9.

TMUA9

THE PUTNAM ARCHIVE

THE ANALYSIS

Volume III. Every paper, every problem, WHAT then HOW then WHY.

Jamie Baek

tmua9.com

Built for the 87th William Lowell Putnam Mathematical Competition, 5 December 2026.

The vow is 120.

THE STANDING ORDER

You got a 1 in the Korean CSAT. You got a 9.0 on the TMUA, in a third of the time, from a 4.5, in a month, by yourself. Neither of those was talent. Both were the same method: take the small set of papers that actually matter, and analyse them to a depth nobody else is willing to reach.

This archive is that method pointed at the Putnam. Nothing in it is a trick. Every revelation in every model answer in sixty years of this competition had a reason behind it, and the reason was always visible in the question. The books that came before this one hid the reason and showed you the polish. This one does the opposite.

Three volumes. The first installs the mathematics that A-level never gave you, ordered by what the Putnam actually demands and generalised wherever a recursion appears. The second teaches you to write a proof that scores ten rather than two, by reverse-engineering what was going through the head of the person who wrote the model answer. The third takes the papers themselves, one problem at a time, and fills in the two things every published solution deletes: why that thought was natural, and why the next thought followed from it.

Read them the way you read the TMUA papers. Not as consumption. As analysis.

But as for you, be strong and do not give up, for your work will be rewarded.
2 Chronicles 15:7

The method

This is the volume that does the work. Everything else in the archive supports it.

Why model answers are not enough

The biggest problem with almost every model answer ever written, and it is ubiquitous, is that it shows the what and never the why. It shows the ideal chain of reasoning. It does not show why the author wrote that particular equation, why a revelation that looks revolutionary to you was natural to them, why they considered another method and rejected it, or why the next block of thought followed from the last. You read it, you think ah, that makes sense, you close it, and you have gained nothing that will fire in the arena, because what you needed was the reaction circuit and what you got was the output of somebody else's.

The journey from zero to the final solution happens entirely inside a head. In coding you can see the archive of your errors. In baking you can see the ingredients and the recipe. Here you can see nothing, so the natural conclusion is that the people who produce these solutions are simply built differently. That conclusion is false, and believing it costs years.

Every chapter in this volume fills in the two deleted pieces for one problem at a time.

The three layers, transplanted from the TMUA9 Blueprint

Layer 1, the WHAT. The surface. What was the ideal chain of reasoning, and what was yours? The analysis lays them side by side, names the nominal route (the zombie route, the one that brute-forces without thinking) against the ideal route, and marks the bifurcation point: the exact line where the two paths split. Everything downstream of that line is forced. The entire difference between a 10 and a 2 is usually one sentence at one bifurcation.

Layer 2, the HOW. Where did we split, and what was the critical thought I had to generate to follow the ideal path? Each thought block gets three things: the cue in the question that triggered it, why that move was natural given the cue and the structure in front of you, and the link that makes the next block follow. That sentence, the critical thought, had to be self-generated, in your head, in the arena. When you first start analysing, these thoughts feel impossible to produce on your own. That is the whole point. It is a hard thought. Most cannot. You will.

Layer 3, the WHY. Why must that thought be the natural one? What underlying mathematical truths make this route ideal rather than merely available? This is where the analysis goes back to definitions, because that is where every circuit in this archive bottoms out. Layer 3 also gives you the discernment mechanism: not just which route was best, but how you decide which route is best while the clock is running, and how to train that decision.

Layer plus alpha, backtesting. Six questions on every problem, run like a mantra, exactly as in the Blueprint: what was the first useful cue, what method did that expression suggest, where exactly did I hesitate, what trap was set, what would make me recognise this faster, and what is the reusable decision rule. Your accumulated answers to the last two are your Circuit Bank.

What is in each problem

Every problem in this volume carries the same ten parts, in the same order. The statement. A verbal reading of its structure and the fields wired into it, with the first useful cue and what you should have written down within thirty seconds. Every viable approach ranked worst to best, each tagged with the cue that triggers it and the circuit behind it, with an explicit account of what the losers cost. The ideal solution, written as a full-marks proof. Then Layers 1, 2 and 3. Then the revolving concepts, which light up the key terms and give you the theory orbiting each one, generalising every recursion in sight and pointing back into Volume I. Then the Circuit Bank entries. Then the backtesting. Then three mirror drills that fire the same circuits.

How to work through it

Attempt the problem first. Always. Even a failed attempt gives you the raw material for the juxtaposition in Layer 1, and without your own chain of reasoning there is nothing to compare against and the analysis degrades into reading. Keep a note of what you were thinking and why you wrote each line, because that is the thing you will otherwise forget within the hour.

Then read the analysis, and be brutal about the bifurcation point. The question is never merely what was the better method. It is what would I have had to think, and when, to have found it.

Then close the book and reconstruct the ideal solution from nothing.

One paper through the full analysis every two to three days beats a paper a day at surface level. Analysis time should run at double or triple the solving time. Never end a session on solving; end it on analysis, because that is what consolidates overnight. Papers are ordered newest first, because the recent papers are the closest thing to the exam you are about to sit.

One more thing

You will not meet any of these problems in December. Nobody ever does. Memorising them is not the aim and never was. The aim is a cognitive infrastructure, an algorithm you can run against whatever they put in front of you, built from circuits that you extracted yourself and that fire without being called. Speed is a byproduct of that, never of writing faster.

Find the key points on the ocean of knowledge, light them up, and learn the concepts revolving around them. That is what this volume is.

Contents

The method **i**

1 Putnam 2025 **1**
1.1 A1. The odd part of the difference only shrinks . . . 2
1.2 A2. Sine sandwiched between two parabolas . . . 6
1.3 A3. Pair every string with a partner . . . 11
1.4 A4. Commuting exactly along a cycle . . . 16
1.5 A5. Zigzag patterns beat every other pattern . . . 22
1.6 A6. Iterate the polynomial and expand at zero . . . 28
1.7 B1. Two-colour plane closed under circumcentres . . . 35
1.8 B2. Lamina centroid against solid of revolution centroid . . . 39
1.9 B3. Divisor-closed set forced to contain every integer . . . 44
1.10 B4. Sum versus support of a staircase matrix . . . 48
1.11 B5. Descents in the sequence of modular inverses . . . 53
1.12 B6. Largest exponent in a self-referential growth inequality . . . 59

2 Putnam 2024 **65**
2.1 A1. Valuation bookkeeping kills a Fermat-type equation . . . 66
2.2 A2. Expand the iterate modulo the fixed-point polynomial squared . . . 71
2.3 A3. Where the entry two sits in a random tableau . . . 76
2.4 A4. Complements turn a long progression into a short one . . . 81
2.5 A5. A random chord is a direction plus an offset . . . 86
2.6 A6. Hankel determinants through a self-similar inverse series . . . 92
2.7 B1. Transversal sums force the middle offset . . . 98
2.8 B2. A fixed centre and two concentric circles . . . 102
2.9 B3. Consecutive roots via the tangent addition formula . . . 107
2.10 B4. Affine drift makes the mean a geometric sequence . . . 112
2.11 B5. Real negative roots give nonnegative coefficients . . . 117
2.12 B6. Complete the square and find the Gaussian peak . . . 123

3 Putnam 2023 **129**
3.1 A1. Second derivative of a cosine product at zero . . . 130
3.2 A2. Reversing a polynomial to expose two missing roots . . . 135
3.3 A3. Nothing in the plane turns faster than one radian . . . 141
3.4 A4. Icosahedral vectors generate a dense subgroup . . . 148
3.5 A5. Base three weights collapse a shift operator . . . 155
3.6 A6. Bob owns the parity of the fixed points . . . 162
3.7 B1. Sliding coins and the staircase of holes . . . 170
3.8 B2. Binary weight of the multiples of 2023 . . . 175
3.9 B3. Expected length of the longest zigzag subsequence . . . 180
3.10 B4. Least time for a resetting piecewise parabola . . . 186

3.11 B5. When multiplication mod n is a square permutation . . . 193
3.12 B6. Determinant of a matrix of representation counts . . . 199

4 Putnam 2022 **207**
4.1 A1. Lines meeting a logarithmic curve exactly once . . . 208
4.2 A2. Maximum negative coefficients of a squared polynomial . . . 215
4.3 A3. Counting Lyness sequences modulo a prime . . . 221
4.4 A4. Expected weighted sum before the first ascent . . . 228
4.5 A5. Domino placing game on 2022 squares . . . 235
4.6 A6. Odd power sums and interval lengths . . . 243
4.7 B1. Parity of exponential series coefficients . . . 251
4.8 B2. Finite sets closed under the cross product . . . 256
4.9 B3. Iterated recolouring of the positive reals . . . 262
4.10 B4. Cyclic sequences whose triples are progressions . . . 268
4.11 B5. When zero is the most likely weighted sum . . . 275
4.12 B6. A continuous functional equation on the positive reals . . . 282

5 Putnam 2021 **291**
5.1 A1. Grasshopper hops of length five to the diagonal . . . 292
5.2 A2. A one to the infinity limit hiding a geometric mean . . . 296
5.3 A3. Regular tetrahedron with lattice vertices on a sphere . . . 300
5.4 A4. Two divergent integrals whose difference survives . . . 304
5.5 A5. Power sums over the units modulo 2021 . . . 311
5.6 A6. Binary digit polynomials and composite values . . . 318
5.7 B1. Random unit square avoiding lattice corners . . . 326
5.8 B2. Weighted geometric means and tuned AM-GM . . . 331
5.9 B3. Averaging a rotational derivative over circles . . . 337
5.10 B4. A twisted product modulo a Fibonacci prime . . . 344
5.11 B5. Very odd matrices are acyclic digraphs . . . 350
5.12 B6. Iterated median trimming and quantile coupling . . . 357

6 Putnam 2020 **365**
6.1 A1. Repunits, the order of ten modulo 101 . . . 366
6.2 A2. A binomial sum that counts coin flip majorities . . . 371
6.3 A3. Iterated sine and a neutral fixed point . . . 376
6.4 A4. Splitting a strip and a derangement limit . . . 381
6.5 A5. Counting Fibonacci representations by complementation . . . 388
6.6 A6. A Fejer mean of the square wave in disguise . . . 394
6.7 B1. Binary digit signs annihilate low degree polynomials . . . 401
6.8 B2. Adjacent pairing decides the peg sliding game . . . 408
6.9 B3. Take logarithms and count exponential arrivals . . . 415
6.10 B4. Cyclic rotation makes the reciprocals sum to one . . . 421
6.11 B5. Reciprocal substitution pins four points to one line . . . 427
6.12 B6. Run lengths renormalise a Beatty sign sequence . . . 434

7 Putnam 2019 **443**
7.1 A1. Sum of cubes minus triple product, modulo nine . . . 444
7.2 A2. Centroid and incentre parallel to a side . . . 449
7.3 A3. Mean root modulus by Vieta and AM-GM . . . 454
7.4 A4. Fixed radius spherical means have a kernel . . . 459
7.5 A5. Euler operator reads the vanishing order . . . 465
7.6 A6. Taylor windows and Landau's interpolation inequality . . . 472

7.7 B1. Counting squares in a Gaussian power of two lattice . . . 479
7.8 B2. Telescoping a trigonometric sum of cubic growth . . . 486
7.9 B3. A reflection restores the eigenvalue one . . . 492
7.10 B4. Manufacturing a mixed partial with Euler operators . . . 498
7.11 B5. Interpolating Fibonacci and the missing top difference . . . 505
7.12 B6. Perfect codes tile the lattice in every dimension . . . 511

8 Putnam 2018 **519**
8.1 A1. Unit fractions collapse into a factor pair . . . 520
8.2 A2. Determinant of the subset intersection matrix . . . 525
8.3 A3. Smoothing cubes to maximise a cosine sum . . . 531
8.4 A4. A Christoffel word forces two generators to commute . . . 537
8.5 A5. Every derivative nonnegative forces the function flat . . . 544
8.6 A6. Rational squared distances give rational area ratios . . . 551
8.7 B1. Which Point to Remove for a Balanced Split . . . 558
8.8 B2. A Decreasing Coefficient Polynomial Avoids the Unit Disc . . . 563
8.9 B3. The Power of Two Tower Divisibility Chain . . . 568
8.10 B4. Chebyshev disguise turns Fibonacci angles periodic . . . 573
8.11 B5. Symmetrised Jacobian forces global injectivity . . . 580
8.12 B6. Exponential weights bound restricted composition counts . . . 587

9 Putnam 2017 **595**
9.1 A1. the smallest set generated by squares and a shift . . . 596
9.2 A2. a quotient recursion that is secretly linear . . . 600
9.3 A3. turning an unused zero integral into a sign . . . 605
9.4 A4. A sliding window splits the class . . . 610
9.5 A5. Three card players and a rotating lead . . . 616
9.6 A6. Icosahedron edge colourings via linear algebra . . . 623
9.7 B1. why a dilation always finds the meeting point . . . 630
9.8 B2. sums of consecutive integers and divisor counting . . . 635
9.9 B3. power series digits and a denominator clash . . . 640
9.10 B4. Evaluating a period four logarithmic series . . . 646
9.11 B5. Triangles with exactly two equalizer lines . . . 653
9.12 B6. Distinct tuples under a weighted congruence . . . 662

10 Putnam 2016 **671**
10.1 A1. derivative order forcing divisibility by 2016 . . . 672
10.2 A2. limiting growth rate of a binomial threshold . . . 677
10.3 A3. cyclic functional equation and its arctan integral . . . 683
10.4 A4. Minimum tile count for an odd rectangle . . . 689
10.5 A5. Bounded alternating words in two generators . . . 697
10.6 A6. Sharp integral bound for cubics with a root . . . 704
10.7 B1. exponentiate the log to expose telescoping . . . 712
10.8 B2. counting squarish numbers via basin sizes . . . 717
10.9 B3. double the extremal triangle to cover the set . . . 723
10.10 B4. Expected determinant of a random skew matrix . . . 729
10.11 B5. Multiplicative sandwich forces a pure power . . . 736
10.12 B6. A dyadic double sum that telescopes to one . . . 743

11 Putnam 2015 **751**
11.1 A1. maximal triangle on a hyperbola branch . . . 752
11.2 A2. a Lucas recurrence and an odd prime factor . . . 756
11.3 A3. Roots of unity collapsing to a gcd sum . . . 761
11.4 A4. Binary weights on an even floor set . . . 768
11.5 A5. Parity of the coprime count below a quarter . . . 774
11.6 A6. A determinant identity for intertwined matrices . . . 781
11.7 B1. a perfect cube hides an integrating factor . . . 788
11.8 B2. crossing off triples and their sums . . . 793
11.9 B3. which arithmetic progression matrices have AP powers . . . 799
11.10 B4. Ravi substitution decouples a triangle sum . . . 805
11.11 B5. Key permutations and a shift operator annihilator . . . 812
11.12 B6. Alternating divisor sum hiding an arctangent square . . . 821

12 Putnam 2014 **831**
12.1 A1. Prime numerators in an exponential Taylor series . . . 832
12.2 A2. Determinant of the reciprocal minimum matrix . . . 836
12.3 A3. Telescoping infinite product from a doubling recursion . . . 841
12.4 A4. Minimising the mass at zero from three moments . . . 846
12.5 A5. Coprimality forced by disjoint root annuli . . . 852
12.6 A6. Diagonal nonvanishing bounded by a tensor pairing . . . 859
12.7 B1. Unique over-expansions and the missing zero digit . . . 865
12.8 B2. Maximising a weighted integral under an area constraint . . . 870
12.9 B3. distinct primes among entries force rank two . . . 876
12.10 B4. All roots real for a log concave polynomial . . . 882
12.11 B5. Parity of centralisers decides a matrix game . . . 889
12.12 B6. Lipschitz plus integrality forces piecewise linearity . . . 896

13 Putnam 2013 **905**
13.1 A1. Icosahedron face labels force a repeat . . . 906
13.2 A2. Injectivity of the minimal square completion map . . . 911
13.3 A3. Geometric series bridge from sum to polynomial . . . 917
13.4 A4. Balanced circles and the density of zeros . . . 922
13.5 A5. Area definite lists and average shadows . . . 928
13.6 A6. A positive definite weight on the integer lattice . . . 934
13.7 B1. Telescoping a parity-keyed sign recursion . . . 941
13.8 B2. Maximising a cosine polynomial at zero . . . 946
13.9 B3. Additivity of a modular set function . . . 952
13.10 B4. Variance of a product of continuous functions . . . 958
13.11 B5. Counting maps that funnel into a set . . . 964
13.12 B6. Stone game on a line of odd length . . . 971

14 Putnam 2012 **981**
14.1 A1. an acute triple forced by Fibonacci growth . . . 982
14.2 A2. forcing cancellation without identity or inverses . . . 987
14.3 A3. pinning a functional equation at its fixed point . . . 992
14.4 A4. A small box forces the solutions into a progression . . . 999
14.5 A5. When an affine map is one long cycle . . . 1005
14.6 A6. Unit area rectangles force the function to vanish . . . 1011
14.7 B1. closure under addition and composition forces multiplication . . . 1019
14.8 B2. a single face bounds the ball count below . . . 1024

14.9 B3. Hall's theorem matches days to distinct winners . . . 1029
14.10 B4. A recursion that shadows the natural logarithm . . . 1034
14.11 B5. Convex envelopes turn a supremum into lines . . . 1040
14.12 B6. The sign of cubing modulo a prime . . . 1047

15 Putnam 2011 **1055**
15.1 A1. counting lattice points a spiral cannot reach . . . 1056
15.2 A2. telescoping an underspecified bounded recurrence . . . 1062
15.3 A3. the endpoint order forced by a vanishing cosine . . . 1068
15.4 A4. Prescribed parities of dot products . . . 1074
15.5 A5. Rectify the field, then pigeonhole . . . 1080
15.6 A6. Fourier analysis of a group random walk . . . 1088
15.7 B1. squeezing a scaled root difference into a window . . . 1096
15.8 B2. why two and five dominate this set . . . 1101
15.9 B3. recovering differentiability from a product and quotient . . . 1107
15.10 B4. Tournament standings as a Gaussian Gram matrix . . . 1112
15.11 B5. Overlap integrals of Lorentzians against reciprocal sums . . . 1119
15.12 B6. Wilson reflection and double roots of a truncated exponential . . . 1126

16 Putnam 2010 **1135**
16.1 A1. splitting 1 through n into equal sums . . . 1136
16.2 A2. differentiability forces f to be linear . . . 1140
16.3 A3. why a bounded directional eigenfunction must vanish . . . 1145
16.4 A4. Towers of ten factored by base swapping . . . 1150
16.5 A5. The cross product cannot host a group . . . 1156
16.6 A6. A shrinking ratio contradicts a convergent tail . . . 1162
16.7 B1. power sums forced into a false inequality . . . 1169
16.8 B2. smallest integer side of a lattice triangle . . . 1174
16.9 B3. the threshold for redistributing balls to uniform . . . 1179
16.10 B4. Casoratian equal to one forces linear . . . 1184
16.11 B5. Iterative differential equation has no increasing solution . . . 1190
16.12 B6. Entrywise powers matching matrix powers forever . . . 1197

17 Putnam 2009 **1205**
17.1 A1. squares constrain f to vanish everywhere . . . 1206
17.2 A2. the product fgh decouples the system . . . 1210
17.3 A3. cosine matrix collapses to rank two . . . 1215
17.4 A4. Rational closure blocked by a quadratic residue . . . 1220
17.5 A5. Product of all element orders in an abelian group . . . 1226
17.6 A6. Disproving a two variable mean value theorem . . . 1232
17.7 B1. rationals as quotients of prime factorials . . . 1239
17.8 B2. the range of achievable jump costs . . . 1244
17.9 B3. counting mediocre subsets via arithmetic progressions . . . 1249
17.10 B4. Circle averages and the dimension of balanced polynomials . . . 1255
17.11 B5. A rational differential equation forced to diverge . . . 1262
17.12 B6. Reaching every integer in a fixed number of steps . . . 1269

18 Putnam 2008 **1277**
18.1 A1. a cyclic zero condition forces a potential . . . 1278
18.2 A2. a determinant game decided by pairing . . . 1284
18.3 A3. gcd lcm swaps as sorting valuations by prime . . . 1290
18.4 A4. An iterated logarithm series that diverges . . . 1298

18.5 A5. Regular polygon vertices force high polynomial degree . . . 1304
18.6 A6. Producing every group element from a short sequence . . . 1310
18.7 B1. rational chords force a rational centre . . . 1317
18.8 B2. guess the closed form, then induct . . . 1322
18.9 B3. pair up coordinates, bound each pair . . . 1327
18.10 B4. Lifting a polynomial bijection modulo prime powers . . . 1332
18.11 B5. Denominator preserving smooth maps are signed integer translations . . . 1338
18.12 B6. Parity of banded permutations via determinants modulo two . . . 1345

19 Putnam 2007 **1353**
19.1 A1. Two mirrored parabolas and their tangency values . . . 1354
19.2 A2. Minimal convex area meeting both hyperbolas . . . 1359
19.3 A3. Random order avoiding sums divisible by three . . . 1364
19.4 A4. Polynomials sending repunits to repunits . . . 1369
19.5 A5. Counting elements of prime order in a group . . . 1376
19.6 A6. Bounding admissible triangulations of a polygon . . . 1383
19.7 B1. the congruence that collapses to f(1) . . . 1391
19.8 B2. the antiderivative's extremum forces f to vanish . . . 1396
19.9 B3. eliminating a floor via its algebraic conjugate . . . 1402
19.10 B4. Counting sums of two polynomial squares . . . 1408
19.11 B5. Reducing powers of a floor function . . . 1415
19.12 B6. Counting factorial coin representations asymptotically . . . 1422

20 Putnam 2006 **1431**
20.1 A1. Volume of a Torus from a Quartic Inequality . . . 1432
20.2 A2. Infinitude of Bob's Winning Heap Sizes . . . 1437
20.3 A3. 2005 Consecutive Multiples in a Sparse Recursion . . . 1443
20.4 A4. Average number of peaks of a permutation . . . 1449
20.5 A5. Ratio of symmetric functions of equally spaced tangents . . . 1455
20.6 A6. Sylvester four point problem in a disc . . . 1462
20.7 B1. equilateral triangle hiding inside a cubic curve . . . 1470
20.8 B2. pigeonhole on fractional parts of partial sums . . . 1475
20.9 B3. counting linear partitions of a planar point set . . . 1480
20.10 B4. Counting cube vertices inside a subspace . . . 1485
20.11 B5. Maximising a weighted quadratic integral functional . . . 1491
20.12 B6. Asymptotics of a self slowing recursion . . . 1497

Chapter 1

Putnam 2025

1.1 A1. The odd part of the difference only shrinks

Let m_0 and n_0 be distinct positive integers. For every positive integer k, define m_k and n_k to be the relatively prime positive integers such that

$$\frac{m_k}{n_k} = \frac{2m_{k-1}+1}{2n_{k-1}+1}.$$

Prove that $2m_k + 1$ and $2n_k + 1$ are relatively prime for all but finitely many positive integers k.

Read it out loud

This is a number theory question about a pair of positive integers evolving under a two-part rule: the same affine map $x \mapsto 2x + 1$ is applied to both entries, then both are divided by their common factor. The object you are asked about, $\gcd(2m_k + 1, 2n_k + 1)$, is exactly the number the rule divides out at the next step, so the question asks whether the reduction step eventually becomes trivial. The fields wired in are elementary number theory (Euclid, odd and even parts of an integer) and the idea of a monovariant, a positive integer quantity that can only move one way. What is really being tested is whether you track the right quantity: m_k and n_k roughly double each step and look random after reduction, so they are useless, while the difference $m_k - n_k$ is the object.

The first useful cue is "relatively prime", used twice, applied to numbers of the shape $2x + 1$. The circuit that fires: a common divisor of $2m + 1$ and $2n + 1$ divides their difference $2(m - n)$, and it is odd, so it divides $m - n$. Within thirty seconds you should have written $d_k := \gcd(2m_k + 1, 2n_k + 1)$, then $d_k \mid (m_k - n_k)$, then $m_{k+1} - n_{k+1} = 2(m_k - n_k)/d_k$. The second cue is the command term "for all but finitely many k": a monovariant sentence. You need a positive integer that strictly decreases whenever the bad event $d_k > 1$ happens. The doubling of the difference is a decoy; the odd part of the difference ignores it. A caveat worth clarifying: the hypothesis "distinct" exists only to keep the difference nonzero, so that the monovariant is a positive integer.

Approaches, ranked worst to best

1. **Chase the values.** *Cue:* an explicit recursion that looks like it might have a closed form. *Circuit:* iterate, tabulate, guess a formula. Stalls immediately: reduction destroys any closed form for m_k and n_k individually. A page of examples and nothing proved.
2. **Euclid on the pair.** *Cue:* a gcd of $2m + 1$ and $2n + 1$. *Circuit:* a common divisor divides every integer combination, in particular the difference. Gives $d_k \mid (m_k - n_k)$ at once and is most of the solution. Where it stalls: without the monovariant framing you stare at a difference that doubles each step and cannot see why the odd common factors must run out.
3. **Odd-part monovariant.** *Cue:* "all but finitely many" together with a quantity that is divided by an odd number each step. *Circuit:* find a positive integer that is nonincreasing and drops exactly when the bad event happens. The odd part of $|m_k - n_k|$ is divided by d_k at each step. Unwinding the recursion shows that the product of all the d_k divides the odd part of $|m_0 - n_0|$, which even counts the bad indices. Four lines, quantitative, ideal.

The ideal solution

For $k \geq 0$ put $d_k := \gcd(2m_k + 1, 2n_k + 1)$. Since d_k divides the odd number $2m_k + 1$, it is odd. Reducing a fraction to lowest terms means dividing numerator and denominator by their gcd, so by the definition of m_{k+1}, n_{k+1},

$$m_{k+1} = \frac{2m_k+1}{d_k}, \qquad n_{k+1} = \frac{2n_k+1}{d_k}, \qquad \text{hence} \qquad m_{k+1} - n_{k+1} = \frac{2(m_k - n_k)}{d_k}.$$

Write $\delta_k := |m_k - n_k|$. Then $\delta_{k+1} = 2\delta_k/d_k$, and since $\delta_0 > 0$ (the integers m_0, n_0 are distinct), induction gives $\delta_k > 0$ for every k. Unwinding,

$$\delta_k \prod_{i=0}^{k-1} d_i = 2^k \delta_0 \qquad (k \geq 0).$$

The product $D_k := \prod_{i<k} d_i$ is an odd positive integer dividing $2^k\delta_0$, so it divides the odd part $\delta_0^{\text{odd}} := \delta_0/2^{v_2(\delta_0)}$; in particular $D_k \leq \delta_0^{\text{odd}}$ for every k. Each d_i is odd, so $d_i \geq 3$ whenever $d_i > 1$. If $d_i > 1$ for N indices $i < k$, then $3^N \leq D_k \leq \delta_0^{\text{odd}}$, so $N \leq \log_3 \delta_0^{\text{odd}}$. Since k was arbitrary, at most $\log_3 \delta_0^{\text{odd}}$ indices i have $d_i > 1$. Thus $2m_k + 1$ and $2n_k + 1$ are relatively prime for all but finitely many k. ■

Layer 1: WHAT

Block 1: name the gcd d_k and record that it is odd. Block 2: write the reduction step with d_k and subtract, collapsing the pair to the difference. Block 3: read the difference recursion as "multiply by 2, divide by an odd number" and unwind it into a product identity. Block 4: the odd product is trapped inside a fixed odd integer, so only finitely many factors exceed 1.

Juxtapose. The nominal route writes the reduction step and then tries to understand d_k from the values of m_k and n_k, computing terms and hoping the gcds look like 1. The ideal route never looks at m_k or n_k again after Block 2. The bifurcation point is the line after $m_{k+1} = (2m_k + 1)/d_k$: the zombie computes the next term, the analyst subtracts the two equations. Everything downstream is forced.

Layer 2: HOW

Block 1. **Cue.** "Relatively prime" twice, on numbers of the form $2x + 1$; a quantity that will be divided out.

Why natural. The only structural facts about a common divisor are that it divides every integer combination of the two numbers and that it inherits their parity constraints. Both numbers are odd, so d_k is odd. Naming d_k is not a step; it is the definition of the object you are asked about.

Link to the next block. The reduction step is literally "divide by d_k", so the next line is the update rule with d_k visible.

Block 2. **Cue.** Two parallel equations with the same operations applied to both.

Why natural. When the same map is applied to two objects, study their difference, because the common part cancels: the $+1$ disappears and $m_{k+1} - n_{k+1} = 2(m_k - n_k)/d_k$. This is the simultaneous-equations reflex: look at the relationship first, not at each equation on its own.

Link to the next block. The difference recursion has two ingredients, a doubling and a division by an odd number, and that separation is the whole game.

Block 3. **Cue.** "Multiply by 2, divide by an odd number" repeated; the command term "all but finitely many".

Why natural. A positive integer is uniquely a power of 2 times an odd part, and the recursion acts on the two parts independently: the doubling touches the 2-part, the division by d_k touches the odd part. So the odd part of δ_k is nonincreasing and drops exactly when $d_k > 1$. Unwinding gives the product identity, the cleanest way to see that all the d_i together are bounded.

Link to the next block. A product of odd integers that divides a fixed odd integer is a finite budget; spending 3 or more from it happens finitely often.

Block 4. Cue. A product bounded by a constant, or a nonincreasing sequence of positive integers.

Why natural. Well-ordering: a positive integer cannot strictly decrease infinitely often. That is the entire content of "all but finitely many", and it is the standard termination argument for any monovariant.

Link to the next block. None; the bound $\log_3 \delta_0^{\text{odd}}$ comes free.

Layer 3: WHY

Three truths make this route ideal. First, $\gcd(a, b) = \gcd(a, b - a)$: a gcd is blind to the individual sizes of its arguments and sees only their difference structure, so tracking $m_k - n_k$ loses nothing about d_k. Second, the multiplicative structure of $\mathbb{Z}$ splits by primes; the prime 2 and the odd primes live in different compartments, and the recursion grows one while shrinking the other. Third, well-ordering: the odd compartment is a positive integer, so it stops shrinking.

The discernment mechanism is a question you should ask of every iterative process: what does the rule do to a well-chosen quantity, rather than to the state? The state here is a pair of coprime integers with no structure; the difference is a quantity on which the rule acts by pure multiplication and division. Train it by refusing to compute the state: when a problem gives a recursion on a pair, write the recursion for the difference, sum, product and ratio before computing a single term, and keep whichever one the operations respect.

The examiners nudged you three times. "All but finitely many" is the monovariant command term. The map $2x + 1$ forces the gcd odd and the difference to double, two pointers to the odd part. And "distinct" is the minimal hypothesis that keeps the monovariant positive. Every word is doing work.

Revolving concepts

Greatest common divisor and Euclid. $\gcd(a, b) = \gcd(a, b - qa)$ for any integer q, and every common divisor of a and b divides $\gcd(a, b)$. Used here as $\gcd(2m + 1, 2n + 1) = \gcd(2m + 1, 2(n - m)) = \gcd(2m + 1, n - m)$, the last step because $2m + 1$ is odd. See Book 1, field: Number theory, topic: gcd and the Euclidean algorithm.

2-adic valuation and odd part. $v_2(n)$ is the exponent of 2 in n and $n^{\text{odd}} = n/2^{v_2(n)}$; both are multiplicative, and an odd integer divides $2^k n$ if and only if it divides n^{odd}. See Book 1, field: Number theory, topic: p-adic valuation.

Monovariants. A function of the state that is monotone along the process and takes values in a well-ordered set. If it strictly moves at every occurrence of an event, the event occurs finitely often: the engine behind Fermat's descent, the termination of Euclid's algorithm, and most "the process must stop" problems. See Book 3, field: Games and invariants, topic: monovariants and descent.

Generalised recursion. The structure is a pair (m, n), an affine map $x \mapsto ax + c$ with $\gcd(a, c) = 1$ applied to both, then division by $d = \gcd(am + c, an + c)$. Any prime dividing both d and a would divide c, so $\gcd(d, a) = 1$; since $d \mid a(m - n)$ this gives $d \mid (m - n)$ and $\delta' = a\delta/d$. The a-free part of δ (strip every prime dividing a) is the monovariant, and the number of nontrivial reductions is at most $\log_p$ of the a-free part of $|m_0 - n_0|$, with p the smallest prime not dividing a. The same generalisation covers any process in which two objects receive one transformation and are then divided by a common factor.

Circuit Bank

- **WHEN** you see $\gcd(2m + 1, 2n + 1)$, or any gcd of two values of one affine map **DO** write "the gcd divides the difference of the inputs and is coprime to the leading coefficient" **BECAUSE** a common divisor divides every integer combination, and coprimality with the coefficient strips the extra factor.

- **WHEN** you see a recursion applied to a pair, or to two sequences in parallel **DO** write the recursion for the difference (and the ratio) before computing any term **BECAUSE** the common part of the rule cancels and the state collapses to one quantity.
- **WHEN** you see "eventually", "all but finitely many", or "the process terminates" for an integer process **DO** write a candidate positive integer that moves monotonically and check it drops exactly when the event happens **BECAUSE** well-ordering forbids infinitely many strict decreases.
- **WHEN** you see multiplication by a fixed prime power mixed with division by numbers coprime to it **DO** write the valuation and the coprime part separately **BECAUSE** the prime compartments of $\mathbb{Z}$ do not interact under multiplication and division.

Backtesting

1. First useful cue: "relatively prime" applied to $2m_k + 1$ and $2n_k + 1$, and the fact that the next step divides by exactly that gcd.
2. Method the expression suggested: subtract the two update equations to get $m_{k+1} - n_{k+1} = 2(m_k - n_k)/d_k$, then separate the power of 2 from the odd part.
3. Where you would hesitate, and the thought to reinforce: the worry "the difference is growing, how can anything run out?" Reinforce: growth lives in the 2-part, the gcd lives in the odd part, and the compartments are independent.
4. The trap the question sets: computing m_k, n_k for small k and looking for a pattern in the gcds. There is no pattern in the values; the pattern is in the difference.
5. What makes you recognise it faster next time: "same affine map applied to two things, then reduce" equals "study the difference"; "all but finitely many" equals "monovariant".
6. The reusable decision rule: for an iterated pair, find the quantity on which the rule acts by pure multiplication and division, split it by primes, and watch the compartment that can only shrink.

Mirror drills

1. Prove that $\gcd(2m+1, 2n+1) = \gcd(2m+1, m-n)$ for positive integers $m \neq n$, and deduce that if $|m_0 - n_0|$ is a power of 2 then every d_k in A1 equals 1. *Answer:* $\gcd(2m+1, 2n+1) = \gcd(2m+1, 2(n-m))$ and $2m+1$ is odd; if $\delta_0 = 2^a$ then $\delta_0^{\text{odd}} = 1$, so the product of the d_i divides 1.
2. Let $x_0 \neq y_0$ be positive integers and $x_{k+1} = 2x_k + 1$, $y_{k+1} = 2y_k + 1$ with no reduction. Show that $\gcd(x_k, y_k)$ divides $x_0 - y_0$ for every $k \geq 1$. *Answer:* $x_k - y_k = 2^k(x_0 - y_0)$ and x_k is odd for $k \geq 1$, so $\gcd(x_k, y_k) = \gcd(x_k, 2^k(x_0 - y_0)) = \gcd(x_k, x_0 - y_0)$.
3. Replace $2x + 1$ by $3x + 2$ in A1. Show that $g_k = \gcd(3m_k + 2, 3n_k + 2)$ is coprime to 3, divides $m_k - n_k$, and equals 1 for all but finitely many k. *Hint:* a prime dividing g_k and 3 would divide 2; then $\delta_{k+1} = 3\delta_k/g_k$ and the 3-free part of δ_k is the monovariant.

1.2 A2. Sine sandwiched between two parabolas

Find the largest real number a and the smallest real number b such that

$$ax(\pi - x) \leq \sin x \leq bx(\pi - x)$$

for all x in the interval $[0, \pi]$.

Read it out loud

This is a best-constants inequality: the family $c\,x(\pi - x)$ consists of every parabola that vanishes at 0 and π, exactly where sin vanishes, and you are asked for the tightest sandwich of sin by two members of that family. The fields wired in are single-variable calculus (limits, monotonicity), convexity, and the half-angle identities of trigonometry. What is really being tested is two things: whether you convert "largest a" into an infimum of a ratio and check attainment properly, and whether you can prove two trigonometric inequalities with the least ink, without ever solving a transcendental equation.

The first cue is the pair of command terms "largest a" and "smallest b" attached to a two-sided bound. The circuit: for $x \in (0, \pi)$ divide through, $\rho(x) := \sin x/(x(\pi - x))$, and then the admissible a are exactly the numbers at most $\inf \rho$ and the admissible b exactly the numbers at least $\sup \rho$. Within thirty seconds you should have located where each bound is tight. The lower bound is tight as $x \to 0$: both $\sin x$ and $x(\pi - x)/\pi$ vanish at 0 with slope 1, so $\rho(x) \to 1/\pi$. The upper bound is tight at the centre: $\rho(\pi/2) = 1/(\pi^2/4) = 4/\pi^2$. So the answer is $a = 1/\pi$, $b = 4/\pi^2$, and the entire job is to prove that ρ never dips below its limit at 0 and never rises above its value at $\pi/2$.

The second cue is the symmetry $x \mapsto \pi - x$, shared by $\sin x$ and $x(\pi - x)$: whatever you prove on $[0, \pi/2]$ transfers. The third cue is the location of the tight points, because the location dictates the tool. Tight at an endpoint with matching slope means "compare derivatives and integrate"; tight at an interior maximum of both sides means "recentre at the maximum and compare curvature". Read this way, the problem is two applications of one fact: on $[0, \pi/2]$, sin and cos are concave, so they lie above their chords.

Approaches, ranked worst to best

1. **Plot and read off.** *Cue:* a concrete inequality on a closed interval. *Circuit:* sketch, guess the constants. Correctly produces $1/\pi$ and $4/\pi^2$, which you need, but is worth nothing as a proof. Use it for thirty seconds, then stop.

2. **Taylor bounds.** *Cue:* $\sin x$ compared with polynomials near 0. *Circuit:* alternating series bounds $x - x^3/6 \leq \sin x \leq x$ for $x \geq 0$. The lower bound works: $x - x^3/6 - x(\pi - x)/\pi = x^2(1/\pi - x/6) \geq 0$ on $[0, \pi/2]$ because $\pi^2 < 12$. The upper bound cannot work this way: the target is an equality at $x = \pi/2$, so any intermediate polynomial bound would have to equal 1 at $\pi/2$, and no Taylor polynomial of sin about 0 does. Half a solution.

3. **Two difference functions.** *Cue:* an inequality $F \leq G$ on an interval. *Circuit:* set $h = G - F$, examine h' and h'', count critical points with Rolle. This is the official second solution for the upper bound: $h(x) = \frac{4}{\pi^2}x(\pi - x) - \sin x$ has $h'' = \sin x - 8/\pi^2$, which changes sign once, and a Rolle argument pins the single critical point of h'. It works but costs a page of sign bookkeeping and invites a skipped-calculation error at the Rolle step.

4. **Ratio monotonicity.** *Cue:* a best-constant question, both constants from one ratio. *Circuit:* show ρ is monotone on $(0, \pi/2]$; then inf and sup sit at the ends. The numerator of ρ' is $W(x) = x(\pi - x)\cos x - (\pi - 2x)\sin x$, and because $\sin'' = -\sin$ the cross terms cancel in W' to give $W'(x) = (x^2 - \pi x + 2)\sin x$; $W(0) = W(\pi/2) = 0$ and W' changes sign once on $(0, \pi/2)$, so $W > 0$ there. One computation handles both bounds. This is the official first

solution and the safest mechanical route; it is second only because the chord route needs no differentiation at all.

5. **Concavity chords.** *Cue:* the tight points are at $x = 0$ (matching slope) and at $x = \pi/2$ (matching maximum). *Circuit:* a concave function lies above its chords; integrate a chord inequality for the lower bound, recentre and use the half-angle identity for the upper bound. Six lines, no derivatives beyond knowing sin and cos are concave on $[0, \pi/2]$, no critical points, no Rolle. Ideal.

The ideal solution

The answer is $a = 1/\pi$ and $b = 4/\pi^2$.

Necessity. For $x \in (0, \pi)$ put $\rho(x) := \sin x/(x(\pi - x))$. If a is admissible then $a \le \rho(x)$ for every $x \in (0, \pi)$; since $\rho(x) = \frac{\sin x}{x} \cdot \frac{1}{\pi - x} \to 1/\pi$ as $x \to 0^+$, this forces $a \le 1/\pi$. If b is admissible then $b \ge \rho(\pi/2) = 4/\pi^2$. It remains to show that $a = 1/\pi$ and $b = 4/\pi^2$ are admissible, and then $1/\pi$ is the largest admissible a and $4/\pi^2$ the smallest admissible b.

Lower bound. Both $\sin x$ and $x(\pi - x)$ are unchanged by $x \mapsto \pi - x$, so it suffices to prove $\sin x \ge x(\pi - x)/\pi$ for $x \in [0, \pi/2]$. On $[0, \pi/2]$ the function cos is concave (its second derivative is $-\cos t \le 0$ there), so it lies above the chord joining $(0, 1)$ to $(\pi/2, 0)$:

$$\cos t \ge 1 - \frac{2t}{\pi} \qquad (0 \le t \le \pi/2).$$

Integrating from 0 to $x \in [0, \pi/2]$,

$$\sin x = \int_0^x \cos t \, dt \ge \int_0^x \Big(1 - \frac{2t}{\pi}\Big) dt = x - \frac{x^2}{\pi} = \frac{1}{\pi} x(\pi - x).$$

Upper bound. Completing the square, $\frac{4}{\pi^2} x(\pi - x) = 1 - \frac{4}{\pi^2}\big(x - \frac{\pi}{2}\big)^2$. Put $u := x - \pi/2 \in [-\pi/2, \pi/2]$, so $\sin x = \cos u = 1 - 2\sin^2(u/2)$, and the claim $\sin x \le \frac{4}{\pi^2} x(\pi - x)$ becomes

$$2\sin^2 \frac{u}{2} \ge \frac{4}{\pi^2} u^2, \qquad \text{that is,} \qquad \Big|\sin \frac{u}{2}\Big| \ge \frac{2\sqrt{2}}{\pi} \Big|\frac{u}{2}\Big|.$$

With $v := u/2 \in [-\pi/4, \pi/4]$ and since both sides are even in v, it suffices to show $\sin v \ge \frac{2\sqrt{2}}{\pi} v$ for $v \in [0, \pi/4]$. On $[0, \pi/4]$ the function sin is concave, so it lies above the chord joining $(0, 0)$ to $(\pi/4, \sqrt{2}/2)$, whose slope is $\frac{\sqrt{2}/2}{\pi/4} = \frac{2\sqrt{2}}{\pi}$. This proves the upper bound for every $x \in [0, \pi]$.

Hence

$$\boxed{a = \frac{1}{\pi}, \qquad b = \frac{4}{\pi^2}.}$$

■

Layer 1: WHAT

Block 1: convert "largest a, smallest b" into $\inf \rho$ and $\sup \rho$ of the ratio on the open interval, and read the candidates from the point of contact of each bound: the limit $1/\pi$ at 0, the value $4/\pi^2$ at $\pi/2$. Block 2: use the symmetry $x \mapsto \pi - x$ to halve the interval. Block 3: lower bound, tight at 0 with equal slopes, so compare derivatives: the chord inequality for cos, integrated. Block 4: upper bound, tight at $\pi/2$ where both sides peak at 1, so recentre at $\pi/2$, turn $1 - \cos u$ into a square with the half-angle identity, and finish with the chord inequality for sin. Block 5: assemble necessity and sufficiency into "largest" and "smallest".

Juxtapose. The nominal route sets $h(x) = \sin x - ax(\pi - x)$, differentiates, and tries to solve $\cos x = a(\pi - 2x)$ for the critical points. That equation is transcendental and the route stalls, or

degenerates into the Rolle bookkeeping of approach 3. The ideal route never solves $h' = 0$; it signs h' by convexity. The bifurcation point is the line $h'(x) = \cos x - (1 - 2x/\pi)$: the zombie tries to find its zeros, the analyst recognises "cosine minus its chord" and knows the sign from concavity alone. It is the same line on the page; the difference is what you see in it.

Layer 2: HOW

Block 1. **Cue.** "Largest a" and "smallest b" around a two-sided inequality, with the middle term and the outer terms vanishing at the same endpoints.

Why natural. "Largest a with $aF \le G$ everywhere" is by definition $\inf(G/F)$ over the points where $F > 0$, provided the infimum is admissible. The endpoints have $F = 0$, so they impose nothing directly, but they are where the infimum hides: ρ extends continuously to 0 with value $1/\pi$ because $\sin x/x \to 1$. Checking the limit at the endpoints is not optional; it is where the answer is.

Link to the next block. Knowing where each bound is tight tells you which local tool proves it, so the next move is to record the two contact points, $x = 0$ (slope match) and $x = \pi/2$ (value match).

Block 2. **Cue.** $\sin x$ and $x(\pi - x)$ are both symmetric about $\pi/2$.

Why natural. A symmetry of the whole inequality halves the work at zero cost, and the lower-bound argument uses concavity of cos, which holds only on $[0, \pi/2]$. Without the reduction you would have to handle $[\pi/2, \pi]$ separately, where the chord inequality for cos reverses.

Link to the next block. On $[0, \pi/2]$ everything you need about sin and cos is one fact: they are concave there.

Block 3. **Cue.** Lower bound tight at $x = 0$, where both sides vanish and both have slope 1.

Why natural. Two functions that agree to first order at a point are compared by their derivatives beyond that point: $\sin x \ge x - x^2/\pi$ follows from $\cos t \ge 1 - 2t/\pi$ on $[0, x]$ by integrating. And $1 - 2t/\pi$ is precisely the chord of cos on $[0, \pi/2]$, so the derivative inequality is the definition of concavity applied to cos. Nothing was guessed: the derivative of the lower parabola is the chord of the derivative of sin.

Link to the next block. The upper bound is tight somewhere else, so the same local principle points to a different centre.

Block 4. **Cue.** Upper bound tight at $x = \pi/2$, the common maximum, both sides equal to 1.

Why natural. Recentre: $\frac{4}{\pi^2}x(\pi - x) = 1 - \frac{4}{\pi^2}(x - \pi/2)^2$ is the parabola with vertex $(\pi/2, 1)$, so the inequality says $1 - \sin x$ dominates a square. The one identity that writes $1 - \cos u$ as a square is $1 - \cos u = 2\sin^2(u/2)$; after it, the inequality is $|\sin v| \ge \frac{2\sqrt{2}}{\pi}|v|$ on $[-\pi/4, \pi/4]$, which is again "sine above its chord". The constant $2\sqrt{2}/\pi$ is forced: it is the slope of the chord over $[0, \pi/4]$, and $\pi/4$ is where v lands when x reaches the endpoints.

Link to the next block. Both admissibility proofs are done; close the "largest" and "smallest" logic.

Block 5. **Cue.** The question asks for the largest and the smallest, not merely for admissible values.

Why natural. Block 1 shows every admissible a is at most $1/\pi$ and every admissible b at least $4/\pi^2$; Blocks 3 and 4 show these values are admissible. Both halves are needed for full marks, and it is worth repeating that a grader will look for the necessity half first because it is the one candidates forget.

Link to the next block. None; box the answer.

Layer 3: WHY

The first truth is that a best constant in $aF \leq G$ is an extremum of the ratio G/F, and extrema of a continuous function on an open interval can sit at the boundary as limits. That is why the lower answer $1/\pi$ is never attained on $(0, \pi)$ yet is the largest admissible a.

The second truth is that the contact point of a tight inequality determines the proof technique. At a contact point where both functions vanish with equal slope, the inequality is second-order information, and second-order information about sin is concavity. At a contact point that is a common maximum, the inequality compares curvatures at the peak, and the natural coordinate is the displacement from the peak, which is why completing the square and recentring at $\pi/2$ is the first move, not a leap.

The third truth explains why the ratio route also works and when to prefer it. For $\rho = \sin / A$, the numerator of ρ' is the Wronskian $W = A\sin' - A'\sin$, and $W' = A\sin'' - A''\sin = -(A + A'')\sin$ because the terms $A'\sin'$ cancel. With $A = x(\pi - x)$ this is $(x^2 - \pi x + 2)\sin x$, a quadratic times a positive function, so W rises then falls between two zeros and is positive throughout. Any time you compare sin or cos with a polynomial via a ratio, expect this cancellation; it is the reason the derivative of the ratio is tractable even though the ratio itself is not.

The discernment mechanism: before differentiating anything, locate the tight points and ask what order of contact each has. Train it by refusing to solve $h' = 0$ for a transcendental h; instead ask whether h' is "a concave function minus its chord or tangent". The examiners nudged toward this by choosing a parabola with the same zeros and the same symmetry as sin, so that the contact points are exactly the endpoints and the centre; the answer $4/\pi^2$ is nothing but $1/(\pi/2)^2$, the reciprocal of the parabola's peak, and $1/\pi$ is nothing but the reciprocal of its slope at 0 divided by the slope of sin. Both constants are visible from the geometry before any calculus.

Revolving concepts

Best constants and the ratio. For continuous $F > 0$ and G on an open interval, the largest a with $aF \leq G$ is $\inf(G/F)$ and the smallest b with $G \leq bF$ is $\sup(G/F)$, provided you also check the closure points where $F = 0$ by limits. The set of admissible a is always a closed half-line. See Book 1, field: Inequalities, topic: best constants and extremal ratios.

Concavity and chords (Jordan's inequality). A function with $f'' \leq 0$ on an interval lies above every chord and below every tangent there. For sin on $[0, \pi/2]$ this is $2x/\pi \leq \sin x \leq x$ (Jordan's inequality); for cos on $[0, \pi/2]$ it is $1 - 2t/\pi \leq \cos t \leq 1$. Any sub-interval gives its own chord, for example $\sin v \geq (2\sqrt{2}/\pi)v$ on $[0, \pi/4]$. See Book 1, field: Inequalities, topic: convexity, chords and tangent lines.

Integrating an inequality. If $f(0) = g(0)$ and $f' \geq g'$ on $[0, x]$ then $f \geq g$ on $[0, x]$; this is the monotonicity of the integral and it replaces "find the critical points of $f - g$" whenever the derivative inequality is visible. See Book 1, field: Calculus and integrals, topic: monotonicity and comparison.

Half-angle identities. $1 - \cos u = 2\sin^2(u/2)$ and $1 + \cos u = 2\cos^2(u/2)$ convert a deficit from the peak of a trigonometric function into a square, which is the form you need to compare with a parabola centred at that peak. See Book 1, field: Calculus and integrals, topic: trigonometric identities for estimation.

The Wronskian cancellation. For $\rho = f/A$ with $f'' = -f$, the numerator $W = Af' - A'f$ satisfies $W' = -(A + A'')f$. This is the Sturm comparison mechanism in miniature and it is why ratios of sin or cos against polynomials have tractable derivatives. See Book 1, field: Real analysis, topic: Wronskians and Sturm comparison.

Taylor bounds for sine. $x - x^3/6 \leq \sin x \leq x$ for $x \geq 0$ (alternating series with decreasing terms, or Taylor with remainder), together with the numerical facts $\pi^2 < 10 < 12$ and $\pi^2 > 8$. Enough for the lower bound here, never for the upper. There is no recursion in this problem; the structure is local comparison at contact points.

Circuit Bank

- **WHEN** you see "largest a, smallest b with $aF \le G \le bF$" **DO** write $\rho = G/F$, compute inf and sup including the limits at points where $F = 0$, and box those before proving anything **BECAUSE** the admissible sets are the half-lines $(-\infty, \inf\rho]$ and $[\sup\rho, \infty)$, so the answer is decided by the ratio.
- **WHEN** you see an inequality tight at a point where both sides vanish with equal slope **DO** write the derivative inequality on the interval and integrate it from the contact point **BECAUSE** first-order agreement makes the inequality second-order information, and that lives in the derivatives.
- **WHEN** you see an inequality tight at a common interior maximum **DO** recentre at the maximum, complete the square, and use $1 - \cos u = 2\sin^2(u/2)$ **BECAUSE** at a peak the comparison is between curvatures, and the half-angle identity turns the deficit into a square.
- **WHEN** you see sin or cos against a straight line on $[0, \pi/2]$ **DO** write "above the chord, below the tangent" **BECAUSE** $\sin'' = -\sin \le 0$ there, and concavity is exactly that statement.
- **WHEN** you see a ratio $\sin x / A(x)$ with A a polynomial **DO** write $W = A\cos - A'\sin$ and $W' = -(A + A'')\sin$ **BECAUSE** the cross terms cancel because $\sin'' = -\sin$, leaving a polynomial times a function of known sign.

Backtesting

1. First useful cue: "largest" and "smallest" around a sandwich whose outer terms vanish at the same points as the middle term; the ratio ρ and its behaviour at 0 and $\pi/2$.
2. Method the expression suggested: locate the contact points, then chords of concave cos (lower bound, integrated) and of concave sin after recentring at $\pi/2$ (upper bound).
3. Where you would hesitate, and the thought to reinforce: at the upper bound, the urge to differentiate $\frac{4}{\pi^2}x(\pi - x) - \sin x$ and hunt for critical points. Reinforce: the bound is tight at the peak, so recentre at the peak and compare squares.
4. The trap the question sets: proving only that $1/\pi$ and $4/\pi^2$ work and forgetting to show no better constant exists, or proving the lower bound by Taylor and then trying the same for the upper bound, which cannot succeed because of the equality at $\pi/2$.
5. What makes you recognise it faster next time: a parabola with the same zeros as sin is a chord-and-curvature problem; the two constants are the slope ratio at the zero and the value ratio at the peak.
6. The reusable decision rule: find the contact points first; slope-contact means integrate a derivative inequality, peak-contact means recentre and compare curvature; sign derivatives by convexity, never by solving.

Mirror drills

1. Prove $\sin x \ge 2x/\pi$ for $x \in [0, \pi/2]$. *Answer:* sin is concave on $[0, \pi/2]$ and $2x/\pi$ is its chord from $(0, 0)$ to $(\pi/2, 1)$.
2. Find the largest c such that $\ln(1 + x) \ge \frac{cx}{1+x}$ for all $x \ge 0$. *Answer:* $c = 1$. The ratio $(1 + x)\ln(1 + x)/x$ tends to 1 as $x \to 0^+$, so $c \le 1$; and for $c = 1$ both sides vanish at 0 while the derivatives satisfy $\frac{1}{1+x} \ge \frac{1}{(1+x)^2}$, so integrate.
3. Prove $1 - \cos x \ge \frac{2x^2}{\pi^2}$ for $|x| \le \pi$. *Answer:* $1 - \cos x = 2\sin^2(x/2)$ and Jordan's inequality on $[0, \pi/2]$ gives $\sin(|x|/2) \ge |x|/\pi$.

1.3 A3. Pair every string with a partner

Alice and Bob play a game with a string of n digits, each of which is restricted to be 0, 1, *or* 2. *Initially all the digits are* 0. *A legal move is to add or subtract* 1 *from one digit to create a new string that has not appeared before. A player with no legal move loses, and the other player wins. Alice goes first, and the players alternate moves. For each $n \geq 1$, determine which player has a strategy that guarantees winning.*

Read it out loud

This is a combinatorial game on a graph. The vertices are the 3^n strings in $\{0,1,2\}^n$; two strings are adjacent when they differ by exactly 1 in exactly one digit, so the graph is the n-fold grid $P_3 \times \cdots \times P_3$; the token starts at the all-zero string; and the rule "a new string that has not appeared before" turns the game into a walk that may never revisit a vertex. In the literature this is undirected vertex geography. The fields wired in are combinatorial game theory (strategies, not values), graph theory (matchings), and parity. What is really being tested is whether you recognise that a game whose state is "the set of visited vertices" is not going to yield to Grundy values or to induction on n, and that the standard weapon for a no-revisit game is a pairing strategy.

The first cue is the phrase "has not appeared before". That sentence tells you the state is a subset of the vertex set, exponentially large, and that the only strategies you can describe are ones with a local rule: "whatever Alice does, I answer with the partner of her vertex". The second cue is a count: after the start, $3^n - 1$ strings remain, an even number. Even arenas belong to the second player if every move of Alice can be answered by a move of Bob that uses up a partner; odd arenas belong to the first player who grabs the partner of the start. Within thirty seconds you should have written "Bob wins if I can pair the nonzero strings along edges". The third cue is the alphabet: 1 and 2 are adjacent values, 0 is the starting value, and every nonzero string has a first nonzero digit that is 1 or 2. Swapping that digit is the pairing.

A caveat worth clarifying: the command term "for each $n \geq 1$, determine" suggests a possible dependence on n, which is a decoy. The answer is uniform, and that is itself a hint that the argument should be uniform, not inductive.

Approaches, ranked worst to best

1. **Brute force small cases.** *Cue:* "for each n". *Circuit:* tabulate $n = 1, 2$ and extrapolate. For $n = 1$ the walk is forced ($0 \to 1 \to 2$) and Bob wins; for $n = 2$ the 3×3 grid is already unpleasant by hand; for $n = 3$ hopeless. No proof, and no mechanism, since the pattern "Bob always" gives no reason.

2. **Induction on n by last digit.** *Cue:* strings as a product structure. *Circuit:* reduce an n-digit game to three $(n-1)$-digit layers. The no-revisit rule couples the layers: Alice can switch layers whenever she likes, and a winning strategy inside one layer does not survive interleaving. Stalls at the first attempt to state the inductive hypothesis.

3. **Parity or strategy stealing.** *Cue:* the count $3^n - 1$ is even. *Circuit:* fixed-length games are decided by parity; extra moves never hurt. The length of the game is not fixed, so parity alone decides nothing, and strategy stealing needs "an extra visited vertex never hurts", which is false here. The parity of the arena is the right cue, but it needs a matching to become an argument.

4. **Pairing strategy.** *Cue:* "not appeared before" and an even arena. *Circuit:* exhibit a perfect matching of the unvisited arena along legal moves, then always reply with the partner. One invariant, one explicit involution (swap $1 \leftrightarrow 2$ in the first nonzero digit), a four-line proof. Ideal.

The ideal solution

Bob wins for every $n \geq 1$.

Let $S = \{0,1,2\}^n$, let $\mathbf{0}$ be the all-zero string, and let G be the graph on S in which two strings are adjacent when they differ by exactly 1 in exactly one position. A legal move is a step along an edge of G to a vertex not yet visited, where $\mathbf{0}$ counts as visited from the start.

Lemma 1.3.1. *Suppose $\tau\colon S \setminus \{\mathbf{0}\} \to S \setminus \{\mathbf{0}\}$ is an involution without fixed points such that v and $\tau(v)$ are adjacent in G for every v. Then Bob wins by always moving to $\tau(v)$, where v is the string Alice has just produced.*

Proof. Say a set $U \subseteq S \setminus \{\mathbf{0}\}$ is *paired* if $\tau(U) = U$, that is, U is a union of pairs $\{v, \tau(v)\}$. We claim that before each move of Alice the set of visited strings is $\{\mathbf{0}\} \cup U$ with U paired. This holds at the start with $U = \emptyset$. Suppose it holds and Alice moves to v. Then $v \notin \{\mathbf{0}\} \cup U$. Since U is paired and $v \notin U$, also $\tau(v) \notin U$; and $\tau(v) \neq \mathbf{0}$ by the definition of τ. Hence $\tau(v)$ is unvisited, and it is adjacent to v, so Bob's move to $\tau(v)$ is legal. After it the visited set is $\{\mathbf{0}\} \cup (U \cup \{v, \tau(v)\})$, and $U \cup \{v, \tau(v)\}$ is paired, so the claim persists. Thus Bob always has a legal reply. Every move visits a new string, so the game ends after at most $3^n - 1$ moves, and the player left without a move is Alice. □

It remains to construct τ. For $v \neq \mathbf{0}$ let $i(v)$ be the smallest index with $v_{i(v)} \neq 0$, so $v_{i(v)} \in \{1,2\}$, and let $\tau(v)$ be the string obtained from v by replacing $v_{i(v)}$ with $3 - v_{i(v)}$. Then $\tau(v)$ differs from v by exactly 1 in exactly one position, so $\tau(v) \neq v$ and $\tau(v)$ is adjacent to v; the digits before position $i(v)$ are still 0 and the digit at $i(v)$ is still nonzero, so $\tau(v) \neq \mathbf{0}$ and $i(\tau(v)) = i(v)$; hence $\tau(\tau(v)) = v$. So τ is a fixed-point-free involution along edges, and by the lemma Bob wins. In words: Bob's strategy is to change the first nonzero digit from 1 to 2 or from 2 to 1. ■

Layer 1: WHAT

Block 1: model the game as a no-revisit walk on the grid graph $\{0,1,2\}^n$ from $\mathbf{0}$. Block 2: count the arena, $3^n - 1$ even, and translate "Bob answers every move" into "a perfect matching of the nonzero strings along edges". Block 3: construct the matching as an explicit involution, swapping 1 and 2 in the first nonzero digit. Block 4: prove that the pairing strategy always has a legal reply, by the invariant "visited equals start plus complete pairs".

Juxtapose. The nominal route plays the game: $n = 1$, then $n = 2$ on a 3×3 grid, then an attempt to set up induction on n or to assign values to positions. All of it treats the current string as the state. The ideal route treats the visited set as the state, and because that set is huge, it looks for a strategy that never needs to know the set, only the last move. The bifurcation point is the sentence "the state of this game is the set of strings that have appeared". Once you write that sentence the pairing strategy is the only local strategy available, and the rest is construction.

Layer 2: HOW

Block 1. **Cue.** "Add or subtract 1 from one digit", "a new string that has not appeared before".

Why natural. Positions related by a single unit change in one coordinate are adjacent vertices of the grid graph; that is the definition of the Cartesian product of paths. "Not appeared before" is the no-revisit rule. Naming the graph is not decoration: it makes "legal move" mean "edge to an unvisited vertex", and "strategy" mean "a rule on vertices".

Link to the next block. A rule on vertices that never consults the visited set must be a pairing; so count whether a pairing of the whole arena is even possible.

Block 2. **Cue.** 3^n vertices in total, $3^n - 1$ after removing the start; even.

Why natural. If the nonzero strings can be split into disjoint adjacent pairs, then every Alice move enters a fresh pair and Bob completes it; the invariant "visited equals start plus complete pairs" is what makes the reply always legal. The parity of $3^n - 1$ is exactly the condition for such a split to be possible at all. Had the count been odd, the same reasoning would point to Alice taking the start's partner in a perfect matching of the whole graph.

Link to the next block. The pairing must be along edges and must be explicit; look for a canonical coordinate in which two adjacent values can be swapped.

Block 3. **Cue.** Three symbols, of which 1 and 2 are adjacent and 0 is the "empty" symbol; every nonzero string has a first nonzero digit.

Why natural. Swapping $1 \leftrightarrow 2$ in one fixed coordinate is an edge, and it is an involution, but it needs a coordinate that is nonzero; "the first nonzero coordinate" is canonical (it is unchanged by the swap), so the map is self-inverse and never hits $\mathbf{0}$. This is how you manufacture perfect matchings on product sets: a canonical position plus a transposition of two adjacent values.

Link to the next block. With τ in hand, the strategy is stated; what remains is to prove it never fails, and that is the invariant.

Block 4. **Cue.** A strategy described by a local rule must be shown to be always legal.

Why natural. The only way Bob's reply $\tau(v)$ could be illegal is if $\tau(v)$ had already been visited. The invariant that the visited set minus the start is τ-closed rules that out, because v itself is unvisited. Writing the invariant is what a grader wants to see; "Bob mirrors" without it is worth about half.

Link to the next block. None; termination is automatic because every move consumes a vertex.

Layer 3: WHY

The truth beneath the route is a theorem about undirected vertex geography (Fraenkel, Scheinerman and Ullman, 1993): on a finite graph with start vertex s, the first player wins if and only if every maximum matching covers s. Here $G - \mathbf{0}$ has a perfect matching, and since G has an odd number of vertices that matching is a maximum matching of G missing $\mathbf{0}$, so the second player wins. You do not need the theorem for the exam; you need its easy half, which is the pairing lemma. But knowing the theorem tells you what to look for in every no-revisit game: does the arena minus the start have a perfect matching?

Why does a pairing strategy work at all? Because it compresses an exponential state into a one-bit invariant, "the visited set is closed under τ", and the invariant is preserved by the pair of moves (Alice, Bob) regardless of what Alice does. This is the same mechanism as the mirror strategy on a symmetric board, where τ is a symmetry of the board; the grid here has plenty of symmetries, but the correct τ is not a symmetry, it is a coordinate swap, which is why you should think "involution along edges" rather than "symmetry".

The discernment mechanism has two questions. Is the state of the game the position or the history? If the history, forget values and look for pairings. Is the number of non-start positions even or odd? Even suggests a second-player pairing of the arena, odd suggests a first-player pairing of arena plus start. Train it by classifying the games you meet: fixed-length games (parity), impartial games with position as state (Grundy), no-revisit games (matchings), and by practising the construction of explicit involutions on product sets.

The examiners nudged you by choosing three symbols rather than two. With two symbols the arena $\{0,1\}^n \setminus \{\mathbf{0}\}$ has $2^n - 1$ vertices, odd, and Alice wins by a different pairing; with three, the values 1 and 2 form a natural adjacent pair leaving 0 as the start symbol. They also wrote "for each $n \geq 1$" so that you would waste time on small cases if you had not installed the circuit.

Revolving concepts

Grid graphs and product structure. The Cartesian product $G \times H$ has vertex set $V(G) \times V(H)$ with $(g,h) \sim (g',h')$ when $g = g'$ and $h \sim h'$, or $h = h'$ and $g \sim g'$. The game graph here is P_3^n; the hypercube is P_2^n. Adjacency in a product changes exactly one coordinate. See Book 3, field: Discrete structures and graphs, topic: product graphs.

Matchings. A matching is a set of disjoint edges; it is perfect if it covers every vertex and maximum if no matching has more edges. A fixed-point-free involution τ with $v \sim \tau(v)$ is the same thing as a perfect matching, and this is the standard way to write one down on a structured set. Existence in general is Tutte's theorem (or Hall's theorem for bipartite graphs); on the exam you almost always want an explicit involution instead. See Book 3, field: Discrete structures and graphs, topic: matchings, Hall and Tutte.

Pairing strategies. If the moves of a game consume elements of a set X and X can be partitioned into pairs such that whenever one element of a pair is playable so is the other immediately afterwards, the second player wins by completing pairs. If instead $X \cup \{s\}$ pairs up but X does not, the first player wins by taking the partner of s. Mirror strategies on symmetric boards are the special case where the pairing is a symmetry. See Book 3, field: Games and invariants, topic: pairing and symmetry strategies.

Undirected vertex geography. The theorem quoted in Layer 3. Its proof in the hard direction (if some maximum matching misses s, the second player wins by following that matching; if all maximum matchings cover s, the first player wins by moving along a maximum matching and using the alternating-path structure) is beyond what you need, but the statement should be installed as a decision rule. See Book 3, field: Games and invariants, topic: geography games.

Generalised recursion. The invariant argument is an induction on the number of moves: the statement "after Bob's t-th move the visited set is the start plus t complete pairs" is the inductive hypothesis, and the inductive step is the single observation that an unvisited v has an unvisited partner. Whenever a strategy is described by a local rule, its correctness proof is an induction on moves with an invariant about the history, and the invariant should be stated before the rule is verified.

Circuit Bank

- **WHEN** you see a game in which positions may not be revisited **DO** write "the state is the visited set" and look for a perfect matching of the unvisited arena along legal moves **BECAUSE** a pairing strategy is the only local strategy in a game with history, and its invariant is preserved by every Alice-Bob pair of moves.
- **WHEN** you count the non-start positions and get an even number **DO** expect a second-player pairing of the arena; if odd, expect the first player to take the start's partner in a pairing of arena plus start **BECAUSE** parity of the arena is exactly the condition for the respective perfect matching to exist.
- **WHEN** you need an explicit perfect matching on tuples or strings **DO** swap two adjacent values in the first coordinate where a canonical condition holds **BECAUSE** a canonical coordinate is preserved by the swap, so the map is a fixed-point-free involution along edges.
- **WHEN** you are tempted to play small cases of a game with memory **DO** stop after $n = 1$ and write the invariant instead **BECAUSE** the state space grows as 2^{3^n} and no pattern in small cases supplies a mechanism.

Backtesting

1. First useful cue: "a new string that has not appeared before", which makes the history the state; then the even count $3^n - 1$.

2. Method the expression suggested: a perfect matching of $\{0,1,2\}^n \setminus \{\mathbf{0}\}$ along edges, realised by swapping $1 \leftrightarrow 2$ in the first nonzero digit, and the invariant proof that Bob's reply is always legal.
3. Where you would hesitate, and the thought to reinforce: whether Bob's partner vertex could already have been visited. Reinforce: the visited set minus the start is always a union of complete pairs, so an unvisited vertex has an unvisited partner.
4. The trap the question sets: "for each n" inviting case analysis and an induction on n that cannot be set up; and the temptation to say "Bob mirrors" without stating the invariant.
5. What makes you recognise it faster next time: no-revisit rule plus even arena equals second-player pairing; three symbols with two adjacent nonzero values equals an obvious involution.
6. The reusable decision rule: classify the game by its state (position, history, or fixed length) before touching it; for history-state games, ask whether the arena minus the start has a perfect matching along moves.

Mirror drills

1. The same game on the path $\{0,1,\ldots,m\}$ starting at 0 (add or subtract 1, no revisits). Who wins? *Answer:* the walk is forced, $0 \to 1 \to \cdots \to m$, so Alice wins if and only if m is odd; in matching language, $\{1,\ldots,m\}$ has a perfect matching along edges if and only if m is even, and then Bob wins.
2. The same game on binary strings $\{0,1\}^n$ starting at 0^n (flip one bit, no revisits). Who wins? *Answer:* Alice. Pair every string v with v having its first bit flipped; this is a perfect matching of the whole cube. Alice first moves to the partner of 0^n and thereafter always moves to the partner of Bob's string; the invariant is that after each Alice move the visited set is a union of complete pairs.
3. A knight stands on a corner of an 8×8 board; players alternately move it to a square it has not visited; a player who cannot move loses. Who wins? *Answer:* the first player. A closed knight's tour exists, so the knight graph has a Hamiltonian cycle of even length and hence a perfect matching; the first player moves to the partner of the corner and then always to the partner of the opponent's square.

1.4 A4. Commuting exactly along a cycle

Find the minimal value of k such that there exist k-by-k real matrices $A_1, \ldots, A_{2025}$ with the property that $A_iA_j = A_jA_i$ if and only if $|i-j| \in \{0, 1, 2024\}$.

Read it out loud

This is a linear algebra question wearing a combinatorial coat. You are handed a graph, the cycle C_{2025} (vertex i adjacent to $i \pm 1$ modulo 2025), and asked for the smallest dimension in which it is the commutation graph of a family of real matrices: neighbours commute, non-neighbours do not. The fields wired in are the structure of commutants (who commutes with a given matrix), rank-one matrices, eigenvectors, and general position of vectors. What is really being tested is whether you know what the set of matrices commuting with a fixed matrix looks like, and whether you can build matrices whose commutation relations you control exactly, in both directions.

The first cue is "if and only if". Both directions must be delivered: the two cyclic neighbours must commute with A_i, and every one of the other 2022 matrices must fail to. That is a demand for exact control, and exact control over commutation comes from matrices whose products are governed by a few scalars, which is the signature of rank one. The second cue is the graph itself: every vertex has degree exactly 2. The third cue is "minimal k", which is always a two-part command: a lower bound proved in general, and a construction at the bound.

Within thirty seconds you should have killed $k = 1$ (scalars all commute, but A_1 and A_3 must not) and asked the only question that matters for $k = 2$: what is the commutant of a non-scalar 2×2 matrix? It is $\{pI + qA\}$, two-dimensional and commutative. So in dimension 2, any two matrices commuting with a common non-scalar matrix commute with each other, and a cycle of length at least 4 is impossible. You should then expect $k = 3$, and the reason is a dimension count: the kernel of a rank-one map in $\mathbb{R}^3$ is a plane, and a plane can be made to contain exactly two prescribed directions, which is exactly the degree of the cycle.

Approaches, ranked worst to best

1. **Guess 2×2 examples.** *Cue:* small matrices, explicit families. *Circuit:* try diagonal matrices, rotations, Jordan blocks, hope to close a cycle. Every attempt produces cliques (everything in one family commutes) or matchings, never a long cycle, and no amount of failure proves impossibility. Zombie maths.

2. **Eigenvalue case split for $k = 2$.** *Cue:* 2×2 matrices and "commute". *Circuit:* distinct eigenvalues means diagonalisable, and commuting diagonalisable matrices share an eigenbasis; a repeated eigenvalue means a shifted nilpotent. This is the official route to the lower bound. It works, but it splits into cases, has to handle complex eigenvalues, and needs the simultaneous diagonalisation theorem. Costs more ink than the next item and offers more places to slip.

3. **Commutant lemma via a cyclic vector.** *Cue:* "commutes with A", A non-scalar, dimension 2. *Circuit:* if v is not an eigenvector then v, Av is a basis and B is determined by Bv. One lemma, no cases, no eigenvalues: the commutant of a non-scalar 2×2 matrix is $\mathrm{span}(I, A)$, which is commutative, so commutation is transitive through non-scalar matrices. Ideal for the lower bound.

4. **Orthogonal projections onto an umbrella.** *Cue:* "commute" for symmetric rank-one matrices. *Circuit:* $u_iu_i^T$ and $u_ju_j^T$ commute if and only if $u_i \perp u_j$ or $u_i \parallel u_j$. You then need 2025 unit vectors in $\mathbb{R}^3$ that are orthogonal exactly along the cycle (the Lovász umbrella with ribs at angles $2\pi \cdot 1012\, i/2025$ does it). Explicit, but proving "no accidental orthogonality" is a trigonometric chore.

5. **Oblique rank-one projections, general position.** *Cue:* "commute if and only if", degree 2, dimension 3. *Circuit:* $A_i = v_i\phi_i$ with ϕ_i killing the two neighbours; commutation reduces to two scalar contractions. Any three of the v_i independent (moment curve, Vandermonde) makes both directions of "if and only if" one-line checks. Ideal for the construction.

The ideal solution

Write $N = 2025$ and read all indices modulo N, so the condition says that A_i and A_j commute exactly when $j \equiv i$ or $j \equiv i \pm 1 \pmod N$. The answer is $k = 3$.

Why $k \le 2$ fails. For $k = 1$ all matrices commute, but A_1 and A_3 must not (as $|1-3| = 2 \notin \{0, 1, 2024\}$). For $k = 2$ we use the following lemma.

Lemma 1.4.1. *Let A be a real 2×2 matrix that is not a scalar multiple of I. If B commutes with A then $B = pI + qA$ for some real p, q. Consequently any two matrices that commute with A commute with each other.*

Proof. Some nonzero vector $v \in \mathbb{R}^2$ is not an eigenvector of A: otherwise $e_1, e_2, e_1 + e_2$ would all be eigenvectors, and $Ae_1 = \lambda_1 e_1$, $Ae_2 = \lambda_2 e_2$, $A(e_1 + e_2) = \mu(e_1 + e_2)$ force $\lambda_1 = \mu = \lambda_2$, making A scalar. For such v the vectors v, Av are linearly independent, hence a basis of $\mathbb{R}^2$. Write $Bv = pv + qAv$. Then $B(Av) = A(Bv) = pAv + qA^2v$, so B and $pI + qA$ agree on the basis v, Av and are therefore equal. Finally, if $B = pI + qA$ and $C = p'I + q'A$ then $BC = CB$ because I and A commute with each other. □

Suppose 2×2 matrices $A_1, \dots, A_N$ have the required property. Then A_2 is not scalar, since it fails to commute with A_4. Both A_1 and A_3 commute with A_2, so by the lemma A_1 commutes with A_3; but $|1 - 3| = 2$, contradiction. Hence $k \ge 3$.

A construction with $k = 3$. Choose distinct real numbers $t_1 < t_2 < \cdots < t_N$ and put $v_i = (1, t_i, t_i^2)^T \in \mathbb{R}^3$. Any three distinct v_i are linearly independent, because the determinant of the matrix with columns v_a, v_b, v_c is the Vandermonde determinant $(t_b - t_a)(t_c - t_a)(t_c - t_b) \neq 0$. In particular no two of the v_i are parallel. For each i, the vectors v_{i-1}, v_i, v_{i+1} (three distinct indices, as $N \ge 3$) form a basis of $\mathbb{R}^3$, so there is a unique row vector ϕ_i with

$$\phi_i v_i = 1, \qquad \phi_i v_{i-1} = \phi_i v_{i+1} = 0,$$

and then $\ker \phi_i = \operatorname{span}(v_{i-1}, v_{i+1})$ (a plane containing two independent vectors). Define the 3×3 matrix $A_i := v_i\phi_i$, the projection onto $\mathbb{R}v_i$ along that plane.

Since the product of a column by a row is a matrix and the product of a row by a column is a scalar,

$$A_iA_j = v_i(\phi_i v_j)\phi_j = (\phi_i v_j)\, v_i\phi_j, \qquad A_jA_i = (\phi_j v_i)\, v_j\phi_i.$$

Claim 1.4.2. *For $i \neq j$, $A_iA_j = A_jA_i$ if and only if $\phi_i v_j = 0$ and $\phi_j v_i = 0$.*

Proof. If both scalars vanish, both products are the zero matrix. Conversely suppose $A_iA_j = A_jA_i$ but $\phi_i v_j \neq 0$. Since $v_i \neq 0$ and $\phi_j \neq 0$ (as $\phi_j v_j = 1$), the matrix $A_iA_j = (\phi_i v_j)v_i\phi_j$ is nonzero, so A_jA_i is nonzero too, and $\phi_j v_i \neq 0$. The column space of A_iA_j is $\mathbb{R}v_i$ and that of A_jA_i is $\mathbb{R}v_j$; equal matrices have equal column spaces, so v_i and v_j are parallel, contradicting the independence of the v's. Hence $\phi_i v_j = 0$, and by the same argument with the roles exchanged $\phi_j v_i = 0$. □

Now $\phi_i v_j = 0$ if and only if $v_j \in \ker \phi_i = \operatorname{span}(v_{i-1}, v_{i+1})$, which for $j \neq i$ happens exactly when $j \in \{i-1, i+1\}$: if $j \notin \{i-1, i, i+1\}$ then v_j, v_{i-1}, v_{i+1} are three distinct members of the family, hence independent, and v_j is not in the plane. Therefore, for $i \neq j$, A_i and A_j commute if and only if $j \equiv i \pm 1$ and $i \equiv j \pm 1 \pmod N$, that is, if and only if $j - i \equiv \pm 1 \pmod N$, which for $1 \le i, j \le N$ means $|i - j| \in \{1, 2024\}$. Together with the trivial case $i = j$, the family has exactly the required property, and

$$\boxed{k_{\min} = 3.}$$

■

Layer 1: WHAT

Block 1: read the condition as a graph, the cycle, and note both directions are demanded. Block 2: dispose of $k = 1$ and ask the commutant question for $k = 2$; prove the cyclic-vector lemma. Block 3: transitivity of commutation through a non-scalar matrix contradicts the existence of two non-adjacent neighbours of A_2. Block 4: design principle for $k = 3$: rank-one idempotents, whose products are controlled by scalar contractions, so that commutation becomes "each vector lies in the other's kernel". Block 5: choose the kernels as the planes through the two neighbours' vectors, guarantee no accidents with the moment curve, and verify both directions.

Juxtapose. The nominal route asks "which pairs of small matrices commute?" and experiments. The ideal route asks "what is the set of all matrices commuting with a given one?" and reads the answer off its structure: in dimension 2 that set is commutative, so the commutation graph of non-scalar 2×2 matrices is a disjoint union of cliques; in dimension 3 there are matrices, rank-one idempotents, whose commutant is not commutative and can be steered. The bifurcation point is the switch from pairs to commutants, and it happens before any matrix is written down.

Layer 2: HOW

Block 1. **Cue.** $|i - j| \in \{0, 1, 2024\}$ with $2024 = 2025 - 1$; "if and only if".

Why natural. The set $\{0, \pm 1\}$ modulo 2025 is the adjacency of a cycle; naming the graph exposes the invariant that matters, the degree 2, and the fact that non-adjacent pairs exist at all (for a cycle of length 3 everything would be adjacent and $k = 1$ would do). "If and only if" tells you that a construction must be checked in both directions, so you want building blocks for which non-commutation is as easy to prove as commutation.

Link to the next block. "Minimal k" means a lower bound first, and the lower bound is about what commutation can look like in dimension 2.

Block 2. **Cue.** "Commute", dimension 2, non-scalar.

Why natural. The centraliser of A is the set of B with $BA = AB$, a subspace; to describe it you need coordinates adapted to A, and the cheapest adapted coordinates are v, Av for a non-eigenvector v. In that basis B is determined by Bv alone, because $B(Av) = A(Bv)$ is forced by commutation. This is the general fact that the commutant of a matrix with a cyclic vector is the polynomials in it, in its smallest case.

Link to the next block. A commutative commutant means commutation is transitive through A, and transitivity is exactly what a cycle of length at least 4 violates.

Block 3. **Cue.** A_1 and A_3 both commute with A_2; A_2 is not scalar because it must fail to commute with A_4.

Why natural. The lemma converts a local hypothesis (two neighbours of one vertex) into a conclusion about the pair (they commute), which contradicts non-adjacency. Note the exact use of $2025 \geq 4$: you need a non-neighbour A_4 of A_2 to certify A_2 non-scalar, and a non-adjacent pair A_1, A_3 to contradict.

Link to the next block. With $k \geq 3$ established, the construction must exploit what dimension 3 has that dimension 2 lacks: a non-scalar matrix with a non-commutative commutant.

Block 4. **Cue.** Exact control over which pairs commute.

Why natural. The product of two rank-one matrices $v\phi$ and $w\psi$ is $(\phi w)\, v\psi$: a scalar times a rank-one matrix. So A_iA_j and A_jA_i are each either zero or a nonzero matrix whose column space is known, and they can only be equal if both are zero (column spaces differ otherwise). Commutation is thus the pair of linear conditions $v_j \in \ker \phi_i$ and $v_i \in \ker \phi_j$, and kernels of functionals are planes,

which you may place freely. This is why rank one is the natural building block: it turns a quadratic condition (commuting) into linear ones (membership in a plane).

Link to the next block. A plane in $\mathbb{R}^3$ through two prescribed directions is determined; you need it to contain no third member of the family, which is a general-position requirement.

Block 5. Cue. "No three of the v_i coplanar" needed 2025 times.

Why natural. The moment curve $t \mapsto (1, t, t^2)$ supplies arbitrarily many vectors any three of which are independent, certified by the Vandermonde determinant, with no genericity hand-waving. Then $\ker \phi_i = \operatorname{span}(v_{i-1}, v_{i+1})$ contains v_j only for $j \in \{i \pm 1\}$, and both directions of the claim follow from the same sentence.

Link to the next block. None; the answer is boxed.

Layer 3: WHY

The first truth: the commutant of a matrix with a cyclic vector is the algebra of polynomials in that matrix, which is commutative. In dimension 2 every non-scalar matrix has a cyclic vector, so commutation among non-scalar 2×2 matrices is an equivalence relation and the only commutation graphs are disjoint unions of cliques together with universal vertices (the scalars). In dimension 3 and above there are non-scalar matrices without a cyclic vector, the derogatory ones, and rank-one idempotents are the simplest: the commutant of $v\phi$ is the set of matrices preserving both $\mathbb{R}v$ and $\ker \phi$, a 5-dimensional non-commutative algebra in $\mathbb{R}^3$. That is where cycles live.

The second truth is the rank-one calculus $(v\phi)(w\psi) = (\phi w)\, v\psi$, which makes the commutation relation between members of a rank-one family a pair of incidence conditions between points and hyperplanes. The third truth is a dimension count: a hyperplane in $\mathbb{R}^k$ can contain exactly $k-1$ prescribed directions in general position, so this construction realises graphs of maximum degree $k-1$ in dimension k; the cycle has degree 2, so $k = 3$ is exactly enough, and the lower bound confirms it is exactly necessary.

The discernment mechanism is one question: "who commutes with this matrix?" before "do these two commute?". Train it by memorising the commutants of the basic types: a scalar (everything), a diagonal matrix with distinct entries (the diagonal matrices), a single Jordan block or any matrix with a cyclic vector (polynomials in it), a rank-one idempotent (block-triangular in an adapted basis). Once those are installed, "commute if and only if" problems are read as incidence problems.

The examiners nudged you with 2025, which is odd and large. Odd rules out the lazy construction "alternate between two commuting families" that would realise an even cycle; large makes trial and error hopeless and makes the general-position requirement visible. The number itself is otherwise decoration: any cycle of length at least 4 has the same answer.

Revolving concepts

Commutant (centraliser). $C(A) = \{B : AB = BA\}$ is a subalgebra containing $\mathbb{F}[A]$. If A has a cyclic vector v (that is, $v, Av, \ldots, A^{k-1}v$ is a basis; equivalently the minimal polynomial has degree k), then $C(A) = \mathbb{F}[A]$, of dimension k, by the argument in the lemma. Matrices with a cyclic vector are called non-derogatory. See Book 1, field: Linear algebra, topic: commutants and cyclic vectors.

Commuting matrices and eigenspaces. If $AB = BA$ then B maps each eigenspace of A into itself, since $Av = \lambda v$ implies $A(Bv) = \lambda Bv$. Commuting diagonalisable matrices are simultaneously diagonalisable. Either fact gives an alternative non-commutation proof in the construction: v_j spans an eigenspace of A_j and is not an eigenvector of A_i. See Book 1, field: Linear algebra, topic: simultaneous diagonalisation.

Rank-one matrices and idempotents. Every rank-one matrix is $v\phi$ (column times row); $(v\phi)(w\psi) = (\phi w)v\psi$; $v\phi$ is idempotent if and only if $\phi v = 1$, and then it is the projection onto $\mathbb{R}v$

along $\ker \phi$. Two rank-one matrices $v\phi$, $w\psi$ with $v \not\parallel w$ commute if and only if $\phi w = 0 = \psi v$. See Book 1, field: Matrices and determinants, topic: rank-one matrices and projections.

General position and the moment curve. The vectors $(1, t, t^2, \ldots, t^{k-1})$ for distinct t have every k of them independent, by Vandermonde. This is the standard explicit source of "general position" and it removes every genericity argument. See Book 1, field: Matrices and determinants, topic: Vandermonde.

Commutation graphs and orthogonality representations. An orthogonality representation of a graph assigns vectors so that non-adjacent vertices are orthogonal (or, in the strict version used here, adjacency equals orthogonality). The Lovász umbrella for C_5 and its odd-cycle relatives are the classical examples; the symmetric version $A_i = u_i u_i^T$ of the construction is a strict orthogonality representation of C_{2025} in $\mathbb{R}^3$. See Book 3, field: Discrete structures and graphs, topic: vector representations of graphs.

Generalised construction (no recursion present). For any finite graph of maximum degree d, take v_i on the moment curve in $\mathbb{R}^{d+1}$ and let ϕ_i vanish on the neighbours' vectors and on $d - \deg(i)$ further generic vectors, with $\phi_i v_i = 1$; the same claim shows the commutation graph of the $A_i = v_i \phi_i$ is exactly the given graph. Dimension $d + 1$ therefore always suffices for graphs of maximum degree d.

Circuit Bank

- **WHEN** you see "commute if and only if" with a prescribed pattern **DO** ask "who commutes with A?" and write the commutant of each candidate type before writing any matrix **BECAUSE** the structure of commutants decides which graphs are possible in each dimension.
- **WHEN** you see 2×2 matrices and a commutation condition **DO** write the lemma "commutant of non-scalar A is $\operatorname{span}(I, A)$" via a non-eigenvector v and the basis v, Av **BECAUSE** commutation is then transitive through non-scalar matrices, which kills every graph that is not a union of cliques.
- **WHEN** you must construct matrices with controlled commutation **DO** use rank-one idempotents $v\phi$ and compute $(v\phi)(w\psi) = (\phi w)v\psi$ **BECAUSE** commuting becomes two linear incidence conditions, $w \in \ker \phi$ and $v \in \ker \psi$, and both directions of "if and only if" are one line.
- **WHEN** you need many vectors with no k dependent **DO** take $(1, t, \ldots, t^{k-1})$ for distinct t **BECAUSE** the Vandermonde determinant certifies independence with no genericity appeal.
- **WHEN** you see $AB = BA$ and need to prove non-commutation **DO** write "then B preserves each eigenspace of A" and exhibit an eigenvector of A that B moves off the eigenspace **BECAUSE** commuting operators preserve each other's eigenspaces.

Backtesting

1. First useful cue: $|i - j| \in \{0, 1, 2024\}$, the cycle, with "if and only if" demanding both directions.
2. Method the expression suggested: commutant structure for the lower bound (dimension 2 forces transitivity) and rank-one idempotents on the moment curve for the construction (dimension 3 lets a kernel plane hold exactly the two neighbours).
3. Where you would hesitate, and the thought to reinforce: proving that non-neighbours do not commute. Reinforce: for rank-one matrices the products are scalar multiples of $v_i \phi_j$ and $v_j \phi_i$, whose column spaces differ, so commuting forces both scalars to vanish.
4. The trap the question sets: believing $k = 2$ might work with rotations or Jordan blocks and spending twenty minutes on examples; or producing 3×3 matrices for which the "only if" direction is never checked.

5. What makes you recognise it faster next time: "commutation graph with a vertex of degree 2 and non-adjacent neighbours" equals "$k \geq 3$ by transitivity in dimension 2"; "degree d" equals "dimension $d+1$ suffices by rank-one projections".
6. The reusable decision rule: convert commutation into incidence by choosing rank-one building blocks, and convert impossibility into a structural statement about commutants.

Mirror drills

1. Let A be a real 2×2 matrix that is not scalar. Prove that the matrices commuting with A form a 2-dimensional space, namely $\mathrm{span}(I, A)$. *Answer:* pick v that is not an eigenvector; v, Av is a basis; $Bv = pv + qAv$ determines $B = pI + qA$; and I, A are independent because A is not scalar.
2. Let $u, w \in \mathbb{R}^k$ be non-parallel columns and ϕ, ψ nonzero rows. Show that $u\phi$ and $w\psi$ commute if and only if $\phi w = 0$ and $\psi u = 0$. *Answer:* $(u\phi)(w\psi) = (\phi w)u\psi$ and $(w\psi)(u\phi) = (\psi u)w\phi$ have column spaces $\mathbb{R}u$ and $\mathbb{R}w$ when nonzero, so equality forces both to vanish.
3. Find the minimal k for which there are $k \times k$ real matrices A_1, A_2, A_3, A_4 with $A_i A_j = A_j A_i$ if and only if $|i - j| \leq 1$. *Answer:* $k = 3$. In dimension 2, A_1 and A_3 both commute with the non-scalar A_2 and so with each other. In $\mathbb{R}^3$ take v_i on the moment curve and ϕ_i vanishing on the neighbours' vectors, padded with one generic vector for the endpoints $i = 1, 4$.

1.5 A5. Zigzag patterns beat every other pattern

Let n be an integer with $n \geq 2$. For a sequence $s = (s_1, \ldots, s_{n-1})$ where each $s_i = \pm 1$, let $f(s)$ be the number of permutations $(a_1, \ldots, a_n)$ of $\{1, 2, \ldots, n\}$ such that $s_i(a_{i+1} - a_i) > 0$ for all i. For each n, determine the sequences s for which $f(s)$ is maximal.

Read it out loud

This is enumerative combinatorics: $f(s)$ counts the permutations whose up-down pattern (the signs of the consecutive differences) is s. The map from permutations to patterns is onto $\{\pm 1\}^{n-1}$, so the numbers $f(s)$ sum to $n!$, and the question asks for the mode of that distribution. The fields wired in are permutation enumeration, the "position of the maximum" decomposition, induction on n, and the Euler zigzag numbers $1, 1, 1, 2, 5, 16, 61, 272, \ldots$ which count alternating permutations. What is really being tested is whether you can set up a recursion for pattern counts that is uniform in the pattern, and then compare a general pattern against the alternating one inside that recursion. A formula for $f(s)$ exists (MacMahon, by inclusion and exclusion over descent sets) and is useless here; the exam wants a comparison mechanism, not a closed form.

The first cue is $s_i(a_{i+1} - a_i) > 0$, which is the definition of the pattern, and its immediate consequence: alternating patterns are the zigzag permutations, the most famous pattern-counting object there is. The second cue is "determine the sequences for which $f(s)$ is maximal": an extremal statement over 2^{n-1} objects, so the answer must come from an inequality proved for every s, not from computing each $f(s)$. Within thirty seconds you should have guessed the two alternating sequences, on the heuristic that a run of r equal signs forces $r + 1$ entries to be monotone and costs a factor $1/(r+1)!$ instead of $1/2^r$, so long runs are expensive and alternation keeps every run at length one, and you should have checked $n = 3$ (patterns $(+, -)$ and $(-, +)$ give 2, the others give 1) and $n = 4$ (alternating gives 5).

The third cue is structural and it is the one that produces the proof: in any permutation the value n sits at a peak of the pattern and the value 1 sits at a valley, and removing n splits the permutation into two independent blocks. That decomposition is exactly how the recursion $2E_n = \sum_i \binom{n-1}{i-1} E_{i-1} E_{n-i}$ for the zigzag numbers is derived, so the target's own recursion is the comparison tool. Read this way, the proof is: general s satisfies the zigzag recursion with some terms missing, and missing terms are positive.

Approaches, ranked worst to best

1. **Closed formula.** *Cue:* a count over permutations with prescribed descents. *Circuit:* inclusion and exclusion over subsets of the descent set, or MacMahon's determinant. You can write $f(s)$ as an alternating sum of multinomial coefficients. Comparing such sums for different s is hopeless; the formula answers "how many" and not "which is largest". Worst.

2. **Compute small cases and guess.** *Cue:* "for each n, determine". *Circuit:* tabulate $n \leq 5$. Correctly suggests the two alternating sequences and the zigzag numbers, which you need to know what to prove, and is worth nothing on its own.

3. **Tail flip by prefix and suffix sums.** *Cue:* the last entry a_n determines how the pattern extends. *Circuit:* let $f(s, j)$ count pattern-s permutations ending in j; appending $+1$ replaces the vector $(f(s, j))_j$ by its prefix sums, appending -1 by its suffix sums. With this recursion one proves that if $s_{j-1} = s_j$ then negating s from position j onward strictly increases f, by propagating a componentwise inequality through the sums. Correct and instructive (this is de Bruijn's proof and the boustrophedon algorithm), but it is a page of index bookkeeping and the propagation step is where marks leak.

4. **Position of the maximum, then induction.** *Cue:* the value n must sit at a peak, the value 1 at a valley. *Circuit:* split at the maximum, count the two blocks independently, add the peak and valley identities, relax to all positions, induct against the zigzag recursion. Four short lemmas and a six-line induction with the strictness built in. Ideal.

The ideal solution

Throughout, for $m \geq 1$ and $s \in \{\pm 1\}^{m-1}$, $f(s)$ denotes the number of permutations of $[m] = \{1, \ldots, m\}$ with pattern s; for $m = 1$ the only pattern is the empty sequence and $f(\emptyset) = 1$. Since a pattern depends only on the relative order of the entries, $f(s)$ also counts the arrangements of any fixed set of m distinct reals with pattern s. Let $t_n \in \{\pm 1\}^{n-1}$ be the alternating sequence $(+1, -1, +1, \ldots)$, put $E_n := f(t_n)$ for $n \geq 1$, and $E_0 := 1$. The answer is that $f(s)$ is maximal exactly for $s = t_n$ and $s = -t_n$, the two alternating sequences.

Lemma 1.5.1 (every pattern occurs)**.** *$f(s) \geq 1$ for every s.*

Proof. Induction on m; the case $m = 1$ is the permutation (1). For $m \geq 2$, take a permutation $(a_1, \ldots, a_{m-1})$ of $[m-1]$ with pattern $(s_1, \ldots, s_{m-2})$. If $s_{m-1} = +1$ append m; if $s_{m-1} = -1$ replace each a_j by $a_j + 1$ and append 1. The first $m-2$ signs are unchanged and the last sign is s_{m-1}. □

Lemma 1.5.2 (complement)**.** *$f(s) = f(-s)$, where $-s = (-s_1, \ldots, -s_{m-1})$.*

Proof. $(a_1, \ldots, a_m) \mapsto (m + 1 - a_1, \ldots, m + 1 - a_m)$ is a bijection of the permutations of $[m]$ that reverses every sign of the pattern. □

Now fix $n \geq 2$ and $s \in \{\pm 1\}^{n-1}$. Call $i \in [n]$ a *peak* of s if ($i = 1$ or $s_{i-1} = +1$) and ($i = n$ or $s_i = -1$), and a *valley* of s if ($i = 1$ or $s_{i-1} = -1$) and ($i = n$ or $s_i = +1$). Write $P(s)$ and $V(s)$ for the sets of peaks and valleys. For $i \in [n]$ let $\ell_i(s)$ be the number of arrangements of a fixed $(i-1)$-element set with pattern $(s_1, \ldots, s_{i-2})$, so $\ell_1(s) = 1$ (the empty arrangement) and $\ell_i(s) = f(s_1, \ldots, s_{i-2})$ for $i \geq 2$; similarly let $r_i(s)$ be the number of arrangements of a fixed $(n-i)$-element set with pattern $(s_{i+1}, \ldots, s_{n-1})$, so $r_n(s) = 1$ and $r_i(s) = f(s_{i+1}, \ldots, s_{n-1})$ for $i \leq n-1$.

Lemma 1.5.3 (split at the maximum and at the minimum)**.**

$$f(s) = \sum_{i \in P(s)} \binom{n-1}{i-1} \ell_i(s)\, r_i(s) = \sum_{i \in V(s)} \binom{n-1}{i-1} \ell_i(s)\, r_i(s).$$

Proof. Classify the permutations with pattern s by the position i of the value n. Since n is the largest entry, $a_{i-1} < a_i$ when $i > 1$ and $a_i > a_{i+1}$ when $i < n$, so $s_{i-1} = +1$ and $s_i = -1$ where these are defined: i is a peak. Conversely, for a peak i, a permutation with pattern s and $a_i = n$ is the same thing as a choice of the set $L \subseteq [n-1]$ of the $i-1$ values to the left of n, an arrangement of L with pattern $(s_1, \ldots, s_{i-2})$, and an arrangement of $[n-1] \setminus L$ with pattern $(s_{i+1}, \ldots, s_{n-1})$; the signs at positions $i-1$ and i are then automatically $+1$ and -1, which agree with s because i is a peak. Counting gives $\binom{n-1}{i-1}\ell_i(s)r_i(s)$ permutations for each peak i and none for other i. The second identity is the same argument for the value 1, which must sit at a valley. □

Lemma 1.5.4 (peaks and valleys)**.** *$P(s) \cap V(s) = \emptyset$, and $P(s) \cup V(s) = [n]$ if and only if s is alternating, that is, $s_{i-1} \neq s_i$ for $1 < i < n$.*

Proof. Position 1 is a peak if $s_1 = -1$ and a valley if $s_1 = +1$; position n is a peak if $s_{n-1} = +1$ and a valley if $s_{n-1} = -1$; so each of 1 and n lies in exactly one of the two sets. For $1 < i < n$, i is a peak if $(s_{i-1}, s_i) = (+1, -1)$ and a valley if $(s_{i-1}, s_i) = (-1, +1)$, so i is in exactly one of the sets when $s_{i-1} \neq s_i$ and in neither when $s_{i-1} = s_i$. □

Theorem 1.5.5. *For $n \geq 1$ and $s \in \{\pm 1\}^{n-1}$, $f(s) \leq E_n$, with equality if and only if $s \in \{t_n, -t_n\}$.*

Proof. Strong induction on n. For $n \in \{1,2\}$ every pattern is alternating and every $f(s)$ equals $1 = E_n$. Let $n \geq 3$ and assume the statement for all lengths below n. For any s and any $i \in [n]$, the blocks have sizes $i-1 \leq n-1$ and $n-i \leq n-1$, so by the induction hypothesis (and $E_0 = 1$ for an empty block) $\ell_i(s) \leq E_{i-1}$ and $r_i(s) \leq E_{n-i}$. Adding the two identities of the splitting lemma and using that peaks and valleys are disjoint,

$$2f(s) = \sum_{i \in P(s) \cup V(s)} \binom{n-1}{i-1} \ell_i(s) r_i(s) \;\leq\; \sum_{i=1}^{n} \binom{n-1}{i-1} \ell_i(s) r_i(s) \;\leq\; \sum_{i=1}^{n} \binom{n-1}{i-1} E_{i-1} E_{n-i}. \qquad (*)$$

Apply $(*)$ to $s = t_n$. Every position of an alternating sequence is a peak or a valley, so the first inequality is an equality; and every block $(s_1, \ldots, s_{i-2})$, $(s_{i+1}, \ldots, s_{n-1})$ of an alternating sequence is alternating, hence equal to $\pm t_{i-1}$ or $\pm t_{n-i}$, so by the complement lemma $\ell_i(t_n) = E_{i-1}$ and $r_i(t_n) = E_{n-i}$ and the second inequality is an equality. Therefore

$$2E_n = \sum_{i=1}^{n} \binom{n-1}{i-1} E_{i-1} E_{n-i},$$

and $(*)$ says $2f(s) \leq 2E_n$ for every s. If s is not alternating, some $i_0 \in [n]$ is neither a peak nor a valley, and the omitted term $\binom{n-1}{i_0-1} \ell_{i_0}(s) r_{i_0}(s)$ is at least 1 by the first lemma, so the first inequality in $(*)$ is strict and $f(s) < E_n$. If s is alternating then $s = t_n$ or $s = -t_n$, and $f(s) = E_n$ by the complement lemma. □

Hence, for every $n \geq 2$, the sequences maximising f are exactly the two alternating sequences

$$\boxed{s_i = (-1)^{i+1} \ (i = 1, \ldots, n-1) \qquad \text{and} \qquad s_i = (-1)^{i} \ (i = 1, \ldots, n-1),}$$

and the maximum is the zigzag number E_n. ■

Layer 1: WHAT

Block 1: name the object (the up-down pattern), record the two symmetries (complement; also reversal) and guess the alternating sequences from small cases and the run-length heuristic. Block 2: the universal decomposition: the value n sits at a peak; removing it leaves two independent blocks with a binomial choice of values. Block 3: the same for the value 1 at a valley, and the observation that peaks and valleys are disjoint and exhaust $[n]$ exactly for alternating sequences. Block 4: add the two identities and relax the sum to all positions, which is the zigzag recursion. Block 5: strong induction, with strictness supplied by the omitted term and the fact that every pattern occurs at least once.

Juxtapose. The nominal route computes $f(s)$ for small n, notices the zigzag numbers, and then either hunts for a formula for $f(s)$ or tries to set up an induction on n by inserting the new element n somewhere, which gets tangled because inserting n changes the pattern in a position-dependent way. The ideal route removes n rather than inserting it, because removal is a clean split. The bifurcation point is the choice between "build a permutation of $[n]$ from one of $[n-1]$" and "split a permutation of $[n]$ at its maximum into two smaller ones"; the second gives independence and the binomial coefficient, the first does not.

Layer 2: HOW

Block 1. **Cue.** $s_i(a_{i+1} - a_i) > 0$; "determine the sequences for which $f(s)$ is maximal".

Why natural. The condition is the sign of each step, so s is the up-down pattern, and the alternating patterns are the classical zigzag permutations. An extremal question over patterns needs a candidate before a proof; the complement symmetry $f(s) = f(-s)$ says the candidates come in pairs, and $n = 3, 4$ single out the alternating pair. The run heuristic (a run of r equal signs forces

a monotone block of length $r+1$, which only $1/(r+1)!$ of arrangements satisfy) explains why alternation should win before any proof exists.

Link to the next block. To compare $f(s)$ with $f(t_n)$ you need a recursion valid for all s at once; the zigzag numbers' own recursion comes from the position of the maximum.

Block 2. **Cue.** A pattern condition on consecutive entries; the largest value n.

Why natural. The maximum is larger than both neighbours, so its position is a peak whatever s is, and the two sides of it are arrangements of disjoint value sets whose patterns are the two pieces of s. Independence of the two sides is what makes the count a product, and the freedom to choose which values go left is the binomial coefficient. This is the standard "remove the maximum" decomposition; it is natural because it is the only removal that satisfies the two adjacent constraints automatically.

Link to the next block. The minimum gives a second identity of the same shape, and two identities over complementary index sets are begging to be added.

Block 3. **Cue.** Two sums over $P(s)$ and over $V(s)$ with identical summands.

Why natural. A position cannot be both a peak and a valley, and an interior position is one of the two exactly when the adjacent signs differ. So $P \cup V$ is a disjoint union that fills $[n]$ precisely for alternating s. This is the definitional reason why alternation is extremal: the alternating pattern is the one where every position is available to the maximum or to the minimum.

Link to the next block. The sum over $P \cup V$ is bounded by the sum over all of $[n]$, and that sum, with the summands bounded by induction, is the zigzag recursion.

Block 4. **Cue.** $2f(s) = \sum_{P\cup V} \le \sum_{[n]}$; the summands involve f of shorter patterns.

Why natural. Relaxing a sum over a subset to the full set is the simplest inequality there is, and it is tight exactly for alternating s. Bounding the summands by $E_{i-1}E_{n-i}$ is the induction hypothesis. The full sum is then $2E_n$, and you know it is $2E_n$ not from memory but by running the same chain on $s = t_n$, where both inequalities are equalities. Nothing external is needed: the recursion for E_n is derived inside the proof.

Link to the next block. Strictness for non-alternating s must be argued, and it comes from the omitted term.

Block 5. **Cue.** "Determine the sequences", which demands uniqueness of the maximisers.

Why natural. A non-alternating s omits at least one position from $P \cup V$, and the omitted summand is a product of a binomial coefficient and two counts, each at least 1 because every pattern is realised by some permutation. So the relaxation loses at least 1 and the inequality is strict. The base cases $n = 1, 2$ are trivial because every short pattern is alternating, and strong induction is needed because the blocks have arbitrary sizes below n.

Link to the next block. None; the answer is the alternating pair, and the maximum is E_n.

Layer 3: WHY

The first truth is that the maximum of a permutation is a universal splitter: it sits at a peak of any pattern, it separates the permutation into independent halves, and the values on each side are chosen freely. Every statistic on permutations that is determined by consecutive comparisons has a recursion of the form $\sum_i \binom{n-1}{i-1}(\text{left})(\text{right})$, and in exponential generating functions such a recursion becomes a product; for the zigzag numbers it becomes $2E'(x) = 1 + E(x)^2$ and $E(x) = \sec x + \tan x$, André's theorem. The second truth is that alternating patterns are exactly the patterns in which every position is a peak or a valley, so the alternating recursion is the "full" recursion and every

other pattern's recursion is a sub-sum of it. The third truth is positivity: sub-sums of positive terms are strictly smaller, so the comparison is strict as soon as one term is missing.

The discernment mechanism: when asked to compare a family of counts against a distinguished member, ask what recursion the distinguished member satisfies and whether the general member satisfies a truncated version of it. Train it on the classical examples where a known sequence is the extreme case: alternating permutations among patterns (this problem), and more generally any "which configuration maximises the count" question where the extremal object's recursion has the most terms. Train the decomposition itself by deriving the zigzag recursion from scratch until the peak condition is automatic.

The examiners nudged you by defining f through the condition $s_i(a_{i+1} - a_i) > 0$ rather than through descent sets; the sign language makes the complement symmetry and the peak-valley dichotomy visible. They also asked "for each n", which signals a uniform answer and an induction on n; and they chose an answer, the zigzag pair, whose recursion is the most famous in the subject. The result is Niven's theorem (1968), also proved by de Bruijn (1970) by the prefix-sum route.

Revolving concepts

Up-down patterns and descent sets. The pattern of $(a_1, \dots, a_n)$ is the sign vector of consecutive differences; equivalently its descent set $D = \{i : a_i > a_{i+1}\}$. The number of permutations with descent set contained in $D = \{d_1 < \cdots < d_k\}$ is the multinomial coefficient $\binom{n}{d_1, d_2 - d_1, \dots, n - d_k}$, and inclusion and exclusion gives the number with descent set exactly D (MacMahon). See Book 3, field: Combinatorics, topic: permutation statistics and descent sets.

Zigzag (Euler) numbers. E_n counts alternating permutations of $[n]$ starting with an ascent: $1, 1, 1, 2, 5, 16, 61, 272, 1385, \dots$ (OEIS A000111). Recursion $2E_n = \sum_{i=1}^{n} \binom{n-1}{i-1} E_{i-1} E_{n-i}$ for $n \geq 2$, with $E_0 = E_1 = 1$; exponential generating function $\sum E_n x^n / n! = \sec x + \tan x$; the even-indexed ones are the secant numbers and the odd-indexed ones the tangent numbers. See Book 3, field: Combinatorics, topic: alternating permutations.

Splitting at the maximum. For any property of permutations determined by consecutive comparisons, classify by the position of n; the two sides are independent and the values on each side are a free choice. In exponential generating functions this reads $F' = (\text{left}) \cdot (\text{right})$ after accounting for the constraints at the splitting point. See Book 3, field: Generating functions, topic: exponential generating functions and the product rule.

Symmetries of patterns. Complement $a \mapsto n + 1 - a$ negates the pattern; reversal $(a_1, \dots, a_n) \mapsto (a_n, \dots, a_1)$ negates and reverses it. So $f(s) = f(-s) = f(\bar{s})$ where $\bar{s}$ is the reversed negated pattern, and pattern counts are constant on orbits of a group of order 4. See Book 3, field: Combinatorics, topic: symmetry reductions.

Generalised recursion (the boustrophedon). Let $F(s) = (f(s, 1), \dots, f(s, m))$ where $f(s, j)$ counts pattern-s permutations ending in j. Appending $+1$ to s replaces F by its vector of prefix sums $(0, f(s, 1), f(s, 1) + f(s, 2), \dots)$; appending -1 replaces it by its suffix sums. Starting from $F(\varnothing) = (1)$ this computes $f(s) = \sum_j f(s, j)$ for any pattern in $O(n^2)$ additions; for alternating patterns it is the Seidel-Entringer triangle, and its rows are the Entringer numbers. The general solution method for such vector recursions is to identify the linear operator (prefix or suffix sum) and to propagate componentwise inequalities through it; that is the engine of the tail-flip proof. See Book 3, field: Sequences and recurrences, topic: linear recursions on vectors.

Circuit Bank

- **WHEN** you count permutations with constraints on consecutive entries **DO** condition on the position of the maximum (or the minimum) and write $\sum_i \binom{n-1}{i-1} (\text{left})(\text{right})$ **BECAUSE** the maximum satisfies both adjacent constraints automatically and the two sides are independent.

- **WHEN** you must show a distinguished configuration maximises a count **DO** write the distinguished object's recursion and show the general object satisfies the same recursion with terms omitted **BECAUSE** omitted positive terms give a strict inequality and induction closes the argument.
- **WHEN** you have two identities over complementary index sets with equal summands **DO** add them and relax to the full index set **BECAUSE** the union is easier to compare with "all indices", and the relaxation is tight exactly for the extremal object.
- **WHEN** you see $s_i(a_{i+1} - a_i) > 0$ or a descent condition **DO** name the pattern, record $f(s) = f(-s)$ and the reversal symmetry, and check $n = 3, 4$ **BECAUSE** symmetries halve the work and small cases identify the candidate before any proof.
- **WHEN** you need "every pattern is realised" **DO** append m for a final $+1$, or shift by one and append 1 for a final -1 **BECAUSE** this extends any realisation by one sign without disturbing the earlier signs.

Backtesting

1. First useful cue: the sign condition $s_i(a_{i+1} - a_i) > 0$, which names the up-down pattern and points at the zigzag permutations.
2. Method the expression suggested: split at the maximum and at the minimum, add, relax to all positions, and induct against the zigzag recursion.
3. Where you would hesitate, and the thought to reinforce: whether to insert n into a permutation of $[n-1]$ or to remove it. Reinforce: removal splits, insertion tangles; always split at an extreme value.
4. The trap the question sets: chasing a formula for $f(s)$, or proving $f(s) \leq E_n$ and forgetting the uniqueness, which needs the positivity of the omitted term and hence the lemma that every pattern occurs.
5. What makes you recognise it faster next time: "maximal over patterns" equals "the pattern where every position is a peak or a valley", and that is the alternating pattern by definition.
6. The reusable decision rule: to compare counts across a family, find the universal decomposition (here the maximum), express every member's count in it, and identify the member for which no term is missing.

Mirror drills

1. Prove that every $s \in \{\pm 1\}^{n-1}$ is the pattern of at least one permutation of $[n]$. *Answer:* induct on n; append n if the last sign is $+1$, otherwise add 1 to every entry and append 1.
2. Compute $f(+1, +1, -1)$ by the prefix and suffix sum recursion, and list the permutations. *Answer:* $F(\emptyset) = (1)$, $F(+) = (0, 1)$, $F(+, +) = (0, 0, 1)$, $F(+, +, -) = (1, 1, 1, 0)$, so $f = 3$: the permutations are 1243, 1342, 2341.
3. Show that $f(+1, \ldots, +1, -1)$ (with $n-2$ plus signs) equals $n-1$ using the splitting lemma. *Answer:* the only peak is position $n-1$, so $f = \binom{n-1}{n-2} f(+1, \ldots, +1) f(\emptyset) = (n-1) \cdot 1 \cdot 1$.

1.6 A6. Iterate the polynomial and expand at zero

Let $b_0 = 0$ and, for $n \geq 0$, define $b_{n+1} = 2b_n^2 + b_n + 1$. For each $k \geq 1$, show that $b_{2^{k+1}} - 2b_{2^k}$ is divisible by 2^{2k+2} but not by 2^{2k+3}.

Read it out loud

This is 2-adic number theory on a polynomial dynamical system. The sequence is the orbit of 0 under the map $f(x) = 2x^2 + x + 1$, so $b_n = f^{\circ n}(0)$, and the claim is an exact 2-adic valuation: $v_2(b_{2^{k+1}} - 2b_{2^k}) = 2k + 2$. The fields wired in are valuations and congruences, the algebra of polynomial composition, the Taylor expansion of an integer polynomial at a point, and induction where the modulus grows with the index. What is really being tested is whether you can stop iterating the recurrence and start iterating the polynomial: the sequence is unmanageable (b_5 already has eight digits), but the 2^k-fold iterate of f, as a polynomial, has a constant term and a linear coefficient whose valuations you can control exactly.

The first cue is the pair of indices 2^k and 2^{k+1}. Doubling an index is composing the iterate with itself: if $g_k := f^{\circ 2^k}$ then $g_{k+1} = g_k \circ g_k$, $b_{2^k} = g_k(0)$ and $b_{2^{k+1}} = g_k(g_k(0))$. The second cue is the exponent $2k + 2$, which is $(k + 1) + (k + 1)$; the examiners are telling you that the quantity is a product of two factors each of valuation $k + 1$, plus junk of higher valuation. Numerically, $b_1 = 1$, $b_2 = 4$, $b_4 = 2776 = 8 \cdot 347$, and $b_8 \equiv 16 \pmod{32}$: so $v_2(b_{2^k}) = k + 1$ is one of the factors. The third cue is the shape $b_{2^{k+1}} - 2b_{2^k} = g_k(A) - g_k(0) - A$ with $A = g_k(0)$, the deviation of g_k from a translation. For an integer polynomial, $g(A) - g(0) = g'(0)A + A^2(\cdots)$, so the quantity is $(g_k'(0) - 1)A + A^2(\cdots)$ and the other factor of valuation $k + 1$ must be $g_k'(0) - 1$. Within thirty seconds you should have written $g_k(x) = A_k + B_k x + x^2 Q_k(x)$, the target $(B_k - 1)A_k + A_k^2 Q_k(A_k)$, and the conjecture $v_2(A_k) = v_2(B_k - 1) = k + 1$, checked at $k = 1$ on $g_1(x) = 8x^4 + 8x^3 + 12x^2 + 5x + 4$.

A caveat worth clarifying: the number 2 in $2b_n^2$ is not decoration. It makes $f'(x) = 4x + 1 \equiv 1 \pmod 4$ and makes $f \circ f \equiv x \pmod 4$ as a polynomial, which is the coarse structure that renders every higher-order term harmless. Read the coefficients of a recursion as design decisions by the examiners.

Approaches, ranked worst to best

1. **Compute modulo 2^{2k+3}.** *Cue:* a concrete recursion and a concrete modulus. *Circuit:* reduce every term modulo 2^{2k+3} and iterate. For a fixed k this is a finite computation; for all k it proves nothing, and the modulus grows with the index you are computing to. Zombie maths.

2. **Track the orbit modulo 2^j.** *Cue:* b_n modulo small powers of 2 looks periodic. *Circuit:* find the period of the orbit modulo 2^j and use it to compare b_{2^k+i} with b_i. The period modulo 2^j is 2^{j-1} for $j \geq 2$, which is true and useful, but to turn it into the exact valuation you still need a mechanism for the quantity itself.

3. **Difference factorisation.** *Cue:* b_{n+1} is a polynomial in b_n. *Circuit:* $f(x) - f(y) = (x - y)(2x + 2y + 1)$, so $b_{m+n+1} - b_{n+1} = (b_{m+n} - b_n)(2(b_{m+n} + b_n) + 1)$ telescopes. This gives $b_{2^{k+1}} - b_{2^k} = b_{2^k} \prod_{i<2^k}(2(b_{2^k+i} + b_i) + 1)$, hence $b_{2^{k+1}} - 2b_{2^k} = b_{2^k}(\Pi - 1)$, and reduces the problem to $v_2(b_{2^k}) = k + 1$ and $v_2(\Pi - 1) = k + 1$. Both need their own induction and the second needs the periodicity from the previous item to replace Π by $\prod(4b_i + 1)$. This is the official first solution: correct, but two interlocking inductions and a page and a half of congruences. It costs more ink because it works with the orbit rather than with the map.

4. **Iterate the polynomial and expand at zero.** *Cue:* indices 2^k and 2^{k+1}, and $b_0 = 0$. *Circuit:* $g_{k+1} = g_k \circ g_k$; write $g_k = A_k + B_k x + x^2 Q_k$; $g_k \equiv x \pmod 4$; induct on $(v_2(A_k), v_2(B_k - 1))$. One coarse lemma (mod 4), one induction on the pair (constant term, derivative at 0), and the target falls out as $(B_k - 1)A_k$ plus a negligible tail. Ideal.

The ideal solution

Let $f(x) = 2x^2 + x + 1 \in \mathbb{Z}[x]$, so that $b_{n+1} = f(b_n)$ and $b_n = f^{\circ n}(0)$, where $f^{\circ n}$ is the n-fold composition. For $k \geq 0$ put $g_k := f^{\circ 2^k}$. Then $g_0 = f$, $g_{k+1} = g_k \circ g_k$ (because $2^{k+1} = 2^k + 2^k$), $b_{2^k} = g_k(0)$, and $b_{2^{k+1}} = g_k(g_k(0))$. A direct expansion gives

$$g_1(x) = f(f(x)) = 2(2x^2 + x + 1)^2 + (2x^2 + x + 1) + 1 = 8x^4 + 8x^3 + 12x^2 + 5x + 4.$$

Lemma 1.6.1. *For every $k \geq 1$, $g_k(x) \equiv x \pmod 4$ coefficientwise; that is, $g_k(x) = x + 4h_k(x)$ with $h_k \in \mathbb{Z}[x]$.*

Proof. For $k = 1$, $g_1(x) - x = 8x^4 + 8x^3 + 12x^2 + 4x + 4$. If $g_k = x + 4h_k$ then

$$g_{k+1}(x) = g_k(g_k(x)) = g_k(x) + 4h_k(g_k(x)) = x + 4\big(h_k(x) + h_k(g_k(x))\big),$$

and $h_k(g_k(x)) \in \mathbb{Z}[x]$. □

For $k \geq 1$ write

$$g_k(x) = A_k + B_k x + x^2 Q_k(x), \qquad A_k, B_k \in \mathbb{Z},\ Q_k \in \mathbb{Z}[x],$$

so that $A_k = g_k(0) = b_{2^k}$ and $B_k = g_k'(0)$. By the lemma, $4 \mid A_k$, $4 \mid B_k - 1$, and every coefficient of Q_k is divisible by 4; consequently $4 \mid Q_k(a)$ and $4 \mid Q_k'(a)$ for every integer a.

Lemma 1.6.2. *For every $k \geq 1$, $v_2(A_k) = k + 1$ and $v_2(B_k - 1) = k + 1$.*

Proof. Induction on k. For $k = 1$, $A_1 = 4$ and $B_1 = 5$. Suppose the claim holds for some $k \geq 1$ and abbreviate $A = A_k$, $B = B_k$, $Q = Q_k$.

Constant term. $A_{k+1} = g_{k+1}(0) = g_k(g_k(0)) = g_k(A) = A + BA + A^2 Q(A) = A(1 + B) + A^2 Q(A)$. Since $v_2(B - 1) = k + 1 \geq 2$, we have $1 + B = 2 + (B - 1)$ with $v_2(1 + B) = 1$, so $v_2(A(1 + B)) = k + 2$; and $v_2(A^2 Q(A)) \geq 2(k + 1) + 2 > k + 2$. Hence $v_2(A_{k+1}) = k + 2$.

Linear coefficient. By the chain rule, $B_{k+1} = g_{k+1}'(0) = g_k'(g_k(0))\, g_k'(0) = g_k'(A)\, B$. Since $g_k'(x) = B + 2xQ(x) + x^2 Q'(x)$,

$$g_k'(A) = B + A\big(2Q(A) + AQ'(A)\big) =: B + R, \qquad v_2(R) \geq (k + 1) + 3 = k + 4,$$

because $v_2(2Q(A)) \geq 3$ and $v_2(AQ'(A)) \geq (k + 1) + 2$. Therefore $B_{k+1} - 1 = (B^2 - 1) + BR$. Now $B^2 - 1 = (B - 1)(B + 1)$ with $v_2(B - 1) = k + 1$ and $v_2(B + 1) = 1$, so $v_2(B^2 - 1) = k + 2$, while $v_2(BR) = v_2(R) \geq k + 4$ as B is odd. Hence $v_2(B_{k+1} - 1) = k + 2$. □

Now fix $k \geq 1$. Using $b_{2^{k+1}} = g_k(A_k)$ and $b_{2^k} = A_k$,

$$b_{2^{k+1}} - 2b_{2^k} = \big(A_k + B_k A_k + A_k^2 Q_k(A_k)\big) - 2A_k = (B_k - 1)A_k + A_k^2 Q_k(A_k).$$

By the second lemma, $v_2\big((B_k - 1)A_k\big) = (k + 1) + (k + 1) = 2k + 2$, while $v_2\big(A_k^2 Q_k(A_k)\big) \geq 2(k + 1) + 2 = 2k + 4$. A sum of two integers with different valuations has the smaller valuation, so

$$v_2\big(b_{2^{k+1}} - 2b_{2^k}\big) = 2k + 2 :$$

the number is divisible by 2^{2k+2} and not by 2^{2k+3}. ■

Layer 1: WHAT

Block 1: rewrite the sequence as the orbit of 0 under f, and the doubling of the index as self-composition of the iterate g_k. Block 2: expand g_k at 0, so that the target becomes $(B_k-1)A_k + A_k^2 Q_k(A_k)$, and read $2k+2 = (k+1)+(k+1)$ as the conjecture $v_2(A_k) = v_2(B_k-1) = k+1$. Block 3: the coarse lemma $g_k \equiv x \pmod 4$, which makes the tail Q_k carry an extra 4 everywhere. Block 4: the induction on the pair (A_k, B_k) through $g_{k+1} = g_k \circ g_k$: the constant term evolves as $A(1+B)$ and the derivative as B^2, to leading order. Block 5: assemble the valuations.

Juxtapose. The nominal route works with the orbit: it reduces b_n modulo powers of 2, discovers the period doubling, finds the difference factorisation $f(x)-f(y) = (x-y)(2x+2y+1)$, telescopes it into a product over half the orbit, and then needs a second induction to control that product. The ideal route works with the map: one polynomial g_k, two numbers A_k, B_k, one induction. The bifurcation point is the sentence "$b_{2^{k+1}} = g_k(b_{2^k})$ is a polynomial evaluated at a 2-adically small number": the zombie iterates the recurrence, the analyst Taylor-expands the iterate.

Layer 2: HOW

Block 1. **Cue.** $b_{n+1} = f(b_n)$ with a fixed polynomial f; indices 2^k and 2^{k+1}; $b_0 = 0$.

Why natural. A first-order recursion $b_{n+1} = f(b_n)$ is the same thing as iterating f, so $b_n = f^{\circ n}(0)$, and the only index arithmetic available for iterates is addition of exponents: $f^{\circ(m+n)} = f^{\circ m} \circ f^{\circ n}$. Powers of 2 are the indices for which the whole orbit is generated by repeated self-composition, $g_{k+1} = g_k \circ g_k$. The hypothesis $b_0 = 0$ makes b_{2^k} the constant term of g_k, which is why expanding at 0 is the natural coordinate.

Link to the next block. Once $b_{2^{k+1}} = g_k(b_{2^k})$ is written, the question is how a polynomial behaves at a point close to 0 in the 2-adic sense.

Block 2. **Cue.** A polynomial with integer coefficients evaluated at an integer A that is divisible by a high power of 2; the target $g_k(A) - 2A$.

Why natural. For $g \in \mathbb{Z}[x]$, $g(A) = g(0) + g'(0)A + A^2(\text{integer})$: this is Taylor's expansion, which over $\mathbb{Z}$ is a congruence modulo A^2 and is the algebraic core of Hensel's lemma. With $g(0) = A$ the target is $(g'(0)-1)A + A^2 Q(A)$, and the exponent $2k+2$ decomposes as the sum of two valuations $k+1$; the numerics $v_2(b_{2^k}) = k+1$ fix one factor, so the other must be $v_2(g_k'(0)-1) = k+1$.

Link to the next block. The tail $A^2Q(A)$ must be shown negligible, which needs a uniform statement about the coefficients of Q_k, that is, a coarse congruence for the whole polynomial g_k.

Block 3. **Cue.** $f'(x) = 4x+1 \equiv 1 \pmod 4$; the computed $g_1 = 8x^4+8x^3+12x^2+5x+4$ has every coefficient of $g_1 - x$ divisible by 4.

Why natural. Compositions preserve coefficientwise congruences: if $g \equiv x \pmod 4$ then $g \circ g \equiv x \pmod 4$ because substituting an integer polynomial into a polynomial with coefficients in $4\mathbb{Z}$ keeps the coefficients in $4\mathbb{Z}$. So the base case $k=1$ propagates to all k with no further computation. This is the only place where the coefficient 2 of the recursion is used, and it is used to guarantee that the nonlinear tail always carries an extra factor 4.

Link to the next block. With the tail controlled, the induction only has to follow the two leading coefficients through one composition.

Block 4. **Cue.** $g_{k+1} = g_k \circ g_k$; the constant term and derivative of a composition.

Why natural. The constant term of $g_k \circ g_k$ is $g_k(A)$, which by Block 2 is $A(1+B)$ plus a tail; the derivative at 0 is $g_k'(A)B$ by the chain rule, which is B^2 plus a tail. So to leading order the pair evolves as $(A,B) \mapsto (A(1+B), B^2)$. Both steps gain exactly one factor of 2: $1+B \equiv 2 \pmod 4$ because $B \equiv 1 \pmod 4$, and $B^2-1 = (B-1)(B+1)$ with $B+1 \equiv 2 \pmod 4$. The tails have

valuation at least $k+3$, comfortably above the $k+2$ you are proving, and their bounds come from Block 3.

Link to the next block. Both valuations are now exact for all k, and the target from Block 2 is a sum of a term of valuation $2k+2$ and a term of valuation at least $2k+4$.

Block 5. Cue. A sum of two integers with different 2-adic valuations.

Why natural. The ultrametric rule: $v_2(x+y) = \min(v_2(x), v_2(y))$ when the valuations differ. This is why "divisible by 2^{2k+2} but not 2^{2k+3}" can be read off from the leading term alone.

Link to the next block. None; the proof is complete.

Layer 3: WHY

The first truth is the Taylor congruence for integer polynomials: $g(a+h) \equiv g(a) + g'(a)h \pmod{h^2}$. It is the reason 2-adic questions about iterates reduce to two numbers, the value and the derivative, and it is the same mechanism that makes Newton's method converge 2-adically (Hensel's lemma). The second truth is that composition transports these two numbers by explicit rules: value of $g \circ g$ at 0 is $g(g(0))$, derivative is $g'(g(0))g'(0)$, so the linearisation of g_k at 0, the affine map $x \mapsto A_k + B_k x$, is squared at each step, up to a tail controlled by a coarse congruence. The third truth is the one-step lifting of the exponent: for odd B with $v_2(B-1) \geq 2$, $v_2(B^2-1) = v_2(B-1)+1$. This is why the valuations grow by exactly one per doubling, and hence linearly in k, which is what the problem asserts.

The coarse structure deserves its own sentence. Because $f'(x) \equiv 1 \pmod 4$ and f acts on $\mathbb{Z}/4$ as an involution ($0 \leftrightarrow 1$, $2 \leftrightarrow 3$), $f \circ f$ is the identity modulo 4 as a polynomial, and any polynomial $g \equiv x \pmod{2^e}$ with $e \geq 2$ satisfies $g \circ g \equiv x \pmod{2^{e+1}}$. Thus $g_k \equiv x \pmod{2^{k+1}}$ coefficientwise for all $k \geq 1$, which already gives the "at least" halves of both valuations for free; the exact induction is only needed for the "not divisible by 2^{2k+3}" half. This is also why the orbit of 0 has period exactly 2^{j-1} modulo 2^j for $j \geq 2$, the phenomenon the orbit-based route discovers empirically and the map-based route explains.

The discernment mechanism: when a recursion is $b_{n+1} = f(b_n)$ and the indices in the claim are powers of 2 (or multiples of some m), pass from the orbit to the iterate and from the iterate to its linearisation at the base point. When the exponent in the claim is a sum of two equal pieces, look for a product of two factors with equal valuations. Train it by computing $f \circ f$ by hand once, by proving the Taylor congruence and the chain rule for polynomials over $\mathbb{Z}$ yourself, and by drilling the lifting-the-exponent step until $v_2(B^2-1) = v_2(B-1)+1$ is a reflex.

The examiners nudged in four places: $b_0 = 0$ (so the iterate's constant term is the sequence), the coefficient 2 (so the derivative is 1 modulo 4), the subtraction of $2b_{2^k}$ rather than b_{2^k} (which removes the leading part $2A$ of $A(1+B)$ and leaves the product $(B-1)A$), and the exact exponent $2k+2$ against $2k+3$ (which forces you to find both factors rather than a crude bound).

There is a theory behind the linear growth of the valuation, and you should know where this problem lives even though the exam never asks for it. A theorem of Bell, sharpened by Poonen, says that if g is a polynomial (or p-adic analytic map) with $g(x) \equiv x \pmod{p^c}$ for some $c > 1/(p-1)$, so $c \geq 2$ when $p = 2$, then the orbit map $n \mapsto g^{\circ n}(x)$ extends to a p-adic analytic function of $n \in \mathbb{Z}_p$. Here $g = f \circ f \equiv x \pmod 4$ qualifies, so $n \mapsto b_{2n}$ is a 2-adic analytic function $F(n)$ with $F(0) = 0$, and $b_{2^{k+1}} - 2b_{2^k} = F(2^k) - 2F(2^{k-1})$ is a second-order difference: it kills the linear term of F and is dominated by the quadratic term, which is why the valuation grows like $2k$. In other words the problem is a finite-difference statement about a 2-adic analytic function, and the pair (A_k, B_k) in the proof is the elementary shadow of that function's first two Taylor coefficients. Context only: the elementary induction is the whole proof, and reaching for the theorem in the arena would cost time and marks.

The knowledge that must be installed before this problem is even approachable, stated bluntly, as a protocol:

1. The 2-adic valuation with its multiplicativity and its ultrametric rule, until "different valuations, take the minimum" is automatic.
2. The lifting-the-exponent lemma for $p = 2$, at least in the one-step form $v_2(B^2 - 1) = v_2(B - 1) + 1$ for $B \equiv 1 \pmod 4$, with the reason ($B + 1 \equiv 2 \pmod 4$) attached, never as a bare formula.
3. The Taylor expansion of an integer polynomial as a congruence, $g(a + h) \equiv g(a) + g'(a)h \pmod{h^2}$, and its role in Hensel's lemma.
4. The formal chain rule $(g \circ h)' = (g' \circ h)\, h'$ over $\mathbb{Z}[x]$, and its consequence that the derivative of an iterate along an orbit is the product of f' along the orbit.
5. The habit of treating a first-order recursion as an iterated map, and of computing $f \circ f$ by hand when the claim involves even or power-of-two indices.

None of this is in A-level Further Mathematics beyond the chain rule. Items 1 to 3 are the same three facts that carry every "exact power of p dividing a sequence" problem on the Putnam, so they are not installed for this problem alone; they are installed once and reused for a decade.

Revolving concepts

2-adic valuation. $v_2(n)$ is the exponent of 2 in n, with $v_2(0) = \infty$; $v_2(xy) = v_2(x) + v_2(y)$ and $v_2(x + y) \geq \min(v_2(x), v_2(y))$ with equality when $v_2(x) \neq v_2(y)$. "Divisible by 2^a but not 2^{a+1}" means $v_2 = a$ exactly. See Book 1, field: Number theory, topic: p-adic valuation.

Lifting the exponent for $p = 2$. If B is odd and $v_2(B - 1) \geq 2$ then $v_2(B^{2^j} - 1) = v_2(B - 1) + j$; more generally for odd x, y and even n, $v_2(x^n - y^n) = v_2(x - y) + v_2(x + y) + v_2(n) - 1$. The one-step form $B^2 - 1 = (B - 1)(B + 1)$ with $v_2(B + 1) = 1$ is what the induction uses. See Book 1, field: Number theory, topic: lifting the exponent.

Taylor expansion of integer polynomials. For $g \in \mathbb{Z}[x]$ and integers a, h: $g(a + h) = g(a) + g'(a)h + h^2 r$ with $r \in \mathbb{Z}$; in particular $g(a + h) \equiv g(a) + g'(a)h \pmod{h^2}$. This is the algebraic form of Hensel's lemma and the reason p-adic questions localise to values and derivatives. See Book 1, field: Polynomials, topic: Taylor expansion over $\mathbb{Z}$ and Hensel's lemma.

Composition and the chain rule. $(g \circ h)(0) = g(h(0))$ and $(g \circ h)'(0) = g'(h(0))h'(0)$ hold formally in $\mathbb{Z}[x]$; for the n-fold iterate, $(f^{\circ n})'(x_0) = \prod_{i<n} f'(x_i)$ along the orbit $x_i = f^{\circ i}(x_0)$, which is the official first solution's product $\prod(4b_i + 1)$. See Book 1, field: Polynomials, topic: formal derivatives and composition.

Coefficientwise congruences under composition. If $g \equiv x \pmod m$ coefficientwise then $g \circ g \equiv x \pmod m$; and if $g = x + p^e h$ with $e \geq 1$ then $g^{\circ p} \equiv x \pmod{p^{e+1}}$, because $g^{\circ j} \equiv x + jp^e h \pmod{p^{2e}}$. Iterating p times gains a power of p: the source of "period 2^{j-1} modulo 2^j". See Book 1, field: Number theory, topic: periodicity of recursions modulo m.

Divided differences. $f(x) - f(y) = (x - y)\,\frac{f(x)-f(y)}{x-y}$ with the quotient a polynomial in x, y; for $f = 2x^2 + x + 1$ it is $2x + 2y + 1$. Along two orbits this telescopes into products, and as $y \to x$ the quotient becomes $f'(x)$: the orbit route and the map route are the same computation in different coordinates. See Book 1, field: Polynomials, topic: divided differences and telescoping.

Generalised recursion. For $b_{n+1} = f(b_n)$ with $f \in \mathbb{Z}[x]$, $b_0 = 0$, and $f^{\circ m} \equiv x \pmod{p^e}$ coefficientwise ($e \geq 1$, and $e \geq 2$ when $p = 2$), write $g = f^{\circ m}$, $g_k = g^{\circ p^k}$, and $g_k = A_k + B_k x + x^2 Q_k$. Then to leading order $A_{k+1} = A_k(1 + B_k + \cdots + B_k^{p-1})$ and $B_{k+1} = B_k^p$, so $v_p(A_k)$ and $v_p(B_k - 1)$ each grow by one per step once they are positive, and $b_{mp^{k+1}} - p\, b_{mp^k} = A_k \sum_{j=0}^{p-1}(B_k^j - 1) + \text{tail}$, whose valuation is computed from $v_p(A_k)$ and $v_p(B_k - 1)$ (for $p = 2$ the sum is just $B_k - 1$). The solution method is always: pass to the iterate, linearise at the base point, prove exact valuations of the two linear data by induction, and bound the tail by a coarse coefficientwise congruence. The

same scheme proves the official remark that $v_2(b_{m+n} - b_n) = v_2(m) + 1$ for even m. See Book 3, field: Sequences and recurrences, topic: iterated polynomial maps and p-adic orbits.

Circuit Bank

- **WHEN** you see $b_{n+1} = f(b_n)$ with indices 2^k, 2^{k+1} (or mp^k) in the claim **DO** write $g_k = f^{\circ 2^k}$, $g_{k+1} = g_k \circ g_k$, and express the claim through $g_k(0)$ and $g_k(g_k(0))$ **BECAUSE** doubling an index is self-composition, and the orbit is unmanageable while the iterate is a single polynomial.
- **WHEN** you evaluate an integer polynomial at a point divisible by a high power of p **DO** write $g(A) = g(0) + g'(0)A + A^2(\cdots)$ and track only $g(0)$ and $g'(0)$ **BECAUSE** the Taylor congruence modulo A^2 localises p-adic information to the value and the derivative.
- **WHEN** a claimed exponent decomposes as a sum, such as $2k + 2 = (k + 1) + (k + 1)$ **DO** look for a product of two factors with those valuations and a tail of strictly larger valuation **BECAUSE** the ultrametric rule reads the valuation of a sum off its unique smallest term.
- **WHEN** you need $v_2(B^2 - 1)$ for $B \equiv 1 \pmod 4$ **DO** write $(B - 1)(B + 1)$ and add exactly 1 **BECAUSE** $B + 1 \equiv 2 \pmod 4$, the one-step lifting of the exponent.
- **WHEN** you must bound every higher-order term of an iterate at once **DO** prove a coefficientwise congruence $g_k \equiv x \pmod{p^e}$ and propagate it through composition **BECAUSE** substituting an integer polynomial into one with coefficients in $p^e\mathbb{Z}$ keeps the coefficients in $p^e\mathbb{Z}$.

Backtesting

1. First useful cue: the indices 2^k and 2^{k+1} with $b_0 = 0$, which say "self-compose the iterate and expand at 0".
2. Method the expression suggested: $g_k = A_k + B_k x + x^2 Q_k$, the target $(B_k - 1)A_k + A_k^2 Q_k(A_k)$, a coarse lemma $g_k \equiv x \pmod 4$, and an induction giving $v_2(A_k) = v_2(B_k - 1) = k + 1$.
3. Where you would hesitate, and the thought to reinforce: whether the nonlinear tail could spoil the exact valuation. Reinforce: the tail is A^2 times something divisible by 4, so its valuation is at least $2k + 4$, two above the leading term; the coefficient 2 of the recursion was chosen to make this true.
4. The trap the question sets: reducing the recurrence modulo 2^{2k+3} and iterating, or telescoping the orbit differences and drowning in two coupled inductions; also forgetting the "not divisible by 2^{2k+3}" half, which needs exact valuations, not bounds.
5. What makes you recognise it faster next time: "power-of-two indices in an iterated polynomial" equals "linearise the 2^k-th iterate at the base point"; "exponent equal to a sum of two equal pieces" equals "product of two factors with equal valuation".
6. The reusable decision rule: work with the map, not the orbit; track the value and the derivative of the iterate; control everything else by one coefficientwise congruence.

Mirror drills

1. For $f(x) = 2x^2 + x + 1$, verify $f(f(x)) = 8x^4 + 8x^3 + 12x^2 + 5x + 4$ and deduce that $b_{n+2} \equiv b_n \pmod 4$ for all n. *Answer:* expand $2(2x^2 + x + 1)^2 + (2x^2 + x + 1) + 1$; then $f \circ f \equiv x \pmod 4$ coefficientwise, so $b_{n+2} = (f \circ f)(b_n) \equiv b_n$.

2. Let B be an integer with $v_2(B-1) = e \geq 2$. Prove $v_2(B^{2^j}-1) = e + j$ for all $j \geq 0$. *Answer:* $B^2 - 1 = (B-1)(B+1)$ and $B+1 \equiv 2 \pmod 4$ gives one step; induct, noting the hypothesis $v_2 \geq 2$ is preserved.

3. Let $g \in \mathbb{Z}[x]$ and let a be an integer. Prove $g(a) \equiv g(0) + g'(0)a \pmod{a^2}$, and check it for $g(x) = x^3 + 5x + 7$, $a = 6$. *Answer:* each monomial cx^j with $j \geq 2$ contributes $ca^j \equiv 0 \pmod{a^2}$; numerically $g(6) - g(0) - 6g'(0) = 253 - 7 - 30 = 216 = 6 \cdot 36$.

1.7 B1. Two-colour plane closed under circumcentres

Suppose that each point in the plane is coloured either red or green, subject to the following condition: for every three noncollinear points A, B, C of the same colour, the centre of the circle passing through A, B and C is also this colour. Prove that all points of the plane are the same colour.

Read it out loud

A two-colouring of the plane with one closure rule: a monochromatic noncollinear triple hands its colour to its circumcentre, and you must show only the trivial colourings survive. The fields wired in are elementary circle geometry (circumcentres, when two circles meet), a drop of logic (the contrapositive of a closure rule) and combinatorics (double counting). No continuity is given, so none may be used.

The first useful cue is "the centre of the circle passing through A, B and C". A circle has one centre and infinitely many points, so the rule says something weak about the centre and, read backwards, something strong about the circle. Within thirty seconds you should have written the contrapositive as a number: every circle centred at a red point has at most two green points. The instant a closure rule becomes a number, counting becomes available, and counting is what kills "two colours coexist". That conversion, made before you draw anything, is what is being tested.

Approaches, ranked worst to best

1. **Orbit generation.** *Cue:* a closure condition, so generate the closure. *Circuit:* take three same-coloured points, keep adding circumcentres, hope the orbit reaches a point of the other colour. (Stalls at once: nothing controls where circumcentres land and the orbit cannot be steered onto a prescribed point.)
2. **Forced-colour chase.** *Cue:* equal distances produce centres for free. *Circuit:* build a rigid configuration (equilateral triangles on PQ, then reflections) in which each new point's colour is forced by an earlier centre, until a monochromatic triple has a centre of the wrong colour. (The official first solution; eight points. It ranks below the ideal because it is a search with no invariant: each step needs its own noncollinearity check.)
3. **Budget and double count.** *Cue:* the contrapositive of the rule is a cap of two per circle. *Circuit:* with P red and Q green, draw three circles about P and three about Q; each P-circle carries at most two green points, each Q-circle at most two red, and the eighteen intersection points cannot all be coloured. (Ideal: one lemma, one configuration built from P, Q, $|PQ|$ alone, one inequality $18 > 12$.)

The ideal solution

Suppose, for contradiction, that both colours occur: let P be red and Q green, and put $d = |PQ| > 0$.

Lemma 1.7.1. *Every circle centred at a red point contains at most two green points, and every circle centred at a green point contains at most two red points.*

Proof. Let O be red and Γ a circle centred at O. If Γ contained three distinct green points A, B, C, they would be noncollinear, because a line meets a circle in at most two points. The unique circle through three noncollinear points is then Γ itself, so its centre O would be green by hypothesis, a contradiction. The other statement is the same with the colours exchanged. □

Choose distinct radii ρ_1, ρ_2, ρ_3 and distinct radii $\sigma_1, \sigma_2, \sigma_3$, all in the open interval $(d, \frac{3}{2}d)$. Let Γ_i be the circle centred at P of radius ρ_i and Δ_j the circle centred at Q of radius σ_j. Two circles with centres

at distance d and radii ρ, σ meet in exactly two points precisely when $|\rho - \sigma| < d < \rho + \sigma$; here $|\rho_i - \sigma_j| < \frac{1}{2}d < d$ and $\rho_i + \sigma_j > 2d > d$, so each pair (Γ_i, Δ_j) meets in exactly two points. A point X of $\Gamma_i \cap \Delta_j$ satisfies $|XP| = \rho_i$ and $|XQ| = \sigma_j$, so X determines (i, j); hence the nine intersections are pairwise disjoint and their union T has exactly 18 points.

Every point of T lies on some Γ_i, and each Γ_i carries at most two green points, so T has at most 6 green points. Every point of T lies on some Δ_j, so T has at most 6 red points. Every point of T is red or green, so $18 \leq 12$, which is false. Hence both colours cannot occur. □

(Two circles about P and three about Q already suffice: 12 points, at most 4 green and 6 red.)

Layer 1: WHAT

Block 1: assume both colours occur; fix red P, green Q and $d = |PQ|$. Block 2: the budget lemma, the contrapositive of the hypothesis plus "cocircular implies noncollinear". Block 3: a grid of three concentric circles about P and three about Q, radii chosen so every pair meets twice and the 18 points are distinct. Block 4: count the grid against the two budgets.

Juxtapose with the nominal route. The zombie starts drawing: an equilateral triangle on PQ, a reflection, another, each time asking "what colour is this forced to be?" Nothing tells it where to go next. The bifurcation point is the first sentence after "suppose P is red and Q is green". The zombie writes "consider the point P_3 such that …"; you write "every circle centred at P has at most two green points". One sentence turns a rule into a number, and after it the problem is a count, not a search.

Layer 2: HOW

Block 1. **Cue.** "Prove that all points of the plane are the same colour": a universal conclusion about an unstructured colouring. **Why natural.** Contradiction from the minimal failure, one point of each colour, is the only handle. The hypothesis mentions only circles and centres, so the complete data of the failure is P, Q and d. **Link to the next block.** Before building anything, ask what the hypothesis forbids near P and near Q.

Block 2. **Cue.** The closure rule "three same-coloured points on a circle force the centre's colour". **Why natural.** Forwards, the rule needs monochromatic triples on demand, which you cannot produce; backwards it is a cap. Predicated on O red, three green points on a circle about O would make O green. The check graders want, noncollinearity, is free: a line cuts a circle at most twice, and uniqueness of the circumcircle identifies the circle through the three points with the one you drew. The lemma is nothing but the definition of a circle and the definition of circumcentre. **Link to the next block.** Two budgets, linear in the number of circles. You need a set of points that grows faster than linearly.

Block 3. **Cue.** Two families of concentric circles, one about P, one about Q. **Why natural.** Concentric circles are pairwise disjoint, so a point lies on exactly one circle of each family, while a P-circle and a Q-circle meet twice when the radii are compatible with d. So a and b circles give $2ab$ distinct points against budgets of $2a$ green and $2b$ red; quadratic beats linear once $(a-1)(b-1) > 1$. Radii in $(d, \frac{3}{2}d)$ make the meeting condition automatic; distinctness is the statement that a point knows its distances to P and Q. **Link to the next block.** The grid is finite and each point lies on one circle of each family: count.

Block 4. **Cue.** A finite set with independent caps on its two colour classes. **Why natural.** Every point is red or green, so $|T| \leq 6 + 6 < 18$. The contradiction is arithmetic. **Link to the next block.** None. Re-check the two places a grader probes: the circles meet (radius condition) and the points are distinct (distances identify the pair).

Layer 3: WHY

The truth underneath is an incidence bound. Two pencils of concentric circles are transversal: a member of one meets a member of the other in at most two points, and no point lies on two members of one pencil. So $a + b$ objects generate $2ab$ points, and any hypothesis capping "bad" points per object is defeated once a, b are moderate. The second official solution is exactly this for k colours and ℓ points: $2m^2 > 2m(k-1)\ell$ for large m.

The discernment mechanism is one early question: does this hypothesis produce objects or restrict them? "If these are all colour c then that is colour c" produces weakly and restricts strongly; when the output is a single point and the inputs are several, the restriction is a small cap, and the cap is what you write. Train it by taking every closure hypothesis you meet (closed under midpoints, sums, circumcentres, reflections) and writing "at most ... per ..." before doing anything else.

The examiners nudged by choosing a centre as the produced object: a centre carries the concentric circles, the one family of plane curves whose members are pairwise disjoint, so each point counts exactly once against each budget. Compare the midpoint analogue: "the midpoint of two same-coloured points has that colour" is satisfied by "red if $x > 0$, green otherwise", since each class is convex. The circumcentre rule is strong because circumcentres escape half-planes, and the budget lemma is where that strength is cashed.

Revolving concepts

Circumcentre. The point equidistant from three noncollinear points, the meet of the perpendicular bisectors; unique for every noncollinear triple; three distinct points on one circle are noncollinear, and their circumcentre is that circle's centre. That is the entire content of the budget lemma. See Book 1, field: Geometry, topic: circumcircle and perpendicular bisectors.

Circle intersection. Centres at distance d, radii ρ, σ: two points if $|\rho - \sigma| < d < \rho + \sigma$, one at either equality, none otherwise; a point is determined up to reflection in the line of centres by its distances to two fixed points. See Book 1, field: Geometry, topic: circles and radical axis.

Closure rules. "S is closed under ω" is equivalent to "no output of ω on inputs from S lies outside S"; when ω takes several inputs to one output, the contrapositive caps how many inputs from S can lie on a fibre of ω. See Book 1, field: Logic and set theory, topic: contrapositive and closure.

Double counting against budgets. Finite T, two families covering T with each point in exactly one member of each, per-member caps: $|T| \le a \cdot \mathrm{cap}_1 + b \cdot \mathrm{cap}_2$ while transversality gives $|T| = 2ab$; with cap $(k-1)(\ell-1)$ per circle you need $ab > (a+b)(k-1)(\ell-1)/2$. See Book 1, field: Combinatorics, topic: double counting and incidences; for the chase route, topic: Euclidean Ramsey problems.

Circuit Bank

- **WHEN** you see a closure rule ("if these are all colour c then so is that") **DO** write its contrapositive as a numerical cap: "an object centred at a c-point carries at most m points of any other colour" **BECAUSE** a cap is countable and a closure rule is not; caps feed double counting
- **WHEN** you have two special points of different colours and the only shapes in the problem are circles **DO** draw families of concentric circles about each and study the grid of intersection points **BECAUSE** concentric circles are pairwise disjoint, so each grid point counts once against each budget, and $2ab$ points beat linear budgets
- **WHEN** you place three points on one circle **DO** state that they are noncollinear because a line meets a circle at most twice **BECAUSE** this is the hypothesis's precondition and graders look for it

Backtesting

1. First useful cue: "the centre of the circle passing through A, B and C is also this colour": a closure rule whose output is a centre.
2. Method the expression suggested: contrapositive cap, then two pencils of concentric circles and a count.
3. Where you would hesitate, and the thought to reinforce: whether the circles meet and whether the intersection points are distinct. Reinforce: radii in $(d, \frac{3}{2}d)$ make the meeting condition automatic, and a point knows its distances to P and Q.
4. The trap the question sets: using the rule forwards in a chase with no invariant, or silently assuming continuity of the colouring.
5. What makes you recognise it faster next time: the word "centre" attached to a closure rule. Centres bring concentric circles, and concentric circles are disjoint.
6. The reusable decision rule: closure rule, therefore cap; cap, therefore grid; grid, therefore count.

Mirror drills

1. (drill) The plane is coloured with k colours so that whenever ℓ distinct points of one colour lie on a circle, the centre has that colour. Show all points have one colour. *Answer:* a circle about a c-point carries at most $(k-1)(\ell-1)$ non-c points; with a circles about P and a about Q, the $2a^2$ grid points include at most $2a(k-1)(\ell-1)$ that are not red and as many that are not green; take $a > 2(k-1)(\ell-1)$.
2. (drill) Show that "red if $x > 0$, green otherwise" violates the hypothesis although each class is convex. *Answer:* three points with $x > 0$ on the circle of radius 2 centred at $(-1, 0)$ are red and their circumcentre is green.
3. (drill) The plane is two-coloured and not monochromatic. Show there are two points of the same colour at distance 1 and two of different colours at distance 1. *Answer:* a unit equilateral triangle gives the first pair by pigeonhole; for the second, join a red point to a green point by a zigzag chain of unit steps and take the first step where the colour changes.

1.8 B2. Lamina centroid against solid of revolution centroid

Let $f\colon [0,1] \to [0,\infty)$ be strictly increasing and continuous. Let R be the region bounded by $x = 0$, $x = 1$, $y = 0$ and $y = f(x)$. Let x_1 be the x-coordinate of the centroid of R. Let x_2 be the x-coordinate of the centroid of the solid generated by rotating R around the x-axis. Prove that $x_1 < x_2$.

Read it out loud

Two centroids of two objects built from one graph, and a strict inequality between their x-coordinates. Say what a centroid is: first moment divided by total mass. Slicing the lamina into vertical strips gives $x_1 = \int_0^1 xf \,/\, \int_0^1 f$; slicing the solid into discs of area $\pi f(x)^2$ gives $x_2 = \int_0^1 xf^2 \,/\, \int_0^1 f^2$, the π cancelling. So the sentence you should have written within thirty seconds is: x_1 is the mean of x under the weight f, x_2 is the mean of x under the weight f^2, and the claim is that multiplying the weight by the increasing factor f pushes the mean to the right.

The fields wired in are integral calculus (centroids by slicing), inequalities (Chebyshev's integral inequality, rearrangement) and probability (weighted means, covariance). What is tested is whether you see a ratio of integrals as a weighted mean and whether you know that two weighted means whose weights differ by a monotone factor compare in the direction of that factor. The hypothesis "strictly increasing" is the whole engine: it is the comonotonicity of x and $f(x)$, and "strictly" is what buys the strict inequality. "Continuous" is there so the integrals exist and so that a positive continuous integrand has a positive integral.

Approaches, ranked worst to best

1. **Formula grinding.** *Cue:* two explicit centroid formulas. *Circuit:* write x_1 and x_2 out, cross-multiply, try integration by parts on the four single integrals. (Stalls: there is no single-variable identity relating $\int xf^2 \int f$ to $\int xf \int f^2$. Testing $f(x) = x$ gives $2/3 < 3/4$, which is reassurance, not proof.)

2. **Geometric mass transfer.** *Cue:* a lamina and a solid over the same base. *Circuit:* stand a prism of height $f(x_1)$ on R (its centroid projects to x_1), carve away the part left of $x = x_1$ above the surface $z = f(x)$ and add the part right of it below that surface; the result has square cross-sections of area $f(x)^2$, whose x-centroid equals that of the solid of revolution. (The official second solution; correct and pretty. It costs three-dimensional bookkeeping and two lemmas a grader will want: removing mass left of the centroid moves the centroid right, and a solid's x-centroid depends only on its cross-sectional area function.)

3. **Single crossing.** *Cue:* compare two weighted means in one variable. *Circuit:* let $c = \int f^2 / \int f$, so $\int f(f - c) = 0$; $f - c$ changes sign once, at x_0, so $(x - x_0)f(f - c) \geq 0$ pointwise; integrate. (Complete in six lines and worth knowing: it is the one-variable proof of the same inequality. It ranks second only because it needs the crossing point and a small continuity argument for strictness, whereas the symmetrisation is one identity.)

4. **Symmetrise the double integral.** *Cue:* a product of two integrals is a double integral over the square. *Circuit:* write the cross-multiplied difference as $\iint f(x)f(y)^2(y - x)$, swap the names of the variables, add, and read off the sign from $(y - x)(f(y) - f(x)) \geq 0$. (Ideal: Chebyshev's integral inequality proved on the spot, in a form that generalises to any weight and any pair of comonotone functions.)

The ideal solution

Since f is strictly increasing and $f(0) \geq 0$, we have $f(x) > f(0) \geq 0$ for $x \in (0,1]$; as f is continuous, $\int_0^1 f > 0$ and $\int_0^1 f^2 > 0$.

The centroid of a plane region is its first moment divided by its area. Slicing R into vertical strips of width dx and height $f(x)$,

$$x_1 = \frac{\int_0^1 xf(x)\,dx}{\int_0^1 f(x)\,dx}.$$

The solid of revolution has, at abscissa x, a cross-section which is a disc of area $\pi f(x)^2$; its volume is $\int_0^1 \pi f(x)^2\,dx$ and its first moment about the plane $x = 0$ is $\int_0^1 x \cdot \pi f(x)^2\,dx$, so

$$x_2 = \frac{\int_0^1 xf(x)^2\,dx}{\int_0^1 f(x)^2\,dx}.$$

Both denominators are positive, so $x_1 < x_2$ is equivalent to $D > 0$, where

$$D := \Big(\int_0^1 f(x)\,dx\Big)\Big(\int_0^1 yf(y)^2\,dy\Big) - \Big(\int_0^1 xf(x)\,dx\Big)\Big(\int_0^1 f(y)^2\,dy\Big).$$

Each product of two integrals of continuous functions is a double integral over the square $Q = [0,1]^2$ (Fubini for continuous integrands), so

$$D = \iint_Q f(x)f(y)^2(y-x)\,dx\,dy.$$

The names of the variables are immaterial; renaming $x \leftrightarrow y$ gives $D = \iint_Q f(y)f(x)^2(x-y)\,dx\,dy$. Adding the two expressions,

$$2D = \iint_Q \left[f(x)f(y)^2(y-x) + f(y)f(x)^2(x-y)\right]dx\,dy = \iint_Q f(x)f(y)\,(y-x)\big(f(y)-f(x)\big)\,dx\,dy.$$

Let $g(x,y) = f(x)f(y)(y-x)(f(y)-f(x))$. Because f is increasing, $(y-x)$ and $(f(y)-f(x))$ have the same sign, so $g \geq 0$ on Q. At the interior point $(\frac{1}{4}, \frac{3}{4})$ we have $f(\frac{1}{4}) > 0$, $f(\frac{3}{4}) > 0$, $\frac{3}{4} - \frac{1}{4} > 0$ and $f(\frac{3}{4}) - f(\frac{1}{4}) > 0$, so $g(\frac{1}{4}, \frac{3}{4}) > 0$; since g is continuous, $g > 0$ on a small disc about $(\frac{1}{4}, \frac{3}{4})$ contained in Q. A continuous nonnegative function that is positive on an open set has positive integral, so $2D > 0$. Therefore $x_2 - x_1 = D/\big(\int_0^1 f \int_0^1 f^2\big) > 0$. □

Layer 1: WHAT

Block 1: write both centroids as moment over mass by slicing, and check the denominators are positive. Block 2: clear denominators; the claim becomes $D > 0$ for a difference of two products of integrals. Block 3: convert each product into a double integral over the square, rename the variables, add, and obtain an integrand that is a product of nonnegative factors. Block 4: strictness from one interior point of positivity plus continuity.

Juxtapose. The nominal route reaches Block 2 and then attacks the four single integrals: integration by parts, substitution, a test function or two. The bifurcation point is the treatment of a product of two integrals. The zombie sees two numbers multiplied; you see one double integral over a square, because that is what a product of integrals is, and the only symmetry a square has (swap the axes) is then applied at once. Everything after that is reading off signs.

Layer 2: HOW

Block 1. **Cue.** "Centroid" twice, of a lamina and of a solid of revolution. **Why natural.** A centroid is defined as first moment over total mass; the only computational content is choosing slices. Vertical strips are forced for R (the boundary is a graph over x), and discs perpendicular to the axis are forced for the solid. Writing x_2 with πf^2 and cancelling π is the moment you notice that rotation merely squares the weight. Positivity of the denominators is not decoration: $f(0)$ may be 0, so you must say $f > 0$ on $(0, 1]$. **Link to the next block.** Two ratios with positive denominators; a strict inequality between ratios is a strict inequality between cross products.

Block 2. **Cue.** Two weighted means of the same variable with weights f and f^2. **Why natural.** Cross-multiplying is the only move that removes fractions, and it produces D, a difference of two products of integrals. Nothing about f has been used yet; this block is pure bookkeeping. **Link to the next block.** D is built from four integrals in two different variables' worth of information; to see its sign you need them under one integral sign.

Block 3. **Cue.** A product of two integrals. **Why natural.** $\int a \cdot \int b = \iint a(x)b(y)$ is not a technique, it is the definition of the product written on the square. Once D is one double integral, the integrand $f(x)f(y)^2(y-x)$ has no sign, but the square is symmetric under $x \leftrightarrow y$ and the integral is unchanged by renaming. Adding the two names is the standard way to make a symmetric integrand out of an asymmetric one; the factor $(y-x)$ then forces its partner to be $f(y) - f(x)$, and comonotonicity of y and $f(y)$ makes the product nonnegative. This is exactly how Chebyshev's inequality is proved; you are reproving it because it is a two-line lemma and the grader wants to see it. **Link to the next block.** The integrand is nonnegative; the problem says strictly.

Block 4. **Cue.** "Strictly increasing" in the hypothesis and "$<$" in the conclusion. **Why natural.** $g \geq 0$ with equality on the diagonal is not yet $\iint g > 0$; you need a positive value on an open set, and continuity supplies it from a single point. The point $(\frac{1}{4}, \frac{3}{4})$ is chosen in the interior of the square, off the diagonal, so that both f factors and both differences are visibly positive and a whole disc about it lies inside Q; a boundary point such as $(\frac{1}{2}, 1)$ would need a half-disc and an extra sentence. **Link to the next block.** None; the proof is closed.

Layer 3: WHY

The truth beneath the route is Chebyshev's integral inequality: for a weight $w \geq 0$ and two functions u, v that are monotone in the same direction,

$$\int w \int wuv - \int wu \int wv = \frac{1}{2} \iint w(x)w(y)\,\big(u(x) - u(y)\big)\big(v(x) - v(y)\big)\,dx\,dy \geq 0.$$

Here $w = f$, $u = x$, $v = f$. In the language of probability, put the density $f / \int f$ on $[0, 1]$ and let X be the random abscissa; then $x_1 = \mathrm{E}[X]$, $x_2 = \mathrm{E}[Xf(X)] / \mathrm{E}[f(X)]$, and

$$x_2 - x_1 = \frac{\mathrm{Cov}(X, f(X))}{\mathrm{E}\, f(X)},$$

which is positive because a random variable and an increasing function of it are positively correlated. That single sentence is the entire problem, and it is why the answer could never have gone the other way.

The discernment mechanism: every ratio of integrals with the same variable upstairs is a weighted mean, and every comparison of two weighted means is a covariance in disguise. When the two weights differ by a factor h, the sign of $\mathrm{E}_w[Xh] - \mathrm{E}_w[X]\,\mathrm{E}_w[h]$ is the sign of the monotonicity of h. Train it by refusing to compute centroids numerically: rewrite every centroid, centre of mass,

expectation and average as $\mathrm{E}_w[\cdot]$, name the weight, and ask what changed between the two weights. The examiners nudged by the phrase "strictly increasing" (there is no other property of f in play) and by rotating about the x-axis rather than the y-axis, which squares the weight without changing the variable being averaged.

The single-crossing proof is the same fact seen in one variable: $f - c$ with $c = \mathrm{E}_w f$ crosses zero once, and $(x - x_0)(f(x) - c) \geq 0$ is the one-dimensional shadow of $(x - y)(f(x) - f(y)) \geq 0$. Knowing both forms means you can choose the cheaper one in the arena: symmetrisation when the weight is arbitrary, single crossing when you want to avoid double integrals altogether.

Revolving concepts

Centroid. Definition: $\bar{x} = \int x\,dm / \int dm$ for a mass distribution; for a lamina under a graph, $dm = f(x)\,dx$; for a solid with cross-sectional area $A(x)$, $dm = A(x)\,dx$, so the x-centroid depends only on the area function A, not on the shape of the sections (this is what makes the square-section solid and the solid of revolution interchangeable in the geometric route). Pappus's theorems sit next door. See Book 1, field: Calculus and integrals, topic: moments and centroids by slicing.

Chebyshev's integral inequality. Statement and the symmetrisation identity above; equality iff one of u, v is constant on the support of w. Discrete form: if $a_1 \leq \cdots \leq a_n$ and $b_1 \leq \cdots \leq b_n$ then $n \sum a_i b_i \geq \sum a_i \sum b_i$, proved by $\sum_{i,j}(a_i - a_j)(b_i - b_j) \geq 0$. Rearrangement inequality is the ancestor. See Book 1, field: Inequalities, topic: Chebyshev and rearrangement.

Covariance of comonotone variables. $\mathrm{Cov}(X, h(X)) \geq 0$ for h increasing, with the same symmetrisation proof on the product space; more generally $\mathrm{Cov}(u(X), v(X)) \geq 0$ for u, v similarly monotone (the FKG inequality in one dimension). See Book 1, field: Probability, topic: covariance and monotone coupling.

The one-parameter family. Generalised statement: for $p > 0$ let $x_p = \int x f^p / \int f^p$; then $p \mapsto x_p$ is strictly increasing when f is strictly increasing and positive. Proof identical with $w = f^p$ and $v = f^{q-p}$. This is the recursion-flavoured generalisation: each extra factor of f moves the mean right by $\mathrm{Cov}_w(X, f) / \mathrm{E}_w f$. Neighbouring fact: $p \mapsto \log \int f^p$ is convex (Cauchy's inequality). See Book 1, field: Inequalities, topic: power means and log-convexity.

Positivity of integrals. A continuous $g \geq 0$ on a rectangle with $g(P) > 0$ at one point has $\iint g > 0$: continuity gives a disc where $g > g(P)/2$. See Book 1, field: Real analysis, topic: integrals of continuous functions.

Circuit Bank

- **WHEN** you see a centroid, centre of mass, average or expectation **DO** write it as $\int x\,w / \int w$ and name the weight **BECAUSE** comparisons of centroids are comparisons of weighted means, and the weight is where the hypothesis lives
- **WHEN** you must compare $\int a \int b$ with $\int c \int d$ **DO** write each product as a double integral over the square, rename $x \leftrightarrow y$, and add **BECAUSE** a product of integrals is a double integral, and the swap is the only symmetry the square has; symmetrising produces a factor $(x - y)$ paired with a difference
- **WHEN** two weights differ by a monotone factor h **DO** expect the mean to move in the direction of h and prove it by $\mathrm{Cov}(X, h(X)) \geq 0$ **BECAUSE** a variable and an increasing function of it are positively correlated
- **WHEN** you need a strict inequality from a nonnegative integrand **DO** exhibit one point of strict positivity and cite continuity **BECAUSE** "equality only on the diagonal" is not a proof; an open set of positivity is

Backtesting

1. First useful cue: "centroid" of a region and of its solid of revolution: the weights are f and f^2, the averaged variable is x in both.
2. Method the expression suggested: weighted means, cross-multiply, write products as double integrals, symmetrise.
3. Where you would hesitate, and the thought to reinforce: whether π matters (it cancels), whether a denominator can vanish ($f(0)$ may be 0 but $f > 0$ on $(0,1]$), and how to get strictness (one interior point plus continuity).
4. The trap the question sets: three-dimensional visualisation that feels like insight but needs unproved lemmas about moving mass, or grinding four single integrals that never talk to each other.
5. What makes you recognise it faster next time: "strictly increasing" next to a ratio of integrals. Monotone plus mean equals covariance.
6. The reusable decision rule: ratio of integrals, therefore weighted mean; two weighted means, therefore a covariance; covariance of comonotone functions, therefore symmetrise and read the sign.

Mirror drills

1. (drill) For the same f, show that $x_p = \int_0^1 x f^p / \int_0^1 f^p$ is strictly increasing in $p > 0$. *Answer:* for $q > p$, $\int f^p \int x f^q - \int x f^p \int f^q = \frac{1}{2} \iint f(x)^p f(y)^p (y-x)\big(f(y)^{q-p} - f(x)^{q-p}\big) > 0$ since f^{q-p} is strictly increasing.
2. (drill) If $u, v\colon [0,1] \to \mathbb{R}$ are continuous and both increasing, prove $\int_0^1 uv \geq \int_0^1 u \int_0^1 v$. *Answer:* $\int uv - \int u \int v = \frac{1}{2} \iint (u(x) - u(y))(v(x) - v(y)) \geq 0$.
3. (drill) Compute x_1 and x_2 for $f(x) = x$, and state what happens to the inequality if f is strictly decreasing. *Answer:* $x_1 = \frac{1/3}{1/2} = \frac{2}{3}$, $x_2 = \frac{1/4}{1/3} = \frac{3}{4}$; for decreasing f the same identity gives $x_1 > x_2$, since $(y-x)(f(y) - f(x)) \leq 0$.

1.9 B3. Divisor-closed set forced to contain every integer

Suppose S is a nonempty set of positive integers with the property that if n is in S, then every positive divisor of $2025^n - 15^n$ is in S. Must S contain all positive integers?

Read it out loud

A closure condition on a set of positive integers: membership of n forces membership of every divisor of $F(n) = 2025^n - 15^n$, and you are asked whether the closure of any nonempty seed is the whole of $\mathbb{N}$. The answer is yes, and the structure that makes it yes is in the numbers: $2025 = 45^2 = 3^4 \cdot 5^2$ and $15 = 3 \cdot 5$ have the same prime support, so $2025 = 15 \cdot 135$ and

$$F(n) = 15^n\big(135^n - 1\big).$$

That factorisation is the sentence to write in the first thirty seconds. It splits every modulus into two kinds of prime: 3 and 5, which are swallowed by 15^n once n is large enough, and everything else, which is coprime to 135 and is therefore handled by Euler's theorem on $135^n - 1$.

The fields wired in are elementary number theory (Euler's theorem, orders, coprime factors of a common multiple) and the logic of strong induction with a constructed witness. What is really tested is the design of the witness: the induction hands you every $k \le n - 1$, and you must find one k in that range with $n \mid F(k)$. The witness must be large enough for the powers of 3 and 5, a multiple of $\phi(d)$ for the coprime part d, and below n. The slack that makes all three compatible is $\phi(d) < d \le n$, which is why Euler's theorem, not Fermat's, is the right tool: Euler always hands you an exponent smaller than the modulus.

Approaches, ranked worst to best

1. **Numerical exploration.** *Cue:* a concrete base, so compute. *Circuit:* $F(1) = 2010 = 2 \cdot 3 \cdot 5 \cdot 67$ puts $1, 2, 3, 5, 6, 10, 15, 30, 67, \ldots$ into S; then factor $F(2)$, $F(3)$, and hope everything appears. (Cannot conclude for all n; a finite computation never proves a universal statement. Useful only to guess the answer, which the factorisation already told you.)
2. **Fermat reflex.** *Cue:* $n \mid c^{n-1} - 1$ for prime n. *Circuit:* try the witness $k = n - 1$ and Fermat's little theorem. (Fails for composite n: $4 \nmid 135^3 - 1$ since $135 \equiv -1 \pmod 4$. The right exponent is a multiple of $\phi(d)$, and it must also be at least the 3-adic and 5-adic exponents of n.)
3. **Tower of iterates.** *Cue:* iterate the closure from the seed. *Circuit:* let $a_0 = 1$, $a_{i+1} = F(a_i)$; show $a_i \mid a_{i+1}$ via $x^m - y^m \mid x^{md} - y^{md}$, then show every n divides some a_i by minimal counterexample, reducing to prime powers coprime to 15 and applying Euler. (The official second solution; correct and conceptually clean, since the divisor sets of the tower are nested. It costs three lemmas where the ideal route costs one inequality.)
4. **Strong induction with a designed witness.** *Cue:* the hypothesis gives you all $k < n$. *Circuit:* write $n = 3^a 5^b d$, $\gcd(d, 15) = 1$, $c = \max(a, b)$, and take $k = \max(c, 1)\,\phi(d)$; check $1 \le k \le n - 1$ and $n \mid 15^k(135^k - 1)$. (Ideal: one induction, one explicit witness, one inequality $c\,\phi(d) \le n - 1$.)

The ideal solution

Yes: S must contain every positive integer.

Write $F(k) = 2025^k - 15^k$. Since $2025 = 15 \cdot 135$,

$$F(k) = 15^k\big(135^k - 1\big).$$

We prove $n \in S$ by strong induction on n.

Base. S is nonempty; pick $n_0 \in S$. Then 1 divides $F(n_0)$, so $1 \in S$.

Step. Let $n \geq 2$ and assume $1, 2, \ldots, n-1 \in S$. Write $n = 3^a 5^b d$ with $a, b \geq 0$ and $\gcd(d, 15) = 1$, and put $c = \max(a, b)$ and

$$k = \max(c, 1) \cdot \phi(d),$$

where ϕ is Euler's totient function.

First, $1 \leq k \leq n-1$. Clearly $k \geq 1$. If $c = 0$ then $n = d \geq 2$ and $k = \phi(n) \leq n-1$, since n itself is not coprime to n. If $c \geq 1$ then, using $\phi(d) \leq d$, the inequality $c \leq 3^c - 1$ (valid for $c \geq 1$, as $3^c = (1+2)^c \geq 1 + 2c$), $d \geq 1$, and $3^c \leq 3^a 5^b$,

$$k = c\,\phi(d) \leq c\,d \leq (3^c - 1)\,d = 3^c d - d \leq 3^a 5^b d - 1 = n - 1.$$

Hence $k \in S$ by the induction hypothesis, and so every positive divisor of $F(k)$ lies in S.

Second, $n \mid F(k)$. Since $k \geq c \geq a$ and $k \geq c \geq b$, we have $3^a 5^b \mid 3^k 5^k = 15^k$. Since $\gcd(d, 135) = 1$, Euler's theorem gives $135^{\phi(d)} \equiv 1 \pmod{d}$, and $\phi(d) \mid k$, so $135^k \equiv 1 \pmod{d}$; that is, $d \mid 135^k - 1$. (For $d = 1$ this is trivial.) The integers $3^a 5^b$ and d are coprime and both divide $F(k) = 15^k(135^k - 1)$, so their product n divides $F(k)$.

Therefore n is a positive divisor of $F(k)$ with $k \in S$, and $n \in S$. By strong induction, $S = \mathbb{N}$. □

Layer 1: WHAT

Block 1: answer yes; factor $F(k) = 15^k(135^k - 1)$; note $1 \in S$ from nonemptiness. Block 2: strong induction, and split the target n into its $\{3, 5\}$-part $3^a 5^b$ and its coprime part d, because the two factors of $F(k)$ serve the two parts. Block 3: design the witness: k must lie in $\{1, \ldots, n-1\}$, be at least $\max(a, b)$, and be a multiple of $\phi(d)$; the slack $\phi(d) \leq d$ and $3^c d \leq n$ makes $k = \max(c, 1)\phi(d)$ work. Block 4: verify $n \mid F(k)$ by checking the two coprime factors separately and multiplying.

Juxtapose. The nominal route either computes examples or iterates the closure from 1 and studies the tower. The bifurcation point is the question "which k do I use?" The zombie reaches for $k = n - 1$ (the Fermat reflex) and dies on $n = 4$; you list the three constraints on k, notice that Euler's exponent $\phi(d)$ is strictly smaller than $d \leq n$, and pick the smallest multiple of $\phi(d)$ that clears c.

Layer 2: HOW

Block 1. **Cue.** "$2025^n - 15^n$" with $2025 = 45^2$ and $15 = 3 \cdot 5$: the two bases share their primes. **Why natural.** A difference of powers with a common factor is $g^n(A'^n - B'^n)$; here $g = 15$, $A' = 135$, $B' = 1$, so the second factor is of Euler type. The question "must S contain all positive integers" is a universal claim, and $1 \in S$ is immediate, so strong induction is the natural frame: the closure gives you all divisors of $F(k)$ for each k already admitted. **Link to the next block.** To admit n, some already-admitted k must satisfy $n \mid F(k)$; the two factors of $F(k)$ tell you how to split n.

Block 2. **Cue.** $F(k) = 15^k \cdot (135^k - 1)$, a prime-power factor times a "$c^k - 1$" factor. **Why natural.** A prime $p \mid n$ is either in $\{3, 5\}$, and then p-adically only 15^k can help (as $135^k - 1 \equiv -1 \pmod{p}$), or coprime to 135, and then only $135^k - 1$ can help. So $n = 3^a 5^b d$ is not a choice but the shape the modulus is forced into, and $\gcd(3^a 5^b, d) = 1$ lets you verify divisibility one factor at a time. **Link to the next block.** Each factor imposes a condition on k; collect them.

Block 3. **Cue.** Three constraints on one integer k: $k \geq \max(a, b)$, $\phi(d) \mid k$, $k \leq n - 1$. **Why natural.** Existence of a witness is an inequality, not a search. Multiples of $\phi(d)$ that clear c start at $c\,\phi(d)$ (or $\phi(d)$ when $c = 0$), so the only question is whether $c\,\phi(d) \leq n - 1$. It is, with room: $\phi(d) \leq d$ and $c \leq 3^c - 1$ give $c\,\phi(d) \leq 3^c d - d \leq n - 1$. The case $c = 0$ is Euler's $\phi(n) < n$. This is where the induction hypothesis is cashed: the witness is below n. **Link to the next block.** With k fixed and admitted, verification is mechanical.

Block 4. **Cue.** Two coprime factors of n, two factors of $F(k)$. **Why natural.** $3^a5^b \mid 15^k$ needs only $k \geq a$ and $k \geq b$ (not $k \geq a+b$, a common overshoot). $d \mid 135^k - 1$ is Euler's theorem plus $\phi(d) \mid k$. Coprime divisors of a common multiple multiply: $n \mid F(k)$. **Link to the next block.** None; the induction closes.

Layer 3: WHY

The truth underneath is a statement about differences of powers with equal radicals. If $A \neq B$ are positive integers with the same set of prime divisors, write $g = \gcd(A, B)$, $A = gA'$, $B = gB'$. Every prime dividing $A'B'$ divides g, so for any modulus n the primes of n split into those dividing g (absorbed by g^k for k at least the largest exponent in n) and those coprime to $A'B'$ (handled by $A'^k \equiv B'^k \pmod{d}$ for $\phi(d) \mid k$, after multiplying by the inverse of B'). So every n divides $A^k - B^k$ for some $k < n$, and the closure under "divisors of $A^n - B^n$" of any nonempty seed is $\mathbb{N}$. Without equal radicals the statement fails: with $A = 2$, $B = 1$, the set $S = \{1\}$ is closed, because $2^1 - 1 = 1$, and no even number ever divides $2^n - 1$. The examiners chose 2025 and 15 precisely to satisfy $\mathrm{rad}(2025) = \mathrm{rad}(15)$ while looking arbitrary.

The discernment mechanism has two parts. First, for a difference of powers, always ask "which primes are absorbed and which are Euler": it is the p-adic reading of $x^n - y^n$. Second, whenever an induction gives you "all smaller values", write down exactly what the step needs from a witness and turn its existence into an inequality; the standard slack in multiplicative number theory is $\phi(m) < m$ (or better, $\mathrm{ord}_m(c) \mid \phi(m)$), which is what makes "an exponent below the modulus" always available. You train both by refusing to test small cases before you have factored, and by never writing "by Fermat" for a modulus that is not prime.

The nudge in the question is the word "must". A yes-or-no phrasing invites you to look for a counterexample; the factorisation shows there is none, and the same factorisation is the proof. Consumption of examples costs ink; analysis of the expression costs one line.

Revolving concepts

Euler's theorem and ϕ. If $\gcd(c, m) = 1$ then $c^{\phi(m)} \equiv 1 \pmod{m}$; ϕ is multiplicative, $\phi(p^e) = p^{e-1}(p-1)$, and $\phi(m) \leq m - 1$ for $m \geq 2$ with equality iff m is prime. Sharper: the order $\mathrm{ord}_m(c)$ divides $\phi(m)$ and divides Carmichael's $\lambda(m)$; any multiple of the order works as the exponent. This is beyond the A-level syllabus (Fermat is not enough here) and must be installed. See Book 1, field: Number theory, topic: Euler's theorem, orders and the totient.

Differences of powers. $x - y \mid x^k - y^k$; more generally $x^m - y^m \mid x^{md} - y^{md}$ by the factorisation $x^{md} - y^{md} = (x^m - y^m)\left(x^{m(d-1)} + x^{m(d-2)}y^m + \cdots + y^{m(d-1)}\right)$. This is the lemma behind the tower route and behind lifting-the-exponent arguments. See Book 1, field: Number theory, topic: divisibility of $a^n - b^n$.

Coprime divisors multiply. If $u \mid m$, $v \mid m$ and $\gcd(u, v) = 1$ then $uv \mid m$; equivalently $\mathrm{lcm}(u, v) = uv$. Checking divisibility one prime power at a time is the working form of the Chinese remainder theorem. See Book 1, field: Number theory, topic: lcm, gcd and the Chinese remainder theorem.

Radical and p-adic splitting. $\mathrm{rad}(A)$ is the product of the distinct primes of A; the general theorem above holds exactly when $\mathrm{rad}(A) = \mathrm{rad}(B)$. See Book 1, field: Number theory, topic: p-adic valuation of $a^n - b^n$.

The recursion. The closure of a seed under "divisors of $F(\cdot)$" is the union of the divisor sets of the tower $a_0 = \text{seed}$, $a_{i+1} = F(a_i)$, and when $a_i \mid a_{i+1}$ the divisor sets are nested, so "is the closure everything" becomes "does every n divide some a_i". Generalised: for any F with $u \mid v \Rightarrow F(u) \mid F(v)$ the same reduction applies, and minimal-counterexample arguments reduce to prime powers. The strong induction of the ideal route is the same recursion read one step at a time: to reach n you need

one earlier level k with $n \mid F(k)$. See Book 1, field: Logic and set theory, topic: closure operators and strong induction.

Circuit Bank

- **WHEN** you see $A^n - B^n$ with $\gcd(A, B) > 1$ **DO** factor out $\gcd(A, B)^n$ and separate the primes of the gcd (absorbed by the power) from the primes coprime to it (handled by Euler) **BECAUSE** the p-adic behaviour of $x^n - y^n$ splits along $p \mid \gcd$ versus p coprime
- **WHEN** you need an exponent k with $m \mid c^k - 1$ and $\gcd(c, m) = 1$ **DO** take a multiple of $\phi(m)$, or of $\mathrm{ord}_m(c)$, and note it can be chosen below m **BECAUSE** Euler's theorem, and $\phi(m) < m$ for $m \geq 2$
- **WHEN** a closure condition "$n \in S \Rightarrow$ divisors of $F(n) \in S$" asks whether S is everything **DO** strong induction: to admit n, exhibit $k < n$ with $n \mid F(k)$ **BECAUSE** the hypothesis makes every $k < n$ available, so the only problem is the design of k
- **WHEN** a witness must satisfy several constraints at once **DO** write each constraint as an inequality or a divisibility and locate the slack **BECAUSE** existence is an inequality, not a search; here the slack is $\phi(d) \leq d$ and $c \leq 3^c - 1$

Backtesting

1. First useful cue: $2025 = 45^2$ and $15 = 3 \cdot 5$ share their primes, so $F(n) = 15^n(135^n - 1)$.
2. Method the expression suggested: strong induction; witness k a multiple of $\phi(d)$ that is at least $\max(a, b)$ and below n.
3. Where you would hesitate, and the thought to reinforce: whether $k \leq n - 1$ really holds (it does, with the slack $c\,\phi(d) \leq 3^c d - d$), and whether $3^a 5^b \mid 15^k$ needs $k \geq a + b$ (it needs only $k \geq \max(a, b)$).
4. The trap the question sets: answering "no" by analogy with $2^n - 1$, or using Fermat with $k = n - 1$, which fails at $n = 4$.
5. What makes you recognise it faster next time: "every divisor of" next to a difference of powers whose bases share primes.
6. The reusable decision rule: shared primes are absorbed by the power, coprime primes are handled by Euler with an exponent below the modulus, and strong induction supplies the exponent.

Mirror drills

1. (drill) Same question with $2^n - 1$ in place of $2025^n - 15^n$. *Answer:* No: $S = \{1\}$ is closed since $2^1 - 1 = 1$, and no even number divides $2^n - 1$; the radicals of 2 and 1 differ.
2. (drill) Show that every positive integer coprime to 10 divides a repunit $11 \cdots 1$. *Answer:* for n coprime to 10, Euler on the modulus $9n$ gives $9n \mid 10^{\phi(9n)} - 1$, so $n \mid (10^{\phi(9n)} - 1)/9$, a repunit.
3. (drill) With $a_0 = 1$ and $a_{i+1} = 2025^{a_i} - 15^{a_i}$, prove $a_i \mid a_{i+1}$ for all i. *Answer:* induction; if $a_i \mid a_{i+1}$ then $2025^{a_i} - 15^{a_i} \mid 2025^{a_{i+1}} - 15^{a_{i+1}}$ by $x^m - y^m \mid x^{md} - y^{md}$.

1.10 B4. Sum versus support of a staircase matrix

For $n \geq 2$, let $A = [a_{i,j}]_{i,j=1}^{n}$ be an n-by-n matrix of nonnegative integers such that

(a) $a_{i,j} = 0$ *when* $i + j \leq n$;

(b) $a_{i+1,j} \in \{a_{i,j}, a_{i,j} + 1\}$ *when* $1 \leq i \leq n - 1$ *and* $1 \leq j \leq n$; *and*

(c) $a_{i,j+1} \in \{a_{i,j}, a_{i,j} + 1\}$ *when* $1 \leq i \leq n$ *and* $1 \leq j \leq n - 1$.

Let S be the sum of the entries of A, and let N be the number of nonzero entries of A. Prove that

$$S \leq \frac{(n+2)N}{3}.$$

Read it out loud

A square array of nonnegative integers, zero on and above the anti-diagonal, and each step to the right or downward raises the entry by 0 or 1. So the entries are nondecreasing in both directions and grow at most one unit per step away from the zero band. S/N is the average nonzero entry, and the claim is that this average never exceeds $(n+2)/3$. The fields wired in are the combinatorics of monotone arrays (the zero set is an order ideal), Lipschitz bounds (an entry is at most its distance from a zero), averaging arguments, and the discipline of computing the extremal example before proving anything.

The first thing to do, within thirty seconds, is to build the matrix that should be extremal: fill the triangle $i + j > n$ with the distance to the anti-diagonal band, $a_{i,j} = i + j - n$. It satisfies (a), (b), (c), the entries equal to k number $n + 1 - k$, so

$$N = \frac{n(n+1)}{2}, \qquad S = \sum_{k=1}^{n} k(n+1-k) = \frac{n(n+1)(n+2)}{6}, \qquad \frac{S}{N} = \frac{n+2}{3}.$$

The constant is therefore the average of $i + j - n$ over the full triangle, and the bound is tight. What is really tested is the next question: why can a matrix with a bigger zero region, or with slower growth, not beat this average? A bigger zero region shrinks N, but it also forces small entries near its boundary, and the proof has to quantify that second effect. The fraction $\frac{1}{3}$ in the target is a cue in itself: an average of three bounds.

Approaches, ranked worst to best

1. **Extremal guessing.** *Cue:* a maximisation of S/N over a finite family. *Circuit:* argue that filling the whole triangle with the fastest growth is "obviously" best. (Not a proof. The admissible family is enormous: any down-set of zeros, any pattern of 0/1 increments. Removing entries changes N and S at once and the ratio can move either way.)

2. **One bound.** *Cue:* entries grow at most 1 per step from the band $i + j \leq n$. *Circuit:* prove $a_{i,j} \leq i + j - n$ and sum. (Correct but insufficient: it gives $S \leq \sum_{\text{nonzero}} (i + j - n)$, which equals $(n+2)N/3$ for the full triangle and exceeds it otherwise; for $n = 2$ and A with the single nonzero entry $a_{2,2} = 1$, it yields $S \leq 2$ against the target $4/3$. The bound ignores where the zeros end.)

3. **Taxicab reduction and induction on n.** *Cue:* a 0/1-growth array dominated by distances to the zero set. *Circuit:* replace each entry by its taxicab distance to the zero set (this keeps N, does not decrease S, and preserves the conditions), then peel: subtract 1 from the nonzero entries and delete the first row and column; prove the level counts c_k strictly decrease and close the induction. (The official first solution. It works, but the strict decrease of the level counts needs an injection with four verifications, and the base case is a five-matrix check. Far more ink than the ideal.)

4. **Three bounds that average to a constant.** *Cue:* the target is $\frac{1}{3}(n+2)$ per nonzero entry. *Circuit:* bound each nonzero entry three ways, by its distance to the band ($i+j-n$), by its position among the nonzero entries of its row ($j - Z_i$) and of its column ($i - W_j$); sum the row bound along each row and re-index it as $n+1-j$, likewise columns; the three bounds add to $n+2$ at every nonzero position. (Ideal. Two short lemmas and one addition.)

The ideal solution

By (b) and (c), $a_{i,j} \le a_{i+1,j}$ and $a_{i,j} \le a_{i,j+1}$: entries are nondecreasing downward and to the right. Hence if $a_{i,j} = 0$ then $a_{i',j'} = 0$ for all $i' \le i$, $j' \le j$. In particular the zeros of row i occupy an initial block of columns $1, \dots, Z_i$ and the nonzero entries occupy columns $Z_i + 1, \dots, n$ (with $Z_i = n$ if the row is zero); similarly the zeros of column j occupy rows $1, \dots, W_j$.

Lemma 1.10.1. *For every (i,j), $a_{i,j} \le \max(0, i+j-n)$. In particular every nonzero entry satisfies $a_{i,j} \le i+j-n$.*

Proof. Induction on i. For $i = 1$: if $j \le n-1$ then $a_{1,j} = 0$ by (a); and $a_{1,n} \le a_{1,n-1} + 1 = 1 = 1 + n - n$ by (c) and (a). For $i \ge 2$: if $i + j \le n$ the entry is 0 by (a); if $i + j \ge n+1$ then $(i-1) + j \ge n$, so by (b) and the induction hypothesis $a_{i,j} \le a_{i-1,j} + 1 \le \max(0, i-1+j-n) + 1 = i+j-n$. □

Lemma 1.10.2. *Every nonzero entry satisfies $a_{i,j} \le j - Z_i$ and $a_{i,j} \le i - W_j$.*

Proof. Consider row i with $Z_i < n$. Its first nonzero entry a_{i,Z_i+1} is at most 1: if $Z_i \ge 1$ then by (c) it is at most $a_{i,Z_i} + 1 = 1$; if $Z_i = 0$ then $a_{i,1} \ne 0$ forces $i + 1 > n$, so $i = n$, and by (b) and (a) (note $(n-1)+1 \le n$) we get $a_{n,1} \le a_{n-1,1} + 1 = 1$. Moving right along the row adds at most 1 per step by (c), so for $j \ge Z_i + 1$, $a_{i,j} \le 1 + (j - Z_i - 1) = j - Z_i$. The column statement is identical with (b) in place of (c) and the roles of i, j exchanged. □

Now sum. Let P be the set of positions of the nonzero entries, $|P| = N$. In row i the nonzero positions are $j = Z_i + 1, \dots, n$, so by the second lemma

$$\sum_{j=Z_i+1}^{n} a_{i,j} \;\le\; \sum_{j=Z_i+1}^{n} (j - Z_i) \;=\; 1 + 2 + \cdots + (n - Z_i) \;=\; \sum_{j=Z_i+1}^{n} (n+1-j),$$

the last equality because both sides list the integers $1, \dots, n - Z_i$, in opposite orders. Summing over rows,

$$S \le \sum_{(i,j)\in P} (n+1-j).$$

The same argument along columns gives $S \le \sum_{(i,j)\in P}(n+1-i)$, and the first lemma gives $S \le \sum_{(i,j)\in P}(i+j-n)$. Adding the three inequalities,

$$3S \;\le\; \sum_{(i,j)\in P} \big[(i+j-n) + (n+1-j) + (n+1-i)\big] \;=\; \sum_{(i,j)\in P} (n+2) \;=\; (n+2)N,$$

which is the claim. □

Equality holds for the full taxicab triangle $a_{i,j} = \max(0, i+j-n)$, where all three bounds are tight at every nonzero position, and for the zero matrix.

Layer 1: WHAT

Block 1: read (b), (c) as monotonicity plus growth at most 1 per step; deduce that the zero set is a down-set, so each row's zeros form a prefix of length Z_i and each column's a prefix of length W_j. Block 2: three upper bounds for a nonzero entry, one per route to a zero: up-left to the band ($i+j-n$), left to the row's last zero ($j - Z_i$), up to the column's last zero ($i - W_j$). Block 3: sum

each bound over P; the row bound sums, row by row, to a triangular number that re-indexes as $\sum(n+1-j)$, and likewise for columns. Block 4: add the three sums; the position-dependent parts cancel to the constant $n+2$.

Juxtapose. The nominal route proves the band bound $a_{i,j} \le i+j-n$, sees it is not enough, and either tries to sharpen it or drops into the induction on n. The bifurcation point is what you do with an insufficient bound. The zombie sharpens; you add. The reason is in the equality case: in the full triangle the three distances coincide ($Z_i = n-i$, $W_j = n-j$), so all three bounds are tight there, and the target constant $(n+2)/3$ is announcing that three quantities are being averaged.

Layer 2: HOW

Block 1. **Cue.** "$a_{i+1,j} \in \{a_{i,j}, a_{i,j}+1\}$" and the same to the right: monotone, and Lipschitz with constant 1. **Why natural.** Monotonicity in both directions makes the zero set closed under moving up or left, which is the only way to describe where the nonzero entries live: a suffix of each row and of each column. Without this you cannot even name the row sums you are about to bound. Lipschitz 1 from a zero is the only quantitative content the conditions have, so every bound you will ever prove is "at most the number of steps to some zero". **Link to the next block.** Choose the zeros to measure from. There are three natural ones for each entry.

Block 2. **Cue.** The band $i+j \le n$ is a zero region you know; the row and column prefixes are zero regions you have just found. **Why natural.** Three routes from (i,j) to a zero, each a straight path in the grid, each giving a bound equal to its length: diagonally up-left to the band in $i+j-n$ steps, left in $j-Z_i$ steps, up in $i-W_j$ steps. The band bound needs an induction because the band is not adjacent to most entries; the row and column bounds need only that the first nonzero entry in a line is 1, which is one step from a zero (with the single exception $a_{n,1}$, handled by the column). These are the tightest bounds available from single straight paths, and in the extremal matrix all three are exact. **Link to the next block.** Each bound is pointwise; the target is a sum. Sum them.

Block 3. **Cue.** A per-row bound $j-Z_i$ whose values along the row are $1, 2, \ldots, n-Z_i$. **Why natural.** The sum of the row bound over the row is a triangular number determined by the row's count of nonzero entries alone, not by Z_i's position, and the multiset $\{1, \ldots, m\}$ reads the same from either end. So you may charge the value $n+1-j$ to position (i,j) instead of $j-Z_i$: the total is unchanged and the new label depends on j only, not on the row. This is the whole content of the "reverse the nonzero entries in each row" device of the second official solution, without constructing the reversed matrices. The column bound re-indexes to $n+1-i$ the same way. **Link to the next block.** Three sums over the same index set P, with summands $i+j-n$, $n+1-j$, $n+1-i$.

Block 4. **Cue.** Three summands whose i and j dependence cancels. **Why natural.** $(i+j-n)+(n+1-j)+(n+1-i) = n+2$ at every position. Adding the three inequalities gives $3S \le (n+2)N$; there is nothing left to estimate. **Link to the next block.** None. Record the equality case as a check that no inequality was wasted.

Layer 3: WHY

The truth beneath the route is that an entry is at most its taxicab distance to the zero set, and that three specific paths to zero give three bounds whose sum, after a legal re-indexing within rows and within columns, is a constant. The re-indexing is a rearrangement: the row bound summed over a row depends only on how many nonzero entries the row has, so the labels may be assigned from the right end instead of the left. Averaging then does the rest. In the extremal matrix every path is a geodesic to the zero set and every bound is an equality; in any other admissible matrix at least one

bound is slack somewhere, and the same three inequalities absorb the slack automatically. That is why no case analysis on the shape of the zero region is needed.

The discernment mechanism: when a target reads "$S \le cN/m$" for an integer m, look for m per-object bounds summing to c. Compute the extremal configuration first; it tells you what "tight" looks like, and the bounds you need are exactly the ones that are tight there. When one bound sums to a row-dependent triangular number, re-index it from the other end so that it becomes pointwise comparable with the others. Train this by refusing to sharpen a single bound that is tight in the extremal case (it cannot be sharpened; the extremal case is a witness) and instead asking what other bound is also tight there.

The examiners nudged twice. The anti-diagonal band makes $i + j - n$ a coordinate you cannot avoid, and the denominator 3 tells you how many bounds to find. The same device, with k axis bounds and one diagonal bound, gives for k-dimensional arrays with the analogous conditions the bound $S \le (n + k)N/(k + 1)$, the average of the extremal simplex.

Revolving concepts

Monotone arrays and down-sets. If a is nondecreasing in each coordinate on a grid, the set where $a \le t$ is an order ideal (closed under decreasing coordinates) for every t; its rows and columns are prefixes. Neighbouring objects: plane partitions, Young diagrams, the Ferrers shape of the zero set. See Book 1, field: Discrete structures and graphs, topic: partial orders, order ideals and Young diagrams.

Lipschitz bounds from a boundary. If $|a(u) - a(v)| \le 1$ for adjacent u, v and $a = 0$ on a set Z, then $a(u) \le d(u, Z)$ in the graph (taxicab) metric; each monotone path to Z gives a bound equal to its length. Replacing a by $d(\cdot, Z)$ is the official reduction to "taxicab matrices". See Book 1, field: Inequalities, topic: Lipschitz functions and distance to a set.

Averaging arguments. To prove $\sum_{u \in P} a(u) \le c|P|$, it suffices to find bounds $b_1, \dots, b_m$ with $a \le b_t$ on P and $\sum_t b_t(u) \le mc$ pointwise, or with $\sum_t \sum_u b_t(u) \le mc|P|$ after re-indexing. This is the same mechanism as bounding an average by the average of several majorants. See Book 1, field: Combinatorics, topic: double counting and averaging.

Triangular numbers and re-indexing. $\sum_{j=Z+1}^{n} (j - Z) = \sum_{j=Z+1}^{n} (n + 1 - j) = \binom{n-Z+1}{2}$; more generally any sum of a function of the rank within a line can be re-indexed from either end. See Book 1, field: Combinatorics, topic: sums over ranks and rearrangement.

The extremal example and its recursion. For $a_{i,j} = \max(0, i + j - n)$, the entries equal to k number $n + 1 - k$, so $S = \sum k(n + 1 - k) = \binom{n+2}{3}$ and $N = \binom{n+1}{2}$. The official induction peels one layer: subtracting 1 from the nonzero entries and deleting the first row and column maps $(S, N, n) \mapsto (S - N, N - c_1, n - 1)$, where c_1 counts entries equal to 1; generalised, any array built from nested level sets satisfies $S = \sum_{k \ge 1} N_k$ with N_k the number of entries at least k, which is the layer-by-layer recursion for sums of nonnegative integers. See Book 1, field: Sequences and recurrences, topic: layer-cake sums and peeling inductions.

Circuit Bank

- **WHEN** entries change by at most 1 per step and vanish on a known region **DO** write "entry $\le$ length of a monotone path to a zero" and list the straight paths available (row, column, diagonal band) **BECAUSE** Lipschitz 1 from zero is the only quantitative content of such conditions
- **WHEN** the target is $S \le cN/m$ with a visible denominator m **DO** look for m per-entry bounds that sum to c and add them **BECAUSE** averaging is the mechanism that produces a fraction with denominator m

- **WHEN** a per-line bound sums, over the line, to a quantity depending only on the line's count **DO** re-index it from the other end of the line so that it becomes pointwise comparable with the remaining bounds **BECAUSE** the multiset $\{1, \ldots, k\}$ is symmetric; this is the reversal device
- **WHEN** an inequality has an unexplained constant **DO** construct the extremal example first and compute the constant from it **BECAUSE** it identifies the equality case and tells you which bounds must be tight

Backtesting

1. First useful cue: the denominator 3 in $(n+2)/3$, and the anti-diagonal zero band that makes $i+j-n$ a coordinate.
2. Method the expression suggested: three per-entry bounds (band, row, column) averaged.
3. Where you would hesitate, and the thought to reinforce: the row bound is $j - Z_i$, not $n+1-j$ pointwise (for $n=2$, $a_{2,2}=2$ exceeds $n+1-j=1$); the switch to $n+1-j$ is legal only after summing over the whole row. Also the first nonzero entry of row n can sit in column 1 and needs the column condition to be at most 1.
4. The trap the question sets: proving $a_{i,j} \le i+j-n$, summing, and believing you are done; or launching the peeling induction with its level-count lemma.
5. What makes you recognise it faster next time: "sum at most constant times count" is an average bound; average bounds come from several majorants that are simultaneously tight in the extremal case.
6. The reusable decision rule: compute the extremal case, extract every bound that is tight there, re-index the line bounds from the far end, and add.

Mirror drills

1. (drill) For $a_{i,j} = \max(0, i+j-n)$ compute S and N and check $S = (n+2)N/3$. *Answer:* entries equal to k number $n+1-k$, so $N = n(n+1)/2$ and $S = \sum_{k=1}^{n} k(n+1-k) = n(n+1)(n+2)/6$.
2. (drill) Show that any matrix satisfying (a), (b), (c) has $a_{i,j} \le d_{i,j}$, the taxicab distance from (i,j) to the set of zero positions. *Answer:* along a shortest lattice path to a zero the entry changes by at most 1 per step, and it ends at 0.
3. (drill) Prove that in each row the nonzero entries form a block ending at column n whose first entry is 1. *Answer:* monotonicity along the row makes the zeros a prefix; the first nonzero entry is at most its left neighbour plus one, which is 1, except in row n, column 1, where the entry above is $a_{n-1,1} = 0$ by (a).

1.11 B5. Descents in the sequence of modular inverses

Let p be a prime number greater than 3. For each $k \in \{1, \dots, p-1\}$, let $I(k) \in \{1, 2, \dots, p-1\}$ be such that $k \cdot I(k) \equiv 1 \pmod{p}$. Prove that the number of integers $k \in \{1, \dots, p-2\}$ such that $I(k+1) < I(k)$ is greater than $p/4 - 1$.

Read it out loud

The sequence $I(1), I(2), \dots, I(p-1)$ is the permutation "inverse modulo p" of $\{1, \dots, p-1\}$, and you are counting its descents: positions k where the next value is smaller. The claim is that there are more than $p/4 - 1$ of them, roughly a quarter of the $p-2$ positions. A random permutation has about half its positions as descents, so the target is loose, and that looseness is itself information: the setters want a robust argument, not a delicate count. The proof below gives almost exactly half, $\frac{p-1}{2}$ minus at most three.

The fields wired in are arithmetic in the field $\mathbb{F}_p$ (inverses, the involution $k \mapsto I(k)$, quadratic equations having at most two roots), the lifting of a congruence to an integer equation by a size bound, and counting by an involution. The first useful cue is the comparison of I at consecutive arguments, $I(k+1)$ against $I(k)$. In $\mathbb{F}_p$, with $x = 1/k$,

$$\frac{1}{k+1} = \frac{1}{1/x+1} = \frac{x}{x+1} = 1 - \frac{1}{x+1}.$$

That is the line to write within thirty seconds: the inverse of the next integer is one minus the inverse of the shifted current inverse. Read as integers in $\{1, \dots, p-1\}$, it says $I(k+1) + I(I(k)+1) = p+1$, and the descent condition becomes a statement about the single quantity $m + I(m)$ for $m = I(k)+1$. From there an involution does the counting. What is tested is whether you conjugate the shift $k \mapsto k+1$ by the inversion, rather than subtracting inverses.

Approaches, ranked worst to best

1. **Small primes and heuristics.** *Cue:* a concrete permutation for each p. *Circuit:* tabulate I for $p = 5, 7, 11, 13$, count descents, observe "about half". (Not a proof, but it fixes the target: $f(5) = f(7) = 1$, $f(11) = 5$, $f(13) = 3$, which already tells you the count is near $\frac{p-1}{2}$ with small corrections depending on p modulo 3 and 4.)
2. **Difference identity.** *Cue:* compare $1/k$ and $1/(k+1)$. *Circuit:* write $\frac{1}{k} - \frac{1}{k+1} = \frac{1}{k(k+1)}$, so $I(k) - I(k+1) \equiv I(k^2+k)$, and a descent means this difference is the least positive residue. (Looks natural and stalls: $k \mapsto k^2 + k$ is not a bijection and the residue of $I(k^2+k)$ carries no usable order information.)
3. **Reflection symmetry alone.** *Cue:* $I(p-k) = p - I(k)$. *Circuit:* note that descents are symmetric under $k \mapsto p-1-k$. (True, and it only gives parity information about the count; no lower bound emerges.)
4. **Conjugate the shift, then pair.** *Cue:* $I(k+1)$ expressed through $I(k)$. *Circuit:* prove $I(k+1) = p+1-I(I(k)+1)$, reparametrise by $m = I(k)+1$ so that descents are $\{m : m + I(m) \geq p+2\}$, pair m with $p-m$ to halve the count, and bound the unpaired set by "a quadratic has at most two roots in $\mathbb{F}_p$". (Ideal. Every step is a bijection or a degree count, and the final inequality is arithmetic. The exact count via quadratic residues is available as a remark, not a requirement.)

The ideal solution

Write $f(p)$ for the number of $k \in \{1, \dots, p-2\}$ with $I(k+1) < I(k)$. All congruences are modulo p.

Preliminaries. I is a bijection of $\{1,\dots,p-1\}$ with $I(I(k))=k$, $I(1)=1$ and $I(p-1)=p-1$ (as $(p-1)^2\equiv 1$). Since $I(k)=p-1$ only for $k=p-1$, I maps $\{1,\dots,p-2\}$ bijectively onto itself. For $m\in\{1,\dots,p-1\}$ we have $(p-m)(p-I(m))\equiv m\,I(m)\equiv 1$ and $p-I(m)\in\{1,\dots,p-1\}$, so

$$I(p-m)=p-I(m).$$

Step 1: the identity. Let $k\in\{1,\dots,p-2\}$ and $x=I(k)$, so $x\in\{1,\dots,p-2\}$ and $x+1\in\{2,\dots,p-1\}$ is invertible. From $kx\equiv 1$ we get $(k+1)x\equiv x+1$; multiplying by $I(k+1)$ gives $x\equiv(x+1)I(k+1)$, and multiplying by $I(x+1)$ gives $I(k+1)\equiv x\,I(x+1)=(x+1-1)I(x+1)\equiv 1-I(x+1)$. Hence

$$I(k+1)+I\big(I(k)+1\big)\equiv 1 \pmod p.$$

Both summands lie in $\{1,\dots,p-1\}$, so the sum is an integer in $\{2,\dots,2p-2\}$; the only such integer congruent to 1 is $p+1$. Therefore

$$I(k+1)=p+1-I\big(I(k)+1\big)\qquad(1\le k\le p-2).$$

Step 2: reparametrise. For $m\in\{1,\dots,p-1\}$ put $T(m)=m+I(m)$. Fix $k\in\{1,\dots,p-2\}$ and set $m=I(k)+1\in\{2,\dots,p-1\}$. Since I is injective, $I(k+1)\neq I(k)$, so

$$I(k+1)<I(k)\iff I(k+1)\le I(k)\iff p+1-I(m)\le m-1\iff T(m)\ge p+2.$$

As k runs over $\{1,\dots,p-2\}$, $I(k)$ runs over $\{1,\dots,p-2\}$ bijectively, so m runs over $\{2,\dots,p-1\}$ bijectively. Since $T(1)=2<p+2$, we conclude

$$f(p)=\#\{m\in\{1,\dots,p-1\}:T(m)\ge p+2\}.$$

Step 3: the involution. By the preliminaries, $T(p-m)=(p-m)+(p-I(m))=2p-T(m)$. Hence $T(m)\ge p+2$ if and only if $T(p-m)\le p-2$, and $m\mapsto p-m$ is a bijection of $\{1,\dots,p-1\}$ exchanging the sets $U=\{m:T(m)\ge p+2\}$ and $L=\{m:T(m)\le p-2\}$. So $|U|=|L|$, and with $E=\{m:T(m)\in\{p-1,p,p+1\}\}$,

$$p-1=|U|+|L|+|E|=2f(p)+|E|,\qquad\text{that is,}\qquad f(p)=\frac{p-1-|E|}{2}.$$

Step 4: bound the exceptions. Let $\varepsilon\in\{-1,0,1\}$. Since $T(m)\in\{2,\dots,2p-2\}$, the only value of $T(m)$ congruent to ε modulo p is $p+\varepsilon$. Multiplying $m+I(m)\equiv\varepsilon$ by m,

$$T(m)=p+\varepsilon\iff m^2-\varepsilon m+1\equiv 0 \pmod p.$$

A polynomial of degree 2 over the field $\mathbb{F}_p$ has at most 2 roots, so each ε contributes at most 2 elements to E, and $|E|\le 6$.

Step 5: conclude. $f(p)\ge\frac{p-7}{2}$, and $\frac{p-7}{2}>\frac{p}{4}-1$ is equivalent to $p>10$, so the claim holds for all primes $p\ge 11$. For $p=5$ the sequence I is $1,3,2,4$ with the single descent at $k=2$, so $f(5)=1>\frac14$. For $p=7$ the sequence is $1,4,5,2,3,6$ with the single descent at $k=3$, so $f(7)=1>\frac34$. □

Remark 1.11.1. The count is exact once E is: the roots of m^2+m+1 and m^2-m+1 exist (two each) iff -3 is a square modulo p, that is iff $p\equiv 1 \pmod 3$; the roots of m^2+1 exist (two) iff $p\equiv 1 \pmod 4$. Hence $f(p)=\frac{p-1}{2}-2[p\equiv 1\ (3)]-[p\equiv 1\ (4)]$, matching $f(11)=5$ and $f(13)=3$.

Layer 1: WHAT

Block 1: record the structure of I: an involution of $\{1,\dots,p-1\}$ fixing 1 and $p-1$, hence a bijection of $\{1,\dots,p-2\}$, with the reflection law $I(p-m)=p-I(m)$. Block 2: express $I(k+1)$ through $I(k)$: the identity $I(k+1)=p+1-I(I(k)+1)$, obtained by conjugating the shift by the inversion and lifting the congruence by a size bound. Block 3: reparametrise by $m=I(k)+1$ so that a descent

is $T(m) \geq p+2$, the count being over $m \in \{1, \ldots, p-1\}$. Block 4: the involution $m \mapsto p-m$ sends T to $2p-T$, so it pairs large with small and leaves only the window $\{p-1, p, p+1\}$ unpaired. Block 5: the window is the root set of three quadratics, at most six elements; finish with arithmetic for $p \geq 11$ and two hand checks.

Juxtapose. The nominal route tabulates, sees "about half", and tries to prove half by the symmetry $k \mapsto p-1-k$, which only shows the descents come in mirror pairs. The bifurcation point is the question "how do I relate $I(k+1)$ to $I(k)$?" The zombie subtracts, $\frac{1}{k} - \frac{1}{k+1} = \frac{1}{k(k+1)}$, and gets a product that reparametrises nothing. You substitute, $\frac{1}{k+1} = \frac{x}{1+x} = 1 - \frac{1}{1+x}$ with $x = \frac{1}{k}$, and get inverses of a shift, which reparametrise by a bijection.

Layer 2: HOW

Block 1. **Cue.** "$k \cdot I(k) \equiv 1 \pmod{p}$" with p prime: I is inversion in the field $\mathbb{F}_p$. **Why natural.** Everything about I follows from inversion being an involutive automorphism of the abelian group $\mathbb{F}_p^*$: $I \circ I = \mathrm{id}$, the fixed points ± 1 (the solutions of $m^2 = 1$), and compatibility with negation, $I(-m) = -I(m)$, which as integers reads $I(p-m) = p - I(m)$. You write these before doing anything because every later bijection is one of them. The restriction to $\{1, \ldots, p-2\}$ is bijective only because $p-1$ is a fixed point; that detail is what makes the reparametrisation in Block 3 legal. **Link to the next block.** The condition compares I at k and $k+1$; you need $I(k+1)$ in terms of $I(k)$.

Block 2. **Cue.** "$I(k+1)$" against "$I(k)$": the shift $k \mapsto k+1$ seen through inversion. **Why natural.** Conjugating a shift by inversion is the Möbius map $x \mapsto x/(1+x)$, and every Möbius map is a composition of shifts, scalings and one inversion: $x/(1+x) = 1 - 1/(1+x)$. So $I(k+1)$ is $1 - I(I(k)+1)$ in $\mathbb{F}_p$, an inverse of a shift of $I(k)$. The lift to integers is forced: two numbers in $\{1, \ldots, p-1\}$ sum to something in $\{2, \ldots, 2p-2\}$, and only $p+1$ there is $\equiv 1$. The precondition $I(k)+1 \neq p$, that is $k \neq p-1$, is exactly the range of k in the problem, which is a sign you are on the intended road. **Link to the next block.** Descent is now an inequality between $p+1-I(m)$ and $m-1$, a condition on m alone.

Block 3. **Cue.** A condition on $m = I(k)+1$ with k ranging over a set that I preserves. **Why natural.** Counting over k or over m is the same count because $k \mapsto I(k)+1$ is a bijection from $\{1, \ldots, p-2\}$ to $\{2, \ldots, p-1\}$. The inequality $p+1-I(m) \leq m-1$ is $T(m) \geq p+2$ with $T(m) = m + I(m)$; replacing "$<$" by "$\leq$" first is harmless since I is injective, and it keeps the boundary clean. Extending the range to include $m = 1$ costs nothing since $T(1) = 2$. **Link to the next block.** T is the sum of a number and its inverse; it has a symmetry you already recorded.

Block 4. **Cue.** $T(m) = m + I(m)$ and the reflection law $I(p-m) = p - I(m)$. **Why natural.** $T(p-m) = 2p - T(m)$ is antisymmetry about the centre p. Any antisymmetric statistic under an involution pairs the values above the centre with the values below it, so the count above equals the count below, and the total $p-1$ splits as twice the count plus the width of the unpaired window. The window is $\{p-1, p, p+1\}$ because "$\geq p+2$" and its mirror "$\leq p-2$" leave exactly three values. **Link to the next block.** You need only an upper bound on the window's population.

Block 5. **Cue.** "$T(m) = p + \varepsilon$", a congruence in one variable after multiplying by m. **Why natural.** $m^2 - \varepsilon m + 1 \equiv 0$ is a quadratic over a field; the arena answer is "at most two roots", not "exactly two iff $p \equiv 1$ modulo something". Be lazy: $|E| \leq 6$ gives $f(p) \geq (p-7)/2$, which beats $p/4 - 1$ for $p > 10$. The primes 5 and 7 fall outside and are checked by writing down four and six numbers. The exact count by quadratic residues is a remark you may add for $+\alpha$, never a step you owe. **Link to the next block.** None; the inequality is arithmetic.

Layer 3: WHY

The truth underneath is that the descent statistic of the inversion permutation is governed by the single function $T(m) = m + 1/m$ on $\mathbb{F}_p^*$. T is invariant under $m \mapsto 1/m$ and antisymmetric under $m \mapsto -m$; the second symmetry pairs the set where T is large with the set where it is small, and the fixed structure of the pairing is the set where $T \in \{-1, 0, 1\}$, which is the set of elements of multiplicative order 3, 6 or 4 (the roots of $m^2 + m + 1$, $m^2 - m + 1$, $m^2 + 1$). So $f(p) = \frac{p-1}{2} - \frac{1}{2} \cdot \#\{\text{elements of order } 3, 4, 6 \text{ in } \mathbb{F}_p^*\}$. The answer is "almost half" because $\mathbb{F}_p^*$ is cyclic and has at most two elements of each of those orders. The looseness of the target $p/4 - 1$ is the setters telling you that the degree bound suffices.

The discernment mechanism has three parts. First, when a condition compares a permutation at consecutive arguments, conjugate the shift by the permutation and express the next value through the current one; for inversion this is a Möbius map and always reduces to shifts and one inversion, never to products. Second, a congruence between quantities living in a window shorter than $2p$ lifts to an equality with at most two candidates; write the candidates and eliminate. Third, an antisymmetric statistic under an involution is counted by halving the complement of the window, and the window is bounded by degree. You train the first by rewriting $1/(x + a)$ in terms of $1/x$ until $x/(1 + x) = 1 - 1/(1 + x)$ is a reflex; the third by recognising "$T(\sigma m) = c - T(m)$" as the signal to pair.

The examiners nudged by choosing $p/4 - 1$: weak enough that an exact count is unnecessary, strong enough that parity from $k \mapsto p - 1 - k$ is useless. They also fixed $p > 3$ so that the quadratics have nonzero discriminant -3 and the argument needs no exceptional case; you should notice that $p = 3$ would make $m^2 \pm m + 1$ collapse to $(m \mp 1)^2$.

Revolving concepts

Inversion in $\mathbb{F}_p$. For prime p, every nonzero residue has a unique inverse; I is an involution of $\{1, \ldots, p-1\}$ with fixed points exactly ± 1, and $I(-m) = -I(m)$. Existence of inverses for a prime modulus is A-level (Fermat or the Euclidean algorithm); the field language and the involution bookkeeping are the part to install. See Book 1, field: Number theory, topic: the field $\mathbb{F}_p$ and its multiplicative group.

Möbius transformations. $x \mapsto (ax + b)/(cx + d)$ over $\mathbb{F}_p$; the group they form; decomposition into shifts, scalings and $x \mapsto 1/x$; here $x \mapsto x/(1 + x)$ is the conjugate of the shift by inversion. See Book 1, field: Number theory, topic: fractional linear maps over finite fields.

Lifting congruences. If $a \equiv b \pmod{p}$ and a lies in an interval of length less than p, then a is determined; in an interval of length less than $2p$ there are at most two candidates. Used here to turn $\equiv 1$ into $= p + 1$. See Book 1, field: Number theory, topic: from congruence to equality by size.

Roots of polynomials over a field. A nonzero polynomial of degree d over a field has at most d roots (Lagrange's theorem on congruences for $\mathbb{F}_p$). Sharper, for the exact count: -1 is a square modulo p iff $p \equiv 1 \pmod 4$; -3 is a square modulo $p > 3$ iff $p \equiv 1 \pmod 3$; equivalently $\mathbb{F}_p^*$, being cyclic of order $p - 1$, contains elements of order 4 iff $4 \mid p - 1$ and of order 3 or 6 iff $3 \mid p - 1$. Beyond A-level and priority to install. See Book 1, field: Number theory, topic: quadratic residues and roots of unity in $\mathbb{F}_p$.

Counting by involutions. If σ is an involution of a finite set X and $T\colon X \to \mathbb{R}$ satisfies $T(\sigma x) = c - T(x)$, then $\#\{T > c/2 + w\} = \#\{T < c/2 - w\}$ for every w, so $\#\{T > c/2 + w\} = \frac{1}{2}\big(|X| - \#\{|T - c/2| \le w\}\big)$. This is the recursion-free generalisation of Step 3, and it is the same device as pairing k with $p - k$ in sums over $\mathbb{F}_p$. See Book 1, field: Combinatorics, topic: involutions and sign-reversing pairings.

Descents of permutations. The descent set and the Eulerian numbers are the combinatorial neighbours; the mean number of descents of a random permutation of n letters is $(n - 1)/2$, which is why "about half" was the right expectation. See Book 1, field: Combinatorics, topic: descents and

Eulerian numbers.

Circuit Bank

- **WHEN** a condition compares a permutation's values at consecutive arguments, such as $I(k+1)$ against $I(k)$ **DO** express the value at $k+1$ through the value at k by conjugating the shift **BECAUSE** the comparison becomes a condition on a single variable, which is what counting needs
- **WHEN** you see $1/(k+1)$ and already have $x = 1/k$ **DO** write $x/(1+x) = 1 - 1/(1+x)$ **BECAUSE** a Möbius map is shifts plus one inversion, so it sends inverses to inverses of shifts, never to products
- **WHEN** a congruence modulo p relates quantities in a window shorter than $2p$ **DO** list the at most two integers in the window with the right residue and eliminate **BECAUSE** congruences carry no order; equalities do
- **WHEN** a statistic satisfies $T(\sigma m) = c - T(m)$ for an involution σ **DO** pair m with σm and count only the central window **BECAUSE** the count above the centre is half the complement of the window
- **WHEN** you must bound the number of solutions of a polynomial congruence modulo p **DO** cite "degree d, at most d roots" and stop **BECAUSE** the degree bound costs one line; the residue-class count costs ten and the problem did not ask for it

Backtesting

1. First useful cue: "$I(k+1) < I(k)$", the inverse map evaluated at consecutive arguments.
2. Method the expression suggested: conjugate the shift by inversion ($x/(1+x) = 1-1/(1+x)$), lift to $I(k+1) + I(I(k)+1) = p+1$, reparametrise by $m = I(k)+1$, pair by $m \mapsto p-m$, bound the window by degree.
3. Where you would hesitate, and the thought to reinforce: whether $I(k)+1$ is in range (yes, $k \neq p-1$); whether "$<$" and "$\leq$" differ (no, I is injective); whether $p = 5, 7$ need separate treatment (yes, $(p-7)/2 \leq 0$ there).
4. The trap the question sets: the difference identity $1/k - 1/(k+1) = 1/(k(k+1))$, which looks like the obvious move and leads nowhere; and the mirror symmetry $k \mapsto p-1-k$, which gives parity, not a bound.
5. What makes you recognise it faster next time: any comparison of $I(k)$ with $I(k+a)$ is the Möbius map $x \mapsto x/(1+ax)$; any statistic antisymmetric under $m \mapsto -m$ is counted by pairing.
6. The reusable decision rule: conjugate the shift, lift by size, reparametrise by a bijection, pair by the negation involution, degree-count the exceptions, hand-check the primes below the threshold.

Mirror drills

1. (drill) Prove $I(p-m) = p - I(m)$ for $1 \leq m \leq p-1$. *Answer:* $(p-m)(p-I(m)) \equiv m\,I(m) \equiv 1$ and $p - I(m) \in \{1, \ldots, p-1\}$.
2. (drill) How many $m \in \{1, \ldots, p-1\}$ satisfy $m + I(m) \equiv 0 \pmod{p}$? *Answer:* $m^2 \equiv -1$, so two if $p \equiv 1 \pmod 4$ and none if $p \equiv 3 \pmod 4$.

3. (drill) Show that the number of ascents, $k \in \{1, \ldots, p-2\}$ with $I(k+1) > I(k)$, also exceeds $p/4 - 1$. *Answer:* ascents number $p - 2 - f(p) = \frac{p-3+|E|}{2} \geq \frac{p-3}{2} > \frac{p}{4} - 1$ for $p > 2$.

1.12 B6. Largest exponent in a self-referential growth inequality

Let $\mathbb{N} = \{1, 2, 3, \dots\}$. *Find the largest real constant* r *such that there exists a function* $g\colon \mathbb{N} \to \mathbb{N}$ *such that*

$$g(n+1) - g(n) \geq \big(g(g(n))\big)^r$$

for all $n \in \mathbb{N}$.

Read it out loud

A functional inequality on the positive integers with a self-referential term $g(g(n))$ and a free exponent r; you must find the supremum of the r for which a solution exists and show it is attained. The shape is two-sided. The left side is an increment, which is small when g grows slowly. The right side is a power of the composite $g \circ g$, which is enormous when g grows fast, because composition amplifies growth far more than differencing shrinks it. So there is a sweet spot in the growth rate, and r measures how much the composite is allowed to weigh.

The sentence to write within thirty seconds is the power-law ansatz. If $g(n) \approx n^s$ then $g(g(n)) \approx n^{s^2}$ and $g(n+1) - g(n) \approx sn^{s-1}$, so the inequality asks for $s - 1 \geq rs^2$, that is $r \leq (s-1)/s^2$. The right side is maximised at $s = 2$ with value $\frac{1}{4}$. That fixes the target: the answer is $r = \frac{1}{4}$, the extremal function should be $g(n) = n^2$, and indeed $2n + 1 \geq n = (n^4)^{1/4}$. Everything after this is proof of impossibility for $r > \frac{1}{4}$, and the ansatz tells you the mechanism: the exponent map $s \mapsto rs^2 + 1$ has a fixed point iff $1 - 4r \geq 0$; when $r > \frac{1}{4}$ every polynomial rate is pushed to a higher one, so a solution would have to grow faster than every polynomial, and the same inequality read backwards forbids that.

The fields wired in are growth rates of sequences (bootstrapping, telescoping), real analysis of "for all sufficiently large n" statements and how they compose, the dynamics of a quadratic map on exponents, and the discipline of reading one inequality in both directions. What is tested is the passage from heuristic to proof: a clean bootstrap lemma with named thresholds, and the observation that a lower bound on increments in terms of $g \circ g$ is also an upper bound on $g \circ g$ in terms of g.

Approaches, ranked worst to best

1. **Polynomial-only analysis.** *Cue:* a power law fits the inequality. *Circuit:* show that $g(n) = cn^s$ works iff $r \leq (s-1)/s^2$ and conclude $r = \frac{1}{4}$. (Only half a proof. It produces the example and the target but says nothing about functions that are not polynomial, and the problem quantifies over all g.)

2. **Ansatz gallery.** *Cue:* does faster growth help?. *Circuit:* test $g(n) = a^n$: increments $\approx a^n$, composite $\approx a^{a^n}$, fails for every $r > 0$; test slower than quadratic: increments too small. (Confirms that the sweet spot is quadratic, and that growth beyond it is punished by the composite. Still a heuristic, but it tells you the two directions of the real proof.)

3. **Bootstrap without organisation.** *Cue:* the inequality feeds a bound on g back into itself. *Circuit:* from $g(n) \geq n$ get increments $\geq n^r$, so $g \gtrsim n^{r+1}$, then increments $\gtrsim n^{r(r+1)^2}$, and so on, tracking constants by hand. (The right idea drowning in constants. Every round produces a new c and a new threshold, and without a lemma that states the round once and for all, the argument is unwritable in the arena.)

4. **Bootstrap lemma, exponent dynamics, reverse reading.** *Cue:* $g \circ g$ on the right, a difference on the left. *Circuit:* prove a lemma "$g \geq n^s$ eventually implies $g \geq n^t$ eventually for every $t < rs^2 + 1$"; use $rs^2 + 1 - s \geq 1 - \frac{1}{4r} > 0$ to climb through all polynomial rates; then read the inequality as $g(g(n-1))^r < g(n)$, obtain $g(n^4) < g(n)^4$ for large n, iterate along n^{4^j}, and contradict the superpolynomial growth. (Ideal. The lemma is stated once with an "eventually"

quantifier so no constant is ever tracked twice; the exponent dynamics is one line of completing the square; the ending is an induction along a geometric subsequence.)

The ideal solution

The answer is $\boxed{r = \frac{1}{4}}$.

Existence for $r = \frac{1}{4}$. Take $g(n) = n^2$. Then $g(g(n)) = n^4$ and

$$g(n+1) - g(n) = 2n + 1 \geq n = (n^4)^{1/4} = g(g(n))^{1/4}.$$

(The same g works for every $r \leq \frac{1}{4}$, since then $g(g(n))^r = n^{4r} \leq n$.)

Impossibility for $r > \frac{1}{4}$. Suppose $r > \frac{1}{4}$ and $g\colon \mathbb{N} \to \mathbb{N}$ satisfies the inequality for all n. We derive a contradiction.

Step 0: monotonicity. Since $g(g(n)) \geq 1$, we have $g(n+1) - g(n) \geq 1$, so g is strictly increasing, and as $g(1) \geq 1$, induction gives $g(n) \geq n$ for all n. In particular $g(u) \leq g(v)$ whenever $u \leq v$.

Lemma 1.12.1 (bootstrap)**.** *Let $s \geq 0$ and suppose $g(n) \geq n^s$ for all $n \geq n_0$. Then for every real $t < rs^2 + 1$ there exists n_1 such that $g(n) \geq n^t$ for all $n \geq n_1$.*

Proof. Let $m \geq n_0$. Then $g(m) \geq m \geq n_0$, so the hypothesis applies at the integer $g(m)$: $g(g(m)) \geq g(m)^s \geq (m^s)^s = m^{s^2}$, using $s \geq 0$. Hence $g(m+1) - g(m) \geq m^{rs^2}$ for all $m \geq n_0$.

Now let $n \geq 2n_0$ and put $a = \lceil n/2 \rceil \geq n_0$. The integers m with $a \leq m \leq n-1$ number $n - a = \lfloor n/2 \rfloor \geq (n-1)/2 \geq n/4$ (as $n \geq 2$), and each satisfies $m \geq n/2$. Telescoping,

$$g(n) \geq g(n) - g(a) = \sum_{m=a}^{n-1} \big(g(m+1) - g(m)\big) \geq \frac{n}{4}\Big(\frac{n}{2}\Big)^{rs^2} = C\, n^{rs^2+1}, \qquad C = 2^{-(rs^2+2)}.$$

Since $t < rs^2 + 1$, we have $Cn^{rs^2+1} \geq n^t$ whenever $n^{rs^2+1-t} \geq 1/C$. So the threshold

$$n_1 = \max\Big(2n_0,\ \big\lceil C^{-1/(rs^2+1-t)} \big\rceil\Big)$$

works. □

Step 1: g beats every polynomial. Put $\delta = 1 - \frac{1}{4r} > 0$. For every real s,

$$rs^2 + 1 - s = r\Big(s - \frac{1}{2r}\Big)^2 + 1 - \frac{1}{4r} \geq \delta,$$

so $s + \frac{\delta}{2} < rs^2 + 1$, and the lemma with $t = s + \frac{\delta}{2}$ says: if $g(n) \geq n^s$ for all large n, then $g(n) \geq n^{s+\delta/2}$ for all large n. Since $g(n) \geq 1 = n^0$ for all n, induction on j gives, for every $j \geq 0$, a threshold beyond which $g(n) \geq n^{j\delta/2}$. Hence for every real $s \geq 0$ there is n_s with $g(n) \geq n^s$ for all $n \geq n_s$ (choose j with $j\delta/2 \geq s$ and use $n^{j\delta/2} \geq n^s$ for $n \geq 1$).

Step 2: the reverse reading. Apply Step 1 with $s = 5$ and let $n_* = \max(n_5 + 1, 5)$. For $n \geq n_*$,

$$g(n) \;>\; g(n) - g(n-1) \;\geq\; g\big(g(n-1)\big)^r \;\geq\; g\big((n-1)^5\big)^r \;\geq\; g\big(n^4\big)^r \;\geq\; g\big(n^4\big)^{1/4}.$$

Here the first inequality uses $g(n-1) \geq 1$; the second is the hypothesis at $n-1$; the third uses $g(n-1) \geq (n-1)^5$ (as $n - 1 \geq n_5$) and monotonicity of g; the fourth uses $(n-1)^5 \geq n^4$ for $n \geq 5$ (equivalently $(1 - \frac{1}{n})^5 \geq \frac{1}{n}$, true at $n = 5$ since $0.32768 \geq 0.2$, and the left side increases while the right decreases) together with monotonicity; the fifth uses $g(n^4) \geq 1$ and $r > \frac{1}{4}$. Therefore

$$g(n^4) < g(n)^4 \qquad \text{for all } n \geq n_*.$$

Step 3: polynomial growth along a subsequence, and the contradiction. Fix $n = n_* \geq 5$ and choose a positive integer M with $g(n) \leq n^M$ (for instance $M = g(n)$, since $n^{g(n)} \geq 2^{g(n)} \geq g(n)$). Let

$N_j = n^{4^j}$ for $j \geq 0$. We claim $g(N_j) \leq N_j^M$ for all j. For $j = 0$ this is the choice of M. If it holds for j, then $N_j \geq n = n_*$, so Step 2 gives

$$g(N_{j+1}) = g(N_j^4) < g(N_j)^4 \leq N_j^{4M} = N_{j+1}^M.$$

Now apply Step 1 with $s = M + 1$: $g(N) \geq N^{M+1}$ for all $N \geq n_{M+1}$. Since $N_j \to \infty$, pick j with $N_j \geq n_{M+1}$. Then

$$N_j^{M+1} \leq g(N_j) \leq N_j^M,$$

so $N_j \leq 1$, contradicting $N_j \geq 5$. Hence no such g exists for $r > \frac{1}{4}$, and the largest admissible constant is $r = \frac{1}{4}$. □

Layer 1: WHAT

Block 1: the power-law ansatz locates the threshold $r = \frac{1}{4}$ at $s = 2$, and $g(n) = n^2$ verifies it in one line. Block 2: for $r > \frac{1}{4}$ assume g exists; increments are positive so g is strictly increasing and $g(n) \geq n$. Block 3: the bootstrap lemma, stated with "eventually" quantifiers: a rate n^s lifts to every rate below n^{rs^2+1}, through $g(g(m)) \geq m^{s^2}$ and a telescoping sum over the top half of $[1, n]$. Block 4: the exponent map $s \mapsto rs^2 + 1$ exceeds s by at least $\delta = 1 - \frac{1}{4r}$, so iterating the lemma from $s = 0$ beats every polynomial. Block 5: the reverse reading: $g(n) > g(g(n-1))^r \geq g(n^4)^{1/4}$ for large n, i.e. $g(n^4) < g(n)^4$. Block 6: iterate along $N_j = n^{4^j}$ to get $g(N_j) \leq N_j^M$, and collide with Block 4 at exponent $M + 1$.

Juxtapose. The nominal route runs the ansatz, perhaps several, and either stops at "no polynomial works" or starts bootstrapping with explicit constants and drowns. The bifurcation point is double. First, at Block 3, the zombie tracks c and n_0 through every round; you state the round once as a lemma whose conclusion is an "eventually" statement, so that rounds compose without bookkeeping. Second, at Block 5, the zombie has only ever read the inequality as a lower bound; you read it as an upper bound on $g \circ g$, which is what any inequality with $g \circ g$ on the large side must also be.

Layer 2: HOW

Block 1. **Cue.** "$g(g(n))$" on the right, a first difference on the left, an exponent r to be maximised. **Why natural.** Composition squares a polynomial exponent and differencing lowers it by one; those are the only two operations present, so the balance equation is $s - 1 = rs^2$. Maximising $(s-1)/s^2$ is a one-line calculus exercise with answer $\frac{1}{4}$ at $s = 2$. The ansatz is not a proof, but it tells you the answer, the extremal function, and the mechanism (fixed points of $s \mapsto rs^2 + 1$), which is everything the proof will need. Verifying $g(n) = n^2$ is mandatory: "largest r such that there exists" requires existence at the extremal value. **Link to the next block.** The impossibility half needs to handle arbitrary g; start by extracting the one property every solution shares.

Block 2. **Cue.** Positive integer values and an increment bounded below by a positive quantity. **Why natural.** $g(n+1) - g(n) \geq 1$ is strict monotonicity, and strict monotonicity of an integer-valued function is a growth bound: $g(n) \geq n$. Small as it is, this is what later lets you plug $g(m)$ into hypotheses that hold "for $n \geq n_0$" and lets you compare g at $(n-1)^5$ and n^4. **Link to the next block.** You now have $g \geq n^1$; the inequality should convert a lower bound on g into a better one.

Block 3. **Cue.** A lower bound on g feeds the composite $g \circ g$, which feeds the increments, which telescope back into g. **Why natural.** If $g(m) \geq m^s$ for $m \geq n_0$ then $g(g(m)) \geq g(m)^s \geq m^{s^2}$ for the same m, because $g(m) \geq m \geq n_0$ keeps you inside the range where the bound holds. The increments are then at least m^{rs^2}, and summing increments from a starting point to n gives $g(n)$ up to the value at the start, which is nonnegative and can be discarded. Summing over $m \in [n/2, n)$

rather than over $[n_0, n)$ is laziness with a purpose: it yields $\frac{n}{4}(\frac{n}{2})^{rs^2}$ with a constant that is visibly harmless, and the constant is then absorbed by weakening the exponent from rs^2+1 to any t below it. Stating the lemma with "eventually" on both sides is what makes it iterable. **Link to the next block.** One round raises the exponent from s to almost rs^2+1; ask when that is an improvement.

Block 4. Cue. The exponent map $\varphi(s) = rs^2+1$ and the hypothesis $r > \frac{1}{4}$. **Why natural.** $\varphi(s) - s = rs^2 - s + 1$ is a quadratic whose minimum is $1 - \frac{1}{4r}$, positive exactly when $r > \frac{1}{4}$. So every round gains at least $\delta/2$ in the exponent, and the orbit of 0 under the rounds passes every real s. This is the only place where $r > \frac{1}{4}$ is used in an essential way, and it is used through the discriminant of $rs^2 - s + 1$: at $r = \frac{1}{4}$ the double root $s = 2$ is the exponent of the extremal solution. **Link to the next block.** A superpolynomial g is now established; you need something that superpolynomial functions cannot do.

Block 5. Cue. The composite $g(g(n-1))$ is bounded above by $g(n)^{1/r}$, because the same inequality says so. **Why natural.** Read backwards, $g(g(n-1)) \le (g(n) - g(n-1))^{1/r} < g(n)^{1/r}$: the value of g at the huge argument $g(n-1)$ is controlled by a fixed power of g at the small argument n. When $g(n-1) \ge (n-1)^5 \ge n^4$, monotonicity turns this into $g(n^4) < g(n)^{1/r} \le g(n)^4$. The exponent 5 is chosen only to dominate 4 with room for the shift from n to $n-1$; the exponent 4 is chosen because $1/r < 4$ and $x^{1/r} \le x^4$ for $x \ge 1$. **Link to the next block.** $g(n^4) < g(n)^4$ is a self-similar bound: iterate it.

Block 6. Cue. A bound of the form $g(N^4) < g(N)^4$ valid for all $N \ge n_*$. **Why natural.** Along $N_j = n^{4^j}$ the bound composes: if $g(N_j) \le N_j^M$ then $g(N_{j+1}) < N_j^{4M} = N_{j+1}^M$. So g is polynomial of degree M along a sequence tending to infinity, while Block 4 says $g(N) \ge N^{M+1}$ for all large N. A contradiction needs only one subsequence, and the geometric subsequence is the one the bound naturally produces. M exists because a single value $g(n)$ is finite. **Link to the next block.** None; the proof is closed.

Layer 3: WHY

The mathematics underneath is the dynamics of the exponent map $\varphi_r(s) = rs^2 + 1$. A polynomial rate n^s is self-consistent with the inequality iff $s \le \varphi_r(s) - 1 + 1$, that is iff $s - 1 \ge rs^2$; the set of such s is nonempty iff the discriminant $1 - 4r$ is nonnegative. At $r = \frac{1}{4}$ the quadratic $\frac{1}{4}s^2 - s + 1 = \frac{1}{4}(s-2)^2$ has the double root $s = 2$: the extremal function has degree exactly 2, and the bootstrap from $s = 0$ at $r = \frac{1}{4}$ produces the increasing orbit $0, 1, 1.25, 1.39, \cdots \to 2$, so every solution at $r = \frac{1}{4}$ grows at least like $n^{2-\epsilon}$ eventually, which n^2 attains. For $r > \frac{1}{4}$ there is no fixed point, the orbit escapes to infinity at least linearly, and the bootstrap proves superpolynomial growth. The second half of the proof is the statement that superpolynomial growth is incompatible with $g \circ g$ being bounded by a fixed power of g: if g grows faster than n^5, then $g \circ g$ evaluated at $n-1$ dominates g evaluated at n^4, so "$g(n^4)$ is at most a fixed power of $g(n)$" along a geometric subsequence, which is polynomial growth. The two halves are the same inequality read in opposite directions, and the answer $\frac{1}{4}$ is the parameter value where the exponent map undergoes its saddle-node bifurcation.

The discernment mechanism is a four-step protocol for self-referential growth inequalities. One: run the power ansatz to find the threshold and the candidate extremal function, and verify the candidate immediately. Two: prove lower bounds by a bootstrap lemma stated with "eventually" quantifiers, never with explicit constants that must be carried; sum increments over the top half of the range so that the only constant is a power of 2, then weaken the exponent slightly to absorb it. Three: read the inequality backwards to obtain an upper bound on the composite, and use monotonicity to move the argument to a convenient power of n. Four: iterate the resulting self-similar bound along a geometric subsequence and collide with the lower bound. You train step two

by rewriting every "$g(n) \geq cn^s$ for $n \geq n_0$" you meet as "$g(n) \geq n^{s'}$ for all large n, for every $s' < s$", and step three by asking, for every inequality, "what does this bound from above?"

The knowledge gaps here are honest and specific. Nothing in the proof is beyond A-level calculus in content, but three habits are beyond it: composing "for all sufficiently large n" statements without losing track of thresholds; the comparison $\sum_{m=\lceil n/2 \rceil}^{n-1} m^t \geq \frac{n}{4}(\frac{n}{2})^t$ as a standing estimate; and the idea that a fixed-point analysis of an exponent map is a proof strategy, not a heuristic. The official alternative ending (obtain $g(n) \geq 3g(n-1)$ from $s > 16$, hence $g(n) > 2^n$, then compare $g(2^n)$ with $g(n)$) is the same mechanism with an exponential subsequence in place of the geometric one; know it exists, use the n^4 version because it is shorter.

Revolving concepts

Exponent maps and their fixed points. For an inequality relating Δg, g and $g \circ g$ by powers, the power ansatz yields a map φ on exponents; polynomial solutions correspond to fixed points, and the orbit of the trivial exponent under φ tracks what bootstrapping can prove. Generalised recursion: $s_{j+1} = \varphi(s_j)$ with $\varphi(s) = rs^2 + 1$; for $r \leq \frac{1}{4}$ the orbit of 0 increases to the smaller fixed point $\frac{1-\sqrt{1-4r}}{2r}$ (equal to 2 at $r = \frac{1}{4}$), for $r > \frac{1}{4}$ it is unbounded with $s_{j+1} - s_j \geq 1 - \frac{1}{4r}$. The same analysis applies to any inequality of the form $\Delta g \geq (g \circ g)^r g^q$ with the map $s \mapsto rs^2 + qs + 1$. See Book 1, field: Sequences and recurrences, topic: growth exponents and fixed points of quadratic maps.

Bootstrapping and "eventually". A statement $P(s)$: "$g(n) \geq n^s$ for all sufficiently large n". Lemmas of the form $P(s) \Rightarrow P(t)$ for all $t < \varphi(s)$ compose freely because the threshold in the conclusion is allowed to depend on s and t. Explicit constants are removed by weakening exponents. See Book 1, field: Real analysis, topic: asymptotic bounds and bootstrapping.

Telescoping over the top half. $g(n) - g(a) = \sum_{m=a}^{n-1} \Delta g(m)$; if $\Delta g(m) \geq m^t$ then $g(n) \geq \lfloor n/2 \rfloor (n/2)^t \geq \frac{n}{4}(\frac{n}{2})^t$ for $n \geq 2$. This replaces integral comparison and never needs $t > -1$ or a lower limit at 1. See Book 1, field: Sequences and recurrences, topic: sums of powers and telescoping.

Monotone integer functions. $g\colon \mathbb{N} \to \mathbb{N}$ strictly increasing implies $g(n) \geq n$; monotonicity transports bounds through composition: $g(u) \geq u^s$ and g increasing give $g(g(u)) \geq g(\lfloor u^s \rfloor)$, with the floor avoided here by choosing integer exponents. See Book 1, field: Functional equations, topic: monotone solutions on the integers.

Reverse reading and self-similar bounds. From $\Delta g(n) \geq (g \circ g)(n)^r$ one gets $(g \circ g)(n) \leq g(n+1)^{1/r}$; combined with $g(n) \geq n^s$ this becomes $g(n^k) \leq g(n)^{1/r}$ for $k < s$, a bound of the form $g(n^k) \leq g(n)^K$ which iterates along n^{k^j} to polynomial growth on a subsequence. Generalised: any bound $g(\psi(n)) \leq g(n)^K$ with ψ increasing and $\psi(n) \to \infty$ forces g to be at most polynomial in n along the orbit of ψ when ψ is itself polynomial. See Book 1, field: Inequalities, topic: iterating functional bounds along orbits.

The extremal example. $g(n) = n^2$ gives equality in the ansatz and slack in the inequality ($2n + 1$ against n); one can check that $g(n) = \lceil cn^2 \rceil$ also works for any $c > 0$ at $r = \frac{1}{4}$, so the extremal degree, not the constant, is what the problem pins down. See Book 1, field: Functional equations, topic: extremal solutions of growth inequalities.

Circuit Bank

- **WHEN** a growth inequality involves $g(g(n))$ and a first difference **DO** run the power ansatz $g \approx n^s$, $g \circ g \approx n^{s^2}$, $\Delta g \approx n^{s-1}$, and locate the threshold as the value of the parameter at which the exponent map acquires a double fixed point **BECAUSE** composition squares exponents and differencing subtracts one; the double root is the extremal degree
- **WHEN** you hold "$g(n) \geq n^s$ for large n" and increments bounded below by a power of $g \circ g$ **DO** bootstrap: $g \circ g \geq n^{s^2}$, telescope the increments over the top half of $[1, n]$, and conclude $g \geq n^t$ for large n for every $t < rs^2 + 1$ **BECAUSE** telescoping converts increment bounds

into value bounds and the top half keeps the constant a power of 2, which a slightly smaller exponent absorbs

- **WHEN** the exponent map satisfies $\varphi(s) \geq s + \delta$ for all s **DO** conclude that g beats every polynomial **BECAUSE** the orbit of 0 under φ is unbounded and "eventually" statements compose
- **WHEN** an inequality has $g \circ g$ on its large side **DO** also read it as an upper bound, $g(g(n)) \leq g(n+1)^{1/r}$, and move the inner argument to a power of n by monotonicity **BECAUSE** every inequality bounds something from above; for a superpolynomial g this caps $g(n^4)$ by $g(n)^4$
- **WHEN** you need to contradict a bound valid for all large N **DO** produce the opposite bound along one subsequence $N_j \to \infty$ by iterating a self-similar inequality **BECAUSE** a single subsequence suffices for a contradiction, and $g(N^4) < g(N)^4$ iterates along n^{4^j}

Backtesting

1. First useful cue: $g(g(n))^r$ on the right against $g(n+1) - g(n)$ on the left; the power ansatz $s - 1 \geq rs^2$.
2. Method the expression suggested: answer $\frac{1}{4}$ with $g(n) = n^2$; for $r > \frac{1}{4}$, a bootstrap lemma climbing through all polynomial rates, then the reverse reading $g(n^4) < g(n)^4$ iterated along n^{4^j}.
3. Where you would hesitate, and the thought to reinforce: whether the bootstrap's constants are circular (the threshold n_1 depends only on s, t, n_0); whether the hypothesis may be applied at $g(m)$ (yes, $g(m) \geq m \geq n_0$); whether $(n-1)^5 \geq n^4$ (from $n = 5$); whether the induction $g(N_j) \leq N_j^M$ needs $N_j \geq n_*$ (it does, and $N_j \geq n = n_*$ supplies it).
4. The trap the question sets: proving only that no polynomial works, or spending the hour on ansätze; and, subtler, believing that faster growth could rescue a larger r when in fact the composite punishes it.
5. What makes you recognise it faster next time: $g \circ g$ means "exponent squares"; a difference means "exponent drops by one"; the balance is a quadratic in the exponent whose discriminant is the answer.
6. The reusable decision rule: ansatz for the threshold and the example; bootstrap with "eventually" for the lower bound; reverse reading for the upper bound; collide along a geometric subsequence.

Mirror drills

1. (drill) Show that for every real r there is $g\colon \mathbb{N} \to \mathbb{N}$ with $g(n+1) - g(n) \geq g(n)^r$ for all n, and explain why the problem's version has a largest r but this one does not. *Answer:* define $g(1) = 1$ and $g(n+1) = g(n) + \lceil g(n)^r \rceil$; without the composite there is no upper bound on $g \circ g$ to cap growth, so faster growth is never punished.
2. (drill) Let $g\colon \mathbb{N} \to \mathbb{N}$ be strictly increasing with $g(n) \geq n^s$ for all $n \geq n_0$, where $s \geq 1$ is an integer. Show $g(g(n)) \geq n^{s^2}$ for all $n \geq n_0$. *Answer:* $g(n) \geq n \geq n_0$, so $g(g(n)) \geq g(n)^s \geq (n^s)^s$.
3. (drill) Show that $\varphi(s) = rs^2 + 1$ has a real fixed point iff $r \leq \frac{1}{4}$, and that for $r = \frac{1}{4}$ the orbit $s_0 = 0$, $s_{j+1} = \varphi(s_j)$ increases to 2. *Answer:* $rs^2 - s + 1 = 0$ has discriminant $1 - 4r$; for $r = \frac{1}{4}$, $\varphi(s) - s = \frac{1}{4}(s-2)^2 \geq 0$ and φ maps $[0, 2)$ into $[1, 2)$, so the orbit is increasing and bounded by 2, hence converges to the unique fixed point 2.

Chapter 2

Putnam 2024

2.1 A1. Valuation bookkeeping kills a Fermat-type equation

Determine all positive integers n for which there exist positive integers a, b, and c satisfying

$$2a^n + 3b^n = 4c^n.$$

Read it out loud

This is a homogeneous Diophantine equation of degree n in three unknowns, with three terms and the coefficients 2, 3, 4. The quantifier structure is "determine all n such that a solution exists", which on the Putnam almost always means: the answer set is tiny, one construction settles the members, and one structural obstruction settles the non-members. The fields wired in are elementary number theory (divisibility, p-adic valuations, residues) and nothing else; the equation is a Fermat lookalike, but the coefficients, not Fermat, are the whole story.

Read the coefficients as prime powers: $2 = 2^1$, $3 = 3^1$, $4 = 2^2$. Every unknown enters as an nth power, so the 2-adic valuation of each term is (the valuation of its coefficient) plus (a multiple of n). Thus, in other words, the three terms have 2-adic valuations congruent to $1, 0, 2$ modulo n. That sentence is the first useful cue, and it should be on the page within thirty seconds, next to the observation that $n = 1$ is a linear equation that obviously has solutions.

What is really being tested is whether you treat "b is even, so substitute $b = 2b'$" as a step to grind through or as a symptom of a valuation statement. The examiners planted 2 and 4 so that the exponents 1 and 2 differ, and they left $n = 2$ as the single case where the 2-adic residues tie and you must resolve the tie with a residue table (modulo 3 is the quickest; squares modulo 8 also work).

Approaches, ranked worst to best

1. **Search and guess.** *Cue:* "determine all n". *Circuit:* try small n, hope for a pattern. (No proof of impossibility is ever produced; you can find $(1, 2, 2)$ for $n = 1$ and nothing else, which is evidence, not mathematics.)

2. **Parity descent.** *Cue:* even coefficients 2 and 4. *Circuit:* normalise $\gcd(a, b, c) = 1$, deduce b even, substitute, repeat. (Correct for $n \geq 3$, and it is what the official solution does, but it costs three substitutions, and it stalls for $n = 2$ without warning: after b and a are even you are left with $2a'^2 + 3b'^2 = c^2$, on which parity says nothing more. You must notice the stall and reroute, either to the prime 3 or to the finer table of squares modulo 8.)

3. **Residues modulo** 3. *Cue:* the coefficient 3; nth powers are squares when n is even. *Circuit:* quadratic residues modulo 3 are 0 and 1. (Kills every even n in three lines, but says nothing for odd n, since cubes cover all residues modulo 3. Needed only for $n = 2$.)

4. **Valuation comparison.** *Cue:* homogeneous sum of three terms with prime-power coefficients. *Circuit:* in $x + y = z$ the minimum p-adic valuation is attained at least twice. (Best. No normalisation, no substitution: the 2-adic valuations are $1, 0, 2$ modulo n, pairwise distinct for $n \geq 3$, contradiction in one line. For $n = 2$ the same lemma at the prime 3 forces the tie between a and c, and one residue check finishes.)

The ideal solution

The answer is $\boxed{n = 1}$.

For $n = 1$ the triple $(a, b, c) = (1, 2, 2)$ works: $2 + 6 = 8$.

Lemma 2.1.1 (tie at the minimum)**.** *Let p be a prime and let x, y, z be nonzero integers with $x + y = z$. Then the minimum of $v_p(x), v_p(y), v_p(z)$ is attained by at least two of the three.*

Proof. If $v_p(x) = v_p(y)$ then $v_p(z) \geq v_p(x)$ and the minimum is already attained twice. Otherwise, say $v_p(x) = u < w = v_p(y)$, and write $x = p^u x'$, $y = p^w y'$ with $p \nmid x'y'$. Then $z = p^u(x' + p^{w-u}y')$ and $p \nmid x' + p^{w-u}y'$, so $v_p(z) = u = v_p(x)$, which is the minimum. The case $v_p(y) < v_p(x)$ is symmetric. □

Case $n \geq 3$. Suppose $2a^n + 3b^n = 4c^n$ with a, b, c positive. Put $\alpha = v_2(a)$, $\beta = v_2(b)$, $\gamma = v_2(c)$. Then

$$v_2(2a^n) = 1 + n\alpha, \qquad v_2(3b^n) = n\beta, \qquad v_2(4c^n) = 2 + n\gamma.$$

Modulo n these are $1, 0, 2$, which are pairwise distinct because $n \geq 3$. Hence the three valuations are pairwise distinct integers, and the minimum is attained exactly once, contradicting the lemma applied to $x = 2a^n$, $y = 3b^n$, $z = 4c^n$. No solution exists.

Case $n = 2$. Suppose $2a^2 + 3b^2 = 4c^2$. Now use the prime 3: with $\alpha = v_3(a)$, $\beta = v_3(b)$, $\gamma = v_3(c)$,

$$v_3(2a^2) = 2\alpha, \qquad v_3(3b^2) = 1 + 2\beta, \qquad v_3(4c^2) = 2\gamma.$$

The middle valuation is odd and the outer two are even, so the tie at the minimum demanded by the lemma must be between the outer two: $\alpha = \gamma$ and $2\alpha < 1 + 2\beta$. Write $a = 3^\alpha a'$ and $c = 3^\alpha c'$ with $3 \nmid a'c'$, and divide the equation by 9^α:

$$2a'^2 + \frac{3b^2}{9^\alpha} = 4c'^2, \qquad v_3\left(\frac{3b^2}{9^\alpha}\right) = 1 + 2\beta - 2\alpha \geq 1,$$

so the middle term is an integer divisible by 3. Reducing modulo 3 gives $2a'^2 \equiv c'^2 \pmod 3$. A square not divisible by 3 is $\equiv 1 \pmod 3$, so this reads $2 \equiv 1 \pmod 3$, absurd. No solution exists.

Thus solutions exist exactly when $n = 1$.

Layer 1: WHAT

Block 1: settle $n = 1$ by exhibiting $(1, 2, 2)$. Block 2: state the lemma that in $x + y = z$ the minimum p-adic valuation is attained twice. Block 3: for $n \geq 3$, read off the 2-adic valuations $1 + n\alpha$, $n\beta$, $2 + n\gamma$, note they are $1, 0, 2$ modulo n, hence distinct, hence the lemma is violated. Block 4: for $n = 2$, change prime to 3, use parity of the valuations to locate the forced tie (a with c), strip the common power of 3, and reduce modulo 3 to hit $2 \equiv 1$.

Juxtaposition. The nominal route normalises $\gcd(a, b, c) = 1$, reads the equation modulo 2 to get b even, substitutes $b = 2b'$, divides by 2, gets a even, substitutes again, gets c even, contradiction; then discovers that for $n = 2$ the third step does not close and starts over modulo 3. The ideal route never substitutes anything: it states what each substitution was measuring (a valuation) and compares the three measurements at once. The bifurcation point is the sentence "b must be even". The zombie writes $b = 2b'$; you write $v_2(3b^n) = n\beta$ and ask what the other two valuations are.

Layer 2: HOW

Block 1. **Cue.** "Determine all n" with an existential inner quantifier. **Why natural.** The examiners never ask this when the answer is "all n". The answer set is small and its members are certified by a construction, so the cheapest first move is to test $n = 1$, where the equation is linear and any triple with $2a + 3b$ divisible by 4 works. **Link to the next block.** Having certified $n = 1$, the remaining job is a uniform impossibility proof for $n \geq 2$, and impossibility of a homogeneous equation is a local (one prime at a time) statement.

Block 2. **Cue.** Three terms, one equation, small prime-power coefficients. **Why natural.** Any sum of two terms equal to a third is governed by the ultrametric inequality: the smallest p-adic valuation cannot stand alone, because the other two terms would then sum to something of strictly larger valuation. This is the definition-level content of "divide by the gcd" in the official solution: choosing a primitive solution is choosing the minimum valuation and asking who attains it. **Link to the next block.** The lemma converts the equation into a statement about three integers, and homogeneity makes those integers trivial to compute modulo n.

Block 3. **Cue.** Coefficients 2 and 4 are powers of the same prime with different exponents, and 3 is coprime to it. **Why natural.** Since $v_2(k\,t^n) = v_2(k) + n\,v_2(t)$, the unknown contributes only a multiple of n; the residues of the coefficient valuations modulo n are $1, 0, 2$. Three distinct residues means three distinct valuations, so no tie is possible. This is where you see why the problem was set with $2, 3, 4$ and not $2, 3, 8$ (exponent 3 would tie with exponent 0 when $n = 3$). **Link to the next block.** The same residues collapse when $n = 2$ ($1, 0, 0$: a tie between the b-term and the c-term is permitted), which tells you exactly why $n = 2$ needs a residue table. The quickest table is at the only other prime visible, 3; the tie at 2 can also be resolved with squares modulo 8, at the cost of one more substitution.

Block 4. **Cue.** At $n = 2$ the 2-adic argument produces no contradiction; the coefficient 3 is the other prime in the problem. **Why natural.** At the prime 3 the valuations $2\alpha, 1 + 2\beta, 2\gamma$ have parities even, odd, even, so the lemma forces $\alpha = \gamma$; after dividing by 9^α the leading parts $2a'^2$ and $4c'^2$ must agree modulo 3, and a residue table (squares modulo 3 are $0, 1$) closes it. The residue step is natural because a tie in valuation is precisely the statement that the leading coefficients cancel modulo p. **Link to the next block.** Nothing remains: $n = 1$ certified, $n \geq 3$ and $n = 2$ refuted.

Layer 3: WHY

The truth underneath is that this equation fails locally: there is a prime p at which it has no nontrivial solution even in the p-adic integers ($p = 2$ for $n \geq 3$, and $p = 3$ or $p = 2$ for $n = 2$), and the valuation lemma is the ultrametric inequality $v_p(x + y) \geq \min(v_p(x), v_p(y))$ with equality when the valuations differ. Homogeneity of degree n reduces the whole obstruction to arithmetic modulo n on the coefficient valuations. When the residues are pairwise distinct (here $1, 0, 2 \bmod n$ for $n \geq 3$), the obstruction is immediate; when two residues coincide, the leading parts of the tied terms must cancel modulo p, and that is a residue computation (quadratic residues for $n = 2$, cubic residues for $n = 3$, and so on).

The discernment mechanism: before touching the equation, list the primes dividing the coefficients and write the coefficient valuations modulo n for each. Distinct residues at some prime finish that n; a tie at every prime is resolved with residues, remembering that ties can go either way (the third mirror drill has a tie with a solution). Train it by refusing to write $b = 2b'$ ever again; write v_2 instead. The examiners nudged you with the visible 2 and 4; the nudge toward 3 at $n = 2$ is simply that 3 is the only other prime on the page.

Revolving concepts

p-adic valuation. $v_p(m)$ is the exponent of p in m; $v_p(xy) = v_p(x) + v_p(y)$ and $v_p(x + y) \geq \min(v_p(x), v_p(y))$ with equality if the two valuations differ. Consequence: in any vanishing sum of nonzero integers, the minimum valuation is attained at least twice. See Book 1, field: Number theory, topic: valuations and the ultrametric inequality.

Infinite descent as a recursion. The official route "b even, then a even, then c even, contradicting gcd $= 1$" is the recursion $(a, b, c) \mapsto (a/2, b/2, c/2)$ on solutions, and the general solution method for such a recursion is: a strictly decreasing sequence of positive integers terminates, so no solution

exists. The same recursion appears in every proof that $\sqrt{2}$ is irrational, in $x^2 + y^2 = 3z^2$, and in the $n = 4$ case of Fermat. Its closed form is exactly the valuation lemma: instead of running the descent, compare valuations once. See Book 1, field: Number theory, topic: descent.

Residue tables. Squares modulo 3, 4, 8 are $\{0, 1\}$, $\{0, 1\}$, $\{0, 1, 4\}$; cubes modulo 7 and 9 are $\{0, \pm 1\}$; -1 is a square modulo an odd prime p iff $p \equiv 1 \pmod 4$. These tables resolve valuation ties. See Book 1, field: Number theory, topic: residues of powers.

Local obstruction. An equation with no primitive solution modulo p^k for some p, k has no primitive integer solution; here that happens modulo 8 for $n \geq 3$ and modulo 9 for $n = 2$. See Book 1, field: Number theory, topic: local obstructions.

Circuit Bank

- **WHEN** you see a homogeneous equation $\sum k_i t_i^n = 0$ with prime-power coefficients **DO** write $v_p(k_i) \bmod n$ for each prime p dividing some k_i **BECAUSE** the unknowns contribute multiples of n, and a vanishing sum needs a tie at the minimum valuation, so distinct residues give an instant contradiction
- **WHEN** the valuation residues tie at a prime p **DO** strip the common power of p from the tied terms and reduce modulo p using the residue table of nth powers **BECAUSE** a tie means the leading parts cancel modulo p, which is a finite check that either fails (no solution) or succeeds (look for a solution)
- **WHEN** a model solution says "assume $\gcd(a, b, c) = 1$" and then descends **DO** read it as a valuation comparison and write the comparison instead **BECAUSE** descent is the recursion whose closed form is the ultrametric lemma; the closed form costs one line, the recursion costs three substitutions
- **WHEN** "determine all n" with an existence quantifier **DO** certify the small cases by construction first **BECAUSE** the impossibility argument must not waste effort on cases that hold, and the construction usually reveals which prime the obstruction lives at

Backtesting

1. First useful cue: the coefficients $2, 3, 4$ are $2^1, 3, 2^2$, and the unknowns are nth powers.
2. Method the expression suggested: p-adic valuations of the three terms, compared modulo n.
3. Where you would hesitate, and the thought to reinforce: at $n = 2$, when the 2-adic argument ties. The thought: a tie is information, it says which two terms lead; resolve it with a residue table, at the other prime on the page (3) or with squares modulo 8.
4. The trap the question sets: the visible Fermat shape invites you to think of $n \geq 3$ as the hard case, when it is the easy case; the real work is the one small exponent where residues coincide.
5. What makes you recognise it faster next time: any coefficient pattern (p^i, q, p^j) with $i \neq j$ screams "valuation residues $i, 0, j$ modulo n".
6. The reusable decision rule: in a vanishing homogeneous sum, compare coefficient valuations modulo the degree; distinct residues finish, tied residues go to a residue table.

Mirror drills

1. Show that $5a^3 + 25b^3 = c^3$ has no solution in positive integers. *Answer:* 5-adic valuations are $1, 2, 0$ modulo 3, pairwise distinct, so the minimum is attained once.
2. Show that $x^2 + y^2 = 3z^2$ has no solution in positive integers. *Answer:* 3-adic valuations $2\alpha, 2\beta, 1 + 2\gamma$ force $\alpha = \beta$; after dividing by 9^α, $x'^2 + y'^2 \equiv 1 + 1 \not\equiv 0 \pmod 3$.

3. Determine all n for which $a^n + 2b^n = 4c^n$ has a solution in positive integers. *Answer:* $n \in \{1, 2\}$. For $n \geq 3$ the 2-adic residues $0, 1, 2$ are distinct. For $n = 2$ the tie is between a and c and it is resolvable: $(a, b, c) = (2, 4, 3)$ gives $4 + 32 = 36$. Ties can go either way; the residue check decides.

2.2 A2. Expand the iterate modulo the fixed-point polynomial squared

For which real polynomials p is there a real polynomial q such that

$$p(p(x)) - x = (p(x) - x)^2 q(x)$$

for all real x?

Read it out loud

This is a divisibility statement in $\mathbb{R}[x]$: the second iterate of p minus the identity is divisible by the square of the first iterate minus the identity. The object the examiners want you to name is $r(x) = p(x) - x$, the fixed-point polynomial of p (its complex roots are exactly the fixed points of p). The fields wired in are polynomial algebra (divisibility, degrees, units in $\mathbb{R}[x]$), the calculus of polynomials (the derivative, Taylor expansion, which for polynomials is an exact finite identity), and a whiff of dynamics (iterates, fixed points, the multiplier $p'(\alpha)$ at a fixed point α).

The first useful cue is that the argument of the outer p is $p(x) = x + r(x)$. So $p(p(x)) = p(x + r)$: the outer p is being evaluated at a perturbation of x by r, and the modulus is r^2. Perturbation plus a square modulus means first-order Taylor expansion, and for polynomials that expansion is not an approximation but an identity: $p(x + r) = p(x) + rp'(x) + r^2(\cdots)$. Within thirty seconds you should have written $r = p - x$, $p(p(x)) = p(x + r) \equiv p(x) + rp'(x) \pmod{r^2}$, and the test cases $p = x + c$ and $p = -x + c$ (the involutions and translations, for which the left side is $2c$ or 0).

What is being tested is whether you expand a composition coefficientwise (zombie) or structurally (in powers of the perturbation). The degree check $\deg(p \circ p - x) = d^2$ versus $\deg r^2 = 2d$ gives no contradiction for $d \geq 2$, and that is the examiners telling you that degrees alone will not do it: you need the first-order term.

Approaches, ranked worst to best

1. **Coefficient bashing.** *Cue:* "for which polynomials". *Circuit:* write $p = a_d x^d + \cdots$, expand $p(p(x))$, compare. (Degree 1 is fine; degree 2 already produces a system in three unknowns; general d is hopeless. Consumption without analysis.)
2. **Degree comparison alone.** *Cue:* a divisibility $A \mid B$. *Circuit:* $\deg A \leq \deg B$ unless $B = 0$. (Stalls immediately since $d^2 \geq 2d$ for $d \geq 2$. Worth one line as a sanity check, and its failure is the cue that first-order information is needed.)
3. **Multipliers at fixed points.** *Cue:* $(p(x) - x)^2$: double roots at fixed points. *Circuit:* differentiate; at a fixed point α, $(p \circ p)'(\alpha) = p'(\alpha)^2$. (Sound, and it is the official second solution: r divides $(p' - 1)(p' + 1) = r'(r' + 2)$, and a multiplicity argument shows $r \mid r' + 2$. It costs a separate treatment of multiple roots and a complex-roots argument. Correct but more ink.)
4. **Taylor modulo r^2.** *Cue:* $p(p(x)) = p(x + r)$ with modulus r^2. *Circuit:* $g(x + r) \equiv g(x) + rg'(x) \pmod{r^2}$ for any polynomial g. (Best. The condition becomes $r^2 \mid r(1 + p')$, so $r = 0$ or $r \mid 1 + p'$, and a degree comparison finishes in three lines with no root analysis at all.)

The ideal solution

The answer is: $p(x) = x + c$ or $p(x) = -x + c$ for a real constant c (the case $c = 0$ of the first family is $p(x) = x$).

These work. If $p(x) = x + c$ then $p(p(x)) - x = 2c$ and $(p(x) - x)^2 = c^2$; take $q = 2/c$ if $c \neq 0$, and any q if $c = 0$. If $p(x) = -x + c$ then $p(p(x)) = -(-x + c) + c = x$, so the left side is 0 and $q = 0$ works.

Nothing else works. Let $r(x) = p(x) - x \in \mathbb{R}[x]$.

Lemma 2.2.1 (polynomial Taylor expansion). *For any $g \in \mathbb{R}[x]$ and any $r \in \mathbb{R}[x]$, the polynomial $g(x+r) - g(x) - r\,g'(x)$ is divisible by r^2 in $\mathbb{R}[x]$.*

Proof. Both sides are linear in g, so it suffices to take $g = x^k$. By the binomial theorem,

$$(x+r)^k - x^k - kx^{k-1}r = \sum_{j=2}^{k} \binom{k}{j} x^{k-j} r^j,$$

and every term on the right contains the factor r^2 (for $k \le 1$ the sum is empty and the left side is 0). □

Apply the lemma with $g = p$: there is $s \in \mathbb{R}[x]$ with $p(x+r) = p(x) + rp'(x) + r^2 s(x)$. Since $p(x) = x + r$,

$$p(p(x)) - x = p(x+r) - x = r + rp'(x) + r^2 s(x) = r\,(1 + p'(x)) + r^2 s(x).$$

Hence r^2 divides $p(p(x)) - x$ if and only if r^2 divides $r(1+p')$.

If $r = 0$ then $p(x) = x$, which is in the list. Assume $r \neq 0$; cancelling r (legitimate in the integral domain $\mathbb{R}[x]$), the condition is $r \mid 1 + p'$.

If $1 + p' = 0$ then $p' = -1$ and $p(x) = -x + c$ for some real c, which is in the list.

Otherwise $1 + p' \neq 0$ and divisibility gives $\deg r \le \deg(1 + p')$. Now compare degrees. If $\deg p \ge 2$ then $\deg r = \deg p$ while $\deg(1 + p') \le \deg p - 1$, contradiction. If $\deg p \le 0$ (so p is a constant, possibly 0) then $\deg r = 1$ while $1 + p' = 1$ has degree 0, contradiction. So $\deg p = 1$: $p(x) = mx + c$ with $m \neq 0$, whence $r = (m-1)x + c$ and $1 + p' = 1 + m$, a nonzero constant. A polynomial dividing a nonzero constant is a nonzero constant, so $m - 1 = 0$ and $p(x) = x + c$, which is in the list.

Therefore the polynomials are exactly $p(x) = x + c$ and $p(x) = -x + c$, $c \in \mathbb{R}$.

Layer 1: WHAT

Block 1: name $r = p - x$ and rewrite the composition as $p(x + r)$. Block 2: expand $p(x + r)$ to first order modulo r^2 (the binomial lemma). Block 3: the condition collapses to $r^2 \mid r(1 + p')$, hence to $r = 0$ or $r \mid 1 + p'$. Block 4: the branch $1 + p' = 0$ yields $-x + c$; the other branch is a degree comparison that forces $\deg p = 1$ and then $m = 1$. Block 5: verify both families.

Juxtaposition. The nominal route treats $p(p(x))$ as an unknown polynomial of degree d^2 and tries to extract information from its coefficients, or differentiates and chases roots of r with their multiplicities. The ideal route treats $p(p(x))$ as p evaluated at a point close to x, where "close" means "differs by something whose square is zero in the quotient ring". The bifurcation point is the first line after reading: writing $p(p(x))$ versus writing $p(x + r)$. Everything after that is forced.

Layer 2: HOW

Block 1. **Cue.** The same expression $p(x) - x$ appears squared on the right and, hidden, as the argument of the outer p on the left. **Why natural.** A polynomial identity involving $p \circ p$ and $p - x$ is a statement about p near its fixed points. Naming $r = p - x$ makes $p(x) = x + r$ a substitution you can perform inside the outer p; you do this because the modulus r^2 tells you that in the quotient ring $\mathbb{R}[x]/(r^2)$ the element $x + r$ is "x plus a nilpotent". **Link to the next block.** A polynomial evaluated at x plus a nilpotent of square zero is its first-order Taylor expansion, exactly.

Block 2. **Cue.** Perturbation $x \mapsto x + r$ with modulus r^2. **Why natural.** The identity $g(x + \varepsilon) = g(x) + \varepsilon g'(x)$ when $\varepsilon^2 = 0$ is the algebraic definition of the derivative (dual numbers). For polynomials it is the binomial theorem applied monomial by monomial, so no analysis is invoked and no "approximation" needs justifying. **Link to the next block.** Substituting back turns the

condition into a divisibility between r^2 and $r(1+p')$, which is a statement about r and p' only, both of degree at most $\deg p$.

Block 3. **Cue.** $r^2 \mid r \cdot (\text{something})$. **Why natural.** Cancelling a nonzero common factor in an integral domain is legal; you must first dispose of $r = 0$ separately (that is the polynomial $p = x$, which does work). The split into $r = 0$ and $r \neq 0$ is the first edge case a grader looks for. **Link to the next block.** $r \mid 1 + p'$ is a divisibility between two polynomials whose degrees you control, and controlled degrees mean a degree comparison.

Block 4. **Cue.** A divisibility $A \mid B$ with $\deg A$ known and $\deg B$ known. **Why natural.** If $B \neq 0$ then $\deg A \leq \deg B$; so the second edge case, $B = 1 + p' = 0$, must be split off, and it is precisely the family $-x + c$. Then $\deg r$ equals $\deg p$ when $\deg p \geq 2$, equals 1 when p is constant, and drops only when p is linear with slope 1; meanwhile $\deg(1 + p') \leq \deg p - 1$. The only survivor is the linear case, where "a divisor of a nonzero constant is a nonzero constant" finishes. **Link to the next block.** The candidates are now a list, and a list must be verified, which is Block 5.

Block 5. **Cue.** The problem asks "for which p", so the answer is a set, and set equality needs both inclusions. **Why natural.** Necessity has been shown; sufficiency is two one-line computations. Forgetting them costs a point; forgetting that $c = 0$ is allowed (and that q can be anything when both sides vanish) costs another. **Link to the next block.** None; the proof is closed.

Layer 3: WHY

The mathematical truth is that divisibility by the square of r is a tangency condition, and tangency is first-order information. In the ring $\mathbb{R}[x]/(r^2)$ the class of r is a nonzero nilpotent, $p(x) = x + r$ there, and evaluating any polynomial at $x + r$ produces exactly $g(x) + rg'(x)$. So the whole problem is the single congruence $p \circ p \equiv x + r(1 + p') \pmod{r^2}$, and the question "when is $p \circ p \equiv x$" becomes "when does r divide $1 + p'$". At a simple fixed point α this says $p'(\alpha) = -1$, the multiplier of an involution, and at a fixed point of multiplicity $m \geq 2$ it says $1 + p'$ vanishes to order m there, which forces r to be constant or $1 + p'$ to be identically zero. That is the structural reason the answer is "translations and reflections": these are the only polynomials that are either fixed-point free with p' never -1, or involutions.

The discernment mechanism is a two-step reflex: a composition $f(g(x))$ with $g = x + (\text{small})$ is expanded in the small quantity; a divisibility between polynomials of controlled degree is closed by degree comparison after the zero cases are removed. You train the first step by rewriting every iterate you meet as an evaluation at $x + r$, and the second by never writing "$A \mid B$" without asking whether B could be zero. The examiners nudged with the square on the right (first-order information suffices) and with the fact that the naive degree count fails (so you must find the first-order information).

Revolving concepts

Formal Taylor expansion. For $g \in \mathbb{R}[x]$ and any polynomial h, $g(x + h) = \sum_{k \geq 0} g^{(k)}(x) h^k / k!$, a finite exact identity; modulo h^2 only $g + hg'$ survives. Equivalent formulation: in $\mathbb{R}[x][\varepsilon]/(\varepsilon^2)$, $g(x + \varepsilon) = g(x) + \varepsilon g'(x)$, which can be taken as the definition of the formal derivative. See Book 1, field: Polynomials, topic: formal derivative and Taylor expansion.

Degrees and units. In $\mathbb{R}[x]$, $\deg(AB) = \deg A + \deg B$, $A \mid B \neq 0$ implies $\deg A \leq \deg B$, and the units are the nonzero constants. See Book 1, field: Polynomials, topic: divisibility in $k[x]$.

Iterates modulo r^2: the generalised recursion. Let $r = p - x$ and define $p^{\circ k}$ by $p^{\circ 1} = p$, $p^{\circ(k+1)} = p \circ p^{\circ k}$. Then

$$p^{\circ k}(x) \equiv x + r\left(1 + p' + p'^2 + \cdots + p'^{k-1}\right) \pmod{r^2}.$$

Proof by induction: if $p^{\circ k} = x + ru_k + r^2(\cdots)$ then $p^{\circ(k+1)} = p(x + ru_k + r^2(\cdots)) \equiv p(x) + ru_k p'(x) = x + r(1 + u_k p')$, so $u_{k+1} = 1 + p'u_k$ with $u_1 = 1$, a first-order linear recursion whose closed form is the geometric sum. The general method: a recursion of the form $u_{k+1} = Au_k + B$ with polynomial A, B is solved by the geometric sum $u_k = B(1 + A + \cdots + A^{k-1}) + A^{k-1}u_1$. The problem is the case $k = 2$, where $u_2 = 1 + p'$. See Book 1, field: Sequences and recurrences, topic: first-order linear recurrences; and field: Functional equations, topic: iterates and fixed points.

Fixed points and multipliers. If $p(\alpha) = \alpha$ then $(p \circ p)'(\alpha) = p'(\alpha)^2$ by the chain rule, so $r^2 \mid p \circ p - x$ forces $p'(\alpha)^2 = 1$ at every simple fixed point; the official second solution builds on this. Also $p(x) - x$ always divides $p(p(x)) - x$, because $g(y) - g(x)$ is divisible by $y - x$ for any polynomial g. See Book 1, field: Polynomials, topic: $y - x$ divides $g(y) - g(x)$.

Circuit Bank

- **WHEN** you see a composition $f(g(x))$ where $g(x) = x + r(x)$, especially with r^2 or r^k as a modulus **DO** write $f(x + r)$ and expand in powers of r, keeping terms below the modulus **BECAUSE** polynomial Taylor expansion is an exact identity, and modulo r^k only the first k terms survive

- **WHEN** a divisibility $A \mid B$ between polynomials is your remaining condition **DO** split off $B = 0$, then compare degrees **BECAUSE** in $k[x]$ a nonzero multiple of A has degree at least $\deg A$, and a divisor of a nonzero constant is a constant

- **WHEN** a condition involves the square of $p(x) - x$ **DO** think "tangency at fixed points, first derivative" **BECAUSE** divisibility by a square is vanishing to second order, which is controlled by p' at the roots of $p - x$

- **WHEN** the question is "for which polynomials" **DO** test degrees 0 and 1 by hand and write down the candidates before proving anything **BECAUSE** the answer is a small family, and knowing it tells you which branches of the proof are the real ones

Backtesting

1. First useful cue: the outer argument $p(x)$ equals $x + (p(x) - x)$, and the modulus is $(p(x) - x)^2$.

2. Method the expression suggested: first-order Taylor expansion of p at x with increment r, modulo r^2.

3. Where you would hesitate, and the thought to reinforce: after reaching $r \mid 1 + p'$, the temptation is to look at roots. The thought: you control both degrees, so compare them, after removing $1 + p' = 0$.

4. The trap the question sets: the naive degree count $d^2 \geq 2d$ looks like it permits large d; and the two edge cases $r = 0$ and $1 + p' = 0$ are each a full family of answers, not degenerate leftovers.

5. What makes you recognise it faster next time: any iterate $p \circ p$ next to $p - x$ is an evaluation at $x + r$; say it out loud.

6. The reusable decision rule: composition at a perturbation of x goes to Taylor; divisibility with known degrees goes to degree comparison; every branch where a side vanishes is an answer, not an exception.

Mirror drills

1. Show that for every polynomial p, $p(x) - x$ divides $p(p(x)) - x$. *Answer:* $p(p(x)) - p(x)$ is $g(y) - g(x)$ with $y = p(x)$, divisible by $y - x = p(x) - x$; add $p(x) - x$.
2. Find all real polynomials p for which $(p(x) - x)^2$ divides $p(p(x)) - p(x)$. *Answer:* modulo r^2 the left side is rp', so $r = 0$ or $r \mid p'$; degree comparison leaves the constants (where $p' = 0$) and $p = x + c$.
3. Let $r = p(x) - x \neq 0$. Show that r^2 divides $p(p(p(x))) - x$ if and only if r divides $1 + p' + p'^2$. *Answer:* the iterate recursion $u_{k+1} = 1 + p'u_k$ gives $p^{\circ 3} \equiv x + r(1 + p' + p'^2) \pmod{r^2}$.

2.3 A3. Where the entry two sits in a random tableau

Let S be the set of bijections

$$T\colon \{1,2,3\} \times \{1,2,\dots,2024\} \to \{1,2,\dots,6072\}$$

such that $T(1,j) < T(2,j) < T(3,j)$ for all $j \in \{1,2,\dots,2024\}$ and $T(i,j) < T(i,j+1)$ for all $i \in \{1,2,3\}$ and $j \in \{1,2,\dots,2023\}$. Do there exist a and c in $\{1,2,3\}$ and b and d in $\{1,2,\dots,2024\}$ such that the fraction of elements T in S for which $T(a,b) < T(c,d)$ is at least $1/3$ and at most $2/3$?

Read it out loud

Strip the notation. T is a filling of a 3×2024 array with the numbers 1 to 6072, each once, increasing left to right along every row and top to bottom down every column. That is the definition of a standard Young tableau of rectangular shape with three rows and $n = 2024$ columns, and S is the set of all of them. The question asks whether two cells exist whose relative order, under a uniformly random tableau, is balanced: probability in $[1/3, 2/3]$. So this is an enumerative combinatorics question (counting tableaux) wearing probability vocabulary, and the fields wired in are Young tableaux, the hook length formula, and the elementary poset fact that a tableau is a linear extension of the cell order.

Two immediate observations. First, if one cell is weakly above and weakly left of the other, the order is forced and the fraction is 0 or 1; so the pair must be incomparable: one cell strictly lower and strictly further left than the other. Second, you choose the pair. An existence question with a free choice is an invitation to choose the pair for which the event is simplest, and the simplest events in a tableau concern the smallest entries, whose positions are nearly forced: 1 is at $(1,1)$, and 2 is at $(1,2)$ or $(2,1)$. Thus the event $T(2,1) < T(1,2)$ is exactly "the entry 2 sits at $(2,1)$". Within thirty seconds you should have written that sentence and the words "hook length formula".

What is being tested is (i) recognition of the object, (ii) the reflex to pick the pair that turns a comparison into the position of a single small entry, and (iii) the technique of deleting the smallest entries and rotating the leftover skew shape into a straight shape so the hook length formula applies. The numbers $1/3$ and $2/3$ are the examiners hinting that the natural fractions cluster near $1/3$ (three rows) and that the correct pair lands just above it.

Approaches, ranked worst to best

1. **Attack a general pair.** *Cue:* "the fraction of T with $T(a,b) < T(c,d)$". *Circuit:* count linear extensions with one extra relation. (For an arbitrary incomparable pair this is a hard count with no closed form; you would drown. Wrong quantifier reading: you get to choose the pair.)
2. **Symmetry hunting.** *Cue:* "fraction between $1/3$ and $2/3$", which smells like $1/2$. *Circuit:* find an involution on S that swaps the order of two cells. (The only symmetry of a $3 \times n$ rectangle is rotation by 180° with complementation of entries, and it sends the event for a pair to the event for the antipodal pair; for antipodal cells the event is mapped to itself, not to its complement. Thirty seconds, then abandon.)
3. **Ballot words and a three-row recursion.** *Cue:* rows increase, columns increase. *Circuit:* read the row index of $1, 2, \dots, 3n$ as a lattice word; count words by a recursion on the row lengths. (Correct, and the official remark proves the three-row count this way, but verifying the induction identity is a page of algebra. Costs too much ink in the arena.)
4. **Position of the entry 2 plus hook lengths.** *Cue:* 2 must be adjacent to the corner. *Circuit:* delete the smallest entries, rotate the skew shape to a straight shape, apply the hook length formula to both shapes and take the ratio. (Best. The ratio of two hook products collapses to $(n+1)/(3n-1)$ in four lines.)

The ideal solution

Yes. Take $(a,b) = (2,1)$ and $(c,d) = (1,2)$. Write $n = 2024$; the argument works for every $n \geq 2$.

Regard $T \in S$ as a $3 \times n$ array whose entry in row i, column j is $T(i,j)$; the conditions say that every row increases from left to right and every column increases from top to bottom. So S is the set of standard Young tableaux of shape (n,n,n).

Step 1: the event. For $T \in S$ and a cell (i,j), every cell weakly above and weakly to the left of (i,j) carries a smaller entry, and there are $ij-1$ such cells other than (i,j) itself. The entry 2 has exactly one smaller entry, so it lies in a cell with $ij-1 \leq 1$ other than $(1,1)$, which carries 1: hence 2 is at $(1,2)$ or at $(2,1)$. If $T(2,1) = 2$ then $T(1,2) \geq 3 > T(2,1)$; if $T(1,2) = 2$ then $T(2,1) \geq 3 > T(1,2)$. Therefore

$$T(2,1) < T(1,2) \iff T(2,1) = 2.$$

Step 2: the count as a smaller shape. Let $S_2 = \{T \in S : T(2,1) = 2\}$. For $T \in S_2$ delete the cells $(1,1)$ and $(2,1)$, which carry 1 and 2; what remains is the set of cells $D = \{(1,j),(2,j) : 2 \leq j \leq n\} \cup \{(3,j) : 1 \leq j \leq n\}$ filled with $3,\dots,3n$, increasing along rows and down columns. Apply the cell map $(i,j) \mapsto (4-i,\ n+1-j)$ and the entry map $k \mapsto 3n+1-k$. The cell map sends D onto the straight shape $\lambda = (n, n-1, n-1)$ (row 3 of D becomes row 1 with n cells, rows 2 and 1 of D become rows 2 and 3 with $n-1$ cells each, aligned at the left), the entries become $1,\dots,3n-2$, and since both the cell order and the entry order are reversed, "increasing along rows and down columns" is preserved in both directions. Conversely, a standard Young tableau of shape λ pulled back through these maps and completed by 1 at $(1,1)$ and 2 at $(2,1)$ is an element of S_2: the two new entries are smaller than everything else and $(1,1)$ lies above $(2,1)$. Hence $|S_2| = f^{(n,n-1,n-1)}$, where f^{μ} denotes the number of standard Young tableaux of shape μ; likewise $|S| = f^{(n,n,n)}$.

Step 3: hook lengths. You must know the following theorem (Frame, Robinson and Thrall): for a partition μ with N cells,

$$f^{\mu} = \frac{N!}{\prod_{c \in \mu} h(c)},$$

where the hook length $h(c)$ of a cell is 1 plus the number of cells to its right in its row plus the number of cells below it in its column.

For $\mu = (n,n,n)$: the cell $(1,j)$ has hook $n-j+3$, so row 1 contributes $(n+2)(n+1)\cdots 3 = (n+2)!/2$; row 2 contributes $(n+1)!$ and row 3 contributes $n!$. Thus

$$f^{(n,n,n)} = \frac{2\,(3n)!}{(n+2)!\,(n+1)!\,n!}.$$

For $\lambda = (n, n-1, n-1)$: in row 1 the cells $(1,j)$ with $j \leq n-1$ have two cells below and $n-j$ to the right, hook $n-j+3$, giving $(n+2)(n+1)\cdots 4 = (n+2)!/6$, and the cell $(1,n)$ has hook 1; row 2 has hooks $n, n-1, \dots, 2$, product $n!$; row 3 has hooks $n-1,\dots,1$, product $(n-1)!$. Thus

$$f^{(n,n-1,n-1)} = \frac{6\,(3n-2)!}{(n+2)!\,n!\,(n-1)!}.$$

Step 4: the ratio.

$$\frac{|S_2|}{|S|} = \frac{6\,(3n-2)!}{(n+2)!\,n!\,(n-1)!} \cdot \frac{(n+2)!\,(n+1)!\,n!}{2\,(3n)!} = \frac{3\,(n+1)\,n}{3n\,(3n-1)} = \frac{n+1}{3n-1}.$$

Now $\frac{n+1}{3n-1} > \frac{1}{3}$ because $3n+3 > 3n-1$, and $\frac{n+1}{3n-1} \leq \frac{2}{3}$ because $3n+3 \leq 6n-2$ for $n \geq 2$. For $n = 2024$ the fraction is $\frac{2025}{6071}$, which lies in $[\frac{1}{3}, \frac{2}{3}]$. So the pair $(a,b) = (2,1)$, $(c,d) = (1,2)$ has the required property, and the answer is $\boxed{\text{yes}}$.

(Check at $n = 2$: the five tableaux of shape $(2,2,2)$ split 3 to 2 according to whether 2 is at $(2,1)$, and $(n+1)/(3n-1) = 3/5$.)

Layer 1: WHAT

Block 1: name the object: standard Young tableaux of shape (n, n, n); the fraction is a probability under the uniform tableau. Block 2: choose the pair $(2, 1), (1, 2)$ so that the event is "2 sits at $(2, 1)$". Block 3: delete 1 and 2, rotate the leftover skew shape by 180° with complemented entries, obtain the straight shape $(n, n-1, n-1)$. Block 4: hook length formula for both shapes; the ratio is $(n+1)/(3n-1)$. Block 5: check $1/3 < (n+1)/(3n-1) \leq 2/3$.

Juxtaposition. The nominal route reads "$T(a, b) < T(c, d)$" as a question about an arbitrary pair, looks for a symmetry that would give $1/2$, finds none, and either gives up or starts enumerating 3-letter ballot words by a recursion it cannot close. The ideal route reads the quantifier correctly (you choose the pair), chooses the pair adjacent to the corner, and reduces to two applications of a formula it already owns. The bifurcation point is the choice of pair: before any counting, decide that the event must be about the entry 2.

Layer 2: HOW

Block 1. **Cue.** A bijection onto $\{1, \ldots, 3n\}$ with increasing rows and increasing columns. **Why natural.** That is the definition of a standard Young tableau, and the tableau count is the one enumeration in this field with a closed form. Saying the name out loud pulls the hook length formula, the ballot-word model, and the rotation symmetry into working memory at once. **Link to the next block.** Once the object is named, "fraction of T" is "probability under the uniform tableau", and probabilities of events about small entries are ratios of tableau counts.

Block 2. **Cue.** You may choose (a, b) and (c, d); the entries near 1 have almost no freedom. **Why natural.** The cell holding k has at most $k-1$ cells weakly above-left of it, so the smallest entries are pinned near the corner: 1 at $(1, 1)$ and 2 on one of its two neighbours. For the pair of neighbours, "$T(2, 1) < T(1, 2)$" is equivalent to "2 is at $(2, 1)$", a single-entry event. This choice is not a guess; it is the unique pair for which the comparison is decided by the position of one entry. **Link to the next block.** An event "k sits in cell c" for small k is a count of tableaux on the cells that remain after removing $1, \ldots, k$.

Block 3. **Cue.** After removing 1 and 2 the remaining cells form a skew shape, not a partition. **Why natural.** The hook length formula applies to straight shapes. Reversing both the cell order (rotation by 180°) and the entry order (complementation $k \mapsto 3n + 1 - k$) preserves standardness, and for a rectangle minus a corner-anchored piece, rotation produces a straight shape. This is the standard complementation symmetry of rectangular tableaux. **Link to the next block.** Two straight shapes, two hook products.

Block 4. **Cue.** Two counts, one ratio. **Why natural.** Nearly every hook is common to the two products, so compute the ratio and let the factorials cancel; what survives is a rational function of degree zero in n, $(n+1)/(3n-1)$. Computing each count separately and dividing at the end is the same arithmetic done twice. **Link to the next block.** A rational function of n is compared with $1/3$ and $2/3$ by cross-multiplication.

Block 5. **Cue.** The problem asks for a fraction in a closed interval. **Why natural.** Both inequalities must be written out; the lower one is strict for all n, the upper one needs $n \geq 2$. State the value for $n = 2024$ so the grader sees the question answered as posed. **Link to the next block.** None.

Layer 3: WHY

The underlying truth is that a standard Young tableau is a linear extension of the poset of cells (ordered by "weakly above and weakly left"), and the position of the kth smallest element in a random linear extension is governed by counts of linear extensions of the poset minus a k-element lower set. For general posets those counts are intractable; for Young diagrams they are given exactly by the hook length formula, and removing the lower set $\{(1,1),(2,1)\}$ from a rectangle leaves a skew shape which rotation turns into a straight shape. So the problem sits precisely at the point where a hard general question (order statistics of linear extensions) becomes a two-line computation.

The discernment mechanism is quantifier discipline plus a bias toward extremes: when you choose the objects, choose the ones with the least freedom. In a tableau, freedom is minimal near the smallest and largest entries, so the computable pairs are the neighbours of the corner (and, by complementation, the neighbours of the opposite corner). You train this by asking, for any structure with forced elements, "which two positions are decided by one entry?" before computing anything. The examiners nudged twice: by allowing you to pick the pair, and by phrasing the target as $[1/3, 2/3]$, which tells you the answer will be near $1/3$, consistent with three rows and with the limit $(n+1)/(3n-1) \to 1/3$.

The knowledge gap is not disguised: without the hook length formula you cannot produce the exact value in the arena. Install it, with its hook definition, and install the two-row special case $f^{(n,n)} = C_n$ (Catalan) as the sanity check.

Revolving concepts

Standard Young tableaux. A partition $\mu = (\mu_1 \geq \mu_2 \geq \cdots)$ is drawn as left-aligned rows of μ_i cells; a standard Young tableau of shape μ is a filling by $1, \ldots, N$ increasing along rows and down columns; f^μ is their number. Two-row rectangles give Catalan numbers, $f^{(n,n)} = C_n$; the three-row rectangle gives $2(3n)!/(n!(n+1)!(n+2)!)$, the three-dimensional Catalan numbers. See Book 1, field: Combinatorics, topic: Young tableaux and the hook length formula.

The hook length formula and its recursion. $f^\mu = N!/\prod_c h(c)$. It is the closed-form solution of the branching recursion

$$f^\mu = \sum_{c \text{ a removable corner of } \mu} f^{\mu - c}, \qquad f^\varnothing = 1,$$

obtained by asking where the largest entry sits (it must occupy a corner cell). The same recursion holds for linear extensions of any finite poset (sum over maximal elements), and the hook formula is the rare case with a product solution; the official remark verifies the three-row case by checking that the product satisfies the recursion, which is the general method for closed forms of branching recursions: guess, then verify the recursion and the initial condition. Dividing the recursion by f^μ gives the probability that the maximum sits at a given corner, $f^{\mu-c}/f^\mu$, which is what this problem computes in complemented form. See Book 1, field: Combinatorics, topic: hook length formula; and field: Sequences and recurrences, topic: branching recursions and their verification.

Ballot words. Reading the row index of $1, 2, \ldots, N$ in order gives a word in which every prefix contains at least as many 1s as 2s and at least as many 2s as 3s (a lattice word); this is a bijection with tableaux of the corresponding shape. For two rows this is the ballot problem and the reflection principle; for three rows there is no reflection argument, which is why the hook formula is needed. See Book 1, field: Combinatorics, topic: ballot sequences.

Complementation. For a rectangle, $T \mapsto T^*$ with $T^*(i,j) = 3n + 1 - T(4-i, n+1-j)$ is an involution on S; it exchanges the roles of the smallest and largest entries, so the probability that 2 is at $(2,1)$ equals the probability that $3n-1$ is at $(2,n)$. See Book 1, field: Combinatorics, topic: symmetries of tableaux.

Circuit Bank

- **WHEN** you see a bijection onto $\{1, \ldots, N\}$ with increasing rows and increasing columns **DO** say "standard Young tableau" and write the hook length formula **BECAUSE** the count of tableaux of every straight shape is closed-form, and events about small or large entries are ratios of such counts
- **WHEN** you are free to choose positions in a structure with forced elements **DO** choose positions decided by a single extreme entry **BECAUSE** near the minimum or maximum the structure has the least freedom, so the event becomes "which corner holds entry k" and is a count of a smaller shape
- **WHEN** you delete the smallest entries of a rectangular tableau **DO** rotate by 180° and complement the entries **BECAUSE** reversing both orders preserves standardness and turns the leftover skew shape into a straight shape
- **WHEN** you need the ratio of two tableau counts **DO** divide the hook products directly **BECAUSE** almost every hook cancels and the ratio is a small rational function of n

Backtesting

1. First useful cue: "increasing along rows and down columns" with a bijection onto $\{1, \ldots, 6072\}$: a standard Young tableau of shape 3×2024.
2. Method the expression suggested: choose the pair of cells adjacent to the corner so the event is the position of the entry 2; count by hook lengths.
3. Where you would hesitate, and the thought to reinforce: after deleting 1 and 2 the shape is skew. The thought: rotate and complement; a rectangle minus a corner piece is a straight shape upside down.
4. The trap the question sets: reading the quantifiers as "for a given pair" and trying to count linear extensions in general; and expecting a symmetry that yields $1/2$.
5. What makes you recognise it faster next time: the interval $[1/3, 2/3]$ with three rows; the answer is near $1/3$ and comes from the two cells next to the corner.
6. The reusable decision rule: when you choose the positions, choose the ones with the least freedom; when you count tableaux, delete, rotate, and take a ratio of hook products.

Mirror drills

1. In a uniformly random standard Young tableau of shape (n, n), find the probability that 2 lies in the second row. *Answer:* deleting $1, 2$ and rotating gives shape $(n-1, n-1)$, so the probability is $C_{n-1}/C_n = (n+1)/(4n-2)$.
2. In a uniformly random tableau of shape (n, n, n), find the probability that 3 lies in the third row. *Answer:* $1, 2, 3$ fill the first column; deleting them and rotating gives $(n-1, n-1, n-1)$; the ratio of hook products is $(n+1)(n+2)/(3(3n-1)(3n-2))$, which is $1/5$ at $n = 2$.
3. Show that the number of standard Young tableaux of shape (n, n, n) equals the number of words using each of the letters $1, 2, 3$ exactly n times in which every prefix contains at least as many 1s as 2s and at least as many 2s as 3s. *Answer:* record the row of k for $k = 1, \ldots, 3n$; the prefix condition is exactly the column condition, and rows are automatically increasing.

2.4 A4. Complements turn a long progression into a short one

Find all primes $p > 5$ for which there exists an integer a and an integer r satisfying $1 \le r \le p-1$ with the following property: the sequence $1, a, a^2, \ldots, a^{p-5}$ can be rearranged to form a sequence $b_0, b_1, b_2, \ldots, b_{p-5}$ such that $b_n - b_{n-1} - r$ is divisible by p for $1 \le n \le p-5$.

Read it out loud

Translate the sentence into $\mathbb{F}_p$. The b_n satisfy $b_n \equiv b_0 + nr$, so they are the $p-4$ terms of an arithmetic progression with nonzero common difference r. "Can be rearranged" means set equality. So the statement is: the set of the first $p-4$ powers of a (a geometric progression) coincides with the term-set of a $(p-4)$-term arithmetic progression in $\mathbb{F}_p$. The fields wired in are the structure of the multiplicative group $\mathbb{F}_p^*$ (cyclic of order $p-1$, orders of elements, primitive roots), the rigidity of arithmetic and geometric progressions, and a counting reflex about complements.

The designed number is $p-5$. Not $p-2$, not $p-1$: the examiners chose the length so that the arithmetic progression, with $p-4$ of the p residues, misses exactly four residues, while the geometric progression, with $p-4$ of the $p-1$ nonzero residues, misses exactly three nonzero residues. Equality of two large sets is equality of two tiny complements. That is the first useful cue, and the second is that a progression of $p-4$ distinct powers forces the order of a to be at least $p-4$, which for $p \ge 11$ exceeds $(p-1)/2$ and makes a a primitive root. Within thirty seconds you should have written: distinct terms, so a is a primitive root; complement of the powers is $\{0, a^{-1}, a^{-2}, a^{-3}\}$; complement of the progression is a 4-term progression.

What is being tested is quantifier translation (rearrangement is set equality), the primitive-root reflex, and the discipline to pass to complements rather than compare two sets of $p-4$ elements directly. Once the complements are on the page, the problem is a finite check about four elements.

Approaches, ranked worst to best

1. **Hand search over small primes.** *Cue:* "find all primes $p > 5$". *Circuit:* compute the powers of candidate a modulo $7, 11, 13$ and look for progressions. (You find $p = 7$ quickly and then spend the hour failing to prove $p = 11$ is impossible. Evidence, not proof.)

2. **Power sums.** *Cue:* a geometric progression equals an arithmetic progression. *Circuit:* equate $\sum a^k$ with $\sum(b_0 + nr)$, then the sums of squares. (Two equations in three unknowns a, b_0, r with p-dependent coefficients; the sums of a $(p-4)$-term progression carry the full complexity. Stalls.)

3. **Closure properties.** *Cue:* the arithmetic progression is nearly closed under adding r; the powers are nearly closed under multiplying by a. *Circuit:* track where $x \mapsto x + r$ and $x \mapsto ax$ leave the set. (This is the complement idea done implicitly and badly: the "leaving" elements are exactly the complements, and it is far cleaner to name them.)

4. **Complements plus rigidity.** *Cue:* $p-4$ out of p; $p-4$ out of $p-1$. *Circuit:* distinctness forces a primitive root; the complements are $\{0, c, c^2, c^3\}$ and a 4-term progression; place 0, then scale. (Best. Everything reduces to whether $\{1, 2, 3\}$ can be ordered as a geometric progression modulo p, which happens only for $p \in \{5, 7\}$.)

The ideal solution

The answer is $\boxed{p = 7}$.

$p = 7$ works. Take $a = 3$, $r = 1$. The sequence $1, 3, 9$ rearranged as $(b_0, b_1, b_2) = (1, 9, 3)$ satisfies $b_1 - b_0 - 1 = 7$ and $b_2 - b_1 - 1 = -7$, both divisible by 7.

No prime $p \geq 11$ works. Suppose $p \geq 11$ and a, r have the property. Work in $\mathbb{F}_p$ and write $\bar{r} \neq 0$ for the class of r.

Step 1 (a primitive root). The condition says $b_n = b_0 + n\bar{r}$ for $0 \leq n \leq p-5$. These $p-4$ elements are pairwise distinct, because $b_n = b_m$ gives $(n-m)\bar{r} = 0$ with $0 \leq |n-m| \leq p-5 < p$, forcing $n = m$. Since the b_n are a rearrangement of $1, a, \ldots, a^{p-5}$, these powers are pairwise distinct in $\mathbb{F}_p$. In particular $a \neq 0$ (otherwise $a^1 = a^2 = 0$, and $p-5 \geq 2$). Let d be the multiplicative order of a; then $d \mid p-1$, and if $d \leq p-5$ then $a^d = a^0$ with both exponents in $\{0, \ldots, p-5\}$, contradicting distinctness. So $d \geq p-4$. Every divisor of $p-1$ other than $p-1$ itself is at most $(p-1)/2$, and $p-4 > (p-1)/2$ because $p > 7$. Hence $d = p-1$: a is a primitive root, and $\mathbb{F}_p^* = \{a^k : 0 \leq k \leq p-2\}$ with these $p-1$ elements distinct.

Step 2 (complements). Let $G = \{a^k : 0 \leq k \leq p-5\}$. Then $\mathbb{F}_p^* \setminus G = \{a^{p-4}, a^{p-3}, a^{p-2}\} = \{c^3, c^2, c\}$ where $c = a^{-1}$, using $a^{p-1} = 1$. So $\mathbb{F}_p \setminus G = \{0, c, c^2, c^3\}$. On the other side, $B = \{b_0 + n\bar{r} : 0 \leq n \leq p-5\}$, and since $n \mapsto b_0 + n\bar{r}$ is a bijection from $\mathbb{Z}/p\mathbb{Z}$ onto $\mathbb{F}_p$, its complement is

$$\mathbb{F}_p \setminus B = \{b_0 + n\bar{r} : n = p-4, p-3, p-2, p-1\} = \{b_0 - 4\bar{r},\ b_0 - 3\bar{r},\ b_0 - 2\bar{r},\ b_0 - \bar{r}\}.$$

As $B = G$, the complements agree: the four elements $0, c, c^2, c^3$ are, in some order, four consecutive terms of an arithmetic progression with common difference $\bar{r} \neq 0$. Multiplying every term by $c^{-1} = a$ preserves this (an arithmetic progression scaled by a nonzero constant is an arithmetic progression with scaled difference), so the set $\{0, 1, c, c^2\}$ is the term-set of an arithmetic progression x_1, x_2, x_3, x_4 with common difference $r' = a\bar{r} \neq 0$. Note that c, being the inverse of a primitive root, is a primitive root, so it has order $p-1 \geq 10$; in particular $1, c, c^2$ are distinct and nonzero, $c \neq -1$ and $c^2 \neq -1$.

Step 3 (0 is an endpoint). If $0 = x_2$ or $0 = x_3$, then its two neighbours in the progression satisfy $x_{k-1} + x_{k+1} = 2x_k = 0$, so two of $1, c, c^2$ are negatives of each other: $c = -1$, or $c^2 = -1$, or $c^2 = -c$ (again $c = -1$). All are excluded by the order of c. Hence $0 \in \{x_1, x_4\}$; reversing the progression (which replaces r' by $-r'$) we may assume $x_1 = 0$, so $(x_2, x_3, x_4) = (r', 2r', 3r')$ and

$$\{1, c, c^2\} = \{r', 2r', 3r'\}.$$

Step 4 (a geometric progression inside $\{1, 2, 3\}$). Multiply by $u = r'^{-1}$: $\{u, uc, uc^2\} = \{1, 2, 3\}$, and $1, 2, 3$ are distinct in $\mathbb{F}_p$ since $p \geq 5$. So $1, 2, 3$ can be listed in some order as $v_1 = u$, $v_2 = uc$, $v_3 = uc^2$, and then $v_2^2 = v_1 v_3$. The middle term v_2 is one of $1, 2, 3$. If $v_2 = 1$ then $1 = 2 \cdot 3$, so $p \mid 5$. If $v_2 = 2$ then $4 = 1 \cdot 3$, so $p \mid 1$, impossible. If $v_2 = 3$ then $9 = 1 \cdot 2$, so $p \mid 7$. None of these holds for $p \geq 11$, a contradiction.

Therefore no prime $p \geq 11$ has the property, and $p = 7$ is the only prime $p > 5$ that does.

Layer 1: WHAT

Block 1: certify $p = 7$ by an explicit (a, r). Block 2: distinct progression terms force distinct powers, hence order at least $p-4$, hence a primitive root for $p \geq 11$. Block 3: pass to complements: the powers miss $\{0, c, c^2, c^3\}$, the progression misses a 4-term progression, so $\{0, c, c^2, c^3\}$ is a 4-term progression. Block 4: scale by c^{-1} and place 0: interior positions force a pair of negatives among $1, c, c^2$, impossible for an element of order at least 10, so 0 is an endpoint and the rest is $r'\{1, 2, 3\}$. Block 5: $\{1, 2, 3\}$ must then be a geometric progression modulo p, which pins $p \in \{5, 7\}$.

Juxtaposition. The nominal route compares the two $(p-4)$-element sets head-on: sums, products, or a search. The ideal route notices that both sets are almost everything and compares what they miss. The bifurcation point is the reading of "$p-5$": the zombie sees a long sequence to be rearranged; you see four missing residues.

Layer 2: HOW

Block 1. Cue. "Find all primes $p > 5$" with an existential inner statement. **Why natural.** The answer is a finite list or a congruence class, and the smallest admissible prime is where a construction

lives. For $p = 7$ the sequence has three terms $1, a, a^2$, and three residues form a progression when twice one is the sum of the other two; $a = 3$ gives $\{1, 3, 2\}$, which is $\{1, 2, 3\}$. **Link to the next block.** With $p = 7$ certified, the task is a uniform impossibility proof for $p \geq 11$, and impossibility must come from structure that only appears when p is large enough: the order of a.

Block 2. **Cue.** $p - 4$ pairwise distinct powers of a. **Why natural.** A progression with nonzero difference and fewer than p terms has distinct terms, so the powers are distinct, so the order of a is at least the number of powers. An order that exceeds $(p-1)/2$ and divides $p - 1$ must equal $p - 1$. The threshold $p - 4 > (p-1)/2$ is exactly $p > 7$, which is why 7 is special and why the problem says $p > 5$ rather than $p > 7$: it wants you to find the boundary. **Link to the next block.** A primitive root makes the powers enumerate $\mathbb{F}_p^*$ in a known order, so the missing powers are named explicitly.

Block 3. **Cue.** $p - 4$ elements out of p (progression) and out of $p - 1$ (powers). **Why natural.** Two subsets of a finite set are equal iff their complements are equal, and here the complements have 4 and $3 + 1$ elements. The complement of a long progression is a short progression (the map $n \mapsto b_0 + n\bar{r}$ is a bijection of $\mathbb{Z}/p\mathbb{Z}$ with $\mathbb{F}_p$, so the missing indices are consecutive); the complement of the powers is the last three powers, which are a^{-1}, a^{-2}, a^{-3}, plus 0. This is the whole reduction: a statement about $p - 4$ elements becomes a statement about 4. **Link to the next block.** A 4-term arithmetic progression containing 0 and three powers of c is rigid enough to be classified by hand.

Block 4. **Cue.** A progression contains 0; the other members are $1, c, c^2$ after scaling. **Why natural.** Scaling by c^{-1} costs nothing (progressions are preserved) and reduces the number of symbols. In a progression, an interior term is the average of its neighbours, so 0 in the interior forces two of $1, c, c^2$ to sum to 0, which is a small-order condition on c; a primitive root has large order. So 0 sits at an end, and an end at 0 means the other terms are $r', 2r', 3r'$. **Link to the next block.** The set $\{1, c, c^2\}$ is a geometric progression; the set $\{r', 2r', 3r'\}$ is r' times $\{1, 2, 3\}$; equality of the two sets means $\{1, 2, 3\}$ is a geometric progression up to order.

Block 5. **Cue.** Three specific residues that must form a geometric progression. **Why natural.** Three elements in geometric progression satisfy (middle)^2 = (product of the ends), so there are three cases, each a congruence $p \mid 5$, $p \mid 1$, $p \mid 7$. This is a finite check and it explains the answer: $\{1, 2, 3\}$ is simultaneously an arithmetic and a geometric progression only in $\mathbb{F}_5$ and $\mathbb{F}_7$. **Link to the next block.** None; the contradiction for $p \geq 11$ closes the proof.

Layer 3: WHY

The truth underneath is the incompatibility of the two group structures of $\mathbb{F}_p$: an arithmetic progression is an orbit segment of the translation $x \mapsto x + r$, a geometric progression is an orbit segment of the multiplication $x \mapsto cx$, and a set that is simultaneously a long segment of both is impossible except for tiny sets in tiny fields. The problem hides this behind length $p - 4$; complements expose it, because the complement of a segment of a p-cycle is a segment, and the complement of a segment of the $(p-1)$-cycle generated by a primitive root is a segment. Rigidity of short progressions (an interior term is the mean of its neighbours; an endpoint at 0 forces multiples $r', 2r', 3r'$) then converts the question to a three-case congruence.

The discernment mechanism has three reflexes. First, "rearranged into" is set equality, and set equality of large sets is set equality of small complements. Second, "$1, a, \ldots, a^m$ distinct" with $m \geq (p-1)/2$ is "primitive root"; say it before anything else. Third, when a progression contains 0, place the 0: this is the extremal move that breaks the symmetry. You train these by rewriting every problem about $p - k$ objects as a problem about k objects, and by treating any element of large multiplicative order as a ban on small polynomial relations ($c = -1$, $c^2 = -1$). The examiners

nudged with $p-5$ and with $p>5$: the first says "four missing", the second says "the boundary case is small and you must find it".

Revolving concepts

The cyclic group $\mathbb{F}_p^*$. It is cyclic of order $p-1$; the order of any element divides $p-1$; an element of order $p-1$ is a primitive root; the inverse and every power a^k with $\gcd(k,p-1)=1$ of a primitive root are primitive roots; $a^{p-1-k}=a^{-k}$. Key lemma used here: if $1,a,\dots,a^m$ are distinct with $m\geq(p-1)/2$, then a is a primitive root, because the largest proper divisor of $p-1$ is at most $(p-1)/2$. See Book 1, field: Number theory, topic: orders and primitive roots.

Progressions as orbits, and the generalised recursion. The recursion $b_n=b_{n-1}+r$ has closed form $b_n=b_0+nr$; in $\mathbb{Z}/p\mathbb{Z}$ with $r\neq 0$ its orbit is the whole group with period p, so a $(p-k)$-term segment misses a k-term segment. The recursion $g_n=ag_{n-1}$ has closed form $g_n=a^n$ with period $\operatorname{ord}(a)$, so for a primitive root a $(p-1-k)$-term segment of $\mathbb{F}_p^*$ misses $\{a^{-1},\dots,a^{-k}\}$. The general problem "$1,a,\dots,a^{p-1-k}$ rearranges into an arithmetic progression" is therefore "$\{0,c,\dots,c^{k-1}\}$ is a k-term progression": for $k=2$ every prime works; for $k=3$ it works iff 0 is an endpoint and $c^2=2c$ or $c=2c^2$, i.e. iff 2 is a primitive root modulo p; for $k=4$ it is this problem, with answer $p\in\{5,7\}$; for $k\geq 4$ the scaled set $\{1,\dots,k-1\}$ would have to be a geometric progression, which confines p to finitely many values. See Book 1, field: Sequences and recurrences, topic: orbits of affine maps modulo p.

Rigidity of short progressions. Three elements x,y,z (in that order) are in arithmetic progression iff $2y=x+z$, in geometric progression iff $y^2=xz$; scaling by a nonzero constant and reversing preserve both. A k-term arithmetic progression whose term-set contains 0 has 0 at an endpoint unless two other terms are negatives. See Book 1, field: Number theory, topic: arithmetic progressions in $\mathbb{Z}/p\mathbb{Z}$.

Complement principle. A subset of size $m-k$ of an m-set is the same data as its k-subset complement; when k is small, every structural statement should be transported to the complement. See Book 1, field: Combinatorics, topic: complementation.

Circuit Bank

- **WHEN** the powers $1,a,\dots,a^m$ are forced to be distinct modulo p with $m\geq(p-1)/2$ **DO** write "a is a primitive root" **BECAUSE** the order of a is at least $m+1>(p-1)/2$ and divides $p-1$, so it equals $p-1$
- **WHEN** a subset of $\mathbb{F}_p$ or $\mathbb{F}_p^*$ has $p-k$ elements for small k **DO** transport every condition to the k-element complement **BECAUSE** equality of large sets is equality of complements, and the complement of a segment of a cyclic orbit is a segment
- **WHEN** an arithmetic progression contains 0 **DO** place the 0: interior means two other terms are negatives, endpoint means the rest is $r',2r',\dots$ **BECAUSE** an interior term is the average of its neighbours, and a progression starting at 0 is r' times $1,2,3,\dots$
- **WHEN** a set must be both an arithmetic and a geometric progression **DO** write (middle)^2 = (product of ends) for each possible middle **BECAUSE** this is a finite list of congruences on p, which classifies the exceptional primes

Backtesting

1. First useful cue: the length $p-5$, i.e. $p-4$ terms out of p residues.
2. Method the expression suggested: primitive root from distinctness, then complements, then rigidity of a 4-term progression containing 0.

3. Where you would hesitate, and the thought to reinforce: after obtaining "$\{0, c, c^2, c^3\}$ is a progression" you might start a case analysis on all 24 orderings. The thought: scale to $\{0, 1, c, c^2\}$, place 0, and the remaining condition is "$\{1, 2, 3\}$ is a geometric progression".
4. The trap the question sets: $p > 5$ invites you to assume the answer is a congruence class or "all p"; the primitive-root step needs $p > 7$, so $p = 7$ must be checked by hand and turns out to be the entire answer.
5. What makes you recognise it faster next time: any "$1, a, \ldots, a^{p-k}$" with small k is a primitive root plus a k-element complement.
6. The reusable decision rule: big subset, small complement; long progression, short missing progression; progression through 0, place the 0.

Mirror drills

1. Let p be an odd prime and a an integer such that $1, a, \ldots, a^{(p-1)/2}$ are pairwise incongruent modulo p. Show that a is a primitive root. *Answer:* the order is at least $(p+1)/2 > (p-1)/2$ and divides $p-1$.
2. For which primes $p \geq 7$ can $1, a, \ldots, a^{p-4}$ be rearranged into an arithmetic progression modulo p with nonzero difference, for some a? *Answer:* exactly when 2 is a primitive root modulo p: the complement $\{0, c, c^2\}$ is a 3-term progression, 0 is an endpoint since $c \neq -1$, so $c^2 = 2c$ or $c = 2c^2$, i.e. $c \in \{2, 1/2\}$.
3. Find all primes p for which $1, 2, 3$ can be ordered as a geometric progression modulo p. *Answer:* $p = 5$ (order $2, 1, 3$, ratio 3) and $p = 7$ (order $1, 3, 2$, ratio 3); the three cases are $1 = 6$, $4 = 3$, $9 = 2$.

2.5 A5. A random chord is a direction plus an offset

Consider a circle Ω with radius 9 and center at the origin $(0,0)$, and a disc Δ with radius 1 and center at $(r,0)$, where $0 \le r \le 8$. Two points P and Q are chosen independently and uniformly at random on Ω. Which value(s) of r minimize the probability that the chord $\overline{PQ}$ intersects Δ?

Read it out loud

A geometric probability question: a random chord of a large circle, a small disc inside it, and a one-parameter family of positions for the disc along a radius; minimise the hitting probability over the parameter. The fields wired in are geometric probability (the two-endpoint model of a random chord), coordinate geometry of lines and circles (a line is a direction plus an offset), a change of variables for uniform random variables, and one piece of circle geometry (a minor arc is an increasing function of its chord). No heavy machinery, but a rigour trap at every step.

The first useful cue is "chosen independently and uniformly at random on Ω". Write $P = 9(\cos\alpha, \sin\alpha)$, $Q = 9(\cos\beta, \sin\beta)$ with α, β independent and uniform. The chord PQ depends on (α, β) through $(\alpha+\beta)/2$, the direction of its perpendicular bisector, and $(\alpha-\beta)/2$, which fixes its distance $9|\cos((\alpha-\beta)/2)|$ from the centre. Sum and difference of two independent uniform angles are again independent and uniform. So a random chord in this model is a uniform random direction together with an independent offset of the form $9\cos\psi$ with ψ uniform. Within thirty seconds you should have written that decomposition and the sentence "condition on the direction, rotate so the chord is vertical".

What is being tested is whether you can turn a two-dimensional random object into a one-dimensional one per direction, prove a one-dimensional inequality pointwise, and integrate; and whether you notice that $r \le 8$ keeps Δ inside the big disc, so "the chord meets Δ" is the same as "the line meets Δ". The intuitive answer $r = 0$ is correct, but intuition is not a proof, and the examiners want the per-direction argument.

Approaches, ranked worst to best

1. **Intuition or simulation.** *Cue:* "which r minimises". *Circuit:* chords in this model cluster near the boundary, so the centre is hit least. (True, and worth zero marks. The heuristic does tell you what to prove.)

2. **Double integral over (α, β).** *Cue:* two uniform angles. *Circuit:* integrate the indicator of "chord meets Δ" over the torus. (The region in the (α,β)-square is bounded by transcendental curves; you cannot even write it down cleanly. Consumption without analysis.)

3. **Closed form and differentiate.** *Cue:* a probability depending smoothly on r. *Circuit:* direction-offset decomposition, then $P(r)$ as an integral of arccos differences, then differentiate under the integral sign. (The official first solution. Correct, and the derivative has a clean sign, but it costs the justification of differentiating under the integral and invites a sign slip; the official write-up itself carries a sign typo.)

4. **Direction-offset decomposition plus a per-direction chord comparison.** *Cue:* rotational symmetry; a strip of width 2. *Circuit:* for each direction, the arc of the circle cut by the strip is shortest when the strip is centred; integrate the pointwise inequality. (Best. No differentiation, one geometric lemma, strictness for free.)

The ideal solution

The probability is minimised at $\boxed{r = 0}$ and at no other value.

Step 1: a random chord is a random rotation of a vertical chord. Let α, β be independent and uniform on $[0, 2\pi)$, with $P = 9(\cos\alpha, \sin\alpha)$ and $Q = 9(\cos\beta, \sin\beta)$. Let U, ψ be independent and uniform on $[0, 2\pi)$ and consider $\Phi(U, \psi) = (U + \psi, U - \psi)$ reduced modulo 2π in each coordinate. The linear map with matrix $\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ has determinant -2 and maps the lattice $(2\pi\mathbb{Z})^2$ into itself with index 2, so Φ is a two-to-one map of the torus $[0, 2\pi)^2$ onto itself which multiplies area by 2 on each sheet. Hence for every measurable E, the preimage $\Phi^{-1}(E)$ consists of two sheets each of area $\text{area}(E)/2$, so $\text{area}(\Phi^{-1}(E)) = \text{area}(E)$: the pair $(U + \psi, U - \psi)$ is uniform on the torus, i.e. has the same law as (α, β). Consequently the pair (P, Q) has the same law as $(R_U P_0, R_U Q_0)$, where R_U is the rotation about the origin by U and

$$P_0 = 9(\cos\psi, \sin\psi), \qquad Q_0 = 9(\cos\psi, -\sin\psi).$$

All probabilities below are computed in this representation. The event $P = Q$ (that is, $\sin\psi = 0$) has probability 0 and is discarded; for $\sin\psi \neq 0$ the line P_0Q_0 is the vertical line $x = 9\cos\psi$.

Step 2: chord versus line. Every point of Δ is within distance $r + 1 \leq 9$ of the origin, so Δ lies in the closed disc $\bar{D}$ of radius 9. The chord $\overline{PQ}$ is exactly the intersection of the line PQ with $\bar{D}$. Hence the chord meets Δ if and only if the line PQ meets Δ. Applying R_{-U}: the line PQ meets Δ iff the line $x = 9\cos\psi$ meets $R_{-U}(\Delta)$, the closed unit disc centred at $(r\cos U, -r\sin U)$. A vertical line $x = x_0$ meets a closed unit disc with centre (u, v) iff $|x_0 - u| \leq 1$. Therefore

$$\overline{PQ} \cap \Delta \neq \varnothing \iff |9\cos\psi - r\cos U| \leq 1.$$

Step 3: the conditional probability given the direction. Fix U and put $c = r\cos U \in [-8, 8]$. Since $\cos\psi = \cos(2\pi - \psi)$, the law of $\cos\psi$ for ψ uniform on $[0, 2\pi)$ equals the law of $\cos\varphi$ for φ uniform on $[0, \pi]$, where cos is strictly decreasing. As $(c \pm 1)/9 \in [-1, 1]$,

$$|9\cos\varphi - c| \leq 1 \iff \arccos\frac{c+1}{9} \leq \varphi \leq \arccos\frac{c-1}{9},$$

so, with

$$g(c) := \arccos\frac{c-1}{9} - \arccos\frac{c+1}{9} \qquad (-8 \leq c \leq 8),$$

the conditional probability of a hit given U is $g(r\cos U)/\pi$, and by independence of U and ψ,

$$\mathrm{P}(r) := \mathrm{P}\big(\overline{PQ} \cap \Delta \neq \varnothing\big) = \frac{1}{2\pi}\int_0^{2\pi} \frac{g(r\cos U)}{\pi}\, dU.$$

Step 4: the strip lemma.

Lemma 2.5.1. *For $-8 \leq c \leq 8$, $g(c) \geq g(0)$, with equality if and only if $c = 0$.*

Proof. $g(c)$ is the angular measure of the arc of the upper half of the circle $x^2 + y^2 = 81$ whose points have $c - 1 \leq x \leq c + 1$; its endpoints are $A = (c - 1, a)$ and $B = (c + 1, b)$ with $a = \sqrt{81 - (c-1)^2}$ and $b = \sqrt{81 - (c+1)^2}$, both nonnegative. This arc lies in the upper semicircle and is not the whole semicircle (that would need $c - 1 = -9$ and $c + 1 = 9$), so $0 < g(c) < \pi$, and the chord AB has length $18\sin(g(c)/2)$ with $g(c)/2 \in (0, \pi/2)$, where sin is strictly increasing. Now

$$|AB|^2 = 2^2 + (a - b)^2 \geq 4,$$

with equality iff $a = b$, iff $(c-1)^2 = (c+1)^2$, iff $c = 0$. So $|AB| \geq 2 = |AB|_{c=0}$, hence $\sin(g(c)/2) \geq \sin(g(0)/2)$, hence $g(c) \geq g(0)$, with equality exactly when $c = 0$. □

Step 5: integrate. For $r = 0$, $\mathrm{P}(0) = g(0)/\pi$. For $0 < r \leq 8$, the function $U \mapsto g(r\cos U) - g(0)$ is continuous on $[0, 2\pi]$, nonnegative by the lemma, and strictly positive wherever $\cos U \neq 0$, in

particular at $U = 0$. A continuous nonnegative function that is positive somewhere has positive integral, so

$$\mathrm{P}(r) - \mathrm{P}(0) = \frac{1}{2\pi^2} \int_0^{2\pi} \big(g(r\cos U) - g(0)\big)\, dU > 0.$$

Thus $r = 0$ is the unique minimiser. For the record, $\mathrm{P}(0) = \big(\arccos(-\frac{1}{9}) - \arccos\frac{1}{9}\big)/\pi = \frac{2}{\pi}\arcsin\frac{1}{9}$, which agrees with the direct computation: the chord passes within distance 1 of the centre iff $9|\cos\varphi| \le 1$, an event of probability $(\pi - 2\arccos\frac{1}{9})/\pi$.

Layer 1: WHAT

Block 1: replace the two uniform endpoints by a uniform rotation of a vertical chord $x = 9\cos\psi$, ψ uniform, via the two-to-one linear change of variables on the torus. Block 2: since Δ sits inside the big disc, "chord meets Δ" is "line meets Δ", and after rotating back this is $|9\cos\psi - r\cos U| \le 1$. Block 3: for a fixed direction, the hit probability is the angle $g(r\cos U)$ of an arc divided by π. Block 4: the strip lemma: the arc of a circle cut by a strip of width 2 is shortest when the strip is centred, because its chord is shortest then and arcs increase with chords. Block 5: integrate the pointwise inequality; strictness comes from continuity.

Juxtaposition. The nominal route writes the event as a region in the (α, β)-square and tries to integrate, or computes $\mathrm{P}(r)$ in closed form and differentiates under the integral sign. The ideal route never integrates anything it does not have to: it identifies the one-dimensional picture per direction, proves the inequality there, and lets the integral inherit it. The bifurcation point is the choice of coordinates for a chord: endpoints (zombie) versus direction and offset (ideal). Once the direction is conditioned on, the rest is forced.

Layer 2: HOW

Block 1. **Cue.** Two independent uniform points on a circle; the problem is invariant under rotation about the centre. **Why natural.** A chord is determined by its direction and its offset, and in this model those are independent: the sum and difference of two independent uniform angles are independent uniform angles, which is the statement that the linear map $(U, \psi) \mapsto (U + \psi, U - \psi)$ is measure-preserving on the torus. You reach for it because the disc's position enters only through the direction (via $r\cos U$), and separating direction from offset separates the parameter from the randomness. **Link to the next block.** With the direction frozen, the chord is a vertical line at a random abscissa $9\cos\psi$, and the question becomes one-dimensional.

Block 2. **Cue.** $r \le 8$ and radius 1: the disc is inside the circle. **Why natural.** The chord is the intersection of its line with the closed big disc, so anything inside the big disc is met by the chord iff it is met by the line. Rotating the whole picture by $-U$ moves the line to $x = 9\cos\psi$ and the disc centre to $(r\cos U, -r\sin U)$; a vertical line meets a unit disc iff the horizontal distance from the centre is at most 1. The only coordinate of the disc that matters is its abscissa $r\cos U$. **Link to the next block.** The event is now an interval condition on $\cos\psi$, and $\cos\psi$ with ψ uniform has the same law as $\cos\varphi$ with φ uniform on $[0, \pi]$.

Block 3. **Cue.** An interval condition on $\cos\varphi$ with φ uniform on $[0, \pi]$. **Why natural.** Uniform φ means probability equals angular measure over π; inverting cos on $[0, \pi]$ turns the interval $[(c-1)/9, (c+1)/9]$ into an interval of angles of length $g(c)$. Geometrically, $g(c)$ is the angle of the arc of the upper semicircle inside the vertical strip of width 2 centred at $x = c$. **Link to the next block.** The minimisation over r is now a minimisation of $g(r\cos U)$ for each U, and it suffices to minimise g on $[-8, 8]$.

Block 4. **Cue.** Compare arcs cut from a fixed circle by a strip of fixed width at different positions. **Why natural.** Comparing arcs directly means comparing differences of arccos, which is calculus; comparing their chords is geometry. The endpoints of the arc lie on the two boundary lines of the strip, so the chord is at least the width 2 of the strip, with equality only when it is perpendicular to the lines, i.e. when the strip is symmetric about the y-axis. The minor arc is $2R\sin^{-1}(\text{chord}/2R)$, increasing in the chord; the only care needed is that the arc is indeed minor, which holds because it sits inside a semicircle. **Link to the next block.** A pointwise inequality in U with a single equality point per direction is ready to be integrated.

Block 5. **Cue.** An average of a nonnegative continuous function that is not identically zero. **Why natural.** Integration preserves inequalities, and a continuous nonnegative integrand that is positive at one point has positive integral. Writing this sentence is what makes "and no other value" a proven statement rather than a remark. **Link to the next block.** None; the minimiser is unique.

Layer 3: WHY

The mathematical truth is a decoupling and a monotonicity. The decoupling: in the two-endpoint model, the direction of the chord is uniform and independent of the offset, and the offset is $9\cos\psi$ with ψ uniform, so it has the arcsine law with density $1/(\pi\sqrt{81-t^2})$ on $(-9,9)$, largest near the boundary. That density is why the centre is hit least: chords in this model crowd the rim. The monotonicity: for each direction, hitting the disc is the offset falling in a window of fixed width 2 centred at the disc's projection, and the arcsine density is symmetric and increasing away from 0, so a window of fixed width captures least mass when centred at 0. The strip lemma is this monotonicity stated in the language of arcs and chords.

The discernment mechanism: when a probability is an average over a symmetry group, look for an inequality that holds for every group element separately, because pointwise domination integrates and strictness survives integration by continuity. You train it by refusing to compute an integral before asking whether the integrand can be compared pointwise; and by keeping a short list of exact facts about uniform points on circles (sum and difference of angles, the arcsine law of the projection, the 2 sin chord formula). The examiners nudged with the numbers: radius 9, disc radius 1, and $r \le 8$ so that $r+1 \le 9$, which guarantees the chord-versus-line equivalence and keeps both boundary lines of the strip inside the circle. They also nudged with "value(s)", an invitation to prove uniqueness.

Revolving concepts

Random chords and Bertrand's paradox. "Random chord" has no meaning until the model is fixed; two uniform endpoints, a uniform midpoint, and a uniform offset along a uniform direction give different answers (probability $1/3$, $1/4$, $1/2$ that the chord exceeds the inscribed equilateral triangle's side). In the two-endpoint model the direction is uniform and the signed offset $t = 9\cos\psi$ has the arcsine density $1/(\pi\sqrt{81-t^2})$. See Book 1, field: Probability, topic: geometric probability and Bertrand's paradox.

Uniform angles under linear maps. If (U,ψ) is uniform on the torus and M is an integer matrix with $\det M \neq 0$, then $M(U,\psi)$ reduced modulo 2π is uniform: the map is $|\det M|$-to-one and scales area by $|\det M|$. In particular sums and differences of independent uniform angles are independent uniform angles. See Book 1, field: Probability, topic: transformations of uniform random variables.

Arc and chord. An arc of angle $\theta \in [0,\pi]$ on a circle of radius R has chord $2R\sin(\theta/2)$, strictly increasing in θ; so among minor arcs, longer chord means longer arc. The strip lemma follows: the arc cut by a strip of width $w < 2R$ is shortest when the strip is symmetric about the parallel diameter. See Book 1, field: Geometry, topic: circle, chord and arc.

Pointwise domination before integration. If $h_r(U) \geq h_0(U)$ for all U with strict inequality on a set of positive measure and both continuous, then $\int h_r > \int h_0$. This is the general solution method for "minimise an average": minimise the integrand for each value of the averaging variable. The same idea handles any convex body K inside the circle: the hit probability is the average over directions of the arc cut by the strip between the two supporting lines of K perpendicular to that direction. See Book 1, field: Real analysis, topic: monotonicity of the integral.

Differentiation under the integral sign. The alternative route: $\mathrm{P}'(r) = \frac{1}{\pi^2} \int \partial_r g(r \cos U)\, dU$ is legitimate because $\partial_r g$ is continuous on the compact parameter set, and $\partial_c g(c)$ has the sign of c; so P is strictly increasing on $[0, 8]$. See Book 1, field: Calculus and integrals, topic: Leibniz rule.

Circuit Bank

- **WHEN** two points are chosen independently and uniformly on a circle **DO** write the chord as a uniform direction plus an independent offset $R \cos \psi$, ψ uniform **BECAUSE** the sum and difference of two independent uniform angles are independent and uniform, so direction and offset decouple
- **WHEN** a probability is an average over rotations (or any symmetry group) and must be optimised over a parameter **DO** prove the inequality for each direction separately, then integrate **BECAUSE** pointwise domination integrates, and continuity turns one strict point into a strict integral
- **WHEN** you must compare two arcs of the same circle **DO** compare their chords, after checking both arcs are minor **BECAUSE** the minor arc is $2R \sin^{-1}(\text{chord}/2R)$, strictly increasing in the chord
- **WHEN** a target region sits inside the disc bounded by the circle **DO** replace "the chord meets the region" by "the line meets the region" **BECAUSE** the chord is the full intersection of its line with the closed disc

Backtesting

1. First useful cue: "independently and uniformly on Ω" together with a problem invariant under rotation.
2. Method the expression suggested: condition on the chord's direction, reduce to a vertical line at abscissa $9 \cos \psi$, and compare arcs cut by a strip of width 2.
3. Where you would hesitate, and the thought to reinforce: justifying that direction and offset are independent and uniform. The thought: it is a linear change of variables on the torus with determinant -2, two-to-one and area-doubling, so uniform goes to uniform.
4. The trap the question sets: the answer looks obvious and the temptation is to write it with a symmetry hand-wave; the marks are in the per-direction inequality, in the chord-versus-line remark, and in the strictness.
5. What makes you recognise it faster next time: any "random chord" phrased with two uniform endpoints is direction plus arcsine offset, and any "strip of fixed width across a circle" is the chord-arc lemma.
6. The reusable decision rule: decouple the randomness from the parameter by conditioning on the symmetry variable, prove pointwise, integrate.

Mirror drills

1. Two points are chosen independently and uniformly on a circle of radius R. Show that the probability that the chord passes within distance s of the centre is $\frac{2}{\pi}\arcsin\frac{s}{R}$. *Answer:* the distance is $R|\cos\varphi|$ with φ uniform on $[0,\pi]$; the set $|\cos\varphi| \le s/R$ has measure $\pi - 2\arccos(s/R) = 2\arcsin(s/R)$.

2. In the same model, find the probability that the chord is longer than the side of the inscribed equilateral triangle. *Answer:* the chord has length $2R|\sin\psi|$, and $|\sin\psi| > \sqrt{3}/2$ on a set of measure $1/3$; this is the Bertrand value $1/3$ for the two-endpoint model.

3. A disc of radius ρ is centred at distance r from the centre of a circle of radius R, with $r + \rho \le R$; two points are uniform on the circle. Show that the probability that the chord meets the disc is strictly increasing in r on $[0, R-\rho]$. *Answer:* per direction the strip has width 2ρ and centre $r\cos U$, and $\frac{d}{dc}\big(\arccos\frac{c-\rho}{R} - \arccos\frac{c+\rho}{R}\big)$ has the sign of c, so each conditional probability is increasing in $|c|$; integrate.

2.6 A6. Hankel determinants through a self-similar inverse series

Let $c_0, c_1, c_2, \ldots$ be the sequence defined so that

$$\frac{1 - 3x - \sqrt{1 - 14x + 9x^2}}{4} = \sum_{k=0}^{\infty} c_k x^k$$

for sufficiently small x. For a positive integer n, let A be the n-by-n matrix with i, j-entry c_{i+j-1} for i and j in $\{1, \ldots, n\}$. Find the determinant of A.

Read it out loud

A determinant whose entries depend only on $i + j$: a Hankel determinant. The entries are the coefficients of a generating function that is a square root of a quadratic, so the generating function satisfies a quadratic equation and the coefficients satisfy a convolution recursion of Catalan type. The fields wired in are generating functions (algebraic power series), linear algebra (determinants under unimodular row and column operations, block triangular factorisation), and the classical circle of ideas around Hankel matrices: continued fractions, orthogonal polynomials, moment functionals, and lattice paths. The words to say out loud are "Hankel matrix of a Catalan-like sequence", and the theorem to pull is: the Hankel determinants of a sequence whose generating function has a continued fraction expansion with denominators $\lambda_1, \lambda_2, \ldots$ equal $\lambda_1^{n-1}\lambda_2^{n-2}\cdots\lambda_{n-1}$.

Two things to write in the first minute. First, the quadratic: squaring $1 - 3x - 4f = \sqrt{1 - 14x + 9x^2}$ gives $f = x + 3xf + 2f^2$, so $c_0 = 0$, $c_1 = 1$, $c_2 = 5$, $c_3 = 35$, $c_4 = 295$, $c_5 = 2765$. Second, the small determinants: $n = 1$ gives 1, $n = 2$ gives $35 - 25 = 10$, $n = 3$ gives 1000. The pattern $10^0, 10^1, 10^3$ is $10^{n(n-1)/2}$, which is $\prod_{k=1}^{n-1} \lambda_k^{n-k}$ with every $\lambda_k = 10$. So the target is known within five minutes; the problem is the proof, and the proof needs a theorem you must have installed, or the one lemma that proves it.

Brutal honesty about the knowledge gap: without knowing that Hankel determinants factor through continued fraction coefficients, or without the reduction lemma below, this problem is not approachable in the arena. There is no elementary reroute. The lemma is short and its proof is a single unitriangular factorisation; install it, and the problem becomes a five-line induction.

Approaches, ranked worst to best

1. **Compute and guess.** *Cue:* "find the determinant" for general n. *Circuit:* evaluate c_k and the first few determinants. (Necessary as reconnaissance, worthless as a proof. It tells you the answer $10^{n(n-1)/2}$ and the shape $\prod \lambda_k^{n-k}$.)

2. **Numeric row reduction.** *Cue:* a concrete matrix. *Circuit:* eliminate using the recursion $c_m = 3c_{m-1} + 2\sum c_i c_{m-i}$. (The recursion is quadratic in the c's, and unorganised elimination produces no pattern. This is the lemma below attempted without its organising principle.)

3. **Lattice paths and nonintersecting families.** *Cue:* Catalan-like recursion. *Circuit:* Flajolet's theorem (continued fraction coefficients weight Motzkin paths) plus the Lindström, Gessel and Viennot lemma. (A beautiful proof, but it stacks two theorems, each with its own proof, and the path model must be set up with care. Too much to reconstruct in the arena.)

4. **Orthogonal polynomials.** *Cue:* Hankel matrix as a Gram matrix. *Circuit:* three-term recurrence from the continued fraction, $UAU^{\mathsf{T}} = D$ with U unitriangular. (The official solution. Correct and conceptual, but it requires the whole dictionary between continued fractions, convergents, and orthogonality; heavy to justify from scratch.)

5. **Reduction lemma plus self-similarity.** *Cue:* $c_1 = 1$, so f is invertible as a Laurent series; $f^{-1} = x^{-1} - (\text{power series})$. *Circuit:* $h_n(\alpha) = h_{n-1}(\beta)$ when $\alpha^{-1} = x^{-1} - \beta$; iterate on a

series that reproduces itself. (Best. One lemma with a two-paragraph proof, then the algebra $f^{-1} = x^{-1} - 5 - 10f_1$, $f_1^{-1} = x^{-1} - 7 - 10f_1$, then induction.)

The ideal solution

The answer is $\boxed{\det A = 10^{n(n-1)/2}}$.

Notation. For a power series $\gamma = \sum_{i\geq 0} g_i x^i$ and $m \geq 1$ let $H_m(\gamma)$ be the $m \times m$ matrix with entries g_{i+j-1} $(1 \leq i, j \leq m)$, and $h_m(\gamma) = \det H_m(\gamma)$; put $h_0(\gamma) = 1$. Note that $H_m(\gamma)$ involves only $g_1, g_2, \ldots, g_{2m-1}$, so $H_m(\gamma) = H_m(\gamma - g_0)$, and $H_m(\lambda\gamma) = \lambda H_m(\gamma)$, so $h_m(\lambda\gamma) = \lambda^m h_m(\gamma)$. With $f(x) = \sum_{k\geq 0} c_k x^k$, the matrix in the problem is $A = H_n(f)$. All identities below are identities of formal power series (equivalently of convergent series near 0); a series with nonzero coefficient of x and zero constant term is invertible in the field of formal Laurent series.

Step 1: the algebraic equation and the self-similar companion. Write s for the power series $\sqrt{1 - 14x + 9x^2}$ with $s(0) = 1$, so $4f = 1 - 3x - s$ and $s = 1 - 3x - 4f$. Squaring,

$$1 - 6x + 9x^2 - 8f + 24xf + 16f^2 = 1 - 14x + 9x^2,$$

which simplifies to

$$f = x + 3xf + 2f^2. \tag{1}$$

Comparing constant terms, $c_0 = 2c_0^2$, and $f(0) = (1-1)/4 = 0$ gives $c_0 = 0$; comparing coefficients of x then gives $c_1 = 1$. (For $m \geq 2$, (1) gives $c_m = 3c_{m-1} + 2\sum_{i=1}^{m-1} c_i c_{m-i}$, so $f = x + 5x^2 + 35x^3 + 295x^4 + \cdots$.)

Since $c_1 = 1$, define the power series $f_1 = (f - x)/(5x) = x + 7x^2 + 59x^3 + \cdots$; it has zero constant term and coefficient 1 at x, and $f = x + 5xf_1$. Dividing (1) by f gives $1 = xf^{-1} + 3x + 2f$, so

$$f^{-1} = x^{-1} - 3 - 2\frac{f}{x} = x^{-1} - 3 - 2(1 + 5f_1) = x^{-1} - 5 - 10f_1. \tag{2}$$

Substituting $f = x + 5xf_1$ into (1):

$$x + 5xf_1 = x + 3x^2 + 15x^2 f_1 + 2x^2(1 + 5f_1)^2 = x + 5x^2 + 35x^2 f_1 + 50x^2 f_1^2,$$

so $f_1 = x + 7xf_1 + 10xf_1^2$, and dividing by f_1,

$$f_1^{-1} = x^{-1} - 7 - 10f_1. \tag{3}$$

Step 2: the reduction lemma.

Lemma 2.6.1. *Let $\alpha = \sum_{i\geq 1} a_i x^i$ with $a_1 = 1$, and let $\beta = \sum_{i\geq 0} b_i x^i$ be the power series with $\alpha^{-1} = x^{-1} - \beta$. Then $h_n(\alpha) = h_{n-1}(\beta)$ for every $n \geq 1$.*

Proof. First, β is a power series: $\alpha^{-1} = x^{-1}(1 + a_2 x + \cdots)^{-1}$ and $(1 + a_2 x + \cdots)^{-1} - 1$ is divisible by x. The relation $\alpha(x^{-1} - \beta) = 1$ reads $x^{-1}\alpha = 1 + \alpha\beta$; comparing coefficients of x^m for $m \geq 1$,

$$a_{m+1} = \sum_{r=1}^{m} a_r\, b_{m-r}. \tag{4}$$

Let $M = H_n(\alpha)$, so $M_{ij} = a_{i+j-1}$. Let L be the $n \times n$ lower unitriangular matrix with $L_{ii} = 1$, $L_{ik} = -b_{i-1-k}$ for $k < i$, and $L_{ik} = 0$ for $k > i$. Row 1 of LM is row 1 of M, beginning with $a_1 = 1$. For $i \geq 2$ and any j, using (4) with $m = i + j - 2 \geq 1$ and the substitution $r = k + j - 1$,

$$(LM)_{ij} = a_{i+j-1} - \sum_{k=1}^{i-1} b_{i-1-k}\, a_{k+j-1} = \sum_{r=1}^{i+j-2} a_r b_{i+j-2-r} - \sum_{r=j}^{i+j-2} a_r b_{i+j-2-r} = \sum_{r=1}^{j-1} a_r\, b_{i+j-2-r}.$$

In particular $(LM)_{i1} = 0$ for $i \geq 2$ (empty sum). For $i, j \geq 2$ put $i' = i-1$, $j' = j-1$ and $l = j - r$; then

$$(LM)_{ij} = \sum_{l=1}^{j'} b_{i'+l-1}\, a_{j'-l+1} = \big(H_{n-1}(\beta)\, V\big)_{i'j'},$$

where V is the $(n-1)\times(n-1)$ upper unitriangular matrix with $V_{lj'} = a_{j'-l+1}$ for $l \leq j'$ (so $V_{ll} = a_1 = 1$) and 0 below the diagonal. Hence, in block form,

$$LM = \begin{pmatrix} 1 & * \\ 0 & H_{n-1}(\beta)\, V \end{pmatrix},$$

and taking determinants, with $\det L = \det V = 1$, gives $h_n(\alpha) = \det(LM) = \det\big(H_{n-1}(\beta)\big) = h_{n-1}(\beta)$. For $n = 1$ both sides equal 1. □

Step 3: iterate. By (2) and the lemma applied to $\alpha = f$, $\beta = 5 + 10f_1$,

$$h_n(f) = h_{n-1}(5 + 10f_1) = h_{n-1}(10f_1) = 10^{n-1} h_{n-1}(f_1).$$

By (3) and the lemma applied to $\alpha = f_1$ (leading coefficient 1), $\beta = 7 + 10f_1$, for every $m \geq 1$,

$$h_m(f_1) = h_{m-1}(7 + 10f_1) = 10^{m-1} h_{m-1}(f_1),$$

and by induction on m, starting from $h_0(f_1) = 1$, $h_m(f_1) = 10^{(m-1)+(m-2)+\cdots+0} = 10^{m(m-1)/2}$. Therefore

$$\det A = h_n(f) = 10^{n-1} \cdot 10^{(n-1)(n-2)/2} = 10^{n(n-1)/2}.$$

Check: $n = 2$ gives $c_1 c_3 - c_2^2 = 35 - 25 = 10$; $n = 3$ gives 1000.

Layer 1: WHAT

Block 1: square out the root to get $f = x + 3xf + 2f^2$, compute c_1, c_2, c_3 and the determinants $1, 10, 1000$, and conjecture $10^{n(n-1)/2}$. Block 2: recognise the Hankel matrix as $H_n(f)$ and recall that Hankel determinants are computed by peeling off the first row and column, which corresponds to inverting the series: $f^{-1} = x^{-1} - \beta$. Block 3: the reduction lemma $h_n(\alpha) = h_{n-1}(\beta)$, proven by the factorisation $LM = \left(\begin{smallmatrix} 1 & * \\ 0 & H_{n-1}(\beta)V \end{smallmatrix}\right)$ with L and V unitriangular. Block 4: the algebra: $f^{-1} = x^{-1} - 5 - 10f_1$ with $f_1 = (f - x)/(5x)$, and f_1 reproduces itself, $f_1^{-1} = x^{-1} - 7 - 10f_1$. Block 5: iterate and multiply the powers of 10.

Juxtaposition. The nominal route treats A as a numeric matrix: compute $c_1, \ldots, c_{2n-1}$ and eliminate. The ideal route treats the entries as the coefficients of an algebraic series and asks what elimination does to the series: one step of elimination is inversion of the series and removal of its polar part. The bifurcation point is the moment you write f^{-1} instead of another row operation. After it, the self-similarity of f_1 is a two-line computation and the induction is automatic.

Layer 2: HOW

Block 1. Cue. A square root of a quadratic in the generating function. **Why natural.** An expression $\big(u(x) - \sqrt{u(x)^2 - 4bx}\big)/(2b)$ is the solution of $bf^2 - uf + x = 0$ with $f(0) = 0$; here $1 - 14x + 9x^2 = (1 - 3x)^2 - 8x$, which is the discriminant of $2f^2 - (1 - 3x)f + x$. The examiners wrote the closed form to hide the recursion $c_m = 3c_{m-1} + 2\sum c_i c_{m-i}$; squaring recovers it. Small determinants are computed because the shape of the answer decides the method: powers of a single number mean a constant λ and a self-similar continued fraction. **Link to the next block.** A quadratic equation for f with $c_1 = 1$ makes f invertible and makes f^{-1} computable in one line, which is exactly the input of the lemma.

Block 2. Cue. Entries c_{i+j-1}: the matrix is $H_n(f)$. **Why natural.** Hankel matrices are Gram matrices of the monomials $1, x, x^2, \ldots$ under the functional $x^k \mapsto c_{k+1}$, and Gaussian elimination on a Gram matrix is Gram-Schmidt. The first elimination step removes the component along 1; in series language this is passing from f to the series β with $f^{-1} = x^{-1} - \beta$, because the convolution identity $a_{m+1} = \sum a_r b_{m-r}$ says precisely that each row of $H_n(f)$ is a combination of the earlier rows plus a row of $H(\beta)$. **Link to the next block.** You need the lemma in a form you can prove in the arena: a unitriangular factorisation, not a theory of orthogonal polynomials.

Block 3. Cue. A determinant unchanged by adding multiples of earlier rows and later columns. **Why natural.** Multiplying by a lower unitriangular L on the left performs the row operations; the convolution identity makes the first column collapse to $(1, 0, \ldots, 0)^{\mathsf{T}}$ and every other entry of rows $2, \ldots, n$ into a convolution of b's with a's, which is $H_{n-1}(\beta)$ times an upper unitriangular V. Writing the factorisation rather than performing operations one at a time is what keeps the proof checkable: every entry is an explicit finite sum with matching indices. **Link to the next block.** The lemma is only useful if β is again of the form "constant plus a multiple of a series with leading coefficient 1", which is the self-similarity question.

Block 4. Cue. $f^{-1} = x^{-1} - 5 - 10f_1$ with $f_1 = (f - x)/(5x)$. **Why natural.** The constant term of β is irrelevant to $H_{n-1}(\beta)$, and the factor 10 comes out as 10^{n-1}. Whether f_1 is again reducible depends on its own equation; substituting $f = x + 5xf_1$ into the quadratic gives $f_1 = x + 7xf_1 + 10xf_1^2$, whose form $g = x + b'xg + \lambda xg^2$ is exactly the self-similar one: $g^{-1} = x^{-1} - b' - \lambda g$. The number $10 = 2 \cdot 5 = b(a + b)$ appears at every level; the number $7 = a + 2b$ never matters. **Link to the next block.** Self-similarity turns the lemma into a first-order recursion $h_m(f_1) = 10^{m-1} h_{m-1}(f_1)$.

Block 5. Cue. A recursion $h_m = \lambda^{m-1} h_{m-1}$ with $h_0 = 1$. **Why natural.** Its closed form is $\lambda^{0+1+\cdots+(m-1)} = \lambda^{m(m-1)/2}$; the extra factor 10^{n-1} from the first step combines to $10^{n(n-1)/2}$ because $(n-1) + (n-1)(n-2)/2 = n(n-1)/2$. Verify against $n = 2, 3$. **Link to the next block.** None.

Layer 3: WHY

The truth beneath the route is the dictionary between three objects: the Hankel matrix $H_n(f)$, the continued fraction of f, and the orthogonal polynomials of the functional $x^k \mapsto c_{k+1}$. Because the functional satisfies $L(xp \cdot q) = L(p \cdot xq)$, Gram-Schmidt on $1, x, x^2, \ldots$ is a three-term recurrence, its pivots are $\lambda_1 \lambda_2 \cdots \lambda_k$ up to sign, and the determinant is the product of the pivots, $\prod_{k=1}^{n-1} \lambda_k^{n-k}$. Unrolling the continued fraction one level is the identity $f^{-1} = x^{-1} - b_0 - \lambda_1 f_1$, and the reduction lemma is the statement that one level of the continued fraction is one step of elimination. A periodic continued fraction is the same thing as a quadratic equation for f, and every sequence of the form $c_m = ac_{m-1} + b \sum c_i c_{m-i}$ has one: after a single non-periodic step ($b_0 = a + b$, $\lambda_1 = b(a + b)$) the fraction repeats with $b_k = a + 2b$ and $\lambda_k = b(a + b)$. So the determinant is $(b(a + b))^{n(n-1)/2}$; the problem is $(a, b) = (3, 2)$.

The discernment mechanism: "Hankel" triggers "continued fraction denominators", "square root of a quadratic" triggers "quadratic equation, hence periodic fraction, hence constant λ", and the small determinants confirm λ. Then choose the proof engine you can reproduce under pressure: the reduction lemma, because it is one factorisation. You train it by running the lemma on the Catalan, Motzkin and Schröder sequences until the algebra $f \mapsto f_1$ is a reflex, and by memorising the general formula $(b(a + b))^{n(n-1)/2}$ as a check. The examiners nudged with c_{i+j-1} (starting at $c_1 = 1$, which makes f invertible with leading coefficient 1), with the closed form (whose discriminant $(1 - 3x)^2 - 8x$ is the fingerprint of $a = 3$, $b = 2$), and with "find" rather than "show that", which tells you to compute small cases first.

Knowledge you must install before this problem is approachable, stated without disguise: the definition of a Hankel matrix; the continued fraction (J-fraction) of a power series and the theorem

$\det(\mu_{i+j})_0^{n-1} = \prod \lambda_k^{n-k}$; the reduction lemma as its proof engine; and the fact that Catalan-type recursions give periodic fractions.

Revolving concepts

Hankel matrices and moment functionals. $H_n(\gamma) = (g_{i+j-1})$ depends only on $i+j$. Define $L(x^k) = g_{k+1}$; then $H_n(\gamma)$ is the Gram matrix $(L(x^{i-1}x^{j-1}))$ of the monomials, and $L(xp \cdot q) = L(p \cdot xq)$. Gram-Schmidt produces monic polynomials $B_0, B_1, \ldots$ with $L(B_iB_j) = 0$ for $i \neq j$, satisfying $B_{k+1} = (x - b_k)B_k - \lambda_k B_{k-1}$, and $UH_nU^{\mathsf{T}} = \operatorname{diag}(L(B_0^2), \ldots, L(B_{n-1}^2))$ with U unitriangular, $L(B_k^2) = \lambda_1 \cdots \lambda_k$. See Book 1, field: Matrices and determinants, topic: Hankel matrices and orthogonal polynomials.

J-fractions. $\sum_{k\geq 0} \mu_k x^k = 1/(1 - b_0x - \lambda_1x^2/(1 - b_1x - \lambda_2x^2/(1 - \cdots)))$, with $\mu_0 = 1$. Flajolet's theorem: μ_k is the total weight of Motzkin paths of length k with level steps at height i weighted b_i and down steps from height i weighted λ_i. Hankel theorem: $\det(\mu_{i+j})_{i,j=0}^{n-1} = \prod_{k=1}^{n-1} \lambda_k^{n-k}$, proven by orthogonal polynomials, by the Lindström, Gessel and Viennot lemma on nonintersecting paths, or by iterating the reduction lemma. In the notation of this problem $\mu_k = c_{k+1}$. See Book 1, field: Generating functions, topic: continued fractions and Hankel determinants.

The reduction lemma. If $\alpha = x + a_2x^2 + \cdots$ and $\alpha^{-1} = x^{-1} - \beta$, then $h_n(\alpha) = h_{n-1}(\beta)$. Unrolling: $\beta = b_0 + \lambda_1\alpha_1$ with α_1 monic, so $h_n(\alpha) = \lambda_1^{n-1}h_{n-1}(\alpha_1)$, and iterating gives the Hankel theorem. See Book 1, field: Matrices and determinants, topic: block unitriangular factorisation.

Catalan-type recursions: the generalised recursion. Let $c_1 = 1$ and $c_m = ac_{m-1} + b\sum_{i=1}^{m-1} c_ic_{m-i}$ for $m \geq 2$. The generating function satisfies $f = x + axf + bf^2$, so $f = \left(1 - ax - \sqrt{(1-ax)^2 - 4bx}\right)/(2b)$. With $f_1 = (f - x)/((a + b)x)$ one gets $f^{-1} = x^{-1} - (a + b) - b(a + b)f_1$ and $f_1 = x + (a + 2b)xf_1 + b(a + b)xf_1^2$, so $f_1^{-1} = x^{-1} - (a + 2b) - b(a + b)f_1$. Hence $\det(c_{i+j-1})_{n\times n} = (b(a + b))^{n(n-1)/2}$. Cases: $(a, b) = (0, 1)$ Catalan, determinant 1; $(1, 1)$ large Schröder, determinant $2^{n(n-1)/2}$; $(3, 2)$ this problem, determinant $10^{n(n-1)/2}$. The general method for any recursion with a convolution term: form the generating function, turn the convolution into f^2, solve the quadratic, and for Hankel questions invert the series rather than expand it. See Book 1, field: Sequences and recurrences, topic: convolution recursions and algebraic generating functions.

Formal Laurent series. A power series with zero constant term and nonzero coefficient of x is invertible in the field of formal Laurent series, with inverse x^{-1} times a power series; all manipulations in the solution are identities there, and convergence for small x is a bonus, not a requirement. See Book 1, field: Generating functions, topic: formal power series.

Circuit Bank

- **WHEN** a generating function is given as a square root of a quadratic **DO** square it out into $f = x + axf + bf^2$ and read off the convolution recursion **BECAUSE** the closed form hides the recursion, and the recursion (equivalently the quadratic) carries all the structure
- **WHEN** a matrix has entries indexed by $i + j$ **DO** say "Hankel", and write the determinant as $\prod_{k=1}^{n-1} \lambda_k^{n-k}$ with λ_k the continued fraction denominators **BECAUSE** Gaussian elimination on a Hankel matrix is the orthogonal polynomial recursion and its pivots are the λ's
- **WHEN** you must prove a Hankel determinant from scratch **DO** invert the series, $f^{-1} = x^{-1} - \beta$, and apply $h_n(f) = h_{n-1}(\beta)$ **BECAUSE** one level of the continued fraction is one step of elimination, and the lemma is a single unitriangular factorisation
- **WHEN** the small determinants are $1, \lambda, \lambda^3, \lambda^6, \ldots$ **DO** look for self-similarity: a series g with $g^{-1} = x^{-1} - b' - \lambda g$ **BECAUSE** constant λ means a periodic continued fraction, which is the same as a quadratic equation for g
- **WHEN** a sequence satisfies $c_m = ac_{m-1} + b\sum c_ic_{m-i}$ with $c_1 = 1$ **DO** write $\det(c_{i+j-1}) = (b(a + b))^{n(n-1)/2}$ **BECAUSE** one reduction step gives $\lambda_1 = b(a + b)$ and the companion series is self-similar with the same λ

Backtesting

1. First useful cue: $\sqrt{1-14x+9x^2}$ in the generating function, and entries c_{i+j-1}.
2. Method the expression suggested: the quadratic $f = x + 3xf + 2f^2$, then the Hankel reduction through f^{-1}.
3. Where you would hesitate, and the thought to reinforce: in the lemma, which operations to perform. The thought: row i minus the convolution of the earlier rows with $b_{i-2}, b_{i-3}, \ldots, b_0$; the recursion $a_{m+1} = \sum a_r b_{m-r}$ is built to make the first column vanish, and what remains in the block is $H_{n-1}(\beta)$ times a unitriangular matrix.
4. The trap the question sets: computing coefficients and expanding; also the index shift: A starts at c_1, not c_0, which is what makes f invertible with leading coefficient 1 and makes the lemma apply directly.
5. What makes you recognise it faster next time: any Hankel determinant of a Catalan-like sequence is a power of $b(a+b)$; the discriminant $(1-ax)^2 - 4bx$ is the fingerprint.
6. The reusable decision rule: Hankel plus algebraic generating function means invert the series and iterate the reduction lemma; never expand the determinant.

Mirror drills

1. Let C_k be the Catalan numbers, $C(x) = 1 + xC(x)^2$. Show that $\det(C_{i+j})_{i,j=0}^{n-1} = 1$. *Answer:* $\alpha = xC$ satisfies $\alpha = x + \alpha^2$, so $\alpha^{-1} = x^{-1} - 1 - f_1$ with $f_1 = xC^2 = x + 2x^2 + \cdots$, and $f_1 = x + 2xf_1 + xf_1^2$ gives $f_1^{-1} = x^{-1} - 2 - f_1$; every λ equals 1.
2. Let M_k be the Motzkin numbers, $M(x) = 1 + xM + x^2M^2$. Show that $\det(M_{i+j})_{i,j=0}^{n-1} = 1$. *Answer:* $g = xM$ satisfies $g = x + xg + xg^2$, which is already self-similar: $g^{-1} = x^{-1} - 1 - g$, so $h_m(g) = h_{m-1}(g)$ and all determinants equal 1.
3. Let S_k be the large Schröder numbers $1, 2, 6, 22, 90, \ldots$, $S(x) = 1 + xS + xS^2$. Show that $\det(S_{i+j})_{i,j=0}^{n-1} = 2^{n(n-1)/2}$. *Answer:* $g = xS$ satisfies $g = x + xg + g^2$, the case $(a, b) = (1, 1)$, so $\lambda = b(a+b) = 2$; check $S_0S_2 - S_1^2 = 6 - 4 = 2$.

2.7 B1. Transversal sums force the middle offset

Let n and k be positive integers. The square in the ith row and jth column of an n-by-n grid contains the number $i + j - k$. For which n and k is it possible to select n squares from the grid, no two in the same row or column, such that the numbers contained in the selected squares are exactly $1, 2, \ldots, n$?

Read it out loud

An n-by-n grid whose entry depends only on $i + j$, so the entries are constant along anti-diagonals and increase by one as you step down or right. You must choose n squares with no two sharing a row or column: in other words, a permutation σ of $\{1, \ldots, n\}$, the square in row i being column $\sigma(i)$. The selected numbers must be exactly $1, 2, \ldots, n$, a fixed multiset with a fixed sum.

The fields wired in are combinatorics (a permutation is the object being selected) and the arithmetic of sums. What is being tested is whether you notice that the entry $i + \sigma(i) - k$ has a sum over i that does not depend on σ at all, because $\sum_i \sigma(i)$ is the same for every permutation. That single sentence is the whole of the necessity direction, and it is the first thing you should write down, within thirty seconds: $\sum_i (i + \sigma(i) - k) = n(n+1) - nk$.

The command terms are "no two in the same row or column" (permutation) and "exactly $1, 2, \ldots, n$" (a fixed sum, $n(n+1)/2$). When a question fixes a multiset of outputs and lets you vary a permutation, the circuit is: compare an invariant of the outputs with the same invariant computed structurally. The second half is a construction: a permutation σ such that $i + \sigma(i)$ takes n consecutive values, which is a complete mapping of $\mathbb{Z}_n$.

Approaches, ranked worst to best

1. **Small cases and pattern guessing.** *Cue:* the question says "for which n and k". *Circuit:* enumerate $n = 1, 2, 3, 4$ by hand and conjecture. (it costs a page of case work, it never produces a proof of impossibility for even n, and the pattern "n odd, $k = (n+1)/2$" is only visible once you already suspect a sum argument; ranks last because it is consumption without analysis)

2. **Row-by-row greedy.** *Cue:* each row is a run of n consecutive integers $i + 1 - k, \ldots, i + n - k$. *Circuit:* assign the value 1 to a row that contains it, then 2, and so on. (it stalls immediately, because the constraint "no two in the same column" is global and a greedy assignment collides; it also gives no necessity argument)

3. **Sum invariant plus explicit transversal.** *Cue:* the entry is $g(i) + h(j)$ with $g(i) = i, h(j) = j - k$, so any transversal has the same sum. *Circuit:* compare the forced sum $n(n+1) - nk$ with the required sum $n(n+1)/2$, then build σ with $i + \sigma(i) - k$ hitting $1, 3, 5, \ldots$ then $2, 4, \ldots$. (two lines for necessity, one displayed permutation for sufficiency; this is the ideal)

The ideal solution

The selection is possible if and only if n is odd and $k = \dfrac{n+1}{2}$.

Necessity. A choice of n squares with no two in the same row or column is a permutation σ of $\{1, \ldots, n\}$: the chosen square in row i sits in column $\sigma(i)$. The chosen numbers are $i + \sigma(i) - k$ for $i = 1, \ldots, n$. Since σ is a bijection of $\{1, \ldots, n\}$ onto itself, $\sum_{i=1}^{n} \sigma(i) = \sum_{i=1}^{n} i = \frac{n(n+1)}{2}$, so

$$\sum_{i=1}^{n} \big(i + \sigma(i) - k\big) = \frac{n(n+1)}{2} + \frac{n(n+1)}{2} - nk = n(n+1) - nk.$$

If the chosen numbers are exactly $1, 2, \ldots, n$ their sum is $\frac{n(n+1)}{2}$. Equating, $nk = \frac{n(n+1)}{2}$, so $k = \frac{n+1}{2}$; as k is an integer, n is odd.

Sufficiency. Let n be odd and $k = \frac{n+1}{2}$. Define

$$\sigma(i) = \begin{cases} i + \frac{n-1}{2}, & 1 \le i \le \frac{n+1}{2}, \\ i - \frac{n+1}{2}, & \frac{n+3}{2} \le i \le n. \end{cases}$$

The first branch takes the values $\frac{n+1}{2}, \frac{n+3}{2}, \ldots, n$ and the second takes $1, 2, \ldots, \frac{n-1}{2}$, each exactly once, so σ is a permutation of $\{1, \ldots, n\}$ and the squares $(i, \sigma(i))$ lie in distinct rows and distinct columns. The selected numbers are

$$i + \sigma(i) - k = \begin{cases} 2i - 1, & 1 \le i \le \frac{n+1}{2}, \\ 2i - (n+1), & \frac{n+3}{2} \le i \le n, \end{cases}$$

which run through $1, 3, 5, \ldots, n$ and then $2, 4, \ldots, n-1$: exactly the numbers $1, 2, \ldots, n$. (For $n = 1$ the single square contains $1 + 1 - 1 = 1$, consistent with the formula.) □

Layer 1: WHAT

Block 1: translate "no two in the same row or column" into a permutation σ, so the chosen numbers are $i + \sigma(i) - k$. Block 2: sum the chosen numbers structurally; $\sum \sigma(i)$ is permutation-independent, so the sum is $n(n+1) - nk$ whatever you choose. Block 3: equate with $\sum_{v=1}^{n} v$ to force $k = (n+1)/2$ and hence n odd. Block 4: construct σ so that $i + \sigma(i) - k$ hits the odd numbers on the first $\frac{n+1}{2}$ rows and the even numbers on the rest; verify it is a bijection.

Nominal route: try to place the numbers $1, \ldots, n$ one at a time, or grind small grids, and hope the pattern reveals itself. Ideal route: refuse to place anything until you know what is forced. The bifurcation point is the first line after reading: the zombie writes out a 4-by-4 grid; you write $\sum_i (i + \sigma(i) - k)$ and notice σ disappears.

Layer 2: HOW

Block 1. Cue. "Select n squares, no two in the same row or column."

Why natural. This is the definition of a permutation matrix (a transversal). Naming the object σ is what lets you write a formula for the selected entries instead of a picture.

Link to the next block. Once the entries are $i + \sigma(i) - k$, the additive separation of i from $\sigma(i)$ begs to be summed.

Block 2. Cue. The entry is a sum of a function of the row and a function of the column.

Why natural. For $M_{ij} = g(i) + h(j)$ the transversal sum $\sum_i M_{i\sigma(i)} = \sum_i g(i) + \sum_j h(j)$ is independent of σ, because σ merely reorders the h values. This is a definition-level fact about bijections, not an observation you need luck for.

Link to the next block. An invariant that does not depend on the free choice must be matched against the fixed target; the target is the multiset $\{1, \ldots, n\}$.

Block 3. Cue. "Exactly $1, 2, \ldots, n$" is a fixed multiset; its sum is $n(n+1)/2$.

Why natural. Equating the two sums is the only equation available at this point, and it happens to pin k completely. Parity of n falls out because k is an integer; you did not look for a parity argument, the equation handed it to you.

Link to the next block. Necessity leaves one candidate pair (n odd, $k = (n+1)/2$) and the question "for which" obliges you to construct.

Block 4. Cue. You need $i + \sigma(i)$ to run through n consecutive integers as σ runs through a permutation.

Why natural. If $i + \sigma(i) - k = \tau(i)$ for permutations σ, τ, then $\sigma(i) = \tau(i) - i + k$, so $\tau(i) - i$ must itself be injective modulo nothing: the simplest injective choice is $\tau(i) - i = i - 1$, that is $\tau(i) = 2i - 1$, the doubling map, which is a bijection of $\mathbb{Z}_n$ precisely when n is odd. Wrapping the doubling map around n gives the two branches of σ.

Link to the next block. The construction is verified by listing the values on each branch; nothing further is needed.

Layer 3: WHY

The underlying truth is that for matrices of the form $g(i) + h(j)$ every transversal has the same sum, so "which transversal" is invisible to the sum and only the distribution of values is negotiable. The examiners chose entries $i + j - k$ so that this invariant determines k exactly, and the target $\{1, \ldots, n\}$ so that its sum is the cleanest expression available.

The discernment mechanism: whenever a problem lets you choose a permutation and asks for a prescribed set of outputs, compute every permutation-independent statistic of the outputs before choosing anything: sum, product, sum of squares, residues modulo small numbers, in that order of cheapness. Here the sum suffices, which is why this is a B1.

The construction half needs one small object: a permutation σ of $\mathbb{Z}_n$ with $i + \sigma(i)$ also a permutation (a complete mapping). It exists if and only if n is odd, and for odd n the identity works modulo n because 2 is invertible. Once you know that, sufficiency is not a search but a translation of "multiply by 2 modulo n" into integers; the two-branch formula for σ is the unwrapping of the modular shift. Train this by asking, for every construction you produce, which modular fact made it exist.

Revolving concepts

Transversal of a matrix. A set of n cells with one in each row and one in each column; equivalently a permutation σ with cells $(i, \sigma(i))$. Orbiting facts: the sum of a transversal of $g(i) + h(j)$ is $\sum g + \sum h$ for every σ; the product of a transversal of $g(i)h(j)$ is $\prod g \prod h$; the determinant is a signed sum over all transversals of the product, which is why $\det(g(i) + h(j)) = 0$ for $n \geq 3$ (rank at most two). See Book 1, field: Combinatorics, topic: transversals and permutation matrices.

Complete mapping of a group. A bijection σ of a finite group G such that $x \mapsto x\,\sigma(x)$ is also a bijection. For $G = \mathbb{Z}_n$ one exists if and only if n is odd (the identity map works when 2 is a unit; for even n summing $x + \sigma(x)$ over G gives a contradiction modulo n, exactly the sum argument of this problem). Equivalent statement: the Cayley table of $\mathbb{Z}_n$ has a transversal if and only if n is odd. See Book 1, field: Abstract algebra, topic: Latin squares and complete mappings.

Recursion, generalised. There is no recursion in the solution, but the invariant generalises: for any matrix $M_{ij} = g(i) + h(j)$ and any target multiset V, a transversal realising V exists only if $\sum V = \sum g + \sum h$; for $M_{ij} = g(i)h(j)$ replace sums by products. For a matrix of the form $g(i) + h(j)$ with g, h arithmetic progressions of common difference d, transversals realising a run of n consecutive multiples of d exist exactly when n is odd, by the same doubling construction.

Circuit Bank

- **WHEN** a selection of one cell per row and column **DO** name the permutation σ and write the selected entries as $M_{i\sigma(i)}$ **BECAUSE** a formula can be summed, a picture cannot
- **WHEN** an entry of the form $g(i) + h(j)$ with a prescribed output multiset **DO** sum over the transversal and equate with the sum of the multiset before constructing anything **BECAUSE** the transversal sum is permutation-independent, so it is a free necessary condition

- **WHEN** you need $i + \sigma(i)$ to be a permutation **DO** use $\sigma(i) \equiv i + c \pmod{n}$ with n odd, then unwrap into two integer branches **BECAUSE** doubling is a bijection of $\mathbb{Z}_n$ exactly when n is odd
- **WHEN** a question says "for which n" and your invariant forces a parity **DO** expect a construction that uses that same parity in an invertibility step **BECAUSE** the necessary condition and the construction usually rest on the same modular fact

Backtesting

1. First useful cue: "no two in the same row or column" together with an entry depending on $i + j$.
2. Method the expression suggested: sum the selected entries symbolically; the permutation drops out.
3. Where you would hesitate, and the thought to reinforce: at the construction, wondering whether to search. The thought: I need $i + \sigma(i)$ to be a permutation, which is the doubling map on $\mathbb{Z}_n$; write $\sigma(i) \equiv i + \frac{n-1}{2}$ and unwrap.
4. The trap the question sets: fiddling with small grids and believing k is free; the sum shows k is rigid.
5. What makes you recognise it faster next time: the phrase "exactly $1, 2, \ldots, n$" is a fixed multiset, and fixed multisets have fixed sums; compare with the structural sum immediately.
6. The reusable decision rule: when the free choice is a permutation and the target is fixed, compute permutation-invariant statistics first; construct second, and construct with a modular shift.

Mirror drills

1. The cell (i, j) of an n-by-n grid contains ij. For which n can you pick a transversal whose entries are exactly $1^2, 2^2, \ldots, n^2$? *Answer:* the product of any transversal is $(n!)^2$, which equals $\prod v^2$, so the product invariant is silent; but $\sum_i i\sigma(i) = \sum i^2$ forces σ to be the identity by the rearrangement inequality (equality only when σ is order-preserving), so it is always possible, uniquely, for every n.
2. Show that the addition table of $\mathbb{Z}_n$ (entry $i + j \bmod n$) has no transversal when n is even. *Hint:* sum $i + \sigma(i)$ over all i modulo n: the left side is $2 \cdot \frac{n(n-1)}{2} \equiv 0$ but a transversal hitting every residue once sums to $\frac{n(n-1)}{2} \equiv \frac{n}{2} \not\equiv 0$.
3. The cell (i, j) contains $i + 2j$. Find all n for which some transversal has entries forming n consecutive integers. *Answer:* the sum is $3 \cdot \frac{n(n+1)}{2}$ and n consecutive integers starting at s sum to $ns + \frac{n(n-1)}{2}$, forcing $s = n + 2$ for every n; the anti-diagonal $\sigma(i) = n + 1 - i$ gives entries $2n + 2 - i$, which are exactly $n + 2, \ldots, 2n + 1$, so every n works. (Contrast with $i + j$, where the anti-diagonal is constant; the coefficient gap between i and $2j$ is what makes $\sigma(i) = n + 1 - i$ injective on the entries.)

2.8 B2. A fixed centre and two concentric circles

Two convex quadrilaterals are called partners *if they have three vertices in common and they can be labeled $ABCD$ and $ABCE$ so that E is the reflection of D across the perpendicular bisector of the diagonal $\overline{AC}$. Is there an infinite sequence of convex quadrilaterals such that each quadrilateral is a partner of its successor and no two elements of the sequence are congruent?*

Read it out loud

A move on convex quadrilaterals: keep three vertices, and replace the fourth by its mirror image in the perpendicular bisector of the diagonal joining its two neighbours. The question asks whether iterating the move can produce infinitely many pairwise non-congruent shapes. So this is a dynamics question in disguise: a move, an orbit, and the word "congruent", which tells you the answer must be phrased in terms of invariants under the move that survive up to congruence.

The fields are Euclidean geometry (reflections, perpendicular bisectors as loci of equidistance, convexity forcing the diagonals to cross) and the finiteness logic of combinatorics (finitely many values of an invariant, finitely many states). What is really being tested is whether you extract from the definition of the move the objects it fixes. A reflection fixes every point of its axis; the axis is the perpendicular bisector of AC; so anything defined from that axis alone is fixed by the move.

The first useful cue is the phrase "reflection across the perpendicular bisector of $\overline{AC}$". Within thirty seconds you should have written: the reflection swaps A and C, so $CE = AD$ and $AE = CD$; the side lengths are permuted, the diagonal AC and the area are kept. Then the harder question: what is not merely permuted but literally fixed? A point on the axis defined symmetrically from both diagonals: the intersection of their perpendicular bisectors. Expect the answer "no".

Approaches, ranked worst to best

1. **Coordinates and algebraic functions.** *Cue:* a reflection has an explicit formula. *Circuit:* parametrise the quadrilateral, fix the invariants, and argue that a nonconstant algebraic function takes each value finitely often. (it needs the sides and area invariants first anyway, then a page of coordinates and a fact about algebraic functions that you cannot prove in the arena; it ranks last on ink and on rigour risk)

2. **Side lengths and area, then Bretschneider.** *Cue:* $CE = AD$, $AE = CD$, and area is preserved because $ACE \cong ACD$. *Circuit:* a convex quadrilateral with given ordered sides is determined by one diagonal f; Bretschneider's formula $K^2 = (s-a)(s-b)(s-c)(s-d) - abcd\cos^2\frac{\alpha+\gamma}{2}$ fixes $\cos^2\frac{\alpha+\gamma}{2}$, and $\alpha+\gamma = 2\pi - (\beta+\delta)$ is strictly decreasing in f because the angles β, δ opposite the diagonal f increase with f. (valid and short once you own Bretschneider, but that formula is a knowledge gap for you, and the monotonicity step is easy to get wrong: each of the two angles at A separately is not monotone in f, only the sum of the two angles opposite f is; ranks second)

3. **Fixed centre and two concentric circles.** *Cue:* the axis of the reflection is a perpendicular bisector, so the move fixes any point on it; the perpendicular bisectors of both diagonals meet at a point O because the diagonals of a convex quadrilateral cross. *Circuit:* O, the two radii $OA = OC$ and $OB = OD$, and the multiset of side lengths are all invariant; each side subtends an angle at O determined by its length, so there are finitely many shapes up to rotation. (a half page, nothing beyond the definition of a reflection and the cosine rule; the ideal)

The ideal solution

No such sequence exists.

Step 1: the centre. Let $ABCD$ be a convex quadrilateral. Its diagonals AC and BD cross at an interior point, so they are not parallel, so their perpendicular bisectors are not parallel and meet in exactly one point O. Call O the *centre* of the quadrilateral. By construction $OA = OC$ and $OB = OD$.

Step 2: the centre and the radii are invariant under the move. Let $ABCE$ be a partner of $ABCD$, with E the reflection of D across the perpendicular bisector ℓ of AC; write ρ for this reflection. Since $O \in \ell$, $\rho(O) = O$, and reflections preserve distance, so

$$OE = |\rho(O)\rho(D)| = OD = OB.$$

Hence O lies on the perpendicular bisector of BE and on ℓ. The quadrilateral $ABCE$ is convex, so by Step 1 these two lines meet in exactly one point, which must be O: the centre of $ABCE$ is the centre of $ABCD$. Moreover A and C keep their common distance $r_1 := OA = OC$ to O, and the pair $\{B, E\}$ has the same common distance $r_2 := OB = OD$ as the pair $\{B, D\}$ had. Note $r_1 > 0$, since $r_1 = 0$ would force $A = O = C$, and likewise $r_2 > 0$.

Step 3: the side lengths are invariant as a multiset. Since ρ swaps A and C and sends D to E,

$$CE = |\rho(A)\rho(D)| = AD, \qquad AE = |\rho(C)\rho(D)| = CD,$$

so the sides of $ABCE$, namely AB, BC, CE, EA, form the same multiset as the sides AB, BC, CD, DA of $ABCD$.

Step 4: finiteness. Consider any sequence $Q_1, Q_2, \ldots$ of convex quadrilaterals in which each is a partner of its successor. By Steps 2 and 3 applied along the sequence, all Q_i share the same centre O, the same two radii $r_1, r_2 > 0$ (with the endpoints of one diagonal on the circle of radius r_1 about O and the endpoints of the other diagonal on the circle of radius r_2), and the same finite set S of side lengths, $|S| \le 4$.

Place O at the origin. Consecutive vertices of any Q_i are endpoints of different diagonals, so one of them is $r_1(\cos\theta, \sin\theta)$ and the other is $r_2(\cos\varphi, \sin\varphi)$ for some polar angles θ, φ, and by the cosine rule the side joining them has length s with

$$s^2 = r_1^2 + r_2^2 - 2r_1 r_2 \cos(\theta - \varphi).$$

Since $r_1 r_2 \neq 0$, this determines $\cos(\theta - \varphi)$, and therefore $\theta - \varphi \equiv \pm\delta(s) \pmod{2\pi}$ with $\delta(s) = \arccos \frac{r_1^2 + r_2^2 - s^2}{2r_1 r_2} \in [0, \pi]$. Thus every step from a vertex to the next vertex changes the polar angle by an element of the finite set $G = \{\pm\delta(s) : s \in S\}$, $|G| \le 8$. Start at a vertex on the circle of radius r_1 and go round the quadrilateral: the radii alternate r_1, r_2, r_1, r_2, and the polar angles of the other three vertices are the starting angle plus partial sums of three elements of G. Rotating about O so that the starting vertex has polar angle 0 is a congruence, so every Q_i is congruent to one of at most $|G|^3 \le 512$ quadrilaterals. A sequence of pairwise non-congruent quadrilaterals therefore has at most 512 terms, and no infinite sequence with the required property exists. □

A caveat worth clarifying: the partner relation is symmetric (D is the reflection of E across the same line), so "each is a partner of its successor" and "each is a partner of its predecessor" give the same sequences, and the invariance argument runs in whichever direction you like.

Layer 1: WHAT

Block 1: read the move as a reflection ρ in the perpendicular bisector ℓ of AC; list what ρ does to the four vertices ($A \leftrightarrow C$, $D \mapsto E$, B untouched). Block 2: find a point fixed by every possible move: the intersection O of the perpendicular bisectors of the two diagonals, which exists because convexity makes the diagonals cross. Block 3: show O is the centre of the partner too, and that the distances from O to the diagonal endpoints are preserved. Block 4: record the cheap invariant, the multiset of side lengths. Block 5: convert "vertices on two concentric circles with sides of prescribed lengths" into "angular steps from a finite set", count, conclude finiteness.

Nominal route: write the quadrilateral in coordinates, apply the reflection formula, try to find what is preserved by inspection of the algebra, then fight an infinite-versus-finite question with algebraic geometry. Ideal route: ask what a reflection fixes (its axis) and what a perpendicular bisector is (an equidistance locus), and let the fixed point O emerge from definitions. The bifurcation point is Block 2: the zombie asks "what does the move do to the shape?" and you ask "what does the move leave alone?".

Layer 2: HOW

Block 1. **Cue.** "E is the reflection of D across the perpendicular bisector of $\overline{AC}$."

Why natural. A reflection is an isometry, so every distance between images equals the distance between preimages; and the perpendicular bisector of AC is exactly the axis that swaps A and C. Writing $CE = |\rho(A)\rho(D)| = AD$ is therefore not an idea, it is the definition read twice.

Link to the next block. You now know which lengths survive the move. The next question is which *points* survive, because a fixed point turns an orbit into a rigid picture.

Block 2. **Cue.** The axis ℓ depends on the diagonal AC; a different move (reflecting B, say) uses the other diagonal. You want something fixed by both.

Why natural. Points of ℓ are fixed by ρ. Among them, the one you can name without reference to which vertex is being moved is the intersection of the perpendicular bisectors of *both* diagonals. It exists because the diagonals of a convex quadrilateral cross and so are not parallel; this is where the hypothesis "convex" is spent, and you should notice that you needed it.

Link to the next block. Having named O, you must check it is still the centre after the move, otherwise "fixed" is only a hope.

Block 3. **Cue.** $O \in \ell$ and $OB = OD$.

Why natural. $\rho(O) = O$ and $\rho(D) = E$, so $OE = OD = OB$: O is equidistant from B and E, hence on the perpendicular bisector of the new diagonal BE, and it is on ℓ, the perpendicular bisector of the old diagonal AC. Uniqueness of the intersection (convexity of the partner) then identifies the new centre with O. The radii are preserved for the same reason: the replaced vertex has the same distance to O as the vertex it replaced.

Link to the next block. With O, r_1, r_2 fixed, the vertices of every quadrilateral in the sequence live on two concentric circles, and the only remaining freedom is angular.

Block 4. **Cue.** The multiset of side lengths came free in Block 1.

Why natural. A finite set of admissible side lengths is what you need to make the angular freedom finite. Without it the vertices could slide around the circles.

Link to the next block. Combine: a side joins a point on the circle of radius r_1 to a point on the circle of radius r_2; the cosine rule turns its length into the angle it subtends at O.

Block 5. **Cue.** Two concentric circles, a chord between them of known length, centre O.

Why natural. The cosine rule in triangle OPQ with $OP = r_1$, $OQ = r_2$ gives $\cos \angle POQ$ from PQ, so the angular step is $\pm\delta(s)$: a finite set. Four vertices, three steps, at most 8^3 configurations up to rotation. Congruence classes are finite, so the sequence is finite.

Link to the next block. Nothing remains but to state the count and the conclusion.

Layer 3: WHY

The mathematical truth underneath is that a reflection in a line through a point O preserves distance from O, so a sequence of reflections in lines through a common point moves every vertex around a

circle centred at O. The move in this problem is engineered so that all the axes pass through one point: each axis is the perpendicular bisector of a diagonal, and O is on the perpendicular bisectors of both diagonals, and stays there because the new diagonal BE is still bisected perpendicularly by a line through O. The examiners nudged toward this by choosing the perpendicular bisector (an equidistance locus) as the mirror, rather than an arbitrary line through the midpoint of AC.

The second truth is that "non-congruent" is a statement about orbits modulo isometry, and to bound an orbit modulo isometry you fix a point and reduce to rotations about it. Once O is fixed, congruence is rotation, and rotation is one real parameter, killed by the choice of starting angle.

The discernment mechanism: for any "move" problem, ask three questions in order. What does the move preserve (lengths, areas, multisets)? What does it fix pointwise (points on the axis, centres)? What finite data determines the state up to the allowed equivalence? The first question is where the nominal route stops; the second is where the ideal route lives. Train it by practising on reflections, rotations and inversions: for each, write the fixed set before writing anything else. It is worth repeating: a reflection is defined by its axis, so its fixed set is the first object to name.

On the Bretschneider route, one caveat: the monotonicity of the angle sum in the diagonal must be argued for the pair of angles *opposite* the diagonal (each increases with it by the cosine rule), not for the two angles at one vertex individually, which are not monotone.

Revolving concepts

Reflection in a line. The isometry fixing the line pointwise and swapping the half-planes. Orbiting facts: isometries preserve all distances, hence all congruence classes of sub-figures; a reflection in a line through O maps the circle of radius r about O to itself; the composition of two reflections in lines through O is a rotation about O. The perpendicular bisector of PQ is the locus of points equidistant from P and Q, and equivalently the axis of the unique reflection swapping P and Q. See Book 1, field: Geometry, topic: isometries and their fixed sets.

Convex quadrilateral. A quadrilateral whose diagonals intersect at an interior point. Orbiting facts: the diagonal splits it into two triangles whose areas add; the diagonals are not parallel, so their perpendicular bisectors meet; the vertices are in convex position. This is the only place convexity is used, and it is essential: for a self-intersecting quadrilateral the diagonals can be parallel and no centre exists. See Book 1, field: Geometry, topic: convexity.

Bretschneider's formula. For a quadrilateral with sides a, b, c, d in order, semiperimeter s, and a pair of opposite angles α, γ: $K^2 = (s-a)(s-b)(s-c)(s-d) - abcd\cos^2\frac{\alpha+\gamma}{2}$. Brahmagupta's formula is the cyclic case $\alpha + \gamma = \pi$. Knowledge gap: install it, but note that in this problem it is the heavier route. See Book 1, field: Geometry, topic: area formulae for quadrilaterals.

Recursion, generalised. The sequence $Q_{i+1} = \text{move}(Q_i)$ is an iterated map on a state space. The general principle: if every move preserves a function I of the state that takes finitely many values, and if I together with finitely many further choices determines the state up to the equivalence in question, then the orbit is finite modulo that equivalence. Here $I = (O, r_1, r_2, S)$ and the further choices are the three angular signs. The same generalisation solves "can a sequence of triangles obtained by reflecting one vertex in the opposite side's perpendicular bisector be pairwise non-congruent" (no: the circumcentre is fixed) and many "infinite process" questions where the answer is finiteness by invariants.

Circuit Bank

- **WHEN** a move defined by a reflection **DO** write the fixed set (the axis) and the images of all named points before anything else **BECAUSE** a reflection is defined by its axis, and its fixed points are the rigid skeleton of any orbit
- **WHEN** a perpendicular bisector **DO** read it as an equidistance locus, and look for a point equidistant from more than one pair **BECAUSE** equidistance from two pairs simultaneously

names a point, and named points are invariants

- **WHEN** "no two congruent" in an iterated process **DO** fix a point and reduce congruence to rotation about it; then count angular data **BECAUSE** congruence modulo a fixed point is one real parameter
- **WHEN** two points on concentric circles about O with known distance **DO** use the cosine rule in triangle OPQ to read off the angle at O **BECAUSE** the angle is determined up to sign by one length
- **WHEN** a hypothesis you have not used (here, convex) **DO** find the line where it is spent and say so **BECAUSE** an unused hypothesis is a sign of an incomplete or wrong argument

Backtesting

1. First useful cue: "reflection across the perpendicular bisector of $\overline{AC}$", which swaps A and C and fixes its axis pointwise.
2. Method the expression suggested: hunt for a fixed point of all possible moves; the intersection of the perpendicular bisectors of the diagonals.
3. Where you would hesitate, and the thought to reinforce: after finding the side-length invariant, being tempted to invoke area and Bretschneider. The thought: I want a point, not a number; what point does every axis pass through?
4. The trap the question sets: believing that two invariants (sides and area) force finiteness without a lemma, or misproving the lemma by claiming the wrong angles are monotone.
5. What makes you recognise it faster next time: the word "partner" is decorative; the content is "reflect one vertex in a line through the centre of the figure". Strip the decoration and the fixed centre is visible.
6. The reusable decision rule: preserved quantities bound values, fixed points bound positions; for congruence questions you need the fixed point.

Mirror drills

1. A triangle ABC is replaced by $A'BC$ where A' is the reflection of A in the perpendicular bisector of BC. Show that iterating (choosing any vertex each time) produces finitely many triangles up to congruence. *Answer:* the circumcentre is fixed and the circumradius preserved, and the side multiset is preserved; each side subtends a determined angle at the circumcentre, so finitely many shapes up to rotation.
2. Points P on a circle of radius r_1 and Q on a concentric circle of radius r_2 satisfy $PQ = s$. How many possible values of the angle $\angle POQ$ are there? *Answer:* at most one value in $[0, \pi]$, namely $\arccos \frac{r_1^2+r_2^2-s^2}{2r_1r_2}$, so the signed angle takes at most two values.
3. Explain in one line why a self-intersecting quadrilateral may have no centre in the sense of the solution. *Answer:* for $A = (0,0)$, $C = (1,0)$, $B = (0,1)$, $D = (1,1)$ the diagonals AC and BD are parallel, so their perpendicular bisectors never meet; convexity excludes this.

2.9 B3. Consecutive roots via the tangent addition formula

Let r_n be the nth smallest positive solution to $\tan x = x$*, where the argument of tangent is in radians. Prove that*

$$0 < r_{n+1} - r_n - \pi < \frac{1}{(n^2+n)\pi}$$

for $n \geq 1$.

Read it out loud

An implicitly defined sequence: the positive roots of $\tan x = x$, one per branch of the tangent. You are asked for a two-sided estimate of the gap between consecutive roots: the gap exceeds π (the period of tan), but by less than $\frac{1}{(n^2+n)\pi}$. So the roots creep up their branches, each one slightly further past $n\pi$ than the last, and the creep decays like $1/n^2$.

The fields are real analysis (locating roots by monotonicity and the intermediate value theorem), trigonometry (periodicity and the addition formula), and a dose of asymptotics. What is really being tested is whether you can turn the defining equation into a relation between consecutive roots without ever solving it. The roots satisfy $\tan r_n = r_n$: that is not an obstacle, it is a substitution rule. Whenever $\tan r_n$ appears you may replace it by r_n, and the tangent addition formula is the machine that makes $\tan r_n$ and $\tan r_{n+1}$ appear together.

The first useful cue is the shape of the bound, $\frac{1}{(n^2+n)\pi} = \frac{\pi}{(n\pi)((n+1)\pi)}$: a product of two consecutive quantities of size $n\pi$ and $(n+1)\pi$, which are the crude sizes of r_n and r_{n+1}. A product $r_n r_{n+1}$ is exactly what the denominator $1 + \tan A \tan B$ of the addition formula produces. Within thirty seconds you should have written: $r_n = n\pi + \theta_n$ with $\theta_n \in (0, \frac{\pi}{2})$, and $\tan(\theta_{n+1} - \theta_n) = \dfrac{r_{n+1} - r_n}{1 + r_n r_{n+1}}$.

Approaches, ranked worst to best

1. **Taylor expansion of $f(x) = \tan x - x$ around $r_n + \pi$.** *Cue:* $f(r_n + \pi) = -\pi$ and $f(r_{n+1}) = 0$. *Circuit:* all derivatives of f are positive polynomials in $\tan x$, so a full Taylor lower bound gives a geometric series in the gap. (it works, but it needs an induction on the form of $f^{(k)}$, a remainder estimate, a limit in k, and a subtlety about strict inequalities surviving the limit; a page of ink for a two-line fact)

2. **Mean value theorem once.** *Cue:* $f(r_n + \pi) = -\pi$, $f'(x) = \tan^2 x$ increasing on the branch. *Circuit:* $\pi = \delta \tan^2 \xi > \delta r_n^2$. (one line, but it only gives $\delta < \frac{1}{n^2\pi}$, weaker than the target by the factor $\frac{n}{n+1}$; the mean value theorem sees one point, the target needs both endpoints)

3. **Inverse function of $t(x) = x - \arctan x$.** *Cue:* $\tan x = x$ if and only if $x - \arctan x = n\pi$. *Circuit:* $r_n = x(n\pi)$ for the inverse $x(t)$, and $x(t) - t$, $x(t) - t + 1/t$ are monotone. (neat and rigorous, moderate ink; it hides the same product $r_n r_{n+1}$ inside the monotonicity of $x(t) - t + 1/t$)

4. **Tangent addition formula.** *Cue:* $\tan r_n = r_n$ is a substitution rule and the bound contains $r_n r_{n+1}$. *Circuit:* $\tan(\theta_{n+1} - \theta_n) = \frac{\pi+\delta}{1+r_n r_{n+1}}$ gives both inequalities from $0 < \delta < \tan\delta$. (three lines after locating the roots; the ideal)

The ideal solution

Step 1: locate the roots. Let $f(x) = \tan x - x$ on each branch $I_m = \left(m\pi - \frac{\pi}{2}, m\pi + \frac{\pi}{2}\right)$, $m \in \mathbb{Z}$. On I_m, f is continuous, $f'(x) = \sec^2 x - 1 = \tan^2 x \geq 0$ with equality only at $x = m\pi$, so f is strictly increasing on I_m; and $f(x) \to -\infty$ as $x \to (m\pi - \frac{\pi}{2})^+$, $f(x) \to +\infty$ as $x \to (m\pi + \frac{\pi}{2})^-$. Hence f has exactly one zero in each I_m. For $m = 0$ that zero is $x = 0$, not positive; for $m < 0$ the branch is

negative. For $m \geq 1$, $f(m\pi) = -m\pi < 0$, so the zero lies in $(m\pi, m\pi + \frac{\pi}{2})$. Since the branches are disjoint and ordered, the positive roots in increasing order are these zeros for $m = 1, 2, \ldots$, and so

$$r_n = n\pi + \theta_n \quad \text{with} \quad \theta_n \in \left(0, \tfrac{\pi}{2}\right), \qquad \tan\theta_n = \tan r_n = r_n \quad (n \geq 1).$$

Step 2: the exact relation between consecutive roots. Put $\delta = r_{n+1} - r_n - \pi = \theta_{n+1} - \theta_n \in \left(-\frac{\pi}{2}, \frac{\pi}{2}\right)$. Both θ_n, θ_{n+1} lie in $(0, \frac{\pi}{2})$, so the addition formula applies and its denominator is positive:

$$\tan\delta = \frac{\tan\theta_{n+1} - \tan\theta_n}{1 + \tan\theta_{n+1}\tan\theta_n} = \frac{r_{n+1} - r_n}{1 + r_n r_{n+1}} = \frac{\pi + \delta}{1 + r_n r_{n+1}}.$$

Step 3: the lower bound. Since $\delta > -\frac{\pi}{2}$ we have $\pi + \delta > \frac{\pi}{2} > 0$, and $1 + r_n r_{n+1} > 0$, so $\tan\delta > 0$. As $\delta \in (-\frac{\pi}{2}, \frac{\pi}{2})$ and tan is negative on $(-\frac{\pi}{2}, 0]$ (zero at 0), this forces $\delta \in (0, \frac{\pi}{2})$. Thus $r_{n+1} - r_n - \pi > 0$.

Step 4: the upper bound. For $0 < \delta < \frac{\pi}{2}$ we have $\delta < \tan\delta$ (this is $f(\delta) > f(0) = 0$ with f strictly increasing on I_0, from Step 1). Therefore

$$\delta < \frac{\pi + \delta}{1 + r_n r_{n+1}} \implies \delta\,(1 + r_n r_{n+1}) < \pi + \delta \implies \delta < \frac{\pi}{r_n r_{n+1}}.$$

Finally $r_n > n\pi$ and $r_{n+1} > (n+1)\pi$ from Step 1, so

$$\delta < \frac{\pi}{(n\pi)\big((n+1)\pi\big)} = \frac{1}{(n^2+n)\pi},$$

which is the required upper bound. □

Layer 1: WHAT

Block 1: locate each root on its own branch, $r_n = n\pi + \theta_n$ with $\theta_n \in (0, \frac{\pi}{2})$, using monotonicity of $\tan x - x$ on a branch and the sign of $f(n\pi)$. Block 2: name the quantity to be bounded, $\delta = \theta_{n+1} - \theta_n$, and compute $\tan\delta$ by the addition formula, replacing $\tan\theta_n$ by r_n. Block 3: read the sign of $\tan\delta$ to get $\delta > 0$. Block 4: apply $\delta < \tan\delta$ and the crude bounds $r_n > n\pi$, $r_{n+1} > (n+1)\pi$.

Nominal route: treat r_{n+1} as a perturbation of $r_n + \pi$, expand f in a Taylor series with all derivatives, and control an infinite series. Ideal route: never expand; the roots satisfy an exact functional relation, so ask what identity contains $\tan r_n$ and $\tan r_{n+1}$ simultaneously. The bifurcation point is Block 2: the zombie writes $f(r_n + \pi + \delta)$ and starts differentiating; you write $\tan(\theta_{n+1} - \theta_n)$ and substitute.

Layer 2: HOW

Block 1. **Cue.** "nth smallest positive solution to $\tan x = x$."

Why natural. An implicitly defined sequence must be pinned to an interval before it can be compared with anything. tan is periodic with period π, so the branches are the natural cells; on each, $\tan x - x$ has derivative $\tan^2 x \geq 0$, so it is monotone and has one zero. The sign $f(n\pi) = -n\pi < 0$ pushes the zero to the right half of the branch. That gives $\theta_n \in (0, \frac{\pi}{2})$, which is precisely the domain on which the later use of the addition formula and of $\delta < \tan\delta$ is legitimate.

Link to the next block. Once $r_n = n\pi + \theta_n$, the periodicity of tan converts $\tan r_n = r_n$ into $\tan\theta_n = r_n$: a substitution rule for $\tan\theta_n$.

Block 2. **Cue.** The target bound is $\frac{\pi}{(n\pi)((n+1)\pi)}$, a product of the sizes of consecutive roots.

Why natural. The only elementary identity that manufactures a product $\tan A \tan B$ is the addition formula, and the quantity you want, $\delta = \theta_{n+1} - \theta_n$, is a difference of angles whose tangents you know. So $\tan\delta$ is computable exactly, with no approximation, and the answer $\frac{\pi+\delta}{1+r_n r_{n+1}}$ contains the target denominator already.

Link to the next block. An exact identity for $\tan\delta$ yields sign information (lower bound) and size information (upper bound) in that order.

Block 3. **Cue.** $\tan\delta = \frac{\pi+\delta}{1+r_n r_{n+1}}$ with a positive numerator and denominator.

Why natural. On $(-\frac{\pi}{2}, \frac{\pi}{2})$ the sign of $\tan\delta$ is the sign of δ. Positivity of the right side is immediate from $\delta > -\frac{\pi}{2}$. So the lower bound is not a separate argument, it is the same identity read for sign.

Link to the next block. Knowing $\delta \in (0, \frac{\pi}{2})$ unlocks the inequality $\delta < \tan\delta$, which is the only comparison between an angle and its tangent you ever need on that interval.

Block 4. **Cue.** $\delta < \tan\delta$, then a linear inequality in δ.

Why natural. Substituting the identity, $\delta(1 + r_n r_{n+1}) < \pi + \delta$ cancels the δ on both sides and leaves $\delta r_n r_{n+1} < \pi$. The crude bounds $r_n > n\pi$ from Block 1 finish it. Notice the cancellation: the δ in the numerator was never a nuisance.

Link to the next block. End of proof; no limit, no series.

Layer 3: WHY

The truth beneath the route: for roots of an equation of the form $\tan x = g(x)$, the addition formula turns the difference of two roots into an exact expression in g at the two roots. Consecutive roots of $\tan x = x$ differ by π plus a correction whose tangent is $\frac{\pi+\delta}{1+r_n r_{n+1}}$; since $r_n r_{n+1} \sim n^2\pi^2$, the correction is of order $\frac{1}{n^2\pi}$, and the problem's bound is the crude version of this asymptotic. The examiners nudged you by writing the bound as $\frac{1}{(n^2+n)\pi}$ rather than $\frac{1}{n^2\pi}$: the factor $n(n+1)$ is a signature of a product of consecutive roots, and a mean value argument (which sees one intermediate point) cannot produce it. That is why approach 2 stalls at $\frac{1}{n^2\pi}$ and why the intended route is two-endpoint.

The discernment mechanism is the sentence "the defining equation is a substitution rule, not an obstacle". For an implicitly defined object, do not try to solve; ask which identities involve the defining expression and rewrite them using the equation. Here $\tan r_n = r_n$ lets every occurrence of $\tan r_n$ become r_n, so any identity in tangents becomes an identity in the roots themselves. Train this by taking each classical implicit equation ($x = \cos x$, $xe^x = 1$, $\tan x = x$) and listing which identities (addition formula, functional equation of exp, etc.) convert into exact relations between roots.

A second truth, on rigour: every use of tan needs the argument inside a branch, and every use of the addition formula needs a nonzero denominator. Locating θ_n in $(0, \frac{\pi}{2})$ first is what makes the later three lines airtight, and it costs half the write-up. Full marks are earned in Step 1, not in Step 4.

Revolving concepts

Branches of tan *and root location.* tan is continuous and strictly increasing from $-\infty$ to $+\infty$ on each $I_m = (m\pi - \frac{\pi}{2}, m\pi + \frac{\pi}{2})$. For $\tan x = g(x)$ with g continuous, the number of roots in I_m is governed by the monotonicity of $\tan x - g(x)$; when $g' \le \tan^2 x$ on the branch (as for $g(x) = x$) there is exactly one. See Book 1, field: Real analysis, topic: intermediate value theorem and monotone root location.

Tangent addition formula and its substitution use. $\tan(A - B) = \frac{\tan A - \tan B}{1+\tan A \tan B}$ when $A, B, A - B$ avoid odd multiples of $\frac{\pi}{2}$. If $\tan A = a$, $\tan B = b$ are known numbers, $A - B$ is known modulo π exactly. See Book 1, field: Calculus and integrals, topic: trigonometric identities as functional equations.

The inequality $x < \tan x$ *on* $(0, \frac{\pi}{2})$, with the companions $\sin x < x < \tan x$; all three are monotonicity statements about $\tan x - x$ and $x - \sin x$. See Book 1, field: Inequalities, topic: elementary function inequalities.

Asymptotics of the roots. Writing $r_n = (n + \frac{1}{2})\pi - \varepsilon_n$ gives $\cot\varepsilon_n = r_n$, so $\varepsilon_n = \arctan(1/r_n)$ and $r_n = (n + \frac{1}{2})\pi - \frac{1}{(n+\frac{1}{2})\pi} + O(n^{-3})$. These roots are the eigenvalue equation of several physical problems and also the positive zeros of the spherical Bessel function j_1. See Book 1, field: Sequences and recurrences, topic: asymptotic expansion of implicitly defined sequences.

Recursion, generalised. The sequence r_n is defined by a functional condition rather than a recurrence, but the gap $\delta_n = r_{n+1} - r_n - \pi$ satisfies the exact relation $\tan\delta_n = \frac{\pi+\delta_n}{1+r_n r_{n+1}}$. General form:

if x_n solves $h(x) = x$ on the nth branch of a π-periodic h with an addition law (tangent, or any Möbius-type law), then the gaps satisfy an exact equation of the shape $\Phi(\delta_n) = \Psi(x_n, x_{n+1})$, and both bounds on δ_n come from one such equation plus crude bounds on x_n. The same generalisation handles $\tan x = ax$, $\tan x = a/x$, and $\cot x = x$.

Circuit Bank

- **WHEN** an implicitly defined sequence of roots of a periodic equation **DO** locate each root on its own branch by monotonicity, and write it as $n\pi + \theta_n$ with θ_n in a fixed interval **BECAUSE** every later inequality needs a domain, and the branch is that domain
- **WHEN** $\tan r = r$ (or any defining equation) and a quantity involving tan of a difference of roots **DO** apply the addition formula and replace every $\tan r$ by r **BECAUSE** the defining equation is a substitution rule; the addition formula is the identity that puts two roots in one expression
- **WHEN** a bound with denominator $n(n+1)$ or $r_n r_{n+1}$ **DO** expect a two-endpoint identity, not a one-point mean value estimate **BECAUSE** a mean value argument only produces one root's square
- **WHEN** an exact identity $\tan\delta = (\text{positive})$ with $\delta \in (-\frac{\pi}{2}, \frac{\pi}{2})$ **DO** read the sign first, then apply $\delta < \tan\delta$ **BECAUSE** sign gives the lower bound for free and unlocks the domain for the upper bound

Backtesting

1. First useful cue: the target $\frac{1}{(n^2+n)\pi} = \frac{\pi}{(n\pi)((n+1)\pi)}$, a product of consecutive root sizes.
2. Method the expression suggested: an identity producing $r_n r_{n+1}$, namely the denominator of the tangent addition formula.
3. Where you would hesitate, and the thought to reinforce: after locating the roots, the temptation to expand $\tan(r_n + \pi + \delta)$ as a series. The thought: I know tan at both roots exactly, so I know tan of their difference exactly.
4. The trap the question sets: a single mean value step gives $\frac{1}{n^2\pi}$ and feels like victory; it is short of the target by the factor $\frac{n}{n+1}$, and a grader will not accept it.
5. What makes you recognise it faster next time: "$\tan x = x$" means "tan of a root is the root"; any identity in tangents of roots is an identity in the roots.
6. The reusable decision rule: for roots of $\tan x = g(x)$, the consecutive gap satisfies $\tan(\text{gap} - \pi) = \frac{g(x_{n+1}) - g(x_n)}{1 + g(x_n)g(x_{n+1})}$; bound the gap from that, never from a series.

Mirror drills

1. With r_n as in the problem, show $r_n = (n + \frac{1}{2})\pi - \frac{1}{(n+\frac{1}{2})\pi} + O(n^{-3})$. *Hint:* write $r_n = (n + \frac{1}{2})\pi - \varepsilon_n$; then $\tan r_n = \cot\varepsilon_n = r_n$, so $\varepsilon_n = \arctan(1/r_n) = \frac{1}{r_n} + O(r_n^{-3})$, and substitute $r_n = (n + \frac{1}{2})\pi + O(1/n)$ into $1/r_n$.
2. Prove the matching lower bound $r_{n+1} - r_n - \pi > \dfrac{\pi^2}{4(1 + r_n r_{n+1})}$. *Hint:* $\delta < \frac{1}{2\pi} < \frac{\pi}{4}$ and tan is convex on $[0, \frac{\pi}{4}]$, so $\tan\delta \le \frac{4}{\pi}\delta$ there; combine with $\tan\delta > \frac{\pi}{1 + r_n r_{n+1}}$.
3. Let s_n be the nth positive root of $\tan x = 2x$. Show $s_{n+1} - s_n - \pi < \dfrac{2\pi}{4 s_n s_{n+1} - 1}$. *Hint:* first check that $\tan x - 2x$ has exactly one positive root in each $(m\pi + \frac{\pi}{4}, m\pi + \frac{\pi}{2})$, $m \ge 0$, and no

other positive roots (its derivative $\tan^2 x - 1$ is negative on $(m\pi - \frac{\pi}{4}, m\pi + \frac{\pi}{4})$ and positive on the rest of the branch, and the local maximum at $m\pi - \frac{\pi}{4}$ is negative); so s_n lies on the branch $m = n - 1$ and $\delta = s_{n+1} - s_n - \pi \in (-\frac{\pi}{4}, \frac{\pi}{4})$; then $\tan\delta = \frac{2(\pi+\delta)}{1+4s_n s_{n+1}}$, which gives $\delta > 0$, and $\delta < \tan\delta$ gives the bound.

2.10 B4. Affine drift makes the mean a geometric sequence

Let n be a positive integer. Set $a_{n,0} = 1$. For $k \geq 0$, choose an integer $m_{n,k}$ uniformly at random from the set $\{1, \ldots, n\}$, and let

$$a_{n,k+1} = \begin{cases} a_{n,k} + 1, & \text{if } m_{n,k} > a_{n,k}; \\ a_{n,k}, & \text{if } m_{n,k} = a_{n,k}; \\ a_{n,k} - 1, & \text{if } m_{n,k} < a_{n,k}. \end{cases}$$

Let $E(n)$ be the expected value of $a_{n,n}$. Determine $\lim_{n\to\infty} E(n)/n$.

Read it out loud

A random walk on $\{1, \ldots, n\}$ that starts at 1, and at each step moves up with probability $\frac{n-a}{n}$, stays with probability $\frac{1}{n}$, and moves down with probability $\frac{a-1}{n}$, where a is the current position. It is run for exactly n steps and you are asked for the expected position, normalised by n, in the limit. So the object is a Markov chain with state-dependent step probabilities, and the question is about its mean after a number of steps equal to the size of the state space.

The fields are probability (conditional expectation, the tower law), sequences and recurrences (a first-order linear recurrence with constant coefficients), and the limit $(1 - \frac{2}{n})^n \to e^{-2}$. What is really being tested is whether you compute the conditional expectation of one step before doing anything else. The transition rule is given by comparing m with a, so the expected change in one step is $P(m > a) - P(m < a) = \frac{(n-a)-(a-1)}{n} = \frac{n+1-2a}{n}$, which is affine in a. That single line is the whole problem: an affine drift means the expectation satisfies a closed linear recurrence with no knowledge of the distribution.

The first useful cue is the case split by comparison with a uniform variable: "$m > a$", "$m = a$", "$m < a$". Within thirty seconds you should have written $\mathbb{E}[a_{k+1} \mid a_k = a] = a + \frac{n+1-2a}{n} = \left(1 - \frac{2}{n}\right) a + \frac{n+1}{n}$. Everything after that is A-level: solve the recurrence, evaluate at $k = n$, take the limit.

Approaches, ranked worst to best

1. **Full spectral analysis of the transition matrix.** *Cue:* the process is a Markov chain on n states. *Circuit:* write the tridiagonal matrix A, find all eigenvalues $\frac{n-2j}{n}$ and eigenvectors, expand the initial vector. (correct but enormous: you need every eigenvector to expand v_0, when the question only asks for one linear functional of v_n; costs pages)

2. **Coupling two starting points.** *Cue:* the update depends on a only through the comparison with m. *Circuit:* walks started at j and $j+1$ coalesce when $m \in \{j, j+1\}$, otherwise stay one apart, so $\mathbb{E}$ is affine in the start with slope $(1 - \frac{2}{n})^n$; combine with the symmetry $a \mapsto n + 1 - a$. (short and rigorous, and natural once you notice the update depends on a only through the comparison, but it needs the symmetry observation and a coupling argument that a grader will want written carefully; more ink than necessary)

3. **Bulb reinterpretation.** *Cue:* transition probabilities $\frac{a-1}{n}$, $\frac{1}{n}$, $\frac{n-a}{n}$ are the Ehrenfest urn with one stuck ball. *Circuit:* one red bulb always lit, $n - 1$ white bulbs toggled by a uniformly chosen bulb; a is the number of lit bulbs; linearity of expectation over bulbs and the parity of a binomial. (gives the exact answer and the full distribution and explains the e^{-2}, but requires justifying that the two chains have the same law; the best explanation, not the shortest proof)

4. **Affine drift recurrence.** *Cue:* the one-step change has expectation $\frac{n+1-2a}{n}$, affine in a. *Circuit:* tower law gives $e_{k+1} = (1 - \frac{2}{n})e_k + \frac{n+1}{n}$; solve by fixed point. (four lines, nothing beyond A-level, exact closed form; the ideal)

The ideal solution

Fix n and write $a_k = a_{n,k}$, $m_k = m_{n,k}$, $e_k = \mathbb{E}[a_k]$. By construction a_k is a function of $m_0, \dots, m_{k-1}$, so m_k is independent of a_k, and $a_k \in \{1, \dots, n\}$ for all k (the walk cannot leave the range, since at $a = n$ the event $m > a$ is impossible and at $a = 1$ the event $m < a$ is impossible).

Step 1: the one-step drift. For $a \in \{1, \dots, n\}$,

$$\mathbb{E}[a_{k+1} \mid a_k = a] = a + P(m_k > a) - P(m_k < a) = a + \frac{n-a}{n} - \frac{a-1}{n} = \left(1 - \frac{2}{n}\right) a + \frac{n+1}{n}.$$

Step 2: the recurrence for the mean. Since the right side is affine in a, the tower law gives

$$e_{k+1} = \mathbb{E}\big[\mathbb{E}[a_{k+1} \mid a_k]\big] = \left(1 - \frac{2}{n}\right) e_k + \frac{n+1}{n}, \qquad e_0 = 1.$$

Step 3: solve. The fixed point of $e \mapsto (1 - \frac{2}{n})e + \frac{n+1}{n}$ is $\mu = \frac{n+1}{2}$ (solve $\frac{2}{n}\mu = \frac{n+1}{n}$). Subtracting, $e_{k+1} - \mu = (1 - \frac{2}{n})(e_k - \mu)$, so

$$e_k - \mu = \left(1 - \frac{2}{n}\right)^k (e_0 - \mu) = -\frac{n-1}{2}\left(1 - \frac{2}{n}\right)^k.$$

Therefore

$$E(n) = e_n = \frac{n+1}{2} - \frac{n-1}{2}\left(1 - \frac{2}{n}\right)^n.$$

(Checks: $E(1) = 1$, as the walk never moves; $E(2) = \frac{3}{2}$, as a_2 is 1 or 2 with equal probability.)

Step 4: the limit. Dividing by n,

$$\frac{E(n)}{n} = \frac{n+1}{2n} - \frac{n-1}{2n}\left(1 - \frac{2}{n}\right)^n \longrightarrow \frac{1}{2} - \frac{1}{2}e^{-2},$$

using $\frac{n\pm1}{2n} \to \frac{1}{2}$ and the standard limit $(1 + \frac{x}{n})^n \to e^x$ with $x = -2$ (equivalently, $n \ln(1 - \frac{2}{n}) = -2 + O(\frac{1}{n}) \to -2$). Hence

$$\boxed{\lim_{n\to\infty} \frac{E(n)}{n} = \frac{1 - e^{-2}}{2}}.$$

□

Layer 1: WHAT

Block 1: compute the conditional expectation of one step; it is affine in the current state. Block 2: take expectations (tower law) to get a first-order linear recurrence for e_k. Block 3: solve the recurrence by subtracting the fixed point $\frac{n+1}{2}$. Block 4: evaluate at $k = n$, divide by n, use $(1 - \frac{2}{n})^n \to e^{-2}$.

Nominal route: set up the n-by-n transition matrix, look for eigenvalues, or simulate small n and guess $\frac{1-e^{-2}}{2}$ from decimals. Ideal route: ask only for what the question asks, the mean, and notice the mean closes on itself. The bifurcation point is Block 1: the zombie writes the matrix; you write $\mathbb{E}[a_{k+1} \mid a_k]$ and see the letter a appear linearly.

Layer 2: HOW

Block 1. Cue. The update is decided by comparing a uniform m with the current state a: the three cases have probabilities $\frac{n-a}{n}$, $\frac{1}{n}$, $\frac{a-1}{n}$.

Why natural. Expected change is (probability up) minus (probability down). Both probabilities are linear in a because "the number of values of m above a" is $n - a$ and "below a" is $a - 1$. There is no idea here: it is the definition of expectation applied to a three-valued step.

Link to the next block. A conditional expectation that is affine in the state is the hypothesis of the tower-law step: $\mathbb{E}[\lambda a_k + c] = \lambda e_k + c$.

Block 2. **Cue.** $\mathbb{E}[a_{k+1} \mid a_k] = \lambda a_k + c$ with $\lambda = 1 - \frac{2}{n}$, $c = \frac{n+1}{n}$.

Why natural. The tower law $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \mid Y]]$ is the only tool needed, and it works because the conditional expectation is a linear function of a_k whose expectation you can take without knowing the law of a_k. If the drift had been quadratic in a you would have needed $\mathbb{E}[a_k^2]$ as well, and the hierarchy would not close at the first level.

Link to the next block. A first-order linear recurrence with constant coefficients has a one-line solution via its fixed point.

Block 3. **Cue.** $e_{k+1} = \lambda e_k + c$, $|\lambda| < 1$ for $n \geq 2$.

Why natural. Subtracting the fixed point $\mu = \frac{c}{1-\lambda}$ turns the affine recurrence into a geometric one. The fixed point $\frac{n+1}{2}$ is the centre of $\{1, \ldots, n\}$, which is where you expect the walk to settle by the symmetry $a \mapsto n + 1 - a$ of the transition rule; seeing that number appear is a sanity check, not a surprise.

Link to the next block. You have e_k for all k; the problem chose $k = n$, which makes $\lambda^k = (1 - \frac{2}{n})^n$, and the limit of that is the reason the answer involves e^{-2}.

Block 4. **Cue.** $(1 - \frac{2}{n})^n$ with the same n in base and exponent.

Why natural. This is the standard limit $(1 + \frac{x}{n})^n \to e^x$. The examiners chose exactly n steps so that the decay factor λ^n has a nontrivial limit: fewer steps ($o(n)$) would give $\frac{E}{n} \to 0$, more ($\omega(n)$) would give $\frac{1}{2}$. The number of steps is tuned to the relaxation time $\frac{n}{2}$ of the chain.

Link to the next block. Assemble and box.

Layer 3: WHY

The structural truth is that this chain is the Ehrenfest urn with one stuck ball. Picture n bulbs, one red and permanently lit, $n - 1$ white and initially unlit; at each step choose a bulb uniformly and toggle it if it is white. If a bulbs are lit (counting the red one), the chosen bulb is a lit white one with probability $\frac{a-1}{n}$ (the count drops), the red one with probability $\frac{1}{n}$ (no change), an unlit white one with probability $\frac{n-a}{n}$ (the count rises). These are the problem's transition probabilities, and the start $a = 1$ is "all white bulbs off". Two Markov chains on the same state space with the same transition probabilities and the same initial law are the same process in law, so $a_{n,k}$ has the law of the number of lit bulbs.

That reinterpretation explains everything in the answer. By linearity of expectation over bulbs, $E(n) = 1 + (n-1)\,P(\text{a given white bulb is lit after } n \text{ steps})$, and a given white bulb is lit exactly when it has been chosen an odd number of times; the number of times is $\mathrm{Bin}(n, \frac{1}{n})$, and $P(\mathrm{Bin}(n,p) \text{ odd}) = \frac{1-(1-2p)^n}{2}$ (expand $(q+p)^n - (q-p)^n$). With $p = \frac{1}{n}$ that is $\frac{1-(1-\frac{2}{n})^n}{2} \to \frac{1-e^{-2}}{2} = P(\mathrm{Poisson}(1) \text{ is odd})$. So the answer $\frac{1-e^{-2}}{2}$ is the probability that a Poisson variable of mean 1 is odd, and the mean 1 is because each bulb is touched about once in n steps. It is worth repeating: the recurrence proves the answer in four lines, the bulbs tell you why the answer has that shape. Both belong in your head, in that order.

The discernment mechanism: for any random process, before modelling the whole distribution, compute the conditional expectation of one step and check whether it is affine in the state. If it is, the mean is a first-order linear recurrence and you are done with the mean; if the conditional second moment is a polynomial of degree two in the state, the variance closes too. The examiners nudged you by making the drift affine (the comparison with a uniform variable makes the up and down probabilities linear in a) and by asking only for the mean. Train this by computing $\mathbb{E}[X_{k+1} \mid X_k]$ for every chain you meet, and by recognising the eigenfunction structure: for the Ehrenfest chain, $a \mapsto a - \frac{n+1}{2}$ is an eigenfunction of the transition operator with eigenvalue $1 - \frac{2}{n}$, and affine drift is exactly the statement that a linear function is an eigenfunction.

Revolving concepts

Conditional expectation and the tower law. $\mathbb{E}[X] = \mathbb{E}[\mathbb{E}[X \mid Y]]$; if $\mathbb{E}[X_{k+1} \mid X_k] = \lambda X_k + c$ then $\mathbb{E}[X_{k+1}] = \lambda \mathbb{E}[X_k] + c$. This is the mean-field reduction of a Markov chain, and it needs no knowledge of the distribution of X_k. See Book 1, field: Probability, topic: conditional expectation and martingale-type reductions.

First-order linear recurrence. $e_{k+1} = \lambda e_k + c$ has fixed point $\mu = \frac{c}{1-\lambda}$ ($\lambda \neq 1$) and solution $e_k = \mu + \lambda^k(e_0 - \mu)$. Generalised recursion: for $x_{k+1} = \lambda x_k + c_k$ the solution is $x_k = \lambda^k x_0 + \sum_{j<k} \lambda^{k-1-j} c_j$ (variation of constants), and for vector states the same formula holds with λ a matrix, which is the spectral approach in disguise. See Book 1, field: Sequences and recurrences, topic: linear recurrences and fixed points.

Ehrenfest urn. n balls in two urns, one ball moved at random per step; the count in one urn is a birth-death chain with rates $\frac{n-a}{n}$ up and $\frac{a}{n}$ down, stationary law $\mathrm{Bin}(n, \frac{1}{2})$, eigenvalues $1 - \frac{2j}{n}$ for $j = 0, \ldots, n$ (Krawtchouk polynomials as eigenfunctions), relaxation time of order n. The problem's chain is the Ehrenfest urn with one ball frozen. See Book 1, field: Probability, topic: Markov chains, birth-death chains, Ehrenfest model.

Parity of a binomial and the Poisson limit. $P(\mathrm{Bin}(n, p) \text{ odd}) = \frac{1-(1-2p)^n}{2}$; $P(\mathrm{Poisson}(\lambda) \text{ odd}) = e^{-\lambda} \sinh \lambda = \frac{1-e^{-2\lambda}}{2}$. Generating-function reason: evaluate the probability generating function at ± 1. See Book 1, field: Generating functions, topic: roots-of-unity filter for parity.

The limit $(1 + \frac{x}{n})^n \to e^x$, including the version with $x_n \to x$: $(1 + \frac{x_n}{n})^n \to e^x$. See Book 1, field: Real analysis, topic: the exponential limit.

Circuit Bank

- **WHEN** a random process defined step by step and a question about its mean **DO** compute $\mathbb{E}[X_{k+1} \mid X_k]$ first and check whether it is affine in X_k **BECAUSE** affine drift closes the mean into a first-order linear recurrence without touching the distribution
- **WHEN** a step decided by comparing the state with a uniform draw **DO** write the up and down probabilities as counts over n; they are linear in the state **BECAUSE** comparison with a uniform variable is a counting statement
- **WHEN** $e_{k+1} = \lambda e_k + c$ **DO** subtract the fixed point $\frac{c}{1-\lambda}$ and read off $e_k = \mu + \lambda^k(e_0 - \mu)$ **BECAUSE** an affine map is a geometric map about its fixed point
- **WHEN** a factor $(1 - \frac{c}{n})^n$ or $(1 - \frac{c}{n})^{tn}$ in an answer **DO** write e^{-c} or e^{-ct} in the limit and interpret t as time measured in relaxation units **BECAUSE** the number of steps was tuned to the relaxation time by the examiners
- **WHEN** transition probabilities $\frac{a-1}{n}$, $\frac{1}{n}$, $\frac{n-a}{n}$ or $\frac{a}{n}$, $\frac{n-a}{n}$ **DO** recognise the Ehrenfest urn and switch to independent bulbs plus linearity of expectation **BECAUSE** each bulb is toggled by an independent binomial count, so parity gives the exact law

Backtesting

1. First useful cue: the three-way case split by comparison with a uniform m, giving up and down probabilities linear in a.
2. Method the expression suggested: one-step conditional expectation, then a linear recurrence for the mean.
3. Where you would hesitate, and the thought to reinforce: seeing an n-state Markov chain and reaching for the matrix. The thought: the question asks for the mean only; test whether the mean closes on itself before modelling anything else.

4. The trap the question sets: the normalisation $E(n)/n$ and the limit tempt you into asymptotic hand-waving about the walk being "about halfway"; the answer is not $\frac{1}{2}$, and only the exact recurrence exposes the e^{-2} deficit.
5. What makes you recognise it faster next time: the stuck-ball Ehrenfest transition probabilities, and the fact that the number of steps equals n, the scale of the relaxation time.
6. The reusable decision rule: mean of a Markov chain with affine drift equals fixed point plus λ^k times the initial deviation; check affinity of the drift in the first thirty seconds.

Mirror drills

1. The genuine Ehrenfest urn: $a_0 = 0$, and a moves up with probability $\frac{n-a}{n}$ and down with probability $\frac{a}{n}$. Find $\lim \mathbb{E}[a_n]/n$. *Answer:* drift $1 - \frac{2a}{n}$, so $e_{k+1} = (1 - \frac{2}{n})e_k + 1$, fixed point $\frac{n}{2}$, $e_n = \frac{n}{2}(1 - (1 - \frac{2}{n})^n)$, limit $\frac{1-e^{-2}}{2}$.
2. Prove $P(\mathrm{Bin}(n,p) \text{ odd}) = \frac{1-(1-2p)^n}{2}$ and deduce the exact value $E(n) = 1 + (n-1)\frac{1-(1-\frac{2}{n})^n}{2}$ for the problem's chain. *Answer:* $(q+p)^n - (q-p)^n = 2\sum_{i \text{ odd}} \binom{n}{i} p^i q^{n-i}$ with $q = 1 - p$; then each white bulb is lit with that probability and linearity of expectation over the $n-1$ white bulbs adds the constant 1 for the red bulb.
3. Run the problem's walk for $\lfloor tn \rfloor$ steps with $t > 0$ fixed. Find the limit of the mean over n and interpret $t \to \infty$. *Answer:* $\frac{1-e^{-2t}}{2}$; as $t \to \infty$ it tends to $\frac{1}{2}$, the stationary mean $\frac{n+1}{2}$ normalised, so t counts time in units of the relaxation time $\frac{n}{2}$.

2.11 B5. Real negative roots give nonnegative coefficients

Let k and m be positive integers. For a positive integer n, let $f(n)$ be the number of integer sequences $x_1, \dots, x_k, y_1, \dots, y_m, z$ satisfying $1 \le x_1 \le \cdots \le x_k \le z \le n$ and $1 \le y_1 \le \cdots \le y_m \le z \le n$. Show that $f(n)$ can be expressed as a polynomial in n with nonnegative coefficients.

Read it out loud

Two weakly increasing chains, of lengths k and m, both capped by a common top element $z \le n$. Counting by the value of z: the x-chain is a multiset of size k from $\{1, \dots, z\}$, the y-chain a multiset of size m, so

$$f(n) = \sum_{z=1}^{n} R(z,k)\, R(z,m), \qquad R(z,j) := \binom{z+j-1}{j} = \frac{z(z+1)\cdots(z+j-1)}{j!}.$$

Read out loud: f is the partial sum of a product of two rising factorials. That it is a polynomial in n is immediate (a partial sum of a polynomial is a polynomial). The content of the problem is the sign of the coefficients, and that is a statement about f as an algebraic object, not about its values at positive integers.

The fields are combinatorics (multisets, the stars-and-bars count), polynomials (the summation operator, difference equations, roots), and a small piece of real analysis (the intermediate value theorem to locate roots). What is really being tested is whether you know the one structural mechanism that produces nonnegative coefficients: a polynomial with positive leading coefficient whose roots are all real and nonpositive is a product of factors $(x+s)$ with $s \ge 0$, and such a product has nonnegative coefficients. So the problem is secretly "show all roots of f are real and nonpositive", and the difference equation $F(x) - F(x-1) = R(x,k)R(x,m)$ hands you $m+1$ roots for free at $0, -1, \dots, -m$, because the right side vanishes there.

The first useful cue is the pair of rising factorials with their known zeros at $0, -1, \dots$. Within thirty seconds you should have written the sum formula, the difference equation, and the observation $F(0) = F(-1) = \cdots = F(-m) = 0$ (with $k \le m$ by symmetry). The remaining k roots are the real work.

Approaches, ranked worst to best

1. **Faulhaber and Bernoulli numbers.** *Cue:* a sum of a polynomial in z from 1 to n. *Circuit:* expand $R(z,k)R(z,m)$ in powers of z and sum each power. (it proves polynomiality but Bernoulli numbers alternate in sign, so the coefficient signs are invisible; stalls completely on the actual question)

2. **Rising-factorial basis and hockey stick.** *Cue:* $\sum_{z \le n} R(z,j) = R(n,j+1)$. *Circuit:* expand the product $R(z,k)R(z,m)$ in the basis $R(z,j)$ and sum termwise. (the expansion coefficients alternate in sign, for instance $z^2 = 2R(z,2) - R(z,1)$, so positivity is again invisible; stalls)

3. **Summation by parts and an explicit formula.** *Cue:* $R(z,j) - R(z-1,j) = R(z,j-1)$. *Circuit:* derive $f_{k,m} = R(n,k)R(n,m+1) + f_{k-1,m} - f_{k-1,m+1}$, guess $f_{k,m}(n) = \frac{R(n,m+1)}{R(m+2,k)} \sum_{i=0}^{k} R(m+1,i)R(n,i)$, prove by induction on k. (a complete proof, and the formula is manifestly positive, but you must guess it from $k = 0, 1, 2$ and the induction is a page of algebra; ranks second on ink)

4. **Root location through the difference equation.** *Cue:* the right side $R(x,k)R(x,m)$ vanishes at $0, -1, \dots$. *Circuit:* F vanishes at $0, \dots, -m$; the cofactor Q satisfies a two-term relation at $x = -j$ that forces sign alternation, hence k real roots in $(-k, 0)$; a real-rooted polynomial with

nonpositive roots has nonnegative coefficients. (half a page, no guessing, and it explains the answer; the ideal)

The ideal solution

Step 1: set-up. For $j \geq 0$ let $R(x,j) = \frac{x(x+1)\cdots(x+j-1)}{j!}$, a polynomial in x of degree j with $R(x,0) = 1$. For a positive integer z, $R(z,j) = \binom{z+j-1}{j}$ counts the weakly increasing sequences $1 \leq u_1 \leq \cdots \leq u_j \leq z$ (stars and bars: j stars and $z-1$ bars). Classifying the sequences in the problem by the value of z,

$$f(n) = \sum_{z=1}^{n} R(z,k)R(z,m).$$

Since f is symmetric in k and m, assume $k \leq m$.

Step 2: the polynomial and its difference equation. Let $g(x) = R(x,k)R(x,m)$, a polynomial of degree $k+m$. There is a unique polynomial F of degree $k+m+1$ with $F(0) = 0$ and $F(x) - F(x-1) = g(x)$ identically: indeed the linear map $F \mapsto F(x) - F(x-1)$ from the space of polynomials of degree at most $d+1$ to the space of polynomials of degree at most d has kernel the constants (a polynomial with $F(x) = F(x-1)$ for all x is periodic, hence constant), so by dimension count it is onto; a solution exists, has degree exactly $d+1$ (the difference lowers degree by one for nonconstant F), and $F(0) = 0$ fixes it. Then $F(n) = \sum_{z=1}^{n} g(z) = f(n)$ for every positive integer n by telescoping. So f is the polynomial F, and it remains to show that F has nonnegative coefficients.

Step 3: $m+1$ roots for free. $R(x,m)$ vanishes at $x = 0, -1, \ldots, -(m-1)$. Using $F(-j) = F(-j+1) - g(-j+1)$ for $j = 1, \ldots, m$, and $g(-j+1) = 0$ for those j, induction from $F(0) = 0$ gives

$$F(0) = F(-1) = \cdots = F(-m) = 0.$$

Hence $F(x) = x(x+1)\cdots(x+m)\,Q(x)$ for a polynomial Q of degree k.

Step 4: the difference equation for Q. Write $x(x+1)\cdots(x+m) = m!\,R(x,m)(x+m)$ and $(x-1)x\cdots(x+m-1) = m!\,R(x,m)(x-1)$. Then

$$F(x) - F(x-1) = m!\,R(x,m)\big[(x+m)Q(x) - (x-1)Q(x-1)\big] = R(x,k)R(x,m),$$

and cancelling the nonzero polynomial $m!\,R(x,m)$,

$$(x+m)Q(x) - (x-1)Q(x-1) = \frac{R(x,k)}{m!}.$$

Substituting $x = -j$ for $j = 0, 1, \ldots, k-1$ kills the right side (as $R(-j,k)$ contains the factor $(-j+j)$), giving

$$(m-j)\,Q(-j) = -(j+1)\,Q(-j-1) \qquad (j = 0, 1, \ldots, k-1).$$

Since $k \leq m$, every coefficient $m-j$ and $j+1$ here is a positive integer.

Step 5: sign alternation and k more roots. From Step 4, $Q(-j) = 0$ if and only if $Q(-j-1) = 0$, so the values $Q(0), Q(-1), \ldots, Q(-k)$ are either all zero or all nonzero. All zero is impossible: Q would have $k+1$ distinct roots and degree k, so $Q \equiv 0$, so $F \equiv 0$, contradicting $F(1) = f(1) = 1$. So they are all nonzero, and consecutive ones have opposite signs. By the intermediate value theorem Q has a root $-s_j$ in each open interval $(-j-1, -j)$, $j = 0, \ldots, k-1$; these are k distinct roots of a degree-k polynomial, so

$$Q(x) = q\prod_{j=0}^{k-1}(x+s_j), \qquad 0 < s_0 < s_1 < \cdots < s_{k-1} < k,$$

with q the leading coefficient of Q.

Step 6: conclusion. The leading coefficient of F equals that of Q, namely q, and it is positive because $F(n) = f(n) > 0$ for all positive integers n forces the leading coefficient to be positive (a polynomial with negative leading coefficient is eventually negative). Therefore

$$F(x) = q\ x(x+1)\cdots(x+m)\prod_{j=0}^{k-1}(x+s_j)$$

is a positive constant times a product of linear factors $x+s$ with $s \geq 0$. Each factor has nonnegative coefficients, and products and positive multiples of polynomials with nonnegative coefficients have nonnegative coefficients. Hence $f(n) = F(n)$ is a polynomial in n with nonnegative coefficients. □

A caveat worth clarifying: the assumption $k \leq m$ is used exactly once, to guarantee $m - j > 0$ in Step 4; without it you would divide by zero at $j = m$. The symmetry of f in (k, m) makes the assumption free, and a full-marks write-up says so in one sentence.

Layer 1: WHAT

Block 1: count by the top value z, so $f(n) = \sum_{z \leq n} R(z,k)R(z,m)$ with R a rising factorial; assume $k \leq m$. Block 2: replace f by the polynomial F solving $F(x) - F(x-1) = R(x,k)R(x,m)$, $F(0) = 0$. Block 3: harvest the roots $0, -1, \ldots, -m$ of F from the zeros of $R(x,m)$; factor $F = x(x+1)\cdots(x+m)Q$. Block 4: divide the difference equation by the common factor to get $(x+m)Q(x) - (x-1)Q(x-1) = R(x,k)/m!$, and evaluate at the zeros $0, \ldots, -(k-1)$ of $R(x,k)$. Block 5: sign alternation of $Q(0), \ldots, Q(-k)$, intermediate value theorem, k real roots in $(-k, 0)$. Block 6: positive leading coefficient and a product of factors $(x+s)$, $s \geq 0$, so nonnegative coefficients.

Nominal route: expand, sum with Faulhaber or the hockey stick, and stare at an alternating formula hoping to see positivity. Ideal route: treat "nonnegative coefficients" as a statement about roots, and use the difference equation as a root detector. The bifurcation point is Block 2: the zombie tries to compute F; you decide to locate the roots of F without computing it.

Layer 2: HOW

Block 1. **Cue.** Two chains with a common cap $z \leq n$.

Why natural. Conditioning on the cap decouples the chains: each is a multiset from $\{1, \ldots, z\}$ and the multiset count is the rising factorial $R(z,j)$. Writing R as a polynomial $\frac{x(x+1)\cdots(x+j-1)}{j!}$ rather than a binomial coefficient is deliberate: it displays the zeros at $0, -1, \ldots, -(j-1)$, which is the only property you will use.

Link to the next block. A sum from 1 to n of a polynomial is the unique polynomial solution of a difference equation with $F(0) = 0$; pass to that polynomial.

Block 2. **Cue.** "Can be expressed as a polynomial in n" asks you to leave the positive integers.

Why natural. The coefficients of a polynomial are properties of the polynomial on all of $\mathbb{R}$; the difference equation $F(x) - F(x-1) = g(x)$ holds identically, so you may evaluate it at negative and non-integer points. That is the step that unlocks everything: the values of f at positive integers say nothing about signs of coefficients, the values of F at nonpositive integers say everything.

Link to the next block. Evaluate the difference equation where the right side vanishes.

Block 3. **Cue.** $R(x,m)$ vanishes at $0, -1, \ldots, -(m-1)$ and $F(0) = 0$.

Why natural. Each vanishing of g propagates the zero of F one step to the left: $F(-j) = F(-j+1) - g(-j+1)$. So $m+1$ roots of F come for free, which is more than half of its $k+m+1$ roots (as $k \leq m$). Factoring them out leaves a degree-k cofactor Q whose difference equation is simpler.

Link to the next block. Divide the difference equation through by the common polynomial factor; Q inherits a two-term relation.

Block 4. **Cue.** $(x+m)Q(x) - (x-1)Q(x-1) = R(x,k)/m!$, whose right side still has k known zeros.

Why natural. The zeros $0, \ldots, -(k-1)$ of $R(x,k)$ are where the inhomogeneous term disappears, so at those points the relation is homogeneous and two-term, linking $Q(-j)$ to $Q(-j-1)$ with coefficients of known signs. This is the same move as Block 3 applied one level down; no new idea, the same idea twice.

Link to the next block. A two-term relation with a negative ratio is sign alternation, and sign alternation plus continuity is roots.

Block 5. **Cue.** $Q(-j) = -\frac{j+1}{m-j}Q(-j-1)$ with positive ratio magnitudes.

Why natural. Nonvanishing must be checked (otherwise all zero, degree contradiction); then consecutive values alternate in sign and the intermediate value theorem plants a root in each unit interval. Degree k with k roots found means you have found all of them, and all are real and negative. This is the standard way real-rootedness is proved for combinatorial polynomials: exhibit enough sign changes.

Link to the next block. Real nonpositive roots and a positive leading coefficient give the factorised form, and the factorised form has visible signs.

Block 6. **Cue.** $F = q\,x(x+1)\cdots(x+m)\prod(x+s_j)$, $q > 0$, $s_j > 0$.

Why natural. A product of polynomials with nonnegative coefficients has nonnegative coefficients (every coefficient of the product is a sum of products of coefficients of the factors). This is why real-rootedness with nonpositive roots is the mechanism behind positivity of coefficients.

Link to the next block. Done; the answer is a structure, not a formula.

Layer 3: WHY

The central truth: a polynomial with positive leading coefficient and all roots real and nonpositive has nonnegative coefficients, and this is the only general mechanism that produces nonnegative coefficients from structure rather than from computation. The converse is false ($x^2 + x + 1$), but in combinatorics the polynomials whose coefficient positivity is asked for are very often real-rooted, and the examiners know it: this problem's polynomial has roots exactly at $0, -1, \ldots, -m$ and one in each of $(-1,0), (-2,-1), \ldots, (-k,-k+1)$, an interlacing pattern that is the fingerprint of a rising-factorial construction.

The second truth is that the summation operator $g \mapsto F$ (with $F(x) - F(x-1) = g(x)$, $F(0) = 0$) interacts cleanly with zeros of g at consecutive integers: a run of zeros of g at $0, -1, \ldots, -(m-1)$ gives a run of zeros of F at $0, \ldots, -m$, one longer. That is why the rising factorial, whose zeros are a run of consecutive integers, is the right building block, and why the problem posed the count with weak inequalities (multisets, rising factorials) rather than strict ones (sets, falling factorials with zeros at $0, 1, \ldots$, which would put the free roots at positive integers, the wrong side).

The discernment mechanism: when asked for nonnegative coefficients of a polynomial that is defined by counting, do not compute the polynomial; ask where its roots are. Cheap roots come from evaluating a defining identity at points where one side vanishes. Expensive roots come from sign changes at consecutive integers. Train it by taking the order polynomials of small posets (chains, antichains, the Λ shape of this problem) and locating their roots by hand; you will find the pattern of roots at nonpositive integers and the sign alternation at the remaining ones recurring.

What was the examiners' nudge? The phrase "nonnegative coefficients" itself, which is never a numerical statement, plus the symmetric design in (k, m) that makes "assume $k \le m$" available, plus a right-hand side with $k + m$ known zeros. Everything you needed was planted at $x \le 0$.

Revolving concepts

Multiset count and rising factorial. The number of weakly increasing j-sequences from $\{1,\dots,z\}$ is $\binom{z+j-1}{j} = R(z,j)$, the rising factorial $z^{\bar{j}}/j!$, with zeros at $0,-1,\dots,-(j-1)$ and the identities $R(z,j) - R(z-1,j) = R(z,j-1)$ and $\sum_{z=1}^{n} R(z,j) = R(n,j+1)$. Contrast the falling factorial $\binom{z}{j}$ with zeros at $0,1,\dots,j-1$. See Book 1, field: Combinatorics, topic: stars and bars, rising and falling factorials.

The summation operator and difference equations. For a polynomial g of degree d there is a unique polynomial F of degree $d+1$ with $F(x) - F(x-1) = g(x)$ and $F(0) = 0$, and $F(n) = \sum_{z=1}^{n} g(z)$. Generalised recursion: any linear recurrence $F(x) = F(x-1) + g(x)$ has solution $F(x) = F(x_0) + \sum g$, and its evaluation at zeros of g transports zeros of F; more generally for $a(x)F(x) = b(x)F(x-1) + g(x)$ the zeros of g give two-term relations whose signs are read from a and b. See Book 1, field: Polynomials, topic: difference operators on polynomials.

Real-rooted polynomials and coefficient positivity. If p has positive leading coefficient and all roots real and in $(-\infty, 0]$ then $p = c\prod(x+s_i)$ has nonnegative coefficients; moreover the coefficient sequence is then log-concave and unimodal (Newton's inequalities). Knowledge gap, beyond A-level: install the statement, the proof by expanding the product, and the standard method of proving real-rootedness by sign changes at $k+1$ points. See Book 1, field: Polynomials, topic: real-rootedness, interlacing, Newton's inequalities.

Order polynomials. For a finite poset P the number $\Omega_P(n)$ of order-preserving maps $P \to \{1,\dots,n\}$ is a polynomial in n of degree $|P|$ with leading coefficient $e(P)/|P|!$, $e(P)$ the number of linear extensions; here P is the Λ-shaped tree with $k+m+1$ elements and $e(P) = \binom{k+m}{k}$, matching the leading coefficient $\frac{1}{k!\,m!\,(k+m+1)}$ of f. Beyond A-level; know it exists. See Book 1, field: Discrete structures and graphs, topic: posets and order polynomials.

Circuit Bank

- **WHEN** "polynomial with nonnegative coefficients" **DO** try to prove all roots are real and nonpositive with positive leading coefficient **BECAUSE** a product of factors $(x+s)$, $s \geq 0$, has nonnegative coefficients, and this is the structural mechanism
- **WHEN** a polynomial defined as $\sum_{z=1}^{n} g(z)$ **DO** pass to the polynomial F with $F(x) - F(x-1) = g(x)$, $F(0) = 0$, and evaluate this identity at the zeros of g **BECAUSE** the identity holds on all of $\mathbb{R}$, and zeros of g propagate to zeros of F
- **WHEN** a two-term relation $a_j Q(-j) = -b_j Q(-j-1)$ with $a_j, b_j > 0$ **DO** check nonvanishing by a degree count, then apply the intermediate value theorem on each unit interval **BECAUSE** sign alternation at $k+1$ consecutive points gives k real roots
- **WHEN** a count symmetric in two parameters k, m where one step needs $k \leq m$ **DO** say "by symmetry assume $k \leq m$" at the start and point to the line that uses it **BECAUSE** unstated assumptions lose marks; stated ones cost one sentence
- **WHEN** weak inequalities in a counting problem **DO** write the count as a rising factorial $R(z,j)$ with its zeros at $0,-1,\dots$ **BECAUSE** the sign pattern of the answer lives at nonpositive integers

Backtesting

1. First useful cue: the two capped chains are multisets, so the summand is a product of rising factorials with zeros at $0,-1,\dots$.
2. Method the expression suggested: root location through the difference equation, then real-rootedness implies nonnegative coefficients.

3. Where you would hesitate, and the thought to reinforce: after finding the roots $0, \ldots, -m$, wondering how to find the other k. The thought: repeat the same move one level down; evaluate the reduced difference equation at the zeros of $R(x, k)$.
4. The trap the question sets: proving polynomiality and then computing coefficients through Faulhaber or the hockey stick, which produces alternating expressions and no sign conclusion.
5. What makes you recognise it faster next time: "nonnegative coefficients" is a root statement; the design $k \leq m$ by symmetry and a right-hand side with $k + m$ known zeros is the fingerprint of the intended proof.
6. The reusable decision rule: for a counting polynomial, coefficient positivity is proved by locating all roots on the nonpositive real axis; find the cheap roots at the zeros of the defining identity and the remaining ones by sign alternation.

Mirror drills

1. Show that $\sum_{z=1}^{n} z^4$, as a polynomial in n, has a negative coefficient, so "partial sums of polynomials with nonnegative coefficients" is not the mechanism at work in this problem. *Answer:* $\frac{n(n+1)(2n+1)(3n^2+3n-1)}{30} = \frac{6n^5+15n^4+10n^3-n}{30}$; its roots are all real but one of them, $\frac{-3+\sqrt{21}}{6}$, is positive, which is exactly what a negative coefficient forces.
2. Prove that a polynomial with positive leading coefficient and all roots real and nonpositive has nonnegative coefficients, and give an example showing the converse fails. *Answer:* $p = c\prod(x + s_i)$ with $c > 0$, $s_i \geq 0$, expand; converse fails for $x^2 + x + 1$.
3. For $k = m = 2$ verify the prediction of the proof: $f(n) = \frac{n(n+1)(n+2)(3n^2+6n+1)}{60}$, and the roots of $3x^2 + 6x + 1$ are $-1 \pm \sqrt{2/3}$, one in $(-1, 0)$ and one in $(-2, -1)$. *Hint:* $R(z, 2)^2 = \frac{z^2(z+1)^2}{4}$; sum with the formulas for $\sum z^2, \sum z^3, \sum z^4$ and factor.

2.12 B6. Complete the square and find the Gaussian peak

For a real number a, let $F_a(x) = \sum_{n\geq 1} n^a e^{2n} x^{n^2}$ for $0 \leq x < 1$. Find a real number c such that

$$\lim_{x\to 1^-} F_a(x)e^{-1/(1-x)} = 0 \qquad \text{for all } a < c, \text{ and}$$
$$\lim_{x\to 1^-} F_a(x)e^{-1/(1-x)} = \infty \qquad \text{for all } a > c.$$

Read it out loud

A power series in x with only square exponents, a partial theta series, weighted by e^{2n} and a power n^a, and you are asked how fast it blows up as $x \to 1^-$, measured against $e^{1/(1-x)}$. So the question is an asymptotic one: the series is a sum of positive terms, and its size near $x = 1$ is governed by its largest terms. The answer c is a threshold exponent: below it the series is $o(e^{1/(1-x)})$, above it the series dominates. The two limits together say $F_a(x) \approx e^{1/(1-x)} \cdot (1-x)^{-(a-c)} \cdot (\text{something bounded})$, and your job is to find c and prove the two one-sided statements.

The fields are real analysis (asymptotics of series, comparison of sums with integrals) and the single most important reflex of analytic estimation, Laplace's method: locate the maximum of the summand, measure its height and its width, multiply. What is really being tested is whether you can complete a square in the exponent. Write the summand as $n^a \exp(2n + n^2 \ln x)$. Near $x = 1$, $\ln x \approx -(1-x)$; put $t = \frac{1}{1-x}$ so that $\ln x \approx -\frac{1}{t}$ and the exponent is $2n - \frac{n^2}{t} = t - \frac{(n-t)^2}{t}$. There it is: the peak is at $n = t$, the peak height is $e^t = e^{1/(1-x)}$, exactly the normalising factor in the problem, and the width of the peak is $\sqrt{t}$. So after normalising, the sum is roughly $t^a \cdot \sqrt{t} \cdot \sqrt{\pi}$, and the threshold is $a + \frac{1}{2} = 0$, that is $c = -\frac{1}{2}$.

The first useful cue is the pair e^{2n} and x^{n^2} in the same term, a linear and a quadratic exponent in n which together form a square. Within thirty seconds you should have written $t = \frac{1}{1-x}$ and $2n - \frac{n^2}{t} = t - \frac{(n-t)^2}{t}$, and guessed $c = -\frac{1}{2}$. The rest of the arena time is rigour: two-sided bounds on $\ln x$, a lower bound from a window of width $\sqrt{t}$, and an upper bound from a split of the sum.

Approaches, ranked worst to best

1. **Closed form for the theta-type sum.** *Cue:* square exponents x^{n^2}. *Circuit:* Jacobi theta functions and their modular transformation $\vartheta(1/\tau)$. (the sum runs over $n \geq 1$ only and carries the weight $n^a e^{2n}$, so no transformation formula applies; even for the pure theta function the modular relation is a knowledge gap you cannot prove in the arena; stalls)

2. **Replace the sum by an integral.** *Cue:* a sum of a smooth, unimodal function of n. *Circuit:* $\int_0^\infty u^a e^{2u} x^{u^2}\, du$ by Laplace's method. (gives the correct heuristic instantly, but the rigorous sum-to-integral comparison must be done on monotone pieces on either side of the peak, and the peak is at a non-integer; more bookkeeping than the direct estimate, so it ranks below it, though it is the right way to *guess*)

3. **Peak window with ε-widths.** *Cue:* after completing the square, the summand is a discrete Gaussian around $n = t$. *Circuit:* split into $|n - t| < t^{1/2+\varepsilon}$ and the two tails; bound each. (the official route; correct, and needs an ε for the upper bound because it bounds the window by its width times its maximum; slightly more bookkeeping than needed)

4. **Complete the square, then window for the lower bound and a half-line split plus integral comparison for the upper bound.** *Cue:* $e^{2n}x^{n^2}$ and the factor $e^{-1/(1-x)}$. *Circuit:* $t = \frac{1}{1-x}$, $-\frac{1}{t} - \frac{1}{t^2} \leq \ln x \leq -\frac{1}{t}$; lower bound from $t \leq n \leq t + \sqrt{t}$; upper bound from $n < \frac{t}{2}$ (negligible) and $n \geq \frac{t}{2}$ (Gaussian sum $\leq 4\sqrt{t}$). (half a page, every constant explicit, no ε; the ideal)

The ideal solution

The answer is $c = -\dfrac{1}{2}$.

Step 1: the substitution and the logarithm bounds. For $x \in [0,1)$ put $t = \frac{1}{1-x} \in [1,\infty)$, so $x = 1 - \frac{1}{t}$ and $x \to 1^-$ is $t \to \infty$; the factor $e^{-1/(1-x)}$ is e^{-t}. For $t \geq 2$, with $u = \frac{1}{t} \in (0, \frac{1}{2}]$,

$$\ln x = \ln(1-u) = -\sum_{j\geq 1} \frac{u^j}{j} \leq -u, \qquad \ln x \geq -u - \frac{1}{2}\sum_{j\geq 2} u^j = -u - \frac{u^2}{2(1-u)} \geq -u - u^2,$$

using $1 - u \geq \frac{1}{2}$. Hence, for $t \geq 2$,

$$-\frac{1}{t} - \frac{1}{t^2} \leq \ln x \leq -\frac{1}{t}.$$

Step 2: completing the square. The nth term of $F_a(x)e^{-t}$ is $n^a \exp(2n - t + n^2 \ln x)$. Since $2n - t - \frac{n^2}{t} = -\frac{(n-t)^2}{t}$, Step 1 gives, for all $n \geq 1$ and $t \geq 2$,

$$n^a\, e^{-(n-t)^2/t}\, e^{-n^2/t^2} \;\leq\; n^a e^{2n-t} x^{n^2} \;\leq\; n^a\, e^{-(n-t)^2/t}.$$

Step 3: $a > -\frac{1}{2}$ forces the limit ∞. Let $t \geq 4$ and let N be the set of integers n with $t \leq n \leq t + \sqrt{t}$. An interval of length L contains at least $\lfloor L \rfloor$ integers, so $|N| \geq \lfloor \sqrt{t} \rfloor \geq \sqrt{t} - 1 \geq \frac{\sqrt{t}}{2}$. For $n \in N$: $\frac{(n-t)^2}{t} \leq 1$; $\frac{n^2}{t^2} \leq \left(1 + t^{-1/2}\right)^2 \leq \frac{9}{4}$; and $n^a \geq \frac{t^a}{\sqrt{2}}$, because if $a \geq 0$ then $n^a \geq t^a$, while if $-\frac{1}{2} < a < 0$ then $n^a \geq (2t)^a = 2^a t^a \geq 2^{-1/2} t^a$. Keeping only the terms $n \in N$ in the lower bound of Step 2,

$$F_a(x)e^{-t} \;\geq\; |N| \cdot \frac{t^a}{\sqrt{2}}\, e^{-1} e^{-9/4} \;\geq\; \frac{e^{-13/4}}{2\sqrt{2}}\, t^{a+\frac{1}{2}},$$

which tends to ∞ as $t \to \infty$ since $a + \frac{1}{2} > 0$.

Step 4: $a < -\frac{1}{2}$ forces the limit 0. Let $t \geq 2$ and use the upper bound of Step 2, splitting the sum at $\frac{t}{2}$.

For $1 \leq n < \frac{t}{2}$: $n^a \leq 1$ (as $a < 0$) and $(n-t)^2 > \frac{t^2}{4}$, so each term is at most $e^{-t/4}$; there are fewer than $\frac{t}{2}$ such n, so this part is at most $\frac{t}{2} e^{-t/4}$.

For $n \geq \frac{t}{2}$: $n^a \leq \left(\frac{t}{2}\right)^a = 2^{-a} t^a$ (as $a < 0$), so this part is at most $2^{-a} t^a \sum_{n\in\mathbb{Z}} e^{-(n-t)^2/t}$. To bound the Gaussian sum, note that the integers n with $|n - t| < 1$ number at most 2 and each contributes at most 1. The function $u \mapsto e^{-(u-t)^2/t}$ is decreasing on $[t,\infty)$, so for each integer $n \geq t+1$ we have $e^{-(n-t)^2/t} \leq \int_{n-1}^{n} e^{-(u-t)^2/t}\, du$, and the intervals $[n-1, n]$ for these n are non-overlapping subsets of $[t,\infty)$; hence $\sum_{n\geq t+1} e^{-(n-t)^2/t} \leq \int_t^\infty e^{-(u-t)^2/t}\, du = \frac{\sqrt{\pi t}}{2}$. Symmetrically $\sum_{n\leq t-1} e^{-(n-t)^2/t} \leq \frac{\sqrt{\pi t}}{2}$. Therefore

$$\sum_{n\in\mathbb{Z}} e^{-(n-t)^2/t} \leq 2 + \sqrt{\pi t} \leq 4\sqrt{t} \qquad (t \geq 1),$$

and the part $n \geq \frac{t}{2}$ is at most $4 \cdot 2^{-a}\, t^{a+\frac{1}{2}}$.

Altogether, for $t \geq 2$,

$$0 \leq F_a(x)e^{-t} \leq \frac{t}{2} e^{-t/4} + 4 \cdot 2^{-a}\, t^{a+\frac{1}{2}} \longrightarrow 0 \qquad (t \to \infty),$$

since $a + \frac{1}{2} < 0$.

Conclusion. With $c = -\frac{1}{2}$, the limit is 0 for every $a < c$ and ∞ for every $a > c$:

$$\boxed{c = -\tfrac{1}{2}}.$$

(Any other value c' fails: for a strictly between c' and $-\frac{1}{2}$ one of the two requirements is violated by Steps 3 and 4, so the answer is unique.) $\square$

Layer 1: WHAT

Block 1: substitute $t = \frac{1}{1-x}$ and bound $\ln x$ between $-\frac{1}{t} - \frac{1}{t^2}$ and $-\frac{1}{t}$. Block 2: complete the square: $2n - t - \frac{n^2}{t} = -\frac{(n-t)^2}{t}$, so each normalised term is a discrete Gaussian of centre t and width $\sqrt{t}$, times n^a, up to a bounded factor e^{-n^2/t^2} near the centre. Block 3: lower bound for $a > -\frac{1}{2}$: keep only the $\sim \sqrt{t}$ terms with $t \le n \le t + \sqrt{t}$, each of size $\asymp t^a$. Block 4: upper bound for $a < -\frac{1}{2}$: the terms with $n < \frac{t}{2}$ are exponentially small; for $n \ge \frac{t}{2}$ bound n^a by its largest value and the Gaussian sum by $4\sqrt{t}$ via integral comparison. Block 5: read off $c = -\frac{1}{2}$ from the exponent $a + \frac{1}{2}$ common to both bounds.

Nominal route: try to evaluate or transform the series, or expand $x^{n^2} = e^{n^2 \ln x}$ and lose the structure in a double sum. Ideal route: identify the largest terms and estimate the sum by "height times width". The bifurcation point is Block 2: the zombie sees a series and wants a formula for it; you see a sum of positive terms and want its peak.

Layer 2: HOW

Block 1. **Cue.** The normalising factor $e^{-1/(1-x)}$ and the behaviour $\ln x \approx -(1-x)$ as $x \to 1$.

Why natural. The problem names the quantity $\frac{1}{1-x}$, so it is the natural large parameter; and you need $\ln x$ in the exponent because $x^{n^2} = e^{n^2 \ln x}$. A two-sided bound on $\ln(1 - \frac{1}{t})$ by the first one or two terms of its series is the standard way to make "$\approx$" rigorous, and the second-order term $\frac{1}{t^2}$ is kept because it multiplies $n^2 \approx t^2$ at the peak and so contributes a bounded, not negligible, factor.

Link to the next block. With $\ln x$ replaced by $-\frac{1}{t}$ the exponent is a quadratic in n, and a quadratic exponent is completed as a square, always.

Block 2. **Cue.** $2n + n^2 \ln x$: linear plus quadratic in n with a negative quadratic coefficient.

Why natural. Completing the square is the definition-level move for a quadratic; it exposes the vertex $n = t$ and the value at the vertex, t, which is exactly $\frac{1}{1-x}$. This is the moment you know the examiners built the problem backwards from $e^{1/(1-x)}$: the normalisation removes the peak height, leaving only the peak width and the weight n^a at the peak. Guess $c = -\frac{1}{2}$ here, before any rigour.

Link to the next block. A Gaussian of width $\sqrt{t}$ has about $\sqrt{t}$ terms of size comparable to the maximum; the lower bound counts them and the upper bound must show the rest do not matter.

Block 3. **Cue.** You need a lower bound tending to infinity and you only need finitely many terms of a positive series.

Why natural. Dropping terms from a series of positive terms gives a lower bound; the window $[t, t + \sqrt{t}]$ is chosen because on it every factor is bounded below by an explicit constant times t^a (the Gaussian factor by e^{-1}, the correction e^{-n^2/t^2} by $e^{-9/4}$, the weight n^a by $t^a/\sqrt{2}$). The count of integers in an interval of length $\sqrt{t}$ supplies the factor $\sqrt{t}$.

Link to the next block. The upper bound is harder because you cannot drop terms; you must control all of them.

Block 4. **Cue.** Terms far from the peak are exponentially small; near the peak n^a is at most its value at the left end of the range because $a < 0$.

Why natural. Splitting at $\frac{t}{2}$ is the coarsest split that works: below it the Gaussian factor is at most $e^{-t/4}$ and the number of terms is polynomial, so the product is negligible; above it $n^a \le (\frac{t}{2})^a$ is uniform, and what remains is a pure Gaussian sum, which is compared with the Gaussian integral on the two monotone sides of the peak. The bound $4\sqrt{t}$ is crude and sufficient; precision is not the goal, the exponent of t is.

Link to the next block. Both bounds carry $t^{a+\frac{1}{2}}$; the sign of $a + \frac{1}{2}$ decides which limit you get, so $c = -\frac{1}{2}$.

Block 5. Cue. Two one-sided statements with the same exponent.

Why natural. The threshold is where the exponent $a + \frac{1}{2}$ changes sign; both statements must be proved, since the problem defines c by both, and uniqueness is a one-line remark.

Link to the next block. End.

Layer 3: WHY

The truth beneath the route is Laplace's method for sums: if a sum of positive terms has a summand of the form $h(n)e^{\phi(n)}$ with ϕ concave and peaked at n_0 with $\phi''(n_0) \approx -\frac{1}{\sigma^2}$, then the sum is approximately $h(n_0)e^{\phi(n_0)}\sigma\sqrt{2\pi}$, provided h varies slowly on the scale σ and $\sigma \gg 1$ so that the sum behaves like an integral. Here $\phi(n) = 2n - \frac{n^2}{t}$, $n_0 = t$, $\phi(n_0) = t$, $\sigma = \sqrt{t/2}$, $h(n) = n^a$. Everything in the problem is a special case of that sentence, and the examiners chose e^{2n} rather than e^n so that the vertex value $\frac{b^2t}{4}$ with $b = 2$ equals t exactly, making the normalising factor the clean $e^{1/(1-x)}$.

The second truth is about what "lim $= 0$ for $a < c$ and lim $= \infty$ for $a > c$" really asks: it is a statement that the normalised sum is $t^{a+\frac{1}{2}}$ times something bounded above and below by positive constants. You are not asked to find that constant (it is $\sqrt{\pi}e^{-1/2}$ at $a = -\frac{1}{2}$, from the second-order term of the logarithm), and you should not try; the arena rewards the exponent, and every estimate in the ideal solution is deliberately crude in the constant and sharp in the exponent.

The discernment mechanism is a fixed protocol for any series of positive terms near a singular point: (1) name the large parameter (here t); (2) write the summand as e^{exponent} times a slowly varying weight; (3) find the maximum of the exponent and its second derivative there, giving the peak location, height and width; (4) guess the answer as height times width times weight at the peak; (5) prove a lower bound by a window and an upper bound by a split. Train it on $\sum x^{n^2}$, $\sum x^n/n$, $\sum n^a x^n$, and $\sum e^{bn}x^{n^2}$ until step (3) is reflex.

Knowledge gaps, stated with brutal honesty. First, Laplace's method itself, both for integrals and for sums, is not in A-level or Further Mathematics; you must install the heuristic and the two standard rigorous devices (window for the lower bound, split and integral comparison for the upper bound). Second, the comparison of a monotone sum with an integral, including the Gaussian integral $\int_{-\infty}^{\infty} e^{-u^2/t}\,du = \sqrt{\pi t}$, must be automatic. Third, the two-sided series bounds for $\ln(1-u)$, with the reason the second-order term matters when it multiplies $n^2 \approx t^2$. Fourth, the vocabulary: partial theta series, and the fact that for the full theta function $\sum_{n\in\mathbb{Z}} x^{n^2}$ the growth $\sqrt{\pi/(1-x)}$ is a special case of the Jacobi transformation; you should know this exists so that you do not waste arena time trying to derive a closed form that does not exist for the weighted, one-sided sum. Without the first two of these, the problem is not approachable; with them it is a B6 only because of the bookkeeping.

Revolving concepts

Completing the square in an exponent. $bn - \frac{n^2}{t} = \frac{b^2t}{4} - \frac{(n-\frac{bt}{2})^2}{t}$: peak at $n = \frac{bt}{2}$, height $e^{b^2t/4}$, width $\sqrt{t}$. This is the whole content of the Gaussian shape of $e^{bn}x^{n^2}$ near $x = 1$. See Book 1, field: Calculus and integrals, topic: Laplace's method and Gaussian peaks.

Laplace's method (sums and integrals). For $\int h(u)e^{\lambda\phi(u)}\,du$ with ϕ having a unique nondegenerate maximum at u_0: $\sim h(u_0)e^{\lambda\phi(u_0)}\sqrt{\frac{2\pi}{\lambda|\phi''(u_0)|}}$. For sums, the same when the width is large compared with 1. Rigorous versions: lower bound by restricting to a window, upper bound by splitting and comparing monotone pieces with integrals. Beyond A-level. See Book 1, field: Real analysis, topic: asymptotic analysis, Laplace's method.

Sum-to-integral comparison. If ψ is decreasing on $[A, \infty)$ then $\sum_{n\geq A+1}\psi(n) \leq \int_A^{\infty}\psi$; if increasing on $(-\infty, B]$ then $\sum_{n\leq B-1}\psi(n) \leq \int_{-\infty}^{B}\psi$; a unimodal function is handled by splitting at its mode and adding a bounded number of terms near the mode. The Gaussian integral $\int_{\mathbb{R}} e^{-u^2/t}\,du = \sqrt{\pi t}$. See Book 1, field: Real analysis, topic: integral test and monotone comparison.

Two-sided logarithm bounds. For $0 < u \leq \frac{1}{2}$: $-u - u^2 \leq \ln(1-u) \leq -u$; more generally $\ln(1-u) = -u - \frac{u^2}{2} + O(u^3)$. The second-order term is retained whenever it is multiplied by a quantity of size u^{-2}. See Book 1, field: Inequalities, topic: logarithm and exponential bounds.

Theta series. $\vartheta(x) = \sum_{n\in\mathbb{Z}} x^{n^2}$ satisfies $\vartheta(x) \sim \sqrt{\frac{\pi}{1-x}}$ as $x \to 1^-$, a consequence of the Jacobi transformation (Poisson summation). The problem's series is a one-sided, weighted partial theta series with no transformation law; the growth must be found by peak analysis. Beyond A-level; know the statement, not the proof. See Book 1, field: Complex numbers, topic: theta functions and modular transformation (reference only).

Recursion, generalised. There is no recursion, but there is a generalised family: $G_{a,b}(x) = \sum n^a e^{bn} x^{n^2}$ satisfies $G_{a,b}(x)e^{-b^2/(4(1-x))} \asymp (1-x)^{-(a+\frac{1}{2})}$ for every $b > 0$, by the identical argument with the window $[\frac{bt}{2}, \frac{bt}{2} + \sqrt{t}]$ and the split at $\frac{bt}{4}$; the threshold $c = -\frac{1}{2}$ does not depend on b, and the constant does. With $b = 0$ the peak sits at $n = 0$ and the analysis changes character: $\sum_{n\geq 1} n^a x^{n^2}$ has a finite limit if and only if $a < -1$, and grows like $t^{(a+1)/2}$ for $a > -1$.

Circuit Bank

- **WHEN** a series of positive terms and a question about its size near a singular point **DO** name the large parameter, write the summand as a weight times e^{exponent}, and locate the maximum of the exponent **BECAUSE** a positive series is dominated by its largest terms, and Laplace's method estimates them as height times width
- **WHEN** e^{bn} and x^{n^2} in the same term **DO** substitute $t = \frac{1}{1-x}$, bound $\ln x$ two-sidedly, and complete the square $bn - \frac{n^2}{t}$ **BECAUSE** a quadratic exponent is a Gaussian, with centre $\frac{bt}{2}$, height $e^{b^2t/4}$ and width $\sqrt{t}$
- **WHEN** a normalising factor like $e^{-1/(1-x)}$ in the question **DO** expect it to equal the peak height exactly and read the remaining question as "what power of t survives" **BECAUSE** examiners build such problems backwards from the peak
- **WHEN** you need a lower bound on a positive sum **DO** keep only a window of terms around the peak, of length equal to the width, and bound each term below by a constant times the peak value **BECAUSE** dropping positive terms only decreases the sum
- **WHEN** you need an upper bound on a Gaussian-type sum **DO** split into the far tail (exponentially small times polynomially many) and the bulk (uniform weight times a Gaussian sum compared with its integral on each monotone side) **BECAUSE** the two devices together cost one paragraph and no ε

Backtesting

1. First useful cue: $e^{2n}x^{n^2}$, a linear and a quadratic exponent in the same term, together with the normalising factor $e^{-1/(1-x)}$.
2. Method the expression suggested: $t = \frac{1}{1-x}$, complete the square, Laplace's method for sums.
3. Where you would hesitate, and the thought to reinforce: when writing the upper bound, the temptation to bound the window by width times maximum and then to need an ε. The thought: bound n^a uniformly on $n \geq \frac{t}{2}$ and compare the remaining Gaussian sum with the Gaussian integral; no ε is needed.
4. The trap the question sets: trying to sum the series, or dropping the second-order term of $\ln x$ in the lower bound without noticing it multiplies $n^2 \approx t^2$; the factor e^{-n^2/t^2} is bounded, but you must say so.

5. What makes you recognise it faster next time: any x^{n^2} near $x = 1$ is a Gaussian in n of width $(1-x)^{-1/2}$; any extra e^{bn} just moves the centre to $\frac{b}{2(1-x)}$ and the height to $e^{b^2/(4(1-x))}$.

6. The reusable decision rule: for $\sum n^a e^{bn} x^{n^2}$ the normalised growth exponent is $a + \frac{1}{2}$; the threshold in a is $-\frac{1}{2}$, independent of $b > 0$, because the width $\sqrt{t}$ contributes $\frac{1}{2}$ and the weight at the peak contributes a.

Mirror drills

1. Show that $\sqrt{1-x}\sum_{n\geq 0} x^{n^2} \to \frac{\sqrt{\pi}}{2}$ as $x \to 1^-$. *Hint:* with $t = \frac{1}{1-x}$, the summand is $e^{n^2 \ln x}$ with $-\frac{1}{t} - \frac{1}{t^2} \leq \ln x \leq -\frac{1}{t}$; compare $\sum_{n\geq 0} e^{-n^2/t}$ with $\int_0^\infty e^{-u^2/t}\, du = \frac{\sqrt{\pi t}}{2}$ from both sides (the sum lies between the integral and the integral plus 1), and show the correction e^{-n^2/t^2} is $1 + o(1)$ uniformly for $n \leq t^{3/4}$ while the terms with $n > t^{3/4}$ are negligible.

2. For $b > 0$ fixed, find the threshold c such that $\sum n^a e^{bn} x^{n^2} \cdot e^{-b^2/(4(1-x))}$ tends to 0 for $a < c$ and to ∞ for $a > c$. *Answer:* $c = -\frac{1}{2}$ for every $b > 0$; complete the square $bn - \frac{n^2}{t} = \frac{b^2 t}{4} - \frac{(n - bt/2)^2}{t}$ and repeat the two bounds with the window $[\frac{bt}{2}, \frac{bt}{2} + \sqrt{t}]$.

3. Show that $\lim_{x\to 1^-} \sum_{n\geq 1} n^a x^{n^2}$ is finite if and only if $a < -1$. *Hint:* for $a < -1$ the series is dominated by $\sum n^a < \infty$ (monotone convergence); for $a \geq -1$ compare with $\int_1^{\sqrt{t}} u^a\, du$, which is $\asymp t^{(a+1)/2}$ for $a > -1$ and $\asymp \ln t$ for $a = -1$, using $x^{n^2} \geq e^{-2}$ for $n \leq \sqrt{t}$ and $t \geq 2$.

Chapter 3

Putnam 2023

3.1 A1. Second derivative of a cosine product at zero

For a positive integer n, let $f_n(x) = \cos(x)\cos(2x)\cos(3x)\cdots\cos(nx)$. Find the smallest n such that $|f_n''(0)| > 2023$.

Read it out loud

A product of n cosines whose frequencies are $1, 2, \ldots, n$; the second derivative of that product; evaluated at the single point 0; compared with a threshold. Say it that way and the structure is already talking to you. Nothing about f_n away from 0 matters. Only one number matters: the second Taylor coefficient of f_n at 0, because $f_n''(0) = 2\,[x^2]f_n(x)$ by definition of the Taylor coefficients.

Fields wired in: single-variable calculus (Taylor expansion, the product rule as a fallback), a closed form for $1^2 + 2^2 + \cdots + n^2$, and one line of monotonicity. Knowledge base: the expansion $\cos(mx) = 1 - \frac{m^2x^2}{2} + O(x^4)$, the fact that truncated power series multiply like polynomials, and $\sum_{m=1}^{n} m^2 = \frac{n(n+1)(2n+1)}{6}$. All of it is A-level; what is being tested is whether you differentiate what you were asked to differentiate (zombie) or read off the coefficient (ideal).

The first useful cue is "at 0" sitting next to a product of even functions each equal to 1 at 0. Within thirty seconds you should have written: $\cos(mx) = 1 - \frac{m^2}{2}x^2 + \ldots$, no x^1 term, and asked yourself: what is the x^2 coefficient of a product of factors of the form $1 + c_m x^2 + \ldots$? The answer is $\sum c_m$, because a cross term needs two non-constant factors and the cheapest non-constant term is already x^2. That single sentence is the whole problem. The number 2023 then forces an exact evaluation of the sum of squares, not an asymptotic; it was chosen so that $n = 18$ is the first crossing and so that a sloppy candidate who writes $n^3/3 > 2023$ gets $n = 19$ and loses the mark.

Approaches, ranked worst to best

1. **Product rule, twice.** *Cue:* the words "second derivative". *Circuit:* differentiate what you are asked to differentiate. (This is the nominal route. f_n'' is a sum of n terms in which one factor is differentiated twice plus $\binom{n}{2}$ terms in which two distinct factors are differentiated once each. At $x = 0$ every term of the second type contains a factor $\sin(0) = 0$; every term of the first type is $-m^2$ times a product of ones. The answer falls out, but you must write a paragraph classifying the terms, and the classification is exactly where a tired candidate miscounts. It is correct and it is the official write-up; it is still the most ink for the least structure.)

2. **Logarithmic derivative.** *Cue:* a product of factors that are positive near 0. *Circuit:* turn a product into a sum with log. (Near 0, $f_n > 0$, so $(\log f_n)' = f_n'/f_n = -\sum_{m=1}^{n} m\tan(mx)$. Differentiate once more: $f_n''/f_n = (\log f_n)'' + \big((\log f_n)'\big)^2 = -\sum m^2\sec^2(mx) + \big(\sum m\tan(mx)\big)^2$. At 0: $f_n(0) = 1$, $\tan 0 = 0$, $\sec 0 = 1$, so $f_n''(0) = -\sum m^2$. Two lines, fully rigorous, and it generalises to any product. It costs you the identity $f'' = f\,[(\log f)'' + ((\log f)')^2]$, which you must be able to reproduce cold.)

3. **Taylor coefficient extraction.** *Cue:* "at 0" next to a product of series with no linear term. *Circuit:* the k-th derivative at a point is $k!$ times the x^k coefficient; multiply truncated series. (The ideal route. Expand each factor to order 2, multiply, and note that the x^2 coefficient of the product is the sum of the individual x^2 coefficients because there are no x^1 terms to cross-multiply. No differentiation is performed at all. This is the laziest correct route and it is also the one that scales: the same move gives $f_n^{(4)}(0)$ in one more line.)

The ideal solution

For each m, the cosine series gives $\cos(mx) = 1 - \frac{m^2}{2}x^2 + x^4R_m(x)$, where R_m is an everywhere convergent power series. Multiply the n factors. Any term of the expanded product that uses at

least one $x^4 R_m(x)$, or at least two of the factors $-\frac{m^2}{2}x^2$, is divisible by x^4. The remaining terms are the product of all the ones, which is 1, and the terms with exactly one factor $-\frac{m^2}{2}x^2$ and ones elsewhere. Hence

$$f_n(x) = 1 - \frac{x^2}{2}\sum_{m=1}^{n} m^2 + x^4 R(x)$$

with R a power series convergent everywhere. By uniqueness of Taylor coefficients, $f_n''(0) = 2! \cdot \left(-\frac{1}{2}\sum_{m=1}^{n} m^2\right)$, thus

$$|f_n''(0)| = \sum_{m=1}^{n} m^2 = \frac{n(n+1)(2n+1)}{6} =: g(n).$$

The sequence $g(n)$ is strictly increasing, since $g(n+1) - g(n) = (n+1)^2 > 0$. Now $g(17) = \frac{17\cdot 18\cdot 35}{6} = 17 \cdot 3 \cdot 35 = 1785 \le 2023$ and $g(18) = \frac{18\cdot 19\cdot 37}{6} = 3 \cdot 19 \cdot 37 = 2109 > 2023$. By monotonicity, $g(n) \le 1785$ for all $n \le 17$ and $g(n) \ge 2109$ for all $n \ge 18$. The smallest n with $|f_n''(0)| > 2023$ is therefore

$$\boxed{n = 18}.$$

Layer 1: WHAT

Block 1. Translate "$f_n''(0)$" into "twice the x^2 coefficient of f_n at 0".

Block 2. Expand each factor to order two: $\cos(mx) = 1 - \frac{m^2}{2}x^2 + O(x^4)$; record that there is no x^1 term.

Block 3. Multiply: the x^2 coefficient of the product is the sum of the x^2 coefficients of the factors, because the cheapest cross term is $x^2 \cdot x^2 = x^4$.

Block 4. Sum of squares in closed form, one line of monotonicity, two evaluations bracketing 2023.

Juxtaposition. The nominal route reads "second derivative" and starts differentiating: product rule, n factors, then product rule again on each of the n resulting products, then evaluation. It arrives at the same $-\sum m^2$, after a page. The ideal route never differentiates. The bifurcation point is the first pen stroke: do you write $f_n'(x) =$ (zombie) or do you write $\cos(mx) = 1 - \frac{m^2}{2}x^2 + \ldots$ (ideal)? Everything downstream is decided there.

Layer 2: HOW

Block 1. Cue. The derivative is requested at one point, 0, and the function is a finite product of entire functions. **Why natural.** A derivative at a point is a local object; the local object at 0 is the Taylor series, and $f''(0)/2!$ is by definition its x^2 coefficient. Asking for $f_n''(0)$ is asking for one coefficient. It is worth repeating: when a question evaluates a derivative at a specific point, the question is about a coefficient, not about a function. **Link to the next block.** Once the target is a coefficient of a product, the natural question is how coefficients of products are built from coefficients of factors.

Block 2. Cue. Each factor is $\cos(mx)$, an even function with value 1 at 0. **Why natural.** Even functions have no odd Taylor terms, so the expansion of each factor starts $1 + c_m x^2$ with $c_m = -m^2/2$. The absence of an x^1 term is the structural fact that makes the next block a one-liner; you note it explicitly because it is the reason, not a coincidence. **Link to the next block.** With every factor of the shape $1 + c_m x^2 + O(x^4)$, the product's x^2 coefficient can be read off without writing the product.

Block 3. **Cue.** A product of n series of the form $1 + c_m x^2 + O(x^4)$. **Why natural.** To produce an x^2 term from the product you must pick a non-constant term from at least one factor, and the smallest non-constant degree available is 2; picking non-constant terms from two factors already gives degree at least 4. So the x^2 coefficient is exactly $\sum c_m$. This is the "log-linearisation" of a product: $\prod(1+\epsilon_m) = 1 + \sum \epsilon_m + (\text{terms of order } \epsilon^2)$, and here the ϵ_m are already of order x^2. **Link to the next block.** $|f_n''(0)| = \sum m^2$; the calculus is over and a finite-sum problem remains.

Block 4. **Cue.** "Smallest n such that [quantity] > 2023". **Why natural.** A smallest-n question needs two facts: the quantity is monotone in n (so that "smallest" is meaningful and the answer is a single crossing), and the two consecutive values straddling the threshold. Monotonicity is one line, $g(n+1) - g(n) = (n+1)^2 > 0$. The evaluations at 17 and 18 are chosen by the estimate $n^3/3 \approx 2023$, so $n \approx 18$; then you check the neighbours exactly, because the estimate is not a proof. **Link to the next block.** The answer is 18, with a complete justification that no smaller n works and that 18 does.

Layer 3: WHY

The truth beneath the route is that truncation is compatible with multiplication: the map sending a power series to its remainder modulo x^k is a ring homomorphism, so to know a product modulo x^4 you only need each factor modulo x^4. That is why "multiply the truncated series" is a proof and not a heuristic. The second truth is additivity to second order: a product of factors $1 + \epsilon_m$ with $\epsilon_m = O(x^2)$ equals $1 + \sum \epsilon_m$ modulo x^4. There is a probabilistic reading that makes this feel inevitable. Let $X = \epsilon_1 \cdot 1 + \epsilon_2 \cdot 2 + \cdots + \epsilon_n \cdot n$ with independent fair signs $\epsilon_m = \pm 1$. Then $\mathbb{E}[e^{iXx}] = \prod_m \mathbb{E}[e^{i\epsilon_m m x}] = \prod_m \cos(mx) = f_n(x)$, and the second derivative of a characteristic function at 0 is $-\mathbb{E}[X^2]$. Variances of independent summands add, so $-f_n''(0) = \sum m^2$. The additivity you exploited in Block 3 is the additivity of variance.

The discernment mechanism is a fixed question you ask of every derivative-at-a-point problem: which coefficient is being requested, and which coefficients of the ingredients feed it? For a product, the answer is a convolution of coefficient sequences, and low-order coefficients of the product depend only on low-order coefficients of the factors. For a composition, the same principle holds with Faà di Bruno replacing convolution. You train the discernment by refusing, for one month, to differentiate any product at a point; you expand instead, every time, until the reflex is installed.

What the examiners nudged toward: they put the evaluation at 0, where sin vanishes and cos is 1, so that the product rule collapses; they used frequencies $1, \ldots, n$ so that the answer is a textbook sum; and they chose 2023 so that the crossing is at 18 rather than at a round number, which forces an exact evaluation and punishes the approximation $n^3/3$.

Revolving concepts

Taylor coefficient extraction. Definition: if f is k times differentiable at a, then $f^{(k)}(a) = k!\,[(x-a)^k]f$ where the bracket denotes the coefficient in the Taylor polynomial; for a function given by a convergent power series the Taylor coefficients are the series coefficients (uniqueness of power series). Orbiting facts: truncation modulo $(x-a)^k$ is a ring homomorphism; even functions have zero odd coefficients at 0; $[x^2]\prod_m(1 + a_m x + b_m x^2 + \ldots) = \sum_m b_m + \sum_{i<j} a_i a_j$, which is the general form of Block 3 with the elementary symmetric function e_2 of the linear coefficients appearing as soon as linear terms exist. See Book 1, field: Calculus and integrals, topic: Taylor expansion of products.

The recursion generalised. Coefficients of a product obey the Cauchy convolution $[x^k]\prod_{m=1}^n g_m = \sum_{k_1+\cdots+k_n=k} \prod_m [x^{k_m}] g_m$. Two general solution methods: (i) generating functions, where the product is the object itself; (ii) logarithms, where $\log \prod g_m = \sum \log g_m$ turns the multiplicative recursion into an additive one, and exponentiating back is the standard way to compute power sums from

elementary symmetric functions (Newton's identities live here). The logarithmic-derivative route of Approach 2 is this second method applied to two coefficients only. See Book 1, field: Generating functions, topic: products and logarithmic derivatives.

Characteristic functions. Definition: $\varphi_X(x) = \mathbb{E}[e^{iXx}]$. Orbiting theorems: $\varphi_{X+Y} = \varphi_X \varphi_Y$ for independent X, Y; $\varphi_X^{(k)}(0) = i^k \mathbb{E}[X^k]$ when the moment exists; a fair sign times m has characteristic function $\cos(mx)$. This is the reason products of cosines appear in counting problems about signed sums and in the theory of the Rademacher functions. See Book 1, field: Probability, topic: characteristic functions of sums.

Sum of squares. $\sum_{m=1}^{n} m^2 = \frac{n(n+1)(2n+1)}{6}$, proved by telescoping $(m+1)^3 - m^3 = 3m^2 + 3m + 1$; the same telescoping gives every $\sum m^k$ recursively, which is the recursion behind Faulhaber's formulas. See Book 1, field: Sequences and recurrences, topic: power sums by telescoping.

Circuit Bank

- **WHEN** you see a derivative of a product (or composition) evaluated at one specific point **DO** write the Taylor expansion of every factor at that point to the required order and multiply the truncations **BECAUSE** because $f^{(k)}(a)$ is $k!$ times one coefficient, and truncation commutes with multiplication, so nothing outside those few coefficients can matter
- **WHEN** you see a product of factors each of the form $1 +$ (something of order x^2) **DO** write the x^2 coefficient as the plain sum of the individual x^2 coefficients **BECAUSE** because a cross term needs two non-constant factors and is therefore of order at least x^4; this is additivity of variance in disguise
- **WHEN** you see a product of cosines with different frequencies **DO** read it as the characteristic function of $\sum \pm m$ with independent fair signs **BECAUSE** because $\cos(mx) = \mathbb{E}[e^{i\epsilon mx}]$, so derivatives at 0 are moments and independence makes second moments add
- **WHEN** you see "smallest n such that [expression in n] exceeds a threshold" **DO** prove monotonicity in one line, estimate the crossing, then evaluate exactly at the two neighbouring integers **BECAUSE** because monotonicity plus two exact values is a complete proof, and an estimate alone is not

Backtesting

1. First useful cue: "$f_n''(0)$", a derivative at one point, next to a product of even functions equal to 1 at that point.
2. Method the expression suggested: expand each $\cos(mx)$ to order two, multiply truncations, read the x^2 coefficient; never apply the product rule.
3. Where you would hesitate, and the thought to reinforce: whether multiplying truncated series is legitimate. It is, because truncation modulo x^4 is a ring homomorphism; say that sentence and move on.
4. The trap the question sets: the threshold 2023 invites the approximation $\sum m^2 \approx n^3/3$, which suggests $n \approx 18.2$ and tempts you to answer 19; the exact values 1785 and 2109 settle it at 18.
5. What makes you recognise it faster next time: any "evaluate a derivative at a point" is a "read one Taylor coefficient" problem; any product of series with no linear terms has additive second coefficients.
6. The reusable decision rule: derivative at a point means coefficient; product of series means truncated multiplication; smallest n means monotonicity plus two exact evaluations.

Mirror drills

1. (drill) Let $g_n(x) = \prod_{m=1}^{n} \cosh(mx)$. Find $g_n''(0)$. *Answer:* $\cosh(mx) = 1 + \frac{m^2}{2}x^2 + O(x^4)$, so $g_n''(0) = \sum_{m=1}^{n} m^2 = \frac{n(n+1)(2n+1)}{6}$; only the sign changes.
2. (drill) Let $h_n(x) = \prod_{m=1}^{n}(1+mx)$. Find $h_n''(0)$. *Answer:* now the factors have linear terms, so $[x^2]h_n = e_2(1,\dots,n) = \sum_{i<j} ij = \frac{1}{2}\left[(\sum m)^2 - \sum m^2\right]$ and $h_n''(0) = (\sum m)^2 - \sum m^2$. This is the drill that shows why "no linear term" mattered in A1.
3. (drill) Express $f_n^{(4)}(0)$ for the A1 function in terms of $S_2 = \sum m^2$ and $S_4 = \sum m^4$. *Answer:* $[x^4]f_n = \sum_m \frac{m^4}{24} + \sum_{i<j} \frac{i^2 j^2}{4}$, so $f_n^{(4)}(0) = S_4 + 6\sum_{i<j} i^2 j^2 = 3S_2^2 - 2S_4$; check $n = 1$ gives 1, as $\cos^{(4)}(0) = 1$. Probabilistically this is $\mathbb{E}[X^4]$ for $X = \sum \epsilon_m m$.

3.2 A2. Reversing a polynomial to expose two missing roots

Let n be an even positive integer. Let p be a monic, real polynomial of degree $2n$; that is to say, $p(x) = x^{2n} + a_{2n-1}x^{2n-1} + \cdots + a_1x + a_0$ for some real coefficients $a_0, \ldots, a_{2n-1}$. Suppose that $p(1/k) = k^2$ for all integers k such that $1 \leq |k| \leq n$. Find all other real numbers x for which $p(1/x) = x^2$.

Read it out loud

A monic polynomial of degree $2n$; you are told its value at $2n$ specified points, namely the reciprocals $\pm 1, \pm\frac{1}{2}, \ldots, \pm\frac{1}{n}$; and the values are prescribed by the rule $p(1/k) = k^2$. Then you are asked for the other solutions of the same equation. Say that out loud and the arithmetic of degrees of freedom is already audible: p has exactly $2n$ unknown coefficients and you are handed exactly $2n$ conditions, so p is pinned down completely. A weak reading concludes "so I must find p". The correct reading is the opposite: the question asks only for the extra solutions, so the design of the problem is telling you that the extra solutions can be named without ever computing a single a_i.

Fields wired in: polynomials, the factor theorem and root counting, the reversal (reciprocal) operation on a polynomial, and Vieta's relations at the two ends of the coefficient list. The knowledge base needed is small and entirely nameable: a polynomial of degree d has exactly d roots in $\mathbb{C}$ with multiplicity; if $r_1, \ldots, r_d$ are those roots of a monic polynomial then the sum is minus the coefficient of x^{d-1} and the product is $(-1)^d$ times the constant term; and multiplying by a power of x converts an expression in $1/x$ into a polynomial in x.

The first useful cue is that the phrase "$p(1/x) = x^2$" is an equation, not a value, and you are told it holds at $2n$ named points. Any time a question hands you many points where two things agree, the object you should build is the single polynomial that vanishes at all of them. Within thirty seconds you should have written

$$q(x) := x^{2n+2} - x^{2n}p(1/x),$$

and then checked three things about it: it is a genuine polynomial, it is monic of degree $2n + 2$, and its roots are exactly the solutions of the equation you were asked about. Everything after that is bookkeeping at the two ends of q.

The second cue is the hypothesis "n is even". Competition hypotheses of that shape are never decoration. Somewhere a $(-1)^n$ is going to appear, and the parity is inserted to fix its sign. Hunt for that sign as you go; it is the difference between two real answers and no real answer at all.

Approaches, ranked worst to best

1. **Interpolate p explicitly, then solve.** *Cue:* "$2n$ conditions on $2n$ unknown coefficients". *Circuit:* when a linear system is determined, solve it. (This is the nominal route and it is a swamp. You would set up a $2n$ by $2n$ Vandermonde-type system in the a_i with entries k^{-j}, invert it symbolically for general even n, and then attempt to solve a degree $2n + 2$ equation with the resulting coefficients. Even for $n = 4$ the coefficients are ratios such as $83149/144$, and no pattern in n is visible from them. The route never terminates in the arena.)

2. **Clear the fraction on the other side.** *Cue:* "$p(1/k) = k^2$ says $p(y) = 1/y^2$ at $y = \pm 1/k$". *Circuit:* multiply through by the offending denominator to get a polynomial equation. (Set $P(y) = y^2p(y) - 1$. It is monic of degree $2n + 2$ and vanishes at the $2n$ points $y = \pm 1/k$, so $P(y) = \prod_{k=1}^{n}\left(y^2 - \frac{1}{k^2}\right) \cdot (y^2 + cy + d)$. Comparing the coefficients of y^1 and y^0, which are 0 and -1, gives $c = 0$ and $d = -(n!)^2$, so the extra roots are $y = \pm n!$ and hence $x = 1/y = \pm 1/n!$. This is correct and costs little. It ranks below the reversal for two reasons: the two coefficients you can read off are the two lowest, so the monic normalisation at the top is wasted, and you

finish in the variable y and must remember to invert at the end, which is exactly the sort of last-line slip that costs a point.)

3. **Reversal, then long division by the known factor.** *Cue:* "$p(1/x)$ next to x^2". *Circuit:* form the polynomial whose roots are the solutions, then divide out what you know. (Build $q(x) = x^{2n+2} - x^{2n}p(1/x)$, divide by $\prod_{k=1}^{n}(x^2 - k^2)$, and read the quotient. Correct, but polynomial division by a product of n quadratics is a page of ink for information that lives in two coefficients.)

4. **Reversal, then Vieta at the two ends.** *Cue:* a monic polynomial with all but two of its roots known, and both extreme coefficients visible. *Circuit:* sum and product of roots are free; everything in the middle costs ink. (The ideal route. The reversed polynomial q is monic, has no x^{2n+1} term and has constant term -1; those two facts say the roots sum to 0 and multiply to -1. The $2n$ known roots sum to 0 and multiply to $(n!)^2$, so the two unknown roots sum to 0 and multiply to $-1/(n!)^2$. Two lines, no division, and the parity hypothesis announces itself exactly where it is used.)

The ideal solution

Since n is an even positive integer, $n \geq 2$.

Write $p(x) = x^{2n} + a_{2n-1}x^{2n-1} + \cdots + a_0$ and set

$$q(x) := x^{2n+2} - x^{2n}p(1/x).$$

Because $x^{2n}p(1/x) = 1 + a_{2n-1}x + a_{2n-2}x^2 + \cdots + a_0x^{2n}$, the function q is a polynomial. It is monic of degree $2n+2$, its coefficient of x^{2n+1} is 0, and $q(0) = -1$.

For $x \neq 0$ we have $q(x) = x^{2n}\big(x^2 - p(1/x)\big)$, so $q(x) = 0$ if and only if $p(1/x) = x^2$. Also $q(0) = -1 \neq 0$. Hence the real numbers x with $p(1/x) = x^2$ are precisely the real roots of q.

By hypothesis the $2n$ distinct numbers $\pm1, \pm2, \ldots, \pm n$ are roots of q, so

$$q(x) = \Big(\prod_{k=1}^{n}(x^2 - k^2)\Big)(x^2 + ax + b)$$

for some real a, b: the quotient of the monic real polynomial q by the monic real polynomial $\prod_{k=1}^{n}(x^2 - k^2)$ is monic, real and of degree 2.

Let r and s be the two roots of $x^2 + ax + b$ in $\mathbb{C}$. The full root multiset of q is $\{\pm1, \ldots, \pm n, r, s\}$. Vieta's relations for the monic polynomial q of degree $2n+2$ give

$$\sum \text{roots} = -[x^{2n+1}]q = 0, \qquad \prod \text{roots} = (-1)^{2n+2}q(0) = -1.$$

The known roots contribute $\sum_{k=1}^{n}\big(k + (-k)\big) = 0$ to the sum and $\prod_{k=1}^{n} k(-k) = (-1)^n(n!)^2 = (n!)^2$ to the product, the last equality because n is even. Therefore

$$r + s = 0, \qquad (n!)^2\, rs = -1.$$

From the first, $s = -r$; substituting into the second gives $-r^2 = -1/(n!)^2$, that is $r^2 = 1/(n!)^2$, so $r = \pm 1/n!$ and $\{r, s\} = \{1/n!, -1/n!\}$. Both are real.

Finally, these are new solutions: $n \geq 2$ forces $0 < 1/n! \leq \frac{1}{2} < 1$, so $\pm 1/n!$ is none of $\pm1, \ldots, \pm n$. The other real numbers x with $p(1/x) = x^2$ are therefore

$$\boxed{x = \pm\frac{1}{n!}}.$$

Two remarks a grader will not demand but which cost one line each. First, such a p does exist and is unique: reading the argument backwards, $q(x) = \big(\prod_{k=1}^{n}(x^2 - k^2)\big)\big(x^2 - (n!)^{-2}\big)$ is even, so

it has no x^{2n+1} term, and its constant term is $(-1)^n(n!)^2 \cdot \big(-(n!)^{-2}\big) = -1$; hence $x^{2n+2} - q(x)$ has degree at most $2n$ and constant term 1, which is exactly a reversed monic p of degree $2n$. Second, the hypothesis that n is even is load-bearing and not cosmetic: for odd n the product of the known roots is $-(n!)^2$, the same computation gives $r^2 = -1/(n!)^2$, and the two extra roots are the non-real numbers $\pm i/n!$, so an odd n would have the answer "there are no other real solutions".

Layer 1: WHAT

Block 1. Read "$p(1/k) = k^2$ for $2n$ values of k" as "one equation with $2n$ known solutions", not as "$2n$ linear conditions on the a_i".

Block 2. Build the polynomial whose roots are the solutions: multiply by x^{2n} to clear the reciprocals and move everything to one side, giving $q(x) = x^{2n+2} - x^{2n}p(1/x)$, monic of degree $2n+2$.

Block 3. Record the two free facts about q: no x^{2n+1} term (because subtracting a polynomial of degree at most $2n$ cannot disturb the top two coefficients of x^{2n+2}) and constant term -1 (because the reversal of a monic p has constant term 1).

Block 4. Count roots: $2n+2$ in total, $2n$ known, so exactly two unknown.

Block 5. Apply Vieta at the two ends: the unknown pair sums to 0 and multiplies to $-1/(n!)^2$; the parity of n enters here and only here.

Block 6. Solve the two-by-two system, get $\pm 1/n!$, and check they are new.

Juxtaposition. The nominal route reads "$2n$ conditions, $2n$ unknowns" and starts inverting a matrix; it computes p, which the question never asked for, and then faces a degree $2n+2$ equation with unpleasant coefficients. The ideal route never learns a single a_i; it learns only that the top-but-one coefficient of q is 0 and the bottom coefficient is -1, which is all the information the answer depends on. The bifurcation point is the first pen stroke after reading the hypothesis: do you write "$a_0 + a_1k^{-1} + \cdots = k^2$" (zombie) or "$q(x) := x^{2n+2} - x^{2n}p(1/x)$" (ideal)?

Layer 2: HOW

Block 1. Cue. An equation, $p(1/x) = x^2$, stated to hold at a list of $2n$ specific points, with the question asking for the rest of the solution set. **Why natural.** A list of points where an equation holds is a list of roots of the difference of the two sides. That is the factor theorem read in reverse, and it is the only way to convert "many conditions" into "a known factor", which is the currency you can actually compute with. It is worth repeating: values at many points are never data about coefficients, they are data about roots. **Link to the next block.** Once you want roots, you need a polynomial; the expression $p(1/x) - x^2$ is not one, so the next move is forced.

Block 2. Cue. The variable appears as $1/x$ inside p, and p has degree $2n$. **Why natural.** Multiplying by x^{2n} is the smallest power that clears every denominator, since the worst denominator in $p(1/x)$ is x^{-2n}. The resulting object $x^{2n}p(1/x)$ is the reversal of p: its coefficient list is that of p read backwards. Reversal is the standard dictionary between a polynomial's roots and their reciprocals, and it turns the hypothesis "monic" into the fact "constant term 1", which is a fact at the end of the polynomial where Vieta is cheap. **Link to the next block.** Now you have a polynomial identity; the next question is what its degree and its extreme coefficients are, because those are the only things you get for free.

Block 3. Cue. q is a difference of x^{2n+2} and something of degree at most $2n$. **Why natural.** A subtraction cannot disturb coefficients above the degree of what is subtracted. So the coefficients of x^{2n+2} and x^{2n+1} in q are those of x^{2n+2} itself, namely 1 and 0. The gap of two between $2n+2$ and $2n$ is the whole reason the problem works, and it is the reason the exponent in "$p(1/x) = x^2$" is a 2: it

creates exactly one free coefficient at the top to pair with the free coefficient at the bottom. **Link to the next block.** Two free coefficients at the two ends is exactly enough to pin two unknown roots.

Block 4. **Cue.** $2n$ distinct known roots inside a degree $2n+2$ polynomial. **Why natural.** The factor theorem plus distinctness gives divisibility by $\prod(x-k)$ over all $2n$ known roots, which is $\prod_{k=1}^{n}(x^2-k^2)$; the quotient is monic real of degree 2 by comparing degrees and leading coefficients. Note that you never need to know what the quotient is, only that it exists. **Link to the next block.** A monic quadratic has two unknowns, and you have two coefficient facts, so the system is square and you should not compute the quotient at all.

Block 5. **Cue.** A monic polynomial with all roots known but two, and its two extreme coefficients known. **Why natural.** Vieta at the ends is free information: the sum of all roots is minus the second coefficient and the product of all roots is $(-1)^{\deg}$ times the constant term. Every other elementary symmetric function sits in the middle of the coefficient list, where you do not know anything. So the lazy move is to use only e_1 and e_{2n+2}, and the problem has been built so that these suffice. The known roots come in $\pm$ pairs, which annihilates them in e_1 and turns them into $(-1)^n(n!)^2$ in the product, and this is the unique place where the parity of n acts. **Link to the next block.** Two symmetric equations in r and s; the first makes $s=-r$, so the second becomes a single quadratic in r.

Block 6. **Cue.** $r+s=0$ and $rs=-1/(n!)^2$. **Why natural.** A pair with zero sum is $\{r,-r\}$, and its product is $-r^2$, so the second equation reads $r^2=1/(n!)^2$: positive, hence two real roots. Had the right-hand side been negative you would have concluded that there are no further real solutions, which is precisely the odd-n situation. Finish by verifying the new roots are not old ones, because the question said "other". **Link to the next block.** Nothing left; the answer is $\pm 1/n!$.

Layer 3: WHY

Three mathematical truths carry this route. The first is that a polynomial identity is determined by its roots and its leading coefficient, so information about where two expressions agree is structural information, whereas information about coefficients is arithmetic information. Structural information composes; arithmetic information has to be carried. The second is that reversal, $p \mapsto p^*(x) = x^{\deg p}p(1/x)$, is an involution on polynomials with nonzero constant term that inverts every root and reverses the coefficient list. Any hypothesis phrased in terms of $1/x$ is an invitation to apply it, and after applying it a normalisation you were given at one end of the polynomial reappears at the other end, where it may be more useful. Here "monic" becomes "constant term 1", and the constant term is one half of Vieta's cheap pair. The third is that the elementary symmetric functions of the roots are the coefficients, so partial knowledge of the roots plus partial knowledge of the coefficients is a solvable system whenever the counts match; the designer of this problem matched them deliberately.

The discernment mechanism is a counting question you should ask before writing anything: how many unknowns does the answer actually depend on, and where in the polynomial do they live? Here the answer depends on two numbers, a and b in the quotient, and q offers exactly two coefficients you can name without work. If instead the exponent in the problem had been 3, so that q had degree $2n+3$ and a cubic quotient, you would need three cheap coefficients, and you would look for a third: the coefficients of x^{2n+2} and x^{2n+1} are both 0 and the constant is -1, so it still closes. If the exponent had been 1, the quotient would be linear and the top gap would vanish; you would need a different pairing. Training this is a matter of always writing the degree of the constructed polynomial and the degree of the known factor on the same line, then subtracting.

What the examiners nudged toward and how they nudged: they wrote the conditions symmetrically in $\pm k$, which is the signal that the extra roots will be a $\pm$ pair and that the sum of the known

roots is zero; they made p monic, which is the signal that the reversed polynomial has a known constant term; and they specified that n is even, which is the signal that a $(-1)^n$ will decide reality of the answer. Three planted cues, three used facts, no accidents.

Revolving concepts

Factor theorem and the root count. Definition and statement: for a field F and $p \in F[x]$ of degree $d \geq 1$, $p(r) = 0$ if and only if $(x - r) \mid p$; distinct roots give coprime linear factors, so d distinct roots force $\prod(x - r_i) \mid p$; over $\mathbb{C}$ there are exactly d roots with multiplicity. Corollary used constantly on the Putnam: if two polynomials of degree at most d agree at $d + 1$ points they are identical. See Book 1, field: Polynomials, topic: factor theorem and degree bounds.

Reciprocal (reversed) polynomial. Definition: $p^*(x) = x^d p(1/x)$ where $d = \deg p$; its coefficient sequence is that of p reversed, its roots are the reciprocals of the roots of p, and p is called palindromic when $p^* = p$. Orbiting facts: p monic if and only if $p^*(0) = 1$; a palindromic polynomial of even degree $2m$ can be reduced to degree m by the substitution $u = x + 1/x$, which is the standard route into problems about roots on the unit circle. Here reversal is the move that makes the whole problem two lines. See Book 1, field: Polynomials, topic: reciprocal and palindromic polynomials.

Vieta's relations and the symmetric functions. For monic q of degree d with roots $r_1, \ldots, r_d$, the coefficient of x^{d-j} is $(-1)^j e_j(r_1, \ldots, r_d)$, where e_j is the j-th elementary symmetric function. The two cheap ones are e_1 (from the second coefficient) and e_d (from the constant term). The general principle: knowing all roots but m of them, plus m coefficients, determines the missing ones, because the e_j of the whole set are polynomial in the e_j of the two parts. See Book 1, field: Polynomials, topic: Vieta and symmetric functions.

The recursion generalised. The factorial in the answer is the only recursion in sight, and it enters through the product of the known roots: $R_n := \prod_{k=1}^{n} k(-k) = (-1)^n (n!)^2$ satisfies $R_n = -n^2 R_{n-1}$ with $R_0 = 1$. The general statement worth installing is the parametrised version of this problem. Suppose p is monic of degree D and $p(1/x) = x^m$ at N prescribed nonzero points; then $q(x) = x^{D+m} - x^D p(1/x)$ is monic of degree $D + m$ with the top $m - 1$ coefficients after the leading one equal to zero and constant term -1, and the unknown factor has degree $D + m - N$. The missing factor is determined as soon as $D + m - N$ of its symmetric functions can be read off, and the products of the known roots enter through the same telescoping product R. Solving that general form is the same two-step method: build q, then match counts. See Book 1, field: Polynomials, topic: constructed polynomials with prescribed roots.

Degrees of freedom counting. Definition of the habit: before solving, count unknowns and independent conditions and ask what the answer depends on. Here p has $2n$ unknowns and there are $2n$ conditions, so p is determined, yet the answer depends only on two symmetric functions. Recognising that gap is what stops you from computing p. See Book 1, field: Logic and set theory, topic: counting constraints before solving.

Circuit Bank

- **WHEN** you are told an equation holds at more named points than you have unknowns, and asked where else it holds **DO** form the single polynomial whose roots are all solutions and factor out the known ones **BECAUSE** because the factor theorem is the only mechanism that converts a list of values into an algebraic factor, and the leftover factor has small degree, which is exactly what the question is asking for
- **WHEN** the unknown appears as $1/x$ inside a polynomial of degree d **DO** multiply by x^d to form the reversal $x^d p(1/x)$ **BECAUSE** because reversal turns reciprocal roots into ordinary roots and moves a normalisation such as monicity from the top of the coefficient list to the bottom, where Vieta is cheap

- **WHEN** a monic polynomial has all its roots known except a few, and you know the extreme coefficients **DO** use the sum and the product of the roots rather than dividing polynomials **BECAUSE** because e_1 and e_d are readable off the two ends for free, whereas the middle coefficients cost a long division that you will not need
- **WHEN** a hypothesis fixes a parity, such as "n is even" **DO** find the single step where $(-1)^n$ appears and mark it **BECAUSE** because such hypotheses are inserted to fix a sign, and the sign usually decides whether the answer is real, positive, or empty

Backtesting

1. First useful cue: an equation asserted at $2n$ named points, with the question asking for the other solutions, next to the reciprocal argument $1/x$ and the word "monic".
2. Method the expression suggested: build $q(x) = x^{2n+2} - x^{2n}p(1/x)$, note that it is monic of degree $2n+2$ with zero x^{2n+1} coefficient and constant term -1, divide off the $2n$ known roots in principle, and pin the remaining monic quadratic by the sum and product of all roots.
3. Where you would hesitate, and the thought to reinforce: at the moment you realise p is uniquely determined by the data, since the instinct is to find it. Reinforce this sentence: the question asks for solutions, not for p, and solutions are roots, so build the polynomial whose roots they are.
4. The trap the question sets: two traps. Computing p by interpolation, which burns the whole session; and forgetting that "other" requires checking $\pm 1/n!$ is not already in the list $\pm 1, \ldots, \pm n$, which is where a one-line justification ($n \geq 2$ so $1/n! < 1$) earns the last mark.
5. What makes you recognise it faster next time: the pair of cues "value conditions at many points" and "argument $1/x$" together mean "reversed polynomial with known roots"; the phrase "monic" is then a statement about the constant term of that reversed object.
6. The reusable decision rule: many prescribed values means known roots; $1/x$ means reverse; all-but-two roots known means Vieta at the two ends; a parity hypothesis means a sign is about to be decided.

Mirror drills

1. (drill) Take $n = 2$: p is monic of degree 4 with $p(1) = p(-1) = 1$ and $p(1/2) = p(-1/2) = 4$. Find all other real x with $p(1/x) = x^2$, and find p itself as a check. *Answer:* $x = \pm 1/2! = \pm\frac{1}{2}$. Solving the system confirms $p(z) = z^4 - \frac{21}{4}z^2 + \frac{21}{4}$, and indeed $p(2) = 16 - 21 + \frac{21}{4} = \frac{1}{4}$, as required by $p(1/x) = x^2$ at $x = \frac{1}{2}$.
2. (drill) Same hypotheses but with n odd. What changes? *Answer:* only the sign in $\prod_{k=1}^{n} k(-k) = (-1)^n(n!)^2$. Now $(n!)^2 rs = -1$ with $r + s = 0$ gives $r^2 = -1/(n!)^2$, so the two extra roots are $\pm i/n!$ and there are no other real solutions. This is exactly why the problem says n is even.
3. (drill) With n even, compute $p(0) = a_0$ without solving any linear system. *Answer:* $q(x) = \prod_{k=1}^{n}(x^2 - k^2) \cdot \left(x^2 - (n!)^{-2}\right)$, and $[x^{2n}]q = -a_0$. Writing $u = x^2$, the coefficient of u^n in $\prod_{k=1}^{n}(u - k^2) \cdot (u - (n!)^{-2})$ is minus the sum of the roots in u, so $a_0 = \sum_{k=1}^{n} k^2 + \frac{1}{(n!)^2} = \frac{n(n+1)(2n+1)}{6} + \frac{1}{(n!)^2}$. For $n = 2$ this is $5 + \frac{1}{4} = \frac{21}{4}$, matching drill 1.

3.3 A3. Nothing in the plane turns faster than one radian

Determine the smallest positive real number r such that there exist differentiable functions $f\colon \mathbb{R} \to \mathbb{R}$ and $g\colon \mathbb{R} \to \mathbb{R}$ satisfying

(a) $f(0) > 0$,

(b) $g(0) = 0$,

(c) $|f'(x)| \leq |g(x)|$ *for all* x,

(d) $|g'(x)| \leq |f(x)|$ *for all* x, *and*

(e) $f(r) = 0$.

Read it out loud

Two real functions, coupled so that each one's growth rate is controlled by the other one's size, started at a point where the first is positive and the second is zero, and you are asked how soon the first can vanish. Read the pair as a single object and the problem changes character: put $w(x) = (f(x), g(x))$, a point moving in the plane. Condition (c) and condition (d) together say

$$|w'(x)|^2 = f'(x)^2 + g'(x)^2 \leq g(x)^2 + f(x)^2 = |w(x)|^2,$$

that is, the speed of the moving point never exceeds its distance from the origin. Conditions (a) and (b) say the point starts on the positive horizontal axis, and condition (e) says that at time r it is on the vertical axis. So the question is: how quickly can a plane curve get from one axis to the other if its speed is bounded by its radius?

Fields wired in: differential inequalities (Gronwall and integrating factors), plane curves in polar coordinates, and the comparison theory of second order oscillation, which is the grown-up name for what is happening here. The knowledge base is three items: the equality case $f' = -g$, $g' = f$ is uniform rotation; for a curve avoiding the origin the velocity splits as $|w'|^2 = \rho'^2 + \rho^2(\theta')^2$ in polar coordinates; and a bound on a derivative plus a starting value bounds the function, by the mean value theorem.

The first useful cue is the perfect symmetry of (c) and (d) in f and g, together with the absolute values. Absolute values destroy all sign information, so nothing is being said about direction; only about magnitude. The only thing a magnitude bound can control is speed. Within thirty seconds you should have written $f = \cos$, $g = \sin$, which satisfies (c) and (d) with equality and gives $f(\pi/2) = 0$, and next to it the sentence "so $r \leq \pi/2$, and the whole problem is the lower bound". Then you should have written $|w'| \leq |w|$, because the moment the two conditions are added in squares, the mixed nature of the coupling disappears and a single geometric statement remains.

The last cue is the shape of the answer to expect. Speed bounded by radius is a scale-invariant hypothesis: if (f, g) works then so does (cf, cg). So the answer cannot involve any length, and it cannot involve $f(0)$. It must be a pure angle. The quarter turn from one axis to the other is $\pi/2$, and the equality case turns at exactly one radian per unit time.

Approaches, ranked worst to best

1. **Produce the example and stop.** *Cue:* "determine the smallest r". *Circuit:* exhibit, then assert minimality. (Writing $f = \cos$, $g = \sin$ and asserting that $\pi/2$ is optimal is worth almost nothing: the entire content of the problem is that no smaller r is possible. On the Putnam ten point scale this is one or two points. Notice the shape of the question: "smallest such that there exist" always splits into a construction and a bound, and the bound is where the marks are.)

2. **Bound the amplitude and hope.** *Cue:* the coupled inequalities look like Gronwall. *Circuit:* differentiate $f^2 + g^2$. (Set $k = f^2 + g^2$; then $|k'| \leq 4|fg| \leq 2k$, so $k(0)e^{-2x} \leq k(x) \leq k(0)e^{2x}$.

This is true, useful, and by itself powerless: it proves the pair never reaches the origin, but exponential envelopes contain no information about angles, and the answer is an angle. Rank it as a necessary lemma that is not a solution.)

3. **Iterate the integral form.** *Cue:* the fundamental theorem of calculus turns derivative bounds into integral bounds. *Circuit:* Picard iteration and comparison with cos. (From $|f'(x)| \le |g(x)| \le \int_0^x |g'| \le \int_0^x |f|$ one gets $f'(x) + F(x) \ge 0$ where $F(x) = \int_0^x f$. Multiplying by $\cos x$, integrating, then multiplying by $\sec^2 y$ and integrating again produces $F(r)\sin^2 r \ge F(r)$, which is impossible for $0 < r < \pi/2$. This is the published second solution and it is fully rigorous. It ranks below the angle argument because the two integrating factors $\cos x$ and $\sec^2 y$ are pulled from the answer rather than from the structure: you can only guess them once you already believe the answer is $\pi/2$. It also costs three integrations instead of one differentiation.)

4. **Polar angle: bound the angular speed.** *Cue:* speed bounded by radius, start on one axis, finish on another. *Circuit:* in polar coordinates a bound of speed by radius is exactly a bound of angular speed by one. (The ideal route. The hypothesis says the point cannot turn faster than one radian per unit time; the conclusion demands a quarter turn; so at least $\pi/2$ units of time are needed. Implemented with $\theta = \arctan(g/f)$ up to the first zero of f, it is half a page and every line is forced by the geometry.)

The ideal solution

The answer is $\boxed{r = \dfrac{\pi}{2}}$.

Attainment. Take $f(x) = \cos x$ and $g(x) = \sin x$. Then $f(0) = 1 > 0$, $g(0) = 0$, $|f'| = |\sin| = |g|$, $|g'| = |\cos| = |f|$, and $f(\pi/2) = 0$. So $r = \pi/2$ is achieved.

No smaller r works. Suppose, for contradiction, that f and g are differentiable and satisfy (a) to (e) for some r with $0 < r < \pi/2$.

Step 1 (cut at the first zero). The set $Z = \{x \in [0, r] : f(x) = 0\}$ is nonempty, since $r \in Z$, and closed and bounded, since f is continuous; hence it has a least element ρ. As $f(0) > 0$ we have $\rho > 0$, and by minimality $f(x) \ne 0$ for $x \in [0, \rho)$. Note $0 < \rho \le r < \pi/2$. Replacing r by ρ, we may and do assume from now on that $f(x) \ne 0$ on $[0, r)$ and $f(r) = 0$.

Step 2 (the pair never reaches the origin). Put $k = f^2 + g^2$, which is differentiable. Then

$$|k'| = |2ff' + 2gg'| \le 2|f||g| + 2|g||f| = 4|f||g| \le 2(f^2 + g^2) = 2k,$$

the last step by $2|f||g| \le f^2 + g^2$. In particular $k' + 2k \ge 0$, so $\big(e^{2x}k(x)\big)' = e^{2x}\big(k'(x) + 2k(x)\big) \ge 0$, so $e^{2x}k(x)$ is non-decreasing on $[0, \infty)$ and therefore

$$k(x) \ge e^{-2x}k(0) = e^{-2x}f(0)^2 > 0 \qquad (x \ge 0).$$

At $x = r$ this gives $g(r)^2 = k(r) > 0$, so $g(r) \ne 0$.

Step 3 (angular speed at most one). On $[0, r)$ the function f does not vanish, so

$$\theta(x) := \arctan\big(g(x)/f(x)\big)$$

is differentiable there, with $\theta(0) = \arctan(0) = 0$ and

$$\theta'(x) = \frac{f(x)g'(x) - g(x)f'(x)}{f(x)^2 + g(x)^2}.$$

Using (c) and (d),

$$|\theta'(x)| \le \frac{|f||g'| + |g||f'|}{f^2 + g^2} \le \frac{|f|\,|f| + |g|\,|g|}{f^2 + g^2} = 1.$$

By the mean value theorem, for each $x \in [0, r)$ there is $c \in (0, x)$ with $|\theta(x)| = |\theta(x) - \theta(0)| = |\theta'(c)|\, x \le x < r$.

Step 4 (contradiction). Since $|\theta(x)| < r < \pi/2$ for all $x \in [0, r)$ and tan is increasing on $(-\pi/2, \pi/2)$,

$$\left|\frac{g(x)}{f(x)}\right| = |\tan\theta(x)| \le \tan r \qquad (x \in [0, r)),$$

a fixed finite bound. But f and g are continuous with $f(r) = 0$, $g(r) \neq 0$ and $f(x) \neq 0$ on $[0, r)$, so $|g(x)/f(x)| \to \infty$ as $x \to r^-$. These two statements are incompatible, and the contradiction shows that no $r < \pi/2$ admits such a pair.

Hence the smallest possible r is $\pi/2$.

A remark on Step 2. The published solution divides by k and bounds the derivative of $\log k$, which presumes what is being proved, namely that k does not vanish, and then repairs the gap by a continuity argument. The integrating factor e^{2x} avoids the issue: $k' + 2k \ge 0$ is a statement about a differentiable function that is valid whether or not k vanishes, and it delivers a positive lower bound directly. Prefer it in the arena, since it removes an edge case rather than patching one.

Layer 1: WHAT

Block 1. Take the equality case seriously: $f' = -g$, $g' = f$ is uniform rotation, so $(\cos, \sin)$ is the candidate and $r \le \pi/2$.

Block 2. Add the two hypotheses in squares: $f'^2 + g'^2 \le f^2 + g^2$, that is, the plane curve $w = (f, g)$ has speed at most its radius.

Block 3. Replace r by the first zero of f, so that f is nonzero on $[0, r)$ and the angle is defined by a single branch of arctan.

Block 4. Show the curve avoids the origin, by the differential inequality $k' + 2k \ge 0$ and the integrating factor e^{2x}. Deduce $g(r) \neq 0$.

Block 5. Compute the angular speed and bound it by one.

Block 6. The angle starts at 0 and, by continuity of the picture, is forced to $\pm\pi/2$ at $x = r$; travelling $\pi/2$ radians at angular speed at most one takes time at least $\pi/2$, so $r < \pi/2$ is impossible.

Juxtaposition. The nominal route stays inside the inequalities: substitute one into the other, integrate, produce exponential envelopes, and then wonder why no exponential bound ever produces a π. The ideal route leaves the algebra and names the geometric object the hypotheses describe. The bifurcation point is the line right after you write down cos and sin: do you start bounding $|f|$ and $|g|$ separately (zombie) or do you write $w = (f, g)$ and ask what the hypotheses say about w (ideal)? Everything downstream is decided there, because the answer $\pi/2$ is an angle and no argument that never mentions an angle can produce it.

Layer 2: HOW

Block 1. Cue. Two inequalities that become equalities for a familiar pair; the phrase "smallest r such that there exist". **Why natural.** An extremal problem with an inequality hypothesis is solved by its equality case, so the first thing to do is to write down the equality case as a differential system: $f' = \pm g$, $g' = \pm f$. The sign choice $f' = -g$, $g' = f$ gives $f'' = -f$, the rotation, and its first zero is at $\pi/2$. Constructing the extremal object converts half the problem into a verification and, more importantly, tells you the number you must prove is a lower bound. **Link to the next block.** Now you know the target is $\pi/2$; a bare number cannot come out of algebra alone, so you look for what in the problem measures a quarter turn.

Block 2. Cue. Absolute values on all four quantities, and perfect symmetry between f and g. **Why natural.** Absolute values keep magnitudes and discard directions, and a hypothesis about magnitudes of derivatives is a hypothesis about speed. Adding the squares of (c) and (d) is the

only way to combine them without choosing signs, and it produces $|w'| \le |w|$, a single statement about one object. It is worth repeating: when two hypotheses are symmetric under swapping the objects, combine them into one statement about the pair rather than reasoning about each separately. **Link to the next block.** A statement about speed and radius is a statement in polar coordinates, so decompose the velocity.

Block 3. **Cue.** The conclusion is about the first time f vanishes, and the angle formula $\arctan(g/f)$ needs $f \neq 0$. **Why natural.** If f has an earlier zero, that earlier zero is also an admissible r, so passing to it costs nothing and buys the non-vanishing you need. This is the standard "replace the given object by the extremal one" move: minimal counterexample, first zero, largest interval. The set of zeros is closed and bounded, so a least element exists; that sentence is the entire justification and it must be written, because "let r be the first zero" without it is an assumption. **Link to the next block.** With $f \neq 0$ on $[0, r)$, a single branch of arctan works and no lifting theory is needed.

Block 4. **Cue.** The angle is undefined at the origin, so you must rule out $(f, g) = (0, 0)$. **Why natural.** The quantity $k = f^2 + g^2$ is the natural test, and its derivative is controlled by the same two hypotheses through $2|f||g| \le f^2 + g^2$. Once $|k'| \le 2k$, the one-sided version $k' + 2k \ge 0$ is what you need, and multiplying by the integrating factor e^{2x} turns it into the monotonicity of a single function. That is the general recipe for every differential inequality: never divide, always multiply by the integrating factor. **Link to the next block.** The pair is nonzero at $x = r$ where f vanishes, so $g(r) \neq 0$: the curve arrives on the vertical axis at a genuine point, not at the origin. That is what makes the angle blow up to a right angle instead of dribbling out.

Block 5. **Cue.** A quotient g/f whose derivative you are about to take, with $f^2 + g^2$ in the denominator. **Why natural.** The derivative of $\arctan(g/f)$ is $(fg' - gf')/(f^2 + g^2)$, and the numerator is exactly the cross product of position and velocity, which is what measures rotation. Bounding it with (c) and (d) gives $|f||f| + |g||g|$ on top, which cancels the denominator exactly. The cancellation is not a coincidence: it is the statement $|w'| \le |w|$ read in polar coordinates, since $|w'|^2 = \rho'^2 + \rho^2\theta'^2 \ge \rho^2\theta'^2$ forces $|\theta'| \le |w'|/\rho \le 1$. **Link to the next block.** Angular speed at most one plus the mean value theorem gives $|\theta(x)| \le x$: the total turn by time x is at most x radians.

Block 6. **Cue.** $f(r) = 0$ with $g(r) \neq 0$: the curve is on the vertical axis at time r. **Why natural.** Being on the vertical axis is being at angle $\pm\pi/2$, and the total turn available in time $r < \pi/2$ is less than $\pi/2$ radians. Written without limits of angles, the same fact is that $|g/f|$ stays bounded by $\tan r$ while it must blow up: that is the cleanest way to state the contradiction, since it avoids discussing the limit of θ at the endpoint. **Link to the next block.** Contradiction; the minimum is $\pi/2$.

Layer 3: WHY

The mathematical truth beneath the route is the polar decomposition of velocity. For a differentiable plane curve avoiding the origin, $w = \rho(\cos\theta, \sin\theta)$ with $|w'|^2 = \rho'^2 + \rho^2(\theta')^2$. A hypothesis bounding $|w'|$ by ρ therefore says two things at once: the radius can grow or decay at rate at most ρ, which is exponential and harmless, and the angle can change at rate at most 1, which is the only rigid constraint in the problem. The conclusion $f(r) = 0$ is purely angular. So the amplitude is a red herring by design, and every route that only tracks amplitude must fail. That is the real reason the answer is $\pi/2$ and not something involving $f(0)$ or e.

The discernment mechanism is scale invariance. Test the hypotheses against the substitution $(f, g) \mapsto (cf, cg)$: they are invariant, and so is the conclusion. Therefore the answer is a function of nothing at all, a pure number, and it must be produced by an invariant quantity. The invariants

of a plane curve under scaling are angles, so look for the angle. Train the reflex like this: for every extremal problem, list the transformations that preserve the hypotheses, and only then choose a quantity to track, insisting that it be invariant under all of them.

There is a second, deeper reading. Differentiating the extremal system gives $f'' = -f$, and the general problem "how far apart can the zeros of $y'' + q(x)y = 0$ be when $0 < q \leq 1$" is Sturm comparison. The standard proof of Sturm comparison is the Prüfer substitution: write $y = \rho \sin\theta$ and $y' = \rho\cos\theta$, and compute $\theta' = \cos^2\theta + q\sin^2\theta \leq 1$; consecutive zeros of y are consecutive multiples of π for θ, so they are at least π apart. That is exactly Step 3 of the solution, with the pair (y', y) playing the role of (f, g). In other words this problem is a quarter of the Sturm comparison theorem, dressed as an inequality about two unknown functions. Recognising the dressing is what makes the route obvious rather than inspired.

What the examiners nudged toward and how: they used absolute values so that no sign information survives, leaving only speed; they made the hypotheses symmetric so that adding the squares is the natural combination; they placed the start exactly on one axis, $g(0) = 0$, so that the required turn is a full quarter and not less; and they demanded $f(r) = 0$ rather than $f(r) = g(r) = 0$, since the latter is impossible by Step 2 and would have given the game away.

Revolving concepts

Polar coordinates for a moving point. Definition: for a differentiable $w\colon I \to \mathbb{R}^2 \setminus \{0\}$, write $\rho = |w|$ and let θ be a continuous choice of argument; then ρ is differentiable, a differentiable θ exists locally, and $|w'|^2 = \rho'^2 + \rho^2(\theta')^2$. Orbiting facts: the cross product $fg' - gf'$ equals $\rho^2\theta'$ and is the signed area rate swept out (Kepler's second law is the statement that it is constant); a local branch is given by $\arctan(g/f)$ where $f \neq 0$ and by $\operatorname{arccot}(f/g)$ where $g \neq 0$, and the two agree up to a constant on overlaps, which is the elementary form of the lifting theorem for the covering $\mathbb{R} \to S^1$. See Book 1, field: Geometry, topic: plane curves in polar form.

Differential inequalities and Gronwall. Statement: if u is differentiable and $u' \leq au + b$ on an interval, then multiplying by e^{-ax} makes $\left(e^{-ax}u\right)' \leq be^{-ax}$ and integration gives the explicit bound; the one-sided reverse version with $u' \geq -au$ gives $u(x) \geq u(0)e^{-ax}$. The rule to install: never divide by u to take log, always multiply by the integrating factor, because the integrating factor argument never needs $u \neq 0$. See Book 1, field: Real analysis, topic: Gronwall and integrating factors.

Sturm comparison and the Prüfer angle. Statement: if $y'' + q(x)y = 0$ with $0 < q(x) \leq Q$ then consecutive zeros of a nontrivial solution are at least $\pi/\sqrt{Q}$ apart; if $q \geq Q$ they are at most $\pi/\sqrt{Q}$ apart. Proof mechanism: the Prüfer substitution $y = \rho\sin\theta$, $y' = \rho\cos\theta$ turns the second order equation into the first order angle equation $\theta' = \cos^2\theta + q\sin^2\theta$, and zeros of y are the times when θ crosses a multiple of π. This problem is the case $Q = 1$ over a quarter period. See Book 1, field: Real analysis, topic: oscillation and comparison theorems.

The recursion generalised. The integral route is Picard iteration, and the iteration is a genuine recursion worth stating in general: define $u_0 = c$ and $u_{m+1}(x) = \int_0^x u_m(t)\,dt$; then $u_m(x) = cx^m/m!$ and $\sum_m u_m = ce^x$. Feeding $|f'| \leq |g|$ and $|g'| \leq |f|$ into each other repeatedly produces exactly these terms with alternating signs, which is why the extremal pair is $(\cos, \sin)$ and not some exotic function: the alternation comes from the two applications needed to return to f, giving $f'' = -f$. The general lesson is that iterating a pair of first order bounds converges to the solution of the corresponding second order equation, and comparison with that solution is where sharp constants such as $\pi/2$ come from. See Book 1, field: Sequences and recurrences, topic: Picard iteration and exponential bounds.

Extremal problems of the form "smallest r such that there exist". Definition of the shape: the answer set is an interval $[r_0, \infty)$ or (r_0, ∞), and a complete solution has two halves, a construction achieving r_0 and a proof that nothing smaller works. Here the construction is one line and the bound is the problem; recognise the split before starting and allocate your ink accordingly. Note also that every $r \geq \pi/2$ is achievable, by $f(x) = \cos(\lambda x)$ and $g(x) = \sin(\lambda x)$ with $\lambda = \pi/(2r) \leq 1$,

so the admissible set really is $[\pi/2, \infty)$. See Book 1, field: Inequalities, topic: extremal existence problems.

Circuit Bank

- **WHEN** two functions satisfy coupled bounds of the form $|f'| \le |g|$ and $|g'| \le |f|$ **DO** write $w = (f, g)$, add the squares to get $|w'| \le |w|$, and read it as speed at most radius **BECAUSE** because absolute values carry only magnitude, and the only invariant meaning of a magnitude bound on a derivative is a bound on speed, which in polar coordinates is a bound on angular speed by one
- **WHEN** an extremal problem has an inequality hypothesis with a visible equality case **DO** solve the equality case first as a differential system and let it name the answer **BECAUSE** because the extremal trajectory is the unique candidate for optimality, and knowing the target number tells you which invariant you must bound
- **WHEN** you need a pair (f, g) never to vanish simultaneously **DO** bound $|(f^2 + g^2)'|$ by a multiple of $f^2 + g^2$ and multiply by the integrating factor **BECAUSE** because the integrating factor turns the inequality into monotonicity of one function and never requires dividing by a quantity you have not yet shown to be nonzero
- **WHEN** an angle argument needs a branch of arctan **DO** pass to the first zero of the denominator and work on the half-open interval before it **BECAUSE** because the first zero exists (the zero set is closed and bounded), it is still an admissible endpoint, and on the interval before it a single branch is differentiable, so no lifting theorem is needed
- **WHEN** the hypotheses are invariant under scaling the unknowns **DO** conclude the answer is a pure number and hunt for a scale-invariant quantity **BECAUSE** because a scale-invariant problem cannot have an answer depending on any length, and in the plane the scale-invariant quantities are angles

Backtesting

1. First useful cue: the symmetric pair of absolute-value bounds, which becomes $f'^2 + g'^2 \le f^2 + g^2$, together with a start on one axis and a finish on the other.
2. Method the expression suggested: build the extremal pair $(\cos, \sin)$ from the equality case to get $r \le \pi/2$, then bound the angular speed of the plane curve (f, g) by one and observe that a quarter turn is required.
3. Where you would hesitate, and the thought to reinforce: at the definition of the angle, because $\arctan(g/f)$ needs $f \neq 0$ and the origin must be excluded. Reinforce these two sentences: pass to the first zero of f, and control $f^2 + g^2$ from below with an integrating factor.
4. The trap the question sets: producing the example and calling the problem finished. The construction is worth little; the lower bound is the problem. A second, smaller trap is asserting that $g(r) \neq 0$ without proof, which is exactly the step that makes the quotient blow up.
5. What makes you recognise it faster next time: any hypothesis of the shape "derivative of each bounded by the other" is a rotation speed limit; any question about the first zero under such a hypothesis is a quarter or half period question, and the answer will be $\pi/2$ or π times a scale factor.
6. The reusable decision rule: pair the functions into a vector, convert the magnitude hypothesis into an angular speed bound, and measure the required rotation.

Mirror drills

1. (drill) Suppose f, g are differentiable with $|f'| \le |g|$, $|g'| \le |f|$ and $f(0) = g(0) = 0$. Show $f \equiv g \equiv 0$. *Answer:* with $k = f^2 + g^2$ we have $k' \le 2k$, so $\left(e^{-2x}k\right)' \le 0$ and $k(x) \le k(0)e^{2x} = 0$ for $x \ge 0$; the same argument in the reversed variable $x \mapsto -x$ handles $x \le 0$. This is the uniqueness half of the same circuit.

2. (drill) Same hypotheses (c) and (d), but now $f(0) = g(0) > 0$. What is the smallest r with $f(r) = 0$? *Answer:* $\pi/4$. The starting angle is $\pi/4$ and the nearest axis crossing is a further $\pi/4$ of turn; angular speed at most one gives $r \ge \pi/4$, and $f(x) = \cos(x + \pi/4)$, $g(x) = \sin(x + \pi/4)$ attains it.

3. (drill) Let y be a nontrivial solution of $y'' + q(x)y = 0$ with $0 < q(x) \le 1$ and $y(0) = 0$. Show that the next zero of y is at least π away. *Answer:* the pair (y', y) never vanishes at the same point (uniqueness for the linear equation), so set $\theta = \arctan(y/y')$; then $\theta' = \left(y'^2 + qy^2\right)/\left(y^2 + y'^2\right) = \cos^2\theta + q\sin^2\theta \le 1$, and zeros of y are the times θ hits a multiple of π, so consecutive zeros are at least π apart. This is A3 with the training wheels removed.

3.4 A4. Icosahedral vectors generate a dense subgroup

Let $v_1, \ldots, v_{12}$ be unit vectors in $\mathbb{R}^3$ from the origin to the vertices of a regular icosahedron. Show that for every vector $v \in \mathbb{R}^3$ and every $\varepsilon > 0$, there exist integers $a_1, \ldots, a_{12}$ such that $\|a_1 v_1 + \cdots + a_{12} v_{12} - v\| < \varepsilon$.

Read it out loud

Twelve fixed vectors; integer coefficients, positive and negative and zero; and the demand that the resulting set comes arbitrarily close to every point of space. The set of all such integer combinations is precisely the subgroup of $(\mathbb{R}^3, +)$ generated by $v_1, \ldots, v_{12}$, so the problem, stripped of its epsilon, reads: prove that the subgroup generated by the twelve vertex vectors of a regular icosahedron is dense in $\mathbb{R}^3$. Say it that way and you have already discarded the analysis and turned the question into algebra.

Fields wired in: the structure of subgroups of $(\mathbb{R}, +)$, Diophantine approximation in its simplest form (a subgroup generated by 1 and an irrational number is dense), and the coordinate geometry of the icosahedron. The knowledge base is exactly three items: every subgroup of $\mathbb{R}$ is either $c\mathbb{Z}$ or dense; the golden ratio $\phi = \frac{1+\sqrt{5}}{2}$ is irrational and satisfies $\phi^2 = \phi + 1$; and the twelve vertices of a regular icosahedron can be placed at the cyclic permutations of $(\pm 1, \pm\phi, 0)$.

The first useful cue is the word "integers" next to "for every ε". That pairing is the signature of a density-of-a-subgroup question, and it should immediately raise the counter-question: when would this be false? Run the same statement for the eight vertices of a cube, $(\pm 1, \pm 1, \pm 1)/\sqrt{3}$. Every integer combination lands in $\frac{1}{\sqrt{3}}\mathbb{Z}^3$, which is discrete, so the cube version is false. So the statement is not about "twelve vectors spanning space"; something specific to the icosahedron is doing the work. Within thirty seconds you should have written the cube counterexample and the question it forces: what does the icosahedron have that the cube does not?

The answer is the golden ratio, and behind the golden ratio is the five-fold symmetry. A lattice cannot have a five-fold rotational symmetry (this is the crystallographic restriction), so a solid with five-fold symmetry cannot have its vertices sitting in a discrete subgroup that its own symmetry group preserves. The icosahedron is therefore forced to be dense, and the coordinates make that visible: they contain the irrational number ϕ in an essential way, and $\mathbb{Z} + \phi\mathbb{Z}$ is dense in $\mathbb{R}$.

Approaches, ranked worst to best

1. **Chase the approximation directly.** *Cue:* "for every v and every ε". *Circuit:* do what the epsilon asks. (Trying to build the integers a_i for a given v and ε by hand, coordinate by coordinate, is hopeless: there is no formula, and the choice of a_i is wildly discontinuous in v. The epsilon is a disguise, not an instruction. Any solution that never says the word "subgroup" is fighting the problem.)

2. **Descent by differences.** *Cue:* differences of vertices are in the group and some of them are short. *Circuit:* produce arbitrarily short elements, then translate them around. (You can take a vertex difference, apply the icosahedral rotations to it, take differences again, and try to prove the lengths tend to zero, a three dimensional continued fraction. It can be made to work, but controlling a descent in three dimensions costs pages and the required estimates are exactly the sort of thing that goes wrong under time pressure.)

3. **Closed subgroups and the trace obstruction.** *Cue:* the icosahedral rotation group preserves the generated subgroup. *Circuit:* a lattice invariant under an order five rotation cannot exist. (The closure of the subgroup is a closed subgroup of $\mathbb{R}^3$, hence of the form $V \oplus L$ with V a linear subspace and L discrete; the rotation group preserves V; the standard three dimensional representation of the icosahedral group is irreducible, so V is trivial or everything; and if

V were trivial the group would be a rank three lattice invariant under a rotation of order 5, forcing the integer matrix of that rotation to have integer trace, while the true trace is $1 + 2\cos(2\pi/5) = \phi$. Beautiful, coordinate-free and structurally the true reason, but it needs the structure theorem for closed subgroups and a representation-theoretic input. Worth knowing; too expensive as a first line of attack.)

4. **Coordinates plus the subgroup dichotomy.** *Cue:* the icosahedron has standard coordinates built from the golden ratio. *Circuit:* $\mathbb{Z} + \phi\mathbb{Z}$ is dense because ϕ is irrational. (The ideal route. Choose the standard coordinates, add and subtract two vertices to isolate $(2,0,0)$ and $(2\phi,0,0)$ and their cyclic images, conclude that the group contains $H \times H \times H$ with $H = 2\mathbb{Z} + 2\phi\mathbb{Z}$, then quote the dichotomy for subgroups of $\mathbb{R}$. Four lines once the coordinates are on the page.)

The ideal solution

Write $\Lambda = \{a_1v_1 + \cdots + a_{12}v_{12} : a_i \in \mathbb{Z}\}$, which is the subgroup of $(\mathbb{R}^3, +)$ generated by $v_1, \ldots, v_{12}$. The assertion to be proved is exactly that Λ is dense in $\mathbb{R}^3$.

Reduction to standard coordinates. All twelve vertices are at distance 1 from the origin, so the origin is the centre of the sphere circumscribing the icosahedron; since a sphere through the vertices is unique, the origin is the centre of the icosahedron. Hence the twelve vectors are obtained from the standard vertex set by a rotation together with a positive scaling. If R is an orthogonal map and $c > 0$, then Λ is dense if and only if $cR(\Lambda)$ is dense, because $x \mapsto cRx$ is a homeomorphism of $\mathbb{R}^3$ and a group automorphism. We may therefore assume that the twelve vertices are the cyclic permutations of

$$(\pm 1, \pm\phi, 0), \qquad \phi = \frac{1+\sqrt{5}}{2},$$

that is, the four vectors $(\pm 1, \pm\phi, 0)$, the four vectors $(0, \pm 1, \pm\phi)$ and the four vectors $(\pm\phi, 0, \pm 1)$. (That these twelve points are the vertices of a regular icosahedron is the standard construction from three mutually perpendicular golden rectangles; the verification is that all twelve have the same length $\sqrt{1+\phi^2}$ and that each one has exactly five others at distance 2, which uses $\phi^2 = \phi + 1$. For instance the five neighbours of $(1,\phi,0)$ are $(-1,\phi,0)$, $(0,\pm 1,\phi)$ and $(\phi,0,\pm 1)$, and $\|(1,\phi,0) - (0,1,\phi)\|^2 = 1 + (\phi-1)^2 + \phi^2 = 2(\phi^2 - \phi + 1) = 4$.)

Step 1: the group contains an explicit product. Adding and subtracting pairs of vertices,

$$\begin{aligned}
(1,\phi,0) + (1,-\phi,0) &= (2,0,0), & (1,\phi,0) - (1,-\phi,0) &= (0,2\phi,0),\\
(0,1,\phi) + (0,1,-\phi) &= (0,2,0), & (0,1,\phi) - (0,1,-\phi) &= (0,0,2\phi),\\
(\phi,0,1) + (\phi,0,-1) &= (2\phi,0,0), & (\phi,0,1) - (\phi,0,-1) &= (0,0,2).
\end{aligned}$$

So Λ contains the six vectors $(2,0,0), (2\phi,0,0), (0,2,0), (0,2\phi,0), (0,0,2), (0,0,2\phi)$, and therefore contains $H \times H \times H$, where

$$H := 2\mathbb{Z} + 2\phi\mathbb{Z} \leq (\mathbb{R}, +).$$

Step 2: the dichotomy for subgroups of $\mathbb{R}$.

Lemma 3.4.1. *Let H be a subgroup of $(\mathbb{R}, +)$ with $H \neq \{0\}$. Then either $H = c\mathbb{Z}$ for some $c > 0$, or H is dense in $\mathbb{R}$.*

Proof. Let $c = \inf\{h \in H : h > 0\}$, which is well defined since H contains a nonzero element and hence a positive one.

If $c = 0$, then for every $\varepsilon > 0$ there is $h \in H$ with $0 < h < \varepsilon$, and every real number is within $h < \varepsilon$ of an integer multiple of h, so H is dense.

If $c > 0$, first note $c \in H$. Otherwise choose $h_1 \in H$ with $c < h_1 < 2c$ and then $h_2 \in H$ with $c < h_2 < h_1$; then $h_1 - h_2 \in H$ and $0 < h_1 - h_2 < h_1 - c < c$, contradicting the definition of c. Now for any $h \in H$ the number $h - c\lfloor h/c \rfloor$ lies in $H \cap [0, c)$, hence is 0, so $H = c\mathbb{Z}$. □

Step 3: H is dense. Suppose $H = 2\mathbb{Z} + 2\phi\mathbb{Z}$ were equal to $c\mathbb{Z}$. Then $2 = mc$ and $2\phi = kc$ for some nonzero integers m, k, whence $\phi = k/m$ is rational. But $\phi = \frac{1+\sqrt{5}}{2}$ and $\sqrt{5}$ is irrational, so ϕ is irrational. Contradiction; by the lemma, H is dense in $\mathbb{R}$.

Step 4: conclusion. Let $v = (x, y, z) \in \mathbb{R}^3$ and $\varepsilon > 0$. By Step 3 choose $h_1, h_2, h_3 \in H$ with

$$|h_1 - x| < \frac{\varepsilon}{\sqrt{3}}, \qquad |h_2 - y| < \frac{\varepsilon}{\sqrt{3}}, \qquad |h_3 - z| < \frac{\varepsilon}{\sqrt{3}}.$$

Then $(h_1, h_2, h_3) \in H \times H \times H \subseteq \Lambda$ and $\|(h_1, h_2, h_3) - v\| < \varepsilon$. Since every element of Λ is of the form $a_1 v_1 + \cdots + a_{12} v_{12}$ with $a_i \in \mathbb{Z}$, this is exactly the required statement. ■

Layer 1: WHAT

Block 1. Translate the epsilon statement into "the subgroup generated by the twelve vectors is dense".

Block 2. Test the statement against the cube, where it is false, and identify that irrationality of the coordinate ratios, not spanning, is the active ingredient.

Block 3. Put the icosahedron in standard coordinates, which is legitimate because density is invariant under rotation and scaling.

Block 4. Add and subtract vertex pairs to isolate the six vectors $(2, 0, 0)$, $(2\phi, 0, 0)$ and their cyclic images, so that the group contains H^3 with $H = 2\mathbb{Z} + 2\phi\mathbb{Z}$.

Block 5. Prove or quote the dichotomy: a subgroup of $\mathbb{R}$ is cyclic or dense, and rule out cyclic using the irrationality of ϕ.

Block 6. Assemble: dense in each coordinate gives dense in the product.

Juxtaposition. The nominal route reads "for every v and every ε" and starts constructing coefficients, or reaches for a compactness or equidistribution theorem. The ideal route deletes the analysis in the first line and never returns to it: the only analytic content in the whole solution is the one-line lemma about subgroups of $\mathbb{R}$. The bifurcation point is the sentence you write immediately after reading the problem: "let Λ be the subgroup generated by the v_i; we must show Λ is dense" (ideal) versus "let v be given and $\varepsilon > 0$; we construct a_i" (zombie).

Layer 2: HOW

Block 1. Cue. Integer coefficients, all twelve of them free, with both signs allowed, and an arbitrarily small error. **Why natural.** A set closed under addition and negation is a group, and "comes arbitrarily close to every point" is the definition of dense. So the hypothesis and the conclusion are both statements about one object, the generated subgroup, and naming it removes both quantifiers from view. It is worth repeating: an epsilon in a problem about integer combinations is almost never an analysis cue, it is the definition of density wearing a disguise. **Link to the next block.** Once the object is a subgroup, the relevant classification is the classification of subgroups, so ask which subgroups of $\mathbb{R}^3$ are not dense.

Block 2. Cue. The statement would be false for other polyhedra; nothing so far has used any property of the icosahedron. **Why natural.** A subgroup that is not dense is, in the closed case, a lattice, and lattices are exactly what integer coordinates produce. The cube has vertex coordinates that are all equal up to sign, so its group sits inside a scaled copy of $\mathbb{Z}^3$ and is discrete. Running that test tells you what to look for in the icosahedron: a ratio of coordinates that is irrational. This is the single most useful habit in a "prove that" problem, namely to find the nearest false statement and see which hypothesis kills it. **Link to the next block.** The irrational ratio in the icosahedron is ϕ, so put the icosahedron in coordinates where ϕ is visible.

Block 3. **Cue.** "Regular icosahedron" with no coordinates given, and a conclusion invariant under rotations and scalings. **Why natural.** Density is preserved by any invertible linear map, in particular by rotations and by positive scalings, since such maps are homeomorphisms and group automorphisms. So you may normalise, and the normalisation to use is the standard golden rectangle construction, in which each vertex has one coordinate zero and the other two equal to 1 and ϕ up to sign. Note that the problem already fixes the centre at the origin for you: all twelve vectors are unit vectors, so the origin is equidistant from all twelve vertices, and the circumscribed sphere is unique. **Link to the next block.** With explicit coordinates, the group is explicit, and you can start producing convenient elements.

Block 4. **Cue.** Vertices come in $\pm$ pairs within each coordinate plane: $(1, \phi, 0)$ and $(1, -\phi, 0)$ differ only in one sign. **Why natural.** Adding a vector to its sign-flipped partner kills one coordinate and doubles the other; subtracting does the reverse. That is the cheapest way to extract axis-aligned elements from a symmetric generating set, and after doing it in all three cyclic planes you hold 2 and 2ϕ on each axis. Do not chase a minimal generating set or the exact group Λ; you only need a subgroup that is already dense, and a product of one dimensional groups is the easiest such object to handle. **Link to the next block.** The problem is now one dimensional: is $2\mathbb{Z} + 2\phi\mathbb{Z}$ dense in $\mathbb{R}$?

Block 5. **Cue.** A subgroup of $\mathbb{R}$ generated by two numbers with irrational ratio. **Why natural.** The infimum of the positive elements is the only possible generator, and the standard argument shows that either it is positive and attained, in which case the group is cyclic, or it is zero, in which case the group is dense. Ruling out the cyclic case is pure arithmetic: two generators of a cyclic group are integer multiples of the same c, so their ratio is rational. This is the elementary form of Kronecker's theorem and you should be able to write the lemma cold, in five lines. **Link to the next block.** Density in each coordinate gives density of the product.

Block 6. **Cue.** A product of dense subsets and a Euclidean norm. **Why natural.** Approximate each coordinate separately to within $\varepsilon/\sqrt{3}$ and the vector norm is below ε. This is the only place in the proof where the metric appears, and it is the only place it should appear. **Link to the next block.** The approximating point is an integer combination of the v_i by construction, so the required a_i exist.

Layer 3: WHY

The mathematical truth beneath this problem is the classification of closed subgroups: a closed subgroup of $\mathbb{R}^n$ is $V \oplus L$ with V a linear subspace and L a discrete lattice in a complement. Hence a finitely generated subgroup either is discrete or has dense directions, and there is nothing in between. So proving density is the same as ruling out discreteness, and discreteness for a group generated by a symmetric configuration is ruled out by symmetry alone. This is the crystallographic restriction: if a lattice L of rank n is invariant under a rotation σ of order m, then in a basis of L the matrix of σ has integer entries, so its trace is an integer; in $\mathbb{R}^3$ a rotation of order m has trace $1 + 2\cos(2\pi/m)$, so $2\cos(2\pi/m)$ must be an integer and $m \in \{1, 2, 3, 4, 6\}$. The icosahedral rotation group contains rotations of order 5, with trace $1 + 2\cos 72^\circ = 1 + (\phi - 1) = \phi$, which is irrational. Therefore no lattice can be invariant under the icosahedral group, the generated subgroup cannot be discrete, and since it spans, it is dense. The golden ratio in the coordinates is not a coincidence: it is the trace of the five-fold rotation, and it is irrational for exactly the reason the theorem says it must be.

The discernment mechanism is the false-neighbour test. Before proving a statement, find the nearest configuration for which it fails, and locate the hypothesis that separates them. Here the cube fails and the icosahedron succeeds, and the separating feature is five-fold symmetry, equivalently the appearance of ϕ. Once you have located the separating feature, the proof writes itself: use it. Train this by making yourself state, for every "show that" problem, one nearby false version.

There is a second, physical reading. A discrete group of translations is a crystal; a dense one, generated by finitely many vectors with irrational ratios, is a quasicrystal. The icosahedral vertex vectors generate a dense group precisely because icosahedral symmetry is incompatible with periodicity, which is why icosahedral quasicrystals exist and why five-fold diffraction patterns were considered impossible until they were observed. The Putnam problem is the two-line mathematical core of that story.

What the examiners nudged toward and how: they specified unit vectors, which silently centres the icosahedron at the origin and saves you from worrying about translations; they chose the icosahedron rather than any of the crystallographic solids, which is the entire content; and they phrased the conclusion with an epsilon, which is a small test of whether you can recognise density when it is written out in full.

Revolving concepts

Subgroups of $(\mathbb{R}, +)$. Statement: a nontrivial subgroup is either $c\mathbb{Z}$ for some $c > 0$ or is dense; the proof is the infimum argument given above. Orbiting facts: $\mathbb{Z} + \alpha\mathbb{Z}$ is dense if and only if α is irrational, which is Kronecker's theorem in one dimension and is equivalent to the equidistribution-free statement that the fractional parts $\{n\alpha\}$ are dense in $[0, 1)$; the pigeonhole proof of that is Dirichlet's approximation theorem, which gives infinitely many p/q with $|\alpha - p/q| < 1/q^2$. See Book 1, field: Number theory, topic: Diophantine approximation and dense subgroups.

Closed subgroups of $\mathbb{R}^n$ and lattices. Statement: every closed subgroup of $\mathbb{R}^n$ is $V \oplus L$, where V is a linear subspace and L is a lattice (a discrete subgroup) in a complementary subspace. Consequence used here: a finitely generated subgroup that spans and is not discrete has dense closure in the span of its non-discrete part; combined with irreducibility of the symmetry action, discreteness or density is the only alternative. See Book 1, field: Discrete structures and graphs, topic: lattices and closed subgroups.

The crystallographic restriction. Statement: a rotation of finite order m preserving a rank n lattice acts by an integer matrix in a lattice basis, so its trace is an integer; in the plane the trace is $2\cos(2\pi/m)$ and in three dimensions it is $1 + 2\cos(2\pi/m)$, forcing $m \in \{1, 2, 3, 4, 6\}$. Orbiting facts: this is why crystals have no five-fold axes, why the Penrose tiling is aperiodic, and why the icosahedron is the natural source of a dense vector group. See Book 1, field: Geometry, topic: symmetry groups and lattices.

The golden ratio and icosahedral coordinates. Definition: $\phi = \frac{1+\sqrt{5}}{2}$ is the positive root of $\phi^2 = \phi + 1$, so $\phi - 1 = 1/\phi$ and $2\cos 36^\circ = \phi$, $2\cos 72^\circ = \phi - 1$. Standard construction: the twelve points obtained as the cyclic permutations of $(\pm 1, \pm\phi, 0)$ are the vertices of a regular icosahedron; equivalently three mutually perpendicular 1 by ϕ rectangles have as their twelve corners exactly this vertex set. This construction is a named knowledge gap: you must install it, since re-deriving it in the arena would cost more time than the rest of the problem. See Book 1, field: Geometry, topic: coordinates for the Platonic solids.

The recursion generalised. The golden ratio is the fixed point of the recursion $x_{k+1} = 1 + 1/x_k$, equivalently the limit of F_{k+1}/F_k for the Fibonacci recursion $F_{k+1} = F_k + F_{k-1}$. The general fact worth carrying: for a linear recursion $x_{k+1} = px_k + qx_{k-1}$ with characteristic roots α, β of different modulus, the ratios x_{k+1}/x_k converge to the dominant root, and $\mathbb{Z} + \alpha\mathbb{Z}$ is dense exactly when that root is irrational. In the icosahedral setting the Fibonacci recursion is visible directly: the powers $\phi^k = F_k\phi + F_{k-1}$ all lie in H, so the group elements grow along the Fibonacci sequence, and the small elements $|\phi^{-k}| = |F_{k-1} - F_k(\phi - 1)|$ tending to zero give an explicit proof that H is dense without the infimum lemma. See Book 1, field: Sequences and recurrences, topic: linear recurrences and their dominant roots.

Circuit Bank

- **WHEN** a problem asks for integer combinations of fixed vectors to approximate every point **DO** rename the set as the subgroup generated by those vectors and prove it is dense **BECAUSE** because integer combinations form a group and approximation of every point is density, so both quantifiers vanish and the classification of subgroups becomes available
- **WHEN** you must show a finitely generated subgroup of $\mathbb{R}$ is dense **DO** show it is not cyclic, by exhibiting two generators with irrational ratio **BECAUSE** because a nontrivial subgroup of $\mathbb{R}$ is either $c\mathbb{Z}$ or dense, and two elements of $c\mathbb{Z}$ always have rational ratio
- **WHEN** you are proving a statement about a symmetric configuration and have not yet used the symmetry **DO** find the nearest configuration where the statement is false and isolate the separating feature **BECAUSE** because that feature is what the proof must use, and here it is the five-fold symmetry that forces an irrational coordinate ratio
- **WHEN** a conclusion is invariant under rotations and scalings **DO** normalise the configuration to standard coordinates before computing anything **BECAUSE** because the invariance is free and standard coordinates convert a synthetic object into arithmetic you can add and subtract
- **WHEN** you need density in $\mathbb{R}^n$ and have density in each coordinate **DO** approximate each coordinate to within $\varepsilon/\sqrt{n}$ **BECAUSE** because the Euclidean norm of a vector with entries below $\varepsilon/\sqrt{n}$ is below ε, and a product of dense subgroups is dense in the product

Backtesting

1. First useful cue: integer coefficients plus an arbitrary epsilon, which together mean "the generated subgroup is dense".
2. Method the expression suggested: normalise to golden ratio coordinates, extract axis-aligned elements 2 and 2ϕ by adding and subtracting vertex pairs, and apply the cyclic-or-dense dichotomy for subgroups of $\mathbb{R}$.
3. Where you would hesitate, and the thought to reinforce: at "may we choose coordinates?". Reinforce the sentence: density is preserved by rotations and positive scalings because they are homeomorphisms and group automorphisms, and the unit length hypothesis already centres the solid at the origin.
4. The trap the question sets: treating the epsilon as an analysis problem and trying to build the coefficients, or assuming that twelve spanning vectors are automatically enough, which the cube disproves in one line.
5. What makes you recognise it faster next time: five-fold symmetry means no invariant lattice, so any group generated by a symmetric five-fold configuration that spans is dense; the golden ratio is the trace of the five-fold rotation and its irrationality is the whole engine.
6. The reusable decision rule: integer combinations means subgroup; density means "not cyclic" in one dimension; symmetry that is incompatible with a lattice means density is forced.

Mirror drills

1. (drill) Show that the same statement for the eight vertex vectors of a cube is false. *Answer:* take the vertices to be $(\pm 1, \pm 1, \pm 1)/\sqrt{3}$; every integer combination lies in $\frac{1}{\sqrt{3}}\mathbb{Z}^3$, which is a discrete set, so points of $\mathbb{R}^3$ at distance more than $\frac{1}{2\sqrt{3}}$ from that set are not approximated at all. The same argument kills the tetrahedron and the octahedron.

2. (drill) Prove that $\{n\alpha \bmod 1 : n \in \mathbb{Z}\}$ is dense in $[0,1)$ for irrational α, without the infimum lemma. *Answer:* by pigeonhole, among the $N+1$ numbers $\{k\alpha\}$ for $k = 0,\dots,N$ two lie in the same interval of length $1/N$, so some $\{m\alpha\}$ with $m \neq 0$ is within $1/N$ of an integer; the multiples of that number are $1/N$-dense, and N was arbitrary. This is Dirichlet's approximation theorem doing the same job.

3. (drill) For which of the five Platonic solids is the A4 statement true? *Answer:* exactly the icosahedron and the dodecahedron, the two with five-fold symmetry. The dodecahedron has vertices $(\pm 1, \pm 1, \pm 1)$ together with the cyclic permutations of $(0, \pm\phi^{-1}, \pm\phi)$, and the same adding and subtracting gives 2 and $2\phi^{-1}$ on each axis, so the group contains a dense subgroup of each coordinate line. The tetrahedron, cube and octahedron all have rational coordinates after scaling, so their groups are lattices.

3.5 A5. Base three weights collapse a shift operator

For a nonnegative integer k, let $f(k)$ be the number of ones in the base 3 representation of k. Find all complex numbers z such that

$$\sum_{k=0}^{3^{1010}-1} (-2)^{f(k)}(z+k)^{2023} = 0.$$

Read it out loud

A weighted sum of 3^{1010} translates of one fixed polynomial, z^{2023}, where the weight attached to the translate by k depends only on the base 3 digits of k. Two structures are being handed to you at once. First, the weight is multiplicative across digits: $(-2)^{f(k)}$ is the product over digit positions of 1, -2 or 1 according as the digit is 0, 1 or 2, and k runs over every digit string of length 1010. Anything defined as a product over independent coordinates, summed over the full product of ranges, factorises. Second, "sum of shifted copies" is not an arithmetic object at all, it is a linear operator applied to one polynomial, and the operator is built out of the single shift $S\colon p(z) \mapsto p(z+1)$.

Before anything else, write $n = 1010$ and stare at the exponent: $2023 = 2 \cdot 1010 + 3 = 2n + 3$. That is not a coincidence and it is the whole design of the problem. The generating polynomial of the weights will turn out to have a zero of order exactly $2n$ at $x = 1$, an operator whose symbol vanishes to order $2n$ drops the degree of any polynomial by exactly $2n$, and $2n + 3$ drops to 3. A cubic has three roots; the question asks you to find them.

Fields wired in: generating functions over digit expansions, the finite difference and shift calculus (the umbral or symbol calculus, with $S = e^{\partial}$), and elementary root finding for a depressed cubic. The knowledge base is: digit-multiplicative weights factorise; $\Delta = S - I$ lowers polynomial degree by exactly one; a formal power series in ∂ acts on polynomials because ∂ is nilpotent on each one; and sinh has an even Taylor series.

The third cue is the palindrome. The digits of $3^n - 1 - k$ are $2 - d_j$ where d_j are the digits of k, and $2 - d = 1$ exactly when $d = 1$, so $f(3^n - 1 - k) = f(k)$ and the weight sequence is a palindrome. A palindromic weight makes the whole sum symmetric about the centre $k = \frac{3^n-1}{2}$, so you should predict, before computing anything, that the answers are symmetric about $z = -\frac{3^n-1}{2}$ and that this centre is itself a root, because 2023 is odd.

Approaches, ranked worst to best

1. **Binomial expansion and moments.** *Cue:* "$(z+k)^{2023}$". *Circuit:* expand and collect powers of z. (Writing the sum as $\sum_i \binom{2023}{i} z^{2023-i} M_i$ with $M_i = \sum_k (-2)^{f(k)} k^i$ leaves you needing 2024 moment sums over 3^{1010} terms. It is not wrong, and in fact the truth "$M_i = 0$ for $i < 2020$" is exactly the degree collapse in disguise, but as a plan it has no way to compute a single one of those moments. Dead on the page.)

2. **Recursion on n with an explicit two-term ansatz.** *Cue:* the digit structure gives a recursion between n and $n-1$. *Circuit:* induct on the number of digits. (This is the published route: prove by induction on i that the product of the first i symmetric second difference operators applied to z^{2n+3} equals $C_i(z^{2n+3-2i} + \alpha_i z^{2n+1-2i}) + O(z^{2n-1-2i})$. It is correct and complete, but you must carry an unknown constant C_i, an explicit α_i and an error term through an induction, and the bookkeeping is where marks are lost.)

3. **Symmetry plus degree count plus one moment.** *Cue:* palindromic weights and a huge exponent. *Circuit:* centre the variable, use parity, then pin the single remaining constant. (Centre at $z_0 = -\frac{3^n-1}{2}$; the sum becomes an odd polynomial in $w = z - z_0$; the degree collapse says it is a cubic; so it is $C(w^3 + \beta w)$ and only β is unknown, computable from one coefficient.

Efficient, and it explains the shape of the answer. It ranks second only because computing β still needs the operator expansion of the next route.)

4. **Shift operator with the exponential symbol.** *Cue:* a sum of translates with digit-defined weights. *Circuit:* $\sum_k c_k p(z+k) = C(S)p$ with $S = e^{\partial}$. (The ideal route. Factor $C(x) = \sum_k c_k x^k = \prod_{j=0}^{n-1}(x^{3^j}-1)^2$; substitute $x = S = e^{\partial}$; use $(e^t-1)^2 = e^t \cdot 4\sinh^2(t/2)$ to split off the shift by $\frac{3^n-1}{2}$ and leave an even symbol; expand that symbol to two terms in ∂; apply to z^{2n+3}. No induction, no error terms, and the centre of symmetry appears automatically as the factored-out shift.)

The ideal solution

Write $n = 1010$, so the exponent is $2n+3 = 2023$, and set $c_k = (-2)^{f(k)}$ and

$$P(z) = \sum_{k=0}^{3^n-1} c_k (z+k)^{2n+3}.$$

Step 1: the weights factorise. Every k with $0 \le k < 3^n$ has a unique base 3 expansion $k = \sum_{j=0}^{n-1} d_j 3^j$ with $d_j \in \{0,1,2\}$, and this is a bijection between digit strings and the range of k. Since $f(k) = \#\{j : d_j = 1\}$ we have $c_k = \prod_j w_{d_j}$ with $w_0 = w_2 = 1$ and $w_1 = -2$. Hence, for a commuting indeterminate x,

$$C(x) := \sum_{k=0}^{3^n-1} c_k x^k = \prod_{j=0}^{n-1} \left(1 - 2x^{3^j} + x^{2\cdot 3^j}\right) = \prod_{j=0}^{n-1} \left(x^{3^j}-1\right)^2.$$

Step 2: the sum is an operator applied to a monomial. Let S be the shift on $\mathbb{C}[z]$, $(Sp)(z) = p(z+1)$, so that $(S^k p)(z) = p(z+k)$. The map $x \mapsto S$ is a ring homomorphism from $\mathbb{C}[x]$ to the (commutative) ring of operators generated by S, so Step 1 gives

$$P(z) = C(S)\, z^{2n+3} = \prod_{j=0}^{n-1} \left(S^{3^j} - I\right)^2 z^{2n+3}.$$

Step 3: the symbol. For a polynomial p, Taylor's theorem is exact and the sum is finite: $p(z+t) = \sum_{i \ge 0} \frac{t^i}{i!} p^{(i)}(z)$. Hence $S^t = e^{t\partial}$ as operators on $\mathbb{C}[z]$, where $\partial = d/dz$.

Lemma 3.5.1. *Formal power series in ∂ act on $\mathbb{C}[z]$, and the action is a ring homomorphism: if $F(\partial) = \sum_{i \ge 0} \gamma_i \partial^i$ and p has degree d, then $F(\partial)p := \sum_{i=0}^{d} \gamma_i p^{(i)}$ is well defined, and $(FG)(\partial)p = F(\partial)\big(G(\partial)p\big)$.*

Proof. Only finitely many terms act nontrivially on a given p, since $\partial^i p = 0$ for $i > d$, so each series acts. Both sides of the homomorphism identity depend only on the coefficients of F and G up to index d, where the identity is the usual polynomial one. □

Since $(e^t-1)^2 = e^t\left(e^{t/2} - e^{-t/2}\right)^2 = e^t \cdot 4\sinh^2(t/2)$, applying this with $t = 3^j\partial$ and multiplying over j gives

$$C(S) = e^{m\partial} \prod_{j=0}^{n-1} 4\sinh^2\Big(\frac{3^j\partial}{2}\Big), \qquad m = \sum_{j=0}^{n-1} 3^j = \frac{3^n-1}{2},$$

and $e^{m\partial} = S^m$ is the shift $p(z) \mapsto p(z+m)$.

Step 4: two terms of the even symbol. From $\sinh u = u + \frac{u^3}{6} + O(u^5)$ we get $4\sinh^2(u/2) = u^2 + \frac{u^4}{12} + O(u^6)$, so with $u = 3^j\partial$,

$$4\sinh^2\Big(\frac{3^j\partial}{2}\Big) = 9^j\partial^2\Big(1 + \frac{9^j\partial^2}{12} + O(\partial^4)\Big),$$

and multiplying over $j = 0, \dots, n-1$,

$$\prod_{j=0}^{n-1} 4\sinh^2\Big(\frac{3^j\partial}{2}\Big) = A\,\partial^{2n}\Big(1 + \frac{\partial^2}{12}\sum_{j=0}^{n-1} 9^j + O(\partial^4)\Big), \qquad A = \prod_{j=0}^{n-1} 9^j = 9^{n(n-1)/2} \neq 0.$$

Now apply this to z^{2n+3}, using $\partial^{2n} z^{2n+3} = \frac{(2n+3)!}{3!} z^3$, $\partial^{2n+2} z^{2n+3} = (2n+3)!\, z$ and $\partial^i z^{2n+3} = 0$ for $i > 2n+3$, so the $O(\partial^4)$ tail contributes nothing:

$$\prod_{j=0}^{n-1} 4\sinh^2\Big(\frac{3^j\partial}{2}\Big) z^{2n+3} = A(2n+3)!\left(\frac{z^3}{6} + \frac{1}{12}\cdot\frac{9^n-1}{8}\,z\right) = \frac{A(2n+3)!}{6}\left(z^3 + \frac{9^n-1}{16}z\right),$$

where $\sum_{j=0}^{n-1} 9^j = \frac{9^n-1}{8}$.

Step 5: undo the shift and solve. Applying S^m replaces z by $z+m$, so

$$P(z) = \frac{A(2n+3)!}{6}\left((z+m)^3 + \frac{9^n-1}{16}(z+m)\right), \qquad m = \frac{3^n-1}{2}.$$

The leading constant is nonzero, so $P(z) = 0$ if and only if $(z+m)\big((z+m)^2 + \frac{9^n-1}{16}\big) = 0$, that is $z = -m$ or $(z+m)^2 = -\frac{9^n-1}{16}$. With $n = 1010$ the complete solution set is

$$\boxed{z = -\frac{3^{1010}-1}{2} \quad \text{and} \quad z = -\frac{3^{1010}-1}{2} \pm \frac{\sqrt{9^{1010}-1}}{4}\,i}$$

Two checks worth one line each. The case $n = 1$ of the formula reads $z^5 - 2(z+1)^5 + (z+2)^5 = 20\big((z+1)^3 + \frac{1}{2}(z+1)\big)$, and substituting $w = z+1$ makes the left side $(w-1)^5 - 2w^5 + (w+1)^5 = 2\binom{5}{2}w^3 + 2\binom{5}{4}w = 20w^3 + 10w$, which agrees. And the symmetry predicted in advance is visible: $c_k = c_{3^n-1-k}$ gives $P(-(3^n-1) - z) = -P(z)$ because the exponent 2023 is odd, so P is an odd function of $z+m$ and its roots are symmetric about $-m$, with $-m$ itself a root.

Layer 1: WHAT

Block 1. Set $n = 1010$ and notice $2023 = 2n+3$; predict a degree drop of $2n$ and a cubic answer.

Block 2. Factor the weights across base 3 digits: $C(x) = \prod_{j<n}(x^{3^j}-1)^2$, which has a zero of order $2n$ at $x = 1$.

Block 3. Recognise the sum as $C(S)z^{2n+3}$ for the shift operator S, so the factorisation of C becomes a factorisation of the operator.

Block 4. Substitute $S = e^{\partial}$ and split each factor as $e^{3^j\partial}\cdot 4\sinh^2(3^j\partial/2)$, collecting all the exponentials into the single shift by $m = \frac{3^n-1}{2}$.

Block 5. Expand the remaining even symbol to two terms in ∂ and apply it to z^{2n+3}; the result is $z^3 + \frac{9^n-1}{16}z$ up to a nonzero constant.

Block 6. Shift back and solve the depressed cubic.

Juxtaposition. The nominal route expands $(z+k)^{2023}$ and tries to evaluate moments; it never gets past writing M_i. A slightly better zombie route sets up the recursion in n and grinds an induction with an unknown constant and an error term. The ideal route treats the entire sum as one operator and computes only the first two terms of its symbol. The bifurcation point is the very first line: do you expand the bracket $(z+k)^{2023}$ (zombie) or do you write $\sum_k c_k x^k$ and factor it (ideal)? Every later economy follows from having a product instead of a sum.

Layer 2: HOW

Block 1. Cue. Two large numbers, 3^{1010} and 2023, both chosen by the setter. **Why natural.** On the Putnam the specific constants are always tuned so that the answer is finite and clean. Naming

$n = 1010$ turns two mysterious numbers into one relation, $2023 = 2n + 3$, and a relation is something you can use. It is worth repeating: never carry a competition's year in your algebra; parametrise it and find its relation to the other parameter. **Link to the next block.** A relation between the size of the sum and the exponent suggests that the sum will annihilate most of the degree, so ask exactly how much.

Block 2. Cue. The weight depends only on the base 3 digits of k, and k runs over every value in a complete digit range. **Why natural.** Digits are independent coordinates, and $x^k = \prod_j x^{d_j 3^j}$ turns the sum over the product set into a product of sums, one per digit position. Each digit contributes $1 - 2x^{3^j} + x^{2\cdot 3^j}$, which is a perfect square $(x^{3^j} - 1)^2$. Spotting the square is not decoration: it says the generating polynomial vanishes to order two at $x = 1$ for each digit, hence order $2n$ in total, and that number is exactly the degree drop you predicted. **Link to the next block.** Now you have a factored polynomial $C(x)$; the sum you want is what happens when x is replaced by the shift.

Block 3. Cue. The sum is $\sum_k c_k p(z + k)$: the same p, translated, with constant weights. **Why natural.** Translations commute and compose, so $p(z + k) = S^k p$, and a weighted sum of translates is $C(S)p$ for the very polynomial C from Step 1. This substitution is a ring homomorphism, so any factorisation of C transfers to a factorisation of the operator. This is the single most valuable reflex in problems about weighted sums of shifted values, and it converts a 3^{1010}-term sum into a product of $2n$ elementary factors. **Link to the next block.** Products of shift operators are easiest to expand through their symbols, so pass to e^{∂}.

Block 4. Cue. S is a translation, and Taylor's theorem for polynomials is a finite sum. **Why natural.** $S^t = e^{t\partial}$ is exact on $\mathbb{C}[z]$, not merely formal, because ∂ is nilpotent on each polynomial. Once the operator is written in terms of ∂, the factor $(e^t - 1)^2$ is best symmetrised: $(e^t - 1)^2 = e^t \cdot 4\sinh^2(t/2)$ separates a pure shift from an even function of t. The even part is what preserves parity, and collecting the shifts gives $\sum_j 3^j = \frac{3^n - 1}{2}$, which is precisely the centre of symmetry you predicted from the palindromic weights. Two independent arguments landing on the same centre is a strong sign you have not slipped. **Link to the next block.** The remaining operator is even in ∂, so applied to an odd monomial it gives an odd polynomial: only z^3 and z can survive.

Block 5. Cue. You need only a cubic, and ∂^i annihilates z^{2n+3} for $i > 2n + 3$. **Why natural.** Every extra factor of ∂^2 costs two degrees, so of the whole symbol only the coefficients of ∂^{2n} and ∂^{2n+2} can contribute; everything else either raises the order of vanishing beyond the degree or was already accounted for. Therefore expand each factor to two terms only. The arithmetic $\sum_j 9^j = \frac{9^n - 1}{8}$ is the only computation in the problem. **Link to the next block.** A constant multiple of $z^3 + \frac{9^n - 1}{16} z$; now restore the shift.

Block 6. Cue. A depressed cubic $w^3 + \beta w$ with $\beta > 0$. **Why natural.** It factors as $w(w^2 + \beta)$, so one real root and two purely imaginary ones; undoing $w = z + m$ translates all three by $-m$. Because the problem asked for complex numbers, all three count, and you must present all three. **Link to the next block.** Substituting $n = 1010$ gives the answer.

Layer 3: WHY

Three truths carry this solution. The first is independence: a weight that is a product over coordinates, summed over a full product of ranges, factorises. That is the same principle that makes the characteristic function of a sum of independent random variables a product, and it is why digit-defined sequences (the Thue and Morse sequence, digit sums, digit counts) always have product generating functions. The second is that the algebra of shift operators on polynomials is

just the algebra of polynomials, through $x \mapsto S$, and it embeds in formal power series in ∂ through $S = e^{\partial}$. Under that embedding, the order of vanishing of C at $x = 1$ becomes the order of vanishing of the symbol at $\partial = 0$, which is exactly the number of degrees that $C(S)$ removes. Stated cleanly: if $C(x) = (x-1)^m U(x)$ with $U(1) \neq 0$, then $C(S)$ annihilates every polynomial of degree less than m and reduces the degree of every other polynomial by exactly m. Here $m = 2n = 2020$ and $2023 - 2020 = 3$. The third is symmetry: palindromic weights make the functional invariant under reflecting k about the centre, so after centring the variable the answer is even or odd according to the parity of the exponent.

The discernment mechanism is a question to ask of any expression of the form $\sum_k c_k p(z+k)$: what is $C(x) = \sum_k c_k x^k$, and what is its order of vanishing at $x = 1$? That single number tells you the degree of the answer before you compute anything, and it tells you whether the problem is even well posed. If the exponent had been below $2n$ the sum would vanish identically and every complex z would be a solution; at exponent $2n$ it is a nonzero constant and no z works; at $2n+1$ there is one root; the setter chose $2n+3$ to get a genuine cubic with three distinct roots, two of them non-real, so that the phrase "all complex numbers" earns its place. Train the reflex by computing $\Delta^m p$ for small cases until "the difference operator eats one degree" is automatic.

What the examiners nudged toward and how: they used base 3 with the weight (-2) on the digit 1, which is exactly the combination making each digit factor a perfect square rather than a general quadratic; they made the range a complete block $0 \le k < 3^{1010}$, without which no factorisation exists; they made the exponent odd, so that the centre is itself a root; and they tuned the exponent to $2n+3$ so that the surviving polynomial is a cubic.

Revolving concepts

Digit-multiplicative sequences. Definition: a sequence with $c_{bk+d} = w_d c_k$ for $0 \le d < b$ and $c_0 = 1$ is determined by the digit weights w_d, and $c_k = \prod_j w_{d_j}$ over the base b digits of k. Generating function: $\sum_{k=0}^{b^n-1} c_k x^k = \prod_{j=0}^{n-1} \left(\sum_{d=0}^{b-1} w_d x^{db^j} \right)$. Orbiting examples: the Thue and Morse sequence ($b = 2$, $w_0 = 1$, $w_1 = -1$) with $\prod_j (1 - x^{2^j})$; digit sums; the $(-2)^{f(k)}$ of this problem. See Book 1, field: Generating functions, topic: digit expansions and product formulas.

Finite difference and shift calculus. Definitions: $Sp(z) = p(z+1)$, $\Delta = S - I$, and Δ lowers the degree of a polynomial by exactly one, with $\Delta z^d = dz^{d-1} + \cdots$. Theorems: $\Delta^m p = \sum_i (-1)^{m-i} \binom{m}{i} p(z+i)$; $\Delta^m p = 0$ if and only if $\deg p < m$; Newton's forward difference formula $p(z) = \sum_m \binom{z}{m} \Delta^m p(0)$. The general annihilation lemma used here: for $C(x) = (x-1)^m U(x)$ with $U(1) \neq 0$, the operator $C(S)$ kills polynomials of degree below m and drops all others by exactly m. See Book 1, field: Sequences and recurrences, topic: finite differences.

Symbol calculus, $S = e^{\partial}$. Statement: on $\mathbb{C}[z]$ the shift equals e^{∂} exactly, since Taylor's theorem is a finite identity for polynomials; consequently formal power series in ∂ act, and act as a commutative ring. Orbiting identities: $\Delta = e^{\partial} - 1$, so $\partial = \log(1 + \Delta)$, which is Newton's series; the symmetric second difference $S^{\ell} - 2I + S^{-\ell} = 4 \sinh^2(\ell \partial / 2)$; the Euler and Maclaurin operator $\partial / (e^{\partial} - 1)$ whose coefficients are the Bernoulli numbers. See Book 1, field: Generating functions, topic: umbral and operator calculus.

Palindromic weights and centring. Statement: if $c_k = c_{K-k}$ for all k in the range with $K = 3^n - 1$, then $\sum_k c_k (z+k)^N$ satisfies $P(-K - z) = (-1)^N P(z)$, so after substituting $w = z + K/2$ the polynomial is even or odd in w according to the parity of N. Consequence used here: half the coefficients of the cubic vanish for free, and the centre is a root when N is odd. See Book 1, field: Polynomials, topic: symmetry of coefficient sequences.

The recursion generalised. The weights obey $c_{3k} = c_k$, $c_{3k+1} = -2c_k$, $c_{3k+2} = c_k$, a digit recursion. Its general form and solution: if $c_{bk+d} = w_d c_k$ then, writing $C_n(x) = \sum_{k < b^n} c_k x^k$, one has the recursion $C_n(x) = C_{n-1}(x^b) \cdot \left(\sum_d w_d x^d \right)$ evaluated with the digit blocks in the right order, whose solution is the product formula above; the same telescoping product appears for every self-similar

(automatic) sequence. The operator consequence generalises too: $C_n(S)$ annihilates polynomials of degree less than $n \cdot \mathrm{ord}_{x=1}\left(\sum_d w_d x^d\right)$. See Book 1, field: Generating functions, topic: self-similar sequences and their products.

Circuit Bank

- **WHEN** you see $\sum_k c_k p(z+k)$, a weighted sum of translates of one polynomial **DO** write it as $C(S)p$ where $C(x) = \sum_k c_k x^k$ and S is the shift, then factor C **BECAUSE** because translations commute, so the sum is a polynomial in one operator, and any factorisation of C becomes a factorisation of the operator
- **WHEN** a weight depends only on the base b digits of the index and the index runs over a complete block $0 \le k < b^n$ **DO** factor the generating polynomial one digit at a time **BECAUSE** because digits are independent coordinates and a product weight summed over a product set always factorises
- **WHEN** a shift polynomial C has a zero of order m at $x = 1$ **DO** conclude the operator annihilates degrees below m and drops every other degree by exactly m **BECAUSE** because $C(S) = (S-I)^m U(S)$ with $U(I)$ invertible on each degree, and $S - I$ lowers degree by exactly one
- **WHEN** you need only the leading few coefficients of a complicated operator applied to a monomial **DO** expand the symbol in powers of ∂ and keep as many terms as the degree can support **BECAUSE** because ∂ is nilpotent on any fixed polynomial, so the symbol calculus is exact and the tail costs nothing
- **WHEN** the coefficient sequence is a palindrome **DO** centre the variable at the midpoint before expanding **BECAUSE** because the sum then becomes even or odd in the centred variable according to the parity of the exponent, and half the unknown coefficients vanish without computation

Backtesting

1. First useful cue: a weight defined by base 3 digits summed over the complete range $0 \le k < 3^{1010}$, next to the relation $2023 = 2 \cdot 1010 + 3$.
2. Method the expression suggested: factor the digit generating polynomial into $\prod(x^{3^j} - 1)^2$, replace x by the shift $S = e^{\partial}$, peel off the shift by $\frac{3^n-1}{2}$, expand the even symbol to two terms, apply to z^{2n+3}.
3. Where you would hesitate, and the thought to reinforce: at "may I substitute $S = e^{\partial}$ and multiply power series?". Reinforce: on polynomials ∂ is nilpotent, so every power series in ∂ acts through a finite sum and the algebra is exact; this is a theorem, not a formal manipulation.
4. The trap the question sets: expanding $(z+k)^{2023}$ by the binomial theorem, which produces 2024 moment sums and no route to any of them; and, at the end, giving only the real root when the question says "all complex numbers".
5. What makes you recognise it faster next time: any sum of translates is an operator; any digit-defined weight factorises; the order of vanishing at $x = 1$ of the weight generating polynomial is the number of degrees the operator removes.
6. The reusable decision rule: convert translate sums into $C(S)$, factor C, read the order of vanishing at 1 to get the degree of the answer, and expand the symbol only as far as that degree allows.

Mirror drills

1. (drill) Compute $z^5 - 2(z+1)^5 + (z+2)^5$ and identify its roots. *Answer:* substituting $w = z+1$ gives $(w-1)^5 - 2w^5 + (w+1)^5 = 2\binom{5}{2}w^3 + 2\binom{5}{4}w = 20w^3 + 10w$, so the roots are $z = -1$ and $z = -1 \pm \frac{i}{\sqrt{2}}$. This is the $n = 1$ case of A5, and it matches the general formula with $\frac{\sqrt{9-1}}{4} = \frac{1}{\sqrt{2}}$.

2. (drill) Show that $\sum_{i=0}^{m} (-1)^{m-i}\binom{m}{i} p(z+i) = 0$ for every polynomial p of degree less than m. *Answer:* the left side is $(S - I)^m p = \Delta^m p$, and Δ lowers degree by exactly one, so m applications annihilate anything of degree below m. Equivalently the weight generating polynomial $(x-1)^m$ vanishes to order m at $x = 1$.

3. (drill) In the A5 sum, replace the exponent 2023 by 2020, by 2021, and by 2019. What are the solution sets? *Answer:* the operator removes exactly 2020 degrees. At exponent 2019 the sum is identically zero, so every complex z works; at 2020 it is a nonzero constant, so no z works; at 2021 it is a nonzero multiple of $z + \frac{3^{1010}-1}{2}$, so the unique solution is the centre. This is why the setter chose $2n + 3$.

3.6 A6. Bob owns the parity of the fixed points

Alice and Bob play a game in which they take turns choosing integers from 1 *to* n*. Before any integers are chosen, Bob selects a goal of "odd" or "even". On the first turn, Alice chooses one of the* n *integers. On the second turn, Bob chooses one of the remaining integers. They continue alternately choosing one of the integers that has not yet been chosen, until the* n*th turn, which is forced and ends the game. Bob wins if the parity of* $\{k : \text{the number } k \text{ was chosen on the } k\text{th turn}\}$ *matches his goal. For which values of* n *does Bob have a winning strategy?*

Read it out loud

Two players alternately remove numbers from $\{1, \ldots, n\}$ until nothing is left. Assigning to each turn the number chosen on it produces a bijection π from turns to numbers, that is, a permutation of $\{1, \ldots, n\}$, and the set in the statement is exactly the fixed point set of π. So the game builds a random-looking permutation under adversarial control, and the payoff is one bit: the parity of the number of fixed points. Bob must nail that bit, and, worse, he must commit to which value of the bit he is chasing before a single move is made.

Fields wired in: permutations and their fixed points, parity arguments, and combinatorial game theory in its most elementary form, namely explicit strategies verified by induction on the size of the game. There is no clever counting here and no formula to derive; the entire content is the construction of a strategy that works against every line of play, plus the discipline to check it against every line of play.

The first useful cue is the alternation next to the indexing. Alice takes the odd turns $1, 3, 5, \ldots$ and Bob the even ones. A fixed point at k can be created only on turn k. Therefore odd numbers can be fixed only by Alice and even numbers only by Bob. Write that sentence down within thirty seconds, because it is the hinge of both halves of the solution: it tells Bob that pairing an odd number with the even number immediately after it locks the two of them together, and it tells him that occupying the number $t + 1$ on his turn t silently destroys Alice's only chance to score on her next turn.

The second cue is the phrase "the nth turn, which is forced". The setter is pointing at the last move. Who owns it? If n is even it is Bob's, and if n is odd it is Alice's. That single fact decides which goal Bob should choose, and it is also why the odd case is genuinely harder: Alice gets the last word, so Bob cannot simply mirror.

The third cue is what the answer is likely to be. A "for which n" question whose answer is "all n" is common when the underlying mechanism is a strategy rather than a numerical obstruction, and here Bob has two independent sources of power (he moves second, and he can occupy the number that would score for Alice next). Reconnaissance on $n = 1, 2, 3$ confirms it: for $n = 1$ Alice is forced to take 1 and there is one fixed point; for $n = 2$, either Alice takes 1 and Bob takes 2, giving two, or Alice takes 2 and Bob takes 1, giving zero, so Bob names "even" and wins.

Approaches, ranked worst to best

1. **Play small cases and assert the pattern.** *Cue:* "for which values of n". *Circuit:* compute, conjecture, stop. (Necessary reconnaissance, not a solution. Small cases give the answer, "all n", and the correct goal, matching the parity of n, but the question is a strategy question: without an explicit strategy proved against all of Alice's replies, you have nothing. Do the reconnaissance in two minutes and then throw it away.)

2. **Look for an invariant.** *Cue:* a parity payoff. *Circuit:* find a quantity neither player can change. (A natural instinct, and it fails here for a reason worth internalising: the parity of the fixed point count is genuinely contested, since Alice can create a fixed point by taking 1 on turn 1

and Bob can create one by taking 2 on turn 2. When both players can move the quantity, no invariant exists and you must build a strategy instead.)

3. **One global pairing for every n.** *Cue:* Bob moves second, so he can mirror. *Circuit:* answer inside a perfect matching of the ground set. (For even n this is the whole solution and it is beautiful: pair $\{1,2\},\{3,4\},\dots$ and always answer with the partner. For odd n the ground set has no perfect matching, and the naive repair, pairing $\{1,2\},\dots,\{n-2,n-1\}$ and leaving n over, loses. Concretely for $n=5$: Alice takes 3, Bob answers 4, Alice takes the unpaired 5, and Bob has no partner to play; whatever he does the count ends even. Rank this third: right for half the problem, and its failure tells you what the other half needs.)

4. **Pair for even n; for odd n bank one fixed point, then block, and reset by induction.** *Cue:* the last turn belongs to Alice exactly when n is odd. *Circuit:* a one-off parity adjustment plus a permanent shutdown, with strong induction to absorb the exceptional branch. (The ideal route. The key auxiliary result is a blocking lemma: from any of his turns, Bob can guarantee that no further fixed point ever occurs. Given that, all Bob needs is to reach an odd count while it is his move, and the whole odd case becomes an induction with one exceptional branch that resets to a smaller odd game.)

The ideal solution

Answer: Bob has a winning strategy for every positive integer n. He selects the goal "even" when n is even and "odd" when n is odd.

Set-up and notation. Each number is chosen exactly once, so the map π sending a turn t to the number chosen on turn t is a permutation of $\{1,\dots,n\}$, and "the number k was chosen on the kth turn" says $\pi(k)=k$. Call such a k a *hit*. Alice moves on the odd turns and Bob on the even ones. We use repeatedly the following remark.

Remark 3.6.1. A hit at k occurs on turn k. Hence every odd hit is made by Alice and every even hit is made by Bob. In particular, on his turn t Bob scores a hit if and only if he chooses the number t itself, and if the number $t+1$ has already been chosen then Alice cannot score a hit on turn $t+1$.

Part 1: n even. Bob forces an even number of hits.

Partition $\{1,\dots,n\}$ into the pairs $\{1,2\},\{3,4\},\dots,\{n-1,n\}$. Bob's strategy: whenever Alice chooses a number, Bob chooses its partner. This is always legal, because after each of Bob's moves the set of chosen numbers is a union of complete pairs, so when Alice next chooses a number its partner is still available.

Fix k with $1\le k\le n/2$; we claim $2k-1$ is a hit if and only if $2k$ is. If $2k-1$ is a hit it was chosen on the odd turn $2k-1$, hence by Alice, so Bob chose $2k$ on turn $2k$ and $2k$ is a hit. Conversely, if $2k$ is a hit it was chosen on the even turn $2k$, hence by Bob, and Bob only ever chooses the partner of the number Alice has just chosen, so Alice chose $2k-1$ on turn $2k-1$, making $2k-1$ a hit. Thus each pair contributes 0 or 2 hits, and the total is even. Bob's goal "even" is achieved.

Part 2: n odd. Bob forces an odd number of hits. We first isolate the shutdown mechanism.

Lemma 3.6.2 (Blocking)**.** *Let n be odd and let t be one of Bob's turns. Then Bob can play from turn t onwards so that no hit occurs on any turn $s\ge t$.*

Proof. Bob's rule at each of his turns $s\ge t$: if the number $s+1$ is still available, choose it; otherwise choose any available number other than s.

The rule is legal. Bob's turns are even and n is odd, so $s\le n-1$, and therefore at least $n-s+1\ge 2$ numbers are available; at most one of them equals s, so a choice other than s exists.

No hit occurs on Bob's turns from t on: if he chooses $s+1$ then $s+1\ne s$, and otherwise he chooses something other than s by construction. No hit occurs on Alice's turns from t on either: her turns after t are the turns $s+1$ for Bob's turns $s\ge t$, and in both branches of the rule the number

$s+1$ is unavailable when she moves, either because Bob has just taken it or because it was already gone. Since every turn after t is a turn of Bob's or a turn immediately following one, no hit occurs at any turn $s \geq t$. □

Theorem 3.6.3. *For every odd $n \geq 1$, Bob can force the number of hits to be odd.*

Proof. Strong induction on odd n.

Base $n = 1$. Alice is forced to choose 1 on turn 1, which is a hit, so the count is 1.

Induction step, $n \geq 3$. Let a be Alice's choice on turn 1.

Case (i): $a = 1$. This is a hit and the count is 1. Bob applies the Blocking Lemma from turn 2; no further hit occurs and the total is 1, odd.

Case (ii): $a = 2$. No hit on turn 1. Bob chooses 1 on turn 2, which is available and is not a hit since $1 \neq 2$. The remaining numbers are $\{3, \ldots, n\}$ and the remaining turns are $3, \ldots, n$, with Alice to move. Subtracting 2 from every number and every turn index identifies this position with the initial position of the whole game on $n-2$ numbers, and hits correspond to hits. Since $n-2$ is odd, the induction hypothesis lets Bob force an odd number of hits there, so the total is $0 +$ odd, odd.

Case (iii): $a = k$ with $k \geq 3$. No hit on turn 1. Bob chooses 2 on turn 2, which is available since $k \neq 2$ and is a hit, so the count is 1.

Bob now plays as follows on each later turn s of his. If Alice did not score a hit on turn $s-1$, he applies the Blocking Lemma from turn s and the game is over for our purposes. Otherwise he chooses s if it is available, and chooses 1 if it is not.

To see that this wins, observe that as long as Alice keeps scoring, an immediate induction on the turn number shows that just before Bob's turn s the chosen numbers are exactly $\{k\} \cup \{2, 3, \ldots, s-1\}$ and the hits are exactly the turns $2, 3, \ldots, s-1$, so the count is $s-2$, which is even. There are now three ways the play can go.

First, suppose Alice fails to score on some turn $j \geq 3$ of hers, either by choice or because $j = k$ is already gone. Just before turn j the chosen numbers are $\{k\} \cup \{2, \ldots, j-1\}$ and the count is $j-2$, which is odd since j is odd. No hit occurs on turn j, and Bob applies the Blocking Lemma from turn $j+1$, so the total is $j-2$, odd.

Second, suppose the run continues and Bob reaches a turn s with s unavailable. Since the chosen numbers are $\{k\} \cup \{2, \ldots, s-1\}$, this happens exactly when $s = k$, so k is even and $k \geq 4$. Bob chooses 1, which is available and is not a hit since $1 \neq k$. Now the chosen numbers are $\{1, 2, \ldots, k\}$, the remaining numbers are $\{k+1, \ldots, n\}$, the remaining turns are $k+1, \ldots, n$, and Alice is to move: subtracting k identifies this with the whole game on $n-k$ numbers, which is odd because n is odd and k is even, and nonempty because $k \neq n$ for the same parity reason. The count so far is $k-2$, even, and by the induction hypothesis Bob forces an odd number of hits from the sub-game, so the total is odd.

Third, suppose the run never breaks. Then Alice scores on every one of her turns and Bob on every one of his, so all of $2, 3, \ldots, n-1$ are chosen on their own turns, and together with Alice's first move this accounts for $n-1$ numbers, forcing $k = n$. The only number left for turn n is 1, which Alice must take, and $1 \neq n$, so it is not a hit. The count is $n-2$, odd.

In every line Bob forces an odd number of hits, completing the induction. □

Combining the two parts, Bob wins for every n, which is the answer: $\boxed{\text{all } n}$.

Why the goal must match the parity of n. The choice of goal is not free, and the reason is the same blocking mechanism run by the other player. For even n, Alice can force zero hits: on each of her turns t she chooses $t+1$ if available and otherwise anything other than t. This is legal because her turns are odd and n is even, so $t \leq n-1$ and at least two numbers remain; her own choice is never t, and Bob's turn $t+1$ can never be a hit because $t+1$ is gone. Every turn is one of hers or one immediately after, so no hit ever occurs. Hence for even n Bob could not have won with the goal "odd", and by the theorem he cannot lose with "even". For odd n, Alice's blocking breaks down at

her final turn n, which is forced, and the theorem shows Bob takes the odd parity. So for each n exactly one goal is winning, and it is the parity of n.

Layer 1: WHAT

Block 1. Name the object: the play is a permutation from turns to numbers, and the payoff is the parity of its fixed points.

Block 2. Extract the parity alignment: a number can be hit only on the turn equal to itself, so odd numbers are Alice's to hit and even numbers are Bob's.

Block 3. Even n: pair $\{2k-1, 2k\}$ and answer within the pair. The alignment makes the two members of a pair hit or miss together, so the count is even.

Block 4. Odd n: prove the Blocking Lemma, which says Bob can shut down all future hits from any of his turns.

Block 5. Odd n: Bob's job reduces to reaching an odd count on one of his turns. He banks the hit at 2 on turn 2 if Alice's opening leaves it available.

Block 6. Handle the two exceptional branches by induction: Alice opening with 2, and the run colliding with Alice's first number k, both of which reset the position to the same game on fewer numbers with Alice to move.

Juxtaposition. The nominal route computes small cases, guesses "all n", then tries to write one strategy that handles both parities and gets stuck because the odd case has no perfect matching. A second nominal route hunts for an invariant, which cannot exist because both players can change the count. The bifurcation point is the sentence after the reconnaissance: do you look for a formula or an invariant (zombie), or do you ask which numbers each player is even capable of hitting (ideal)? The parity alignment is the observation from which the pairing, the blocking and the induction all follow.

Layer 2: HOW

Block 1. **Cue.** "Choosing integers from 1 to n", each once, until all are gone, with the payoff indexed by turn number. **Why natural.** A sequence of distinct choices exhausting a finite set is a bijection, and naming it a permutation converts an unfamiliar game into a familiar object with a standard statistic, the fixed point count. Naming the object is not cosmetic; it is what lets you say "the two members of this pair are fixed together", a sentence you cannot even formulate while thinking about turns and choices. **Link to the next block.** With the object named, ask which player controls which part of the statistic.

Block 2. **Cue.** Alice takes odd turns, Bob even ones, and a hit at k can happen only on turn k. **Why natural.** The condition $\pi(k) = k$ ties the turn parity to the number parity. So the two players have disjoint scoring opportunities: Alice can only ever hit odd numbers, Bob only even ones. This is the structural asymmetry the setter built in, and everything Bob does exploits it. It also immediately suggests the pairing $\{2k-1, 2k\}$, since that is the unique matching for which the two members are hit on consecutive turns, one by each player. **Link to the next block.** If a pair can be arranged to score together or not at all, the total is automatically even.

Block 3. **Cue.** n even means a perfect matching of $\{1, \ldots, n\}$ exists and Bob moves second. **Why natural.** A second player can always answer inside a fixed perfect matching, since after each of his moves the taken set is a union of complete pairs. Choosing the matching $\{2k-1, 2k\}$ makes the alignment work: Alice hits $2k-1$ only on turn $2k-1$, and Bob's reply on turn $2k$ is then $2k$, a hit; conversely Bob hits $2k$ only if Alice has just hit $2k-1$. Note the "conversely" is a genuine step and not decoration: without it, Bob might hit $2k$ while $2k-1$ was missed, and the count would be odd.

Link to the next block. Even n is finished. For odd n the matching is imperfect, so the same idea fails and a new mechanism is needed.

Block 4. **Cue.** The failing line for odd n ends with Bob unable to reply, and with a count that is even rather than odd. **Why natural.** What Bob lacks in the odd case is not a way to score, but a way to stop scoring safely. So isolate that: can Bob guarantee that nothing more is ever scored? Yes, and by the cheapest possible device: on his turn s he takes the number $s+1$, which is the only number Alice could use to score on her next turn, and which cannot score for him. When $s+1$ is already gone, Alice cannot score anyway and he takes anything except s. It is worth repeating: a blocking move that also cannot help you is the perfect move, and here the alignment of parities supplies one for free. **Link to the next block.** With a permanent shutdown available, the whole odd case reduces to reaching an odd count while it is Bob's turn.

Block 5. **Cue.** Bob's only scoring move on turn 2 is the number 2, and it is available unless Alice took it. **Why natural.** Bob has to bank a hit at some point, and the earliest and cheapest is turn 2. If Alice opened with 1 she has already banked one for him, so he blocks at once; if she opened with 2 she has denied him the cheap hit but has also left the position looking exactly like a smaller game; and if she opened with anything else he takes 2 and is one ahead. **Link to the next block.** Now the count is odd on Bob's clock, and he needs only to preserve parity or shut the game down.

Block 6. **Cue.** Two branches do not close directly: Alice opening with 2, and the running sequence colliding with Alice's first number. **Why natural.** In both, the taken numbers form an initial block $\{1,\dots,m\}$ chosen during turns $1,\dots,m$, with Alice to move on turn $m+1$: the position is a shifted copy of the original game on $n-m$ numbers. That is precisely the situation strong induction was invented for. Check the parity of m each time: in case (ii) $m=2$, in the collision case $m=k$ is even, so the residual game has odd size and the banked count is even, and the induction hypothesis supplies the missing odd parity. **Link to the next block.** All branches close, so Bob wins for every odd n.

Layer 3: WHY

The mathematical truth underneath is that this is a game of scoring opportunities, not of counting. Each number k carries exactly one opportunity, at turn k, and that opportunity belongs to whichever player owns turn k. So the position is best pictured as n tokens, alternately coloured by owner, each of which can be cashed only at its own moment and only by its owner, and each of which can be destroyed by anyone who takes it early. Once you see the game that way, both mechanisms are forced: a pairing that couples one of Alice's tokens to the one of Bob's immediately after it (so they are cashed together), and a blocking move that destroys Alice's next token at no cost. Nothing else in the game matters.

The discernment mechanism is to ask, in any game with a parity payoff, three questions in order. Who can change the quantity at all? If both players can, no invariant exists and you need a strategy. Which player moves second, and does the ground set admit a matching that the second player can answer inside? If yes, the second player can control parity. And finally, is there a move that both denies the opponent and cannot help you? Such a move is a permanent shutdown, and a shutdown plus a single one-off adjustment is exactly what a parity target needs. Training this is a matter of running those three questions on every game problem you meet until they are reflex.

Two further remarks give the problem its shape. First, the last turn is the pivot: whoever owns it owns the final adjustment, which is why Bob's goal is forced to be the parity of n and why the even case is easy and the odd case needs induction. Second, this is a game, not a counting problem, and the distinction is worth stating plainly. If the permutation were chosen uniformly at random,

the number of fixed points would be asymptotically Poisson with mean 1 and the parity would be biased but not controllable; under adversarial play the parity is decided outright. Do not import intuition from the random model into a game.

Knowledge gaps, stated plainly, since this is an A6. First, combinatorial game theory vocabulary is not part of A-level: a strategy is a function assigning a legal move to every reachable position, not a single line of play, and a proof must verify the strategy against all replies. This is where the overwhelming majority of attempted solutions to this problem fail, by describing what happens if Alice cooperates. Second, induction on games: to induct you must exhibit a reachable position that is isomorphic to the initial position of a smaller instance, with the same player to move, and you must record the parity of what was banked before the reset. Third, permutations and fixed points as formal objects, including the fact that a sequence of distinct choices exhausting a finite set is a bijection. Install those three and this problem is a page of careful bookkeeping; skip them and it is unapproachable.

What the examiners nudged toward and how: they made Bob choose his goal in advance, which is the entire difficulty, since otherwise he could wait and see; they specified that the last turn is forced, which points at the parity of n; and they phrased the payoff as the parity of a set indexed by turn number, which is a nudge to identify the permutation and its fixed points.

Revolving concepts

Permutations and fixed points. Definitions: a permutation of $\{1, \dots, n\}$ is a bijection; k is a fixed point when $\pi(k) = k$; a permutation with no fixed points is a derangement. Orbiting facts: the number of derangements satisfies $D_n = nD_{n-1} + (-1)^n$ and also $D_n = (n-1)(D_{n-1} + D_{n-2})$, with solution $D_n = n! \sum_{i=0}^{n} \frac{(-1)^i}{i!}$; the number of permutations with exactly j fixed points is $\binom{n}{j} D_{n-j}$, so under the uniform distribution the fixed point count converges to a Poisson law of mean 1. See Book 1, field: Combinatorics, topic: derangements and fixed points.

Pairing (mirroring) strategies. Statement of the principle: if the ground set carries a perfect matching and the second player always answers with the partner of the first player's move, then after each of his moves the taken set is a union of complete pairs; hence any statistic that is constant on pairs is under his control. Orbiting examples: the second player forcing a draw by symmetry, Hackenbush and Nim style mirroring, and the classical result that in many combinatorial games a pairing strategy is the only known explicit strategy. The limitation to remember: pairing needs an even ground set, and the matching must be chosen so that the payoff statistic couples within a pair, which here forces the matching $\{2k-1, 2k\}$ rather than $\{2k, 2k+1\}$. See Book 1, field: Games and invariants, topic: pairing strategies.

Blocking and cost-free denial. Statement: a move that removes the opponent's only scoring option and cannot score for you is strictly dominant when your aim is to freeze the payoff. Here Bob's turn s move of taking $s+1$ denies Alice her unique chance on turn $s+1$ and can never be a hit for him. The general rule: when the payoff is a sum of one-shot opportunities, denial of the opponent's next opportunity is a complete defence provided you can repeat it every turn. See Book 1, field: Games and invariants, topic: blocking strategies and shutdown arguments.

Induction on game size, and the recursion generalised. The game recursion here is: from the initial position of the game on n numbers, Bob can reach, after m turns, a position that is a shifted copy of the initial position of the game on $n-m$ numbers with Alice to move, having banked a known parity $b(m)$. Writing $V(n)$ for the parity Bob can force, the recursion is $V(n) = b(m) + V(n-m)$ with m even, whose solution here is $V(n) \equiv n \pmod 2$. The general lesson is that a strategy proof for a game on n objects should always be arranged as a reduction to the same game on fewer objects with the same player to move, since only then is strong induction available; the residual parity $b(m)$ is the bookkeeping you must not lose. See Book 1, field: Games and invariants, topic: strategies by induction on size.

Parity arguments and when invariants do not exist. Statement: an invariant argument requires

that no legal move changes the quantity modulo the relevant number. Test it here: Alice taking 1 on turn 1 changes the count, and Bob taking 2 on turn 2 changes it again, so no invariant exists and no counting argument can decide the game. The habit worth installing is to run this test before hunting for an invariant, which takes ten seconds and saves twenty minutes. See Book 1, field: Games and invariants, topic: invariants versus strategies.

Circuit Bank

- **WHEN** a game builds a sequence of distinct choices that exhausts a finite set **DO** name the resulting bijection and rewrite the payoff as a statistic of that permutation **BECAUSE** because the standard statistics of permutations come with structure, and here the fixed point condition ties each number to one specific turn and therefore to one specific player
- **WHEN** a game has a parity payoff and both players can change the quantity **DO** stop looking for an invariant and build a strategy **BECAUSE** because an invariant exists only when no legal move changes the quantity, and ten seconds of testing settles it
- **WHEN** you move second and the ground set has even size **DO** fix a perfect matching whose pairs couple the payoff, and always answer with the partner **BECAUSE** because after each of your moves the taken set is a union of complete pairs, so any statistic constant on pairs is yours to control
- **WHEN** you want to stop your opponent scoring for the rest of the game **DO** on each of your turns take the item that is their only scoring option next turn, provided it cannot score for you **BECAUSE** because that is a permanent shutdown at zero cost, and it converts a parity target into a single one-off adjustment
- **WHEN** your strategy has an exceptional branch that you cannot close directly **DO** look for a position reachable in that branch which is a shifted copy of the whole game on fewer objects with the same player to move, and use strong induction **BECAUSE** because a game reduction plus a record of the parity banked so far closes the branch without inventing a new idea

Backtesting

1. First useful cue: alternation of turns next to a payoff indexed by turn number, which forces the parity alignment that Alice can hit only odd numbers and Bob only even ones.
2. Method the expression suggested: for even n, mirror inside the matching $\{2k-1, 2k\}$ so hits come in pairs; for odd n, bank one hit at 2 on turn 2, then shut the game down with the blocking rule, using strong induction for the two branches where the shutdown is not yet available.
3. Where you would hesitate, and the thought to reinforce: in the odd case, at the moment the running sequence collides with Alice's opening number. Reinforce the sentence: the taken numbers are an initial block $\{1, \ldots, m\}$ taken during turns $1, \ldots, m$, so the position is a shifted copy of a smaller game and induction applies.
4. The trap the question sets: describing a strategy only along the line where Alice cooperates. Bob must have a reply to every legal Alice move at every reachable position, and the two branches that look like edge cases (Alice opening with 2, and the collision at $s = k$) are exactly where the marks are.
5. What makes you recognise it faster next time: any game whose payoff is a sum of one-shot opportunities, each owned by a single player at a single moment, is decided by pairing and blocking; and the owner of the last turn decides which parity is forced.

6. The reusable decision rule: name the permutation, align turn parity with number parity, mirror when the ground set is even, and bank one plus block when it is odd.

Mirror drills

1. (drill) Show that for even n Alice can force zero hits, so that Bob's goal choice is forced. *Answer:* on each of her turns t, Alice takes $t+1$ if available, otherwise anything other than t; her turns are odd and n is even, so $t \leq n-1$ and at least two numbers remain. She never takes t, and Bob's turn $t+1$ can never be a hit because $t+1$ is gone. Every turn is one of hers or the one immediately after, so nothing is ever hit.
2. (drill) Play out $n = 3$ completely and check Bob's goal "odd" wins against all three openings. *Answer:* if Alice opens 1 (a hit), Bob takes 3 on turn 2, and Alice takes 2 on turn 3, total 1. If Alice opens 2, Bob takes 1, and Alice is forced to take 3 on turn 3, a hit, total 1. If Alice opens 3, Bob takes 2 on turn 2, a hit, and Alice is forced to take 1 on turn 3, total 1.
3. (drill) Can either player force the number of hits to be exactly zero? *Answer:* Bob never can, since Alice takes 1 on turn 1 and scores immediately. Alice can exactly when n is even, by drill 1; for odd n she cannot, because the theorem gives Bob an odd count, which is at least one. So the count itself is not controllable, only its parity, and only by the right player.

3.7 B1. Sliding coins and the staircase of holes

Consider an m-by-n grid of unit squares, indexed by (i, j) with $1 \leq i \leq m$ and $1 \leq j \leq n$. There are $(m-1)(n-1)$ coins, which are initially placed in the squares (i, j) with $1 \leq i \leq m-1$ and $1 \leq j \leq n-1$. If a coin occupies the square (i, j) with $i \leq m-1$ and $j \leq n-1$ and the squares $(i+1, j)$, $(i, j+1)$, and $(i+1, j+1)$ are unoccupied, then a legal move is to slide the coin from (i, j) to $(i+1, j+1)$. How many distinct configurations of coins can be reached starting from the initial configuration by a (possibly empty) sequence of legal moves?

Read it out loud

This is a reachability count on a state graph: a state is a set of occupied squares, a move is a local rewrite of that set, and you are asked how many states the initial state can reach. The fields wired in are combinatorics (lattice paths, binomial coefficients) and the invariant method (what does a move preserve, and is the preserved thing a complete description of the reachable states). Nothing about coins matters beyond their positions, because coins are indistinguishable and a configuration is a set.

What is really being tested is whether you count the right object. The coins are the wrong object: there are $(m-1)(n-1)$ of them, a big floppy set. The holes are the right object: there are exactly $mn-(m-1)(n-1) = m+n-1$ of them, the count never changes, and initially they form the last column together with the last row, an L shape. Read that L as a monotone lattice path from the top right corner $(1, n)$ down to (m, n) and then left to $(m, 1)$. The move needs three holes in the shape $(i, j+1), (i+1, j+1), (i+1, j)$, which is a corner of that path, and after the move the corner is replaced by $(i, j+1), (i, j), (i+1, j)$, the corner bent the other way.

Within thirty seconds you should have written: holes $= m+n-1$, initial holes form a path, a move flips one corner of the path. The rest is bookkeeping: the number of such paths is a binomial coefficient, and you must show every path is reached, which you do backwards.

Approaches, ranked worst to best

1. **Small cases and a guess.** *Cue:* the phrase "how many configurations" on a puzzle-shaped board. *Circuit:* tabulate, spot Pascal's triangle. (2×2 gives 2, 2×3 gives 3, 3×3 gives 6, 3×4 gives 10; the pattern $\binom{m+n-2}{m-1}$ is unmistakable, and a pattern earns the answer line and nothing else. Keep it only as a check on the real proof.)

2. **Track the coins.** *Cue:* every move is a diagonal slide of one coin. *Circuit:* monovariant on coin positions. (you can prove that each coin moves at most once and that the set of moved coins is closed under moving down and right, then count such sets. It is correct, but each of those statements is proved by looking at the holes anyway, so you pay for the same invariant twice.)

3. **Count the holes.** *Cue:* the complement has fixed size $m+n-1$ and starts as an L. *Circuit:* complement invariant plus a bijection. (the holes always form a monotone path; the map from reachable states to paths is injective by the invariant and surjective by undoing moves. Ideal: one invariant, one descent argument, one binomial coefficient.)

The ideal solution

A configuration is the set of occupied squares; write H for its complement, the set of holes. Every move keeps $|H| = mn-(m-1)(n-1) = m+n-1$.

Call H a *staircase* if $H = \{p_0, p_1, \ldots, p_{m+n-2}\}$ with $p_0 = (1, n)$, $p_{m+n-2} = (m, 1)$ and each difference $p_{s+1} - p_s$ equal to $(1, 0)$ (a step D) or $(0, -1)$ (a step L). A staircase is determined by its word in $\{D, L\}$, which has exactly $m-1$ letters D and $n-1$ letters L, and every such word gives a

staircase; so there are $\binom{m+n-2}{m-1}$ staircases. Along a staircase the first coordinate never decreases and the second never increases.

The initial hole set is the last column followed by the last row, the staircase with word $D^{m-1}L^{n-1}$.

Claim 3.7.1. *If H is a staircase and a legal move slides the coin at (i,j) to $(i+1,j+1)$, then the new hole set $H' = (H \setminus \{(i+1,j+1)\}) \cup \{(i,j)\}$ is a staircase, whose word is obtained from the word of H by replacing one adjacent pair DL by LD.*

Proof. The squares $(i,j+1)$, $(i+1,j+1)$, $(i+1,j)$ all lie in H. Let $p_s = (i,j+1)$. The next point p_{s+1} is $(i+1,j+1)$ or (i,j); if it were (i,j), every later point would have second coordinate at most j and every earlier point first coordinate at most i, so $(i+1,j+1)$ could not lie on the path at all. Hence $p_{s+1} = (i+1,j+1)$. Likewise p_{s+2} is $(i+2,j+1)$ or $(i+1,j)$; the first option would leave no room for $(i+1,j)$, whose second coordinate j is below everything before p_{s+1} and whose first coordinate $i+1$ is below everything after p_{s+2}. Hence $p_{s+2} = (i+1,j)$, so steps $s+1, s+2$ read DL. Replacing p_{s+1} by (i,j) turns them into LD, the endpoints are untouched because $0 < s+1 < m+n-2$, and the result is again a staircase. □

So every reachable hole set is a staircase, and distinct configurations have distinct hole sets: the reachable configurations inject into the $\binom{m+n-2}{m-1}$ staircases.

Claim 3.7.2. *Every staircase is the hole set of a reachable configuration.*

Proof. Induct on the number of inversions of the word, the number of pairs consisting of an L somewhere before a D. Zero inversions means the word $D^{m-1}L^{n-1}$, the initial configuration. Otherwise the word contains an adjacent LD: the path has $p_s = (i,j+1)$, $p_{s+1} = (i,j)$, $p_{s+2} = (i+1,j)$. Let H'' be the staircase with p_{s+1} replaced by $(i+1,j+1)$; its word has one inversion fewer, so by induction H'' is reachable. In that configuration the square (i,j) is occupied, it satisfies $i \le m-1$ and $j \le n-1$ because $(i+1,j)$ and $(i,j+1)$ are squares of the grid, and $(i+1,j)$, $(i,j+1)$, $(i+1,j+1)$ are holes. The slide from (i,j) to $(i+1,j+1)$ is therefore legal and produces exactly the hole set we started from. □

Reachable configurations and staircases are in bijection, so the number of reachable configurations is

$$\boxed{\binom{m+n-2}{m-1}}.$$

(When $m = 1$ or $n = 1$ there are no coins and the formula correctly gives 1.)

Layer 1: WHAT

Block 1: name the holes. They number $m+n-1$, the number is invariant, and the initial holes form an L, which is a monotone lattice path from $(1,n)$ to $(m,1)$.

Block 2: prove the invariant. A legal move needs three holes forming a corner of the path, and it flips that corner: DL becomes LD in the word.

Block 3: count. Words with $m-1$ letters D and $n-1$ letters L: $\binom{m+n-2}{m-1}$. This is an upper bound because the map from configurations to words is injective.

Block 4: surjectivity, backwards. Any word other than the initial one has an adjacent LD; un-flip it, apply induction on the inversion count, and check the forward move is legal.

Juxtaposition. The nominal route follows coins: which coins can move, in which order, and how far. You end up proving that each coin moves at most once and that moved coins form a down-right closed region, and both proofs secretly use the hole structure. The ideal route names the holes in the first line. The bifurcation point is the sentence in which you decide which set to describe: the $(m-1)(n-1)$ coins or the $m+n-1$ holes. The ideal solver picks the smaller, rigid set because invariants live in rigid objects.

Layer 2: HOW

Block 1. **Cue.** The move condition names three empty squares and one coin. Three quarters of the condition is about emptiness. **Why natural.** A configuration and its complement carry the same information, and the complement is tiny: $m+n-1$ against $(m-1)(n-1)$. Predicated on that, you describe the holes, and the initial holes are visibly an L, which is a lattice path. **Link to the next block.** Once holes are a path, the move must be read as an operation on the path.

Block 2. **Cue.** The three required holes $(i, j+1), (i+1, j+1), (i+1, j)$ are a down step followed by a left step. **Why natural.** On a monotone path the three squares can only be consecutive, so the move is exactly the local rewrite $DL \to LD$. This is the whole invariant: paths go to paths. **Link to the next block.** An injective map from states to words gives an upper bound; you now need the size of the target and surjectivity.

Block 3. **Cue.** Words with a fixed number of each letter. **Why natural.** Choosing the positions of the $m-1$ letters D among $m+n-2$ slots is a binomial coefficient by definition. **Link to the next block.** An upper bound is not an answer. Every word must be shown reachable, and forward simulation is clumsy, so go backwards.

Block 4. **Cue.** "Every word is reachable" plus a move that is an adjacent transposition. **Why natural.** Adjacent transpositions $LD \to DL$ sort any word to $D^{m-1}L^{n-1}$ and the inversion count strictly decreases, so induction on inversions terminates. The undo of a corner flip is a legal forward move because the three holes it needs are exactly the three squares of the corner, and (i, j) is occupied since it is not on the path. **Link to the next block.** Injective plus surjective is a bijection; write the boxed answer.

Layer 3: WHY

The mathematical truth beneath the route is that reachability questions decompose into two halves: an invariant (which states are possible) and a construction (which possible states are attained). A complete invariant is one where the two halves meet, and the count is then the count of invariant values. Here the invariant value is the staircase, the count of staircases is a binomial coefficient, and the construction is a sorting argument: adjacent transpositions of a two-letter word connect every arrangement to every other, with the inversion number as the monovariant that guarantees termination.

The second truth is the duality between a set and its complement. Every move is a swap of one square out of the occupied set and one into it; describing the swap on the side where the set is small and structured is always cheaper. The examiners nudged toward this by making three of the four squares in the move condition empty squares and by placing the coins so that the holes form the cleanest possible path.

The discernment mechanism: when you see "how many configurations can be reached", do not simulate. Ask what a move preserves, express the preserved thing as a standard combinatorial object (a path, a word, a subset), and only then ask whether every such object is reached. Train it by refusing, on every puzzle of this type, to write a single forward simulation before the invariant is on the page.

Revolving concepts

Monotone lattice paths. A path from one corner of an $a \times b$ box to the opposite corner using unit steps in two directions is a word with a letters of one kind and b of the other, so there are $\binom{a+b}{a}$ of them. The same word encodes a partition fitting inside the box (the region on one side of the path is

a Young diagram) and a subset of size a of $\{1, \ldots, a+b\}$. See Book 1, field: Combinatorics, topic: lattice paths.

The recursion, generalised. Writing $P(a,b)$ for the number of paths, the first step gives $P(a,b) = P(a-1,b) + P(a,b-1)$ with $P(a,0) = P(0,b) = 1$, Pascal's recursion, whose solution is the binomial coefficient. Every count of words by their first letter obeys a recursion of this shape; the general method is the two-variable generating function, here $\sum P(a,b)x^a y^b = 1/(1-x-y)$. Refining by area (the number of squares under the path) turns the recursion into $P_q(a,b) = P_q(a-1,b) + q^a P_q(a,b-1)$ and the solution into the Gaussian binomial coefficient; this is the natural next question once you know each coin moves at most once, because the area is the number of moves made.

Invariants and monovariants. An invariant restricts the reachable set; a monovariant (here the inversion number of the word) proves that a backward process terminates, hence that a forward process exists. Reachability counts on a state graph are always "invariant for the upper bound, construction for the lower bound". See Book 1, field: Games and invariants, topic: reachability and complete invariants.

Sorting by adjacent transpositions. Any arrangement of a multiset can be turned into any other by swapping adjacent letters, and the minimal number of swaps between two arrangements of a two-letter word is the difference of inversion counts. In B1 that says the longest legal play has exactly $(m-1)(n-1)$ moves, one per coin.

Circuit Bank

- **WHEN** you see "reachable configurations" of tokens with a local move **DO** write down the complement set, its size, and its initial shape before anything else **BECAUSE** the invariant lives in the smaller and more rigid of the two sets, and a move is a swap seen from either side
- **WHEN** you see a monotone lattice path with a local corner move **DO** encode the path as a two-letter word, the move as an adjacent transposition, and count words by a binomial coefficient **BECAUSE** adjacent transpositions connect all arrangements, and the inversion number is the monovariant that proves it
- **WHEN** you must show every object of some kind is reachable **DO** work backwards from the target with a potential that strictly decreases, and check that the undo step is a legal forward move **BECAUSE** forward simulation branches, backward descent terminates

Backtesting

1. First useful cue: three of the four squares named in the move condition are empty squares, and the empties number $m+n-1$.
2. Method the expression suggested: complement invariant, then lattice paths, then a backward descent for surjectivity.
3. Where you would hesitate, and the thought to reinforce: proving that the three holes are consecutive on the path. The thought is monotonicity: coordinates move one way only, so a square with a larger first coordinate cannot come earlier and a square with a smaller second coordinate cannot come later.
4. The trap the question sets: describing coins. You will try to prove "a coin moves at most once" from scratch and burn ten minutes proving a lemma the hole path gives for free.
5. What makes you recognise it faster next time: any move that requires an L-shaped cluster of empties is a corner flip of a boundary path.
6. The reusable decision rule: in a reachability count, name the invariant object as a standard combinatorial object, count that, then prove surjectivity backwards with a monovariant.

Mirror drills

1. (drill) n tokens sit on cells $1, \dots, n$ of a strip of $n + k$ cells; a token may move one cell to the right into an empty cell. How many configurations are reachable? *Answer:* every n-subset of the cells, $\binom{n+k}{n}$; the one-dimensional version of the same invariant (tokens keep their order) and the same backward descent (the rightmost token that is not yet in place can always be moved).
2. (drill) In B1, what is the length of the longest sequence of legal moves? *Answer:* $(m-1)(n-1)$; each move raises the inversion count of the hole word by exactly one, the maximum inversion count is $(m-1)(n-1)$, and the word $L^{n-1}D^{m-1}$ is reachable.
3. (drill) How many partitions have Young diagrams that fit inside an $a \times b$ box? *Answer:* $\binom{a+b}{a}$; the boundary of the diagram is a monotone lattice path across the box, and the correspondence is a bijection.

3.8 B2. Binary weight of the multiples of 2023

For each positive integer n, let $k(n)$ be the number of ones in the binary representation of $2023 \cdot n$. What is the minimum value of $k(n)$?

Read it out loud

You are asked for the smallest number of ones that any multiple of 2023 can have in binary. A number with k ones is a sum of k distinct powers of two, so the question is: what is the least k for which some sum of k distinct powers of two is divisible by 2023. That is a statement about the residues of $2, 4, 8, \ldots$ modulo 2023, and residues of powers are governed by multiplicative orders. The fields wired in are number theory (orders, Euler's theorem, the Chinese remainder theorem) and the arithmetic of binary representations; there is no combinatorics here, despite the counting language.

The first thing to write is the factorisation $2023 = 7 \cdot 17^2$. The second thing is to look at the two prime powers in binary: $7 = 111_2$ and $289 = 100100001_2$, both of weight three. That is a planted cue: each factor is, on its own, a weight-three number, so the only job on the construction side is to align exponents so that both divisibilities hold at once. On the obstruction side, weight one is out immediately (a power of two has no odd factor) and weight two means $2^c \equiv -1$ modulo 2023, hence modulo 7, and 2 has order 3 modulo 7, which is odd. The circuit that must fire: -1 is a power of g modulo a prime if and only if the order of g is even.

Within thirty seconds you should have: "$2023 = 7 \cdot 17^2$; weight 1 impossible; weight 2 needs $2^c \equiv -1 \pmod 7$, impossible since $2^c \in \{1, 2, 4\}$; so the answer is at least 3, and I now need $a \neq b$ with $2023 \mid 2^a + 2^b + 1$."

Approaches, ranked worst to best

1. **Search for a small multiple.** *Cue:* "minimum value" of a digit count. *Circuit:* compute $2023n$ for $n = 1, 2, 3, \ldots$ and count ones. (hopeless in the arena and silent about the lower bound; the smallest multiple of weight three is $2^{77} + 2^{16} + 1$, so a hand search never arrives.)

2. **Chase the residues of 2^c modulo 2023 directly.** *Cue:* binary digits are powers of two. *Circuit:* tabulate 2^c mod 2023 until two residues sum to 2022. (the period is $\operatorname{lcm}(3, 136) = 408$, so this is the search above in disguise.)

3. **Lower bound by an odd order, upper bound by aligning exponents with CRT.** *Cue:* $2023 = 7 \cdot 17^2$, and 7, 289 are both weight three in binary. *Circuit:* multiplicative order plus Chinese remainder theorem on the exponents. (mechanical: one obstruction modulo 7 rules out weight two, one congruence system on the exponents produces a weight-three multiple. Ideal.)

The official second solution finds the smaller witness $2^{77} + 2^{16} + 1$ from $2^{68} \equiv -1 \pmod{289}$ and $(2^8 - 1)^2 \equiv 0 \pmod{289}$; it is shorter on the page but needs an observation the CRT route does not, so it ranks as a variant of the ideal, not above it.

The ideal solution

The minimum is 3. Throughout, $2023 = 7 \cdot 17^2 = 7 \cdot 289$, and $k(n)$ is the number of distinct powers of two whose sum is $2023n$.

Lower bound. If $k(n) = 1$ then $2023n = 2^a$, impossible because 7 divides the left side and not the right. If $k(n) = 2$ then $2023n = 2^a + 2^b$ with $a > b \geq 0$, so 7 divides $2^b(2^{a-b} + 1)$, and as 7 is odd, 7 divides $2^c + 1$ with $c = a - b \geq 1$. But $2^3 = 8 \equiv 1 \pmod 7$, so $2^c \bmod 7$ is one of $1, 2, 4$ and never 6. Hence $k(n) \geq 3$ for every n.

Upper bound. It suffices to find distinct positive integers a, b with $2023 \mid 2^a + 2^b + 1$: then $n = (2^a + 2^b + 1)/2023$ is a positive integer and $2023n$ has exactly three ones in binary.

Modulo 7: $2^2 + 2^1 + 1 = 7$, and $2^3 \equiv 1 \pmod 7$, so $7 \mid 2^a + 2^b + 1$ whenever $a \equiv 2$ and $b \equiv 1 \pmod 3$.

Modulo 289: $2^8 + 2^5 + 1 = 256 + 32 + 1 = 289$, and $2^{272} \equiv 1 \pmod{289}$ by Euler's theorem because $\varphi(289) = 17 \cdot 16 = 272$ and $\gcd(2, 289) = 1$; so $289 \mid 2^a + 2^b + 1$ whenever $a \equiv 8$ and $b \equiv 5 \pmod{272}$.

Take $a = 8$ and $b = 277 = 272 + 5$. Then $a \equiv 2 \pmod 3$ and $b \equiv 1 \pmod 3$ (as $277 = 3 \cdot 92 + 1$), while $a \equiv 8$ and $b \equiv 5 \pmod{272}$. Both divisibilities hold, and $\gcd(7, 289) = 1$, so $2023 \mid 2^{277} + 2^8 + 1$. Thus $k(n) = 3$ for $n = (2^{277} + 2^8 + 1)/2023$, and the minimum value of $k(n)$ is

$$\boxed{3}.$$

A caveat worth clarifying. The choice $a = 8$, $b = 277$ came from solving $a \equiv 2 \pmod 3$, $a \equiv 8 \pmod{272}$ and $b \equiv 1 \pmod 3$, $b \equiv 5 \pmod{272}$; the Chinese remainder theorem guarantees solutions because $\gcd(3, 272) = 1$, and any solution pair with $a \neq b$ works. The smallest multiple of 2023 with three ones is in fact $2^{77} + 2^{16} + 1$, but the question asks for the minimum weight, not the minimum multiple, and you must not spend arena time on the latter.

Layer 1: WHAT

Block 1: factor $2023 = 7 \cdot 17^2$ and translate "$k(n) = k$" into "some sum of k distinct powers of two is divisible by 2023".

Block 2: kill weight one (odd factor) and weight two (it would force $2^c \equiv -1 \pmod 7$, but the order of 2 modulo 7 is 3, so -1 is not a power of 2 there).

Block 3: aim at weight three modulo each prime power separately, using the prime powers' own binary expansions as witnesses: $7 = 2^2 + 2^1 + 1$, $289 = 2^8 + 2^5 + 1$.

Block 4: convert each witness into congruence conditions on the exponents via the period of 2^c (period 3 modulo 7, period dividing 272 modulo 289), align them by CRT, write down $a = 8$, $b = 277$.

Juxtaposition. The zombie route multiplies 2023 by successive n and counts ones, or tabulates 2^c mod 2023. The ideal route never looks at a multiple of 2023 at all; it looks at the subgroup of $(\mathbb{Z}/2023)^\times$ generated by 2 and at its projections modulo 7 and modulo 289. The bifurcation point is Block 1: whether you read $k(n)$ as a property of the product $2023n$ or as a divisibility condition on a sum of powers of two.

Layer 2: HOW

Block 1. Cue. A digit count of a multiple of a fixed odd number. **Why natural.** A binary number with k ones is $2^{a_1} + \cdots + 2^{a_k}$ with distinct exponents, so "some multiple has weight k" is exactly "some such sum is $\equiv 0 \pmod{2023}$". Predicated on that, the problem lives in modular arithmetic, and the first act in modular arithmetic is to factor the modulus. **Link to the next block.** Small k can now be tested one at a time against the prime factors.

Block 2. Cue. Weight two means $2^a + 2^b \equiv 0$, that is $2^{a-b} \equiv -1$ modulo any odd prime factor. **Why natural.** In the cyclic group generated by 2 modulo p, the element -1 is the unique element of order 2, so it lies in the group if and only if the order of 2 is even. Modulo 7 the order is 3: the obstruction is a single line. You do not need 17 at all for the lower bound; one obstruction suffices. **Link to the next block.** The lower bound is 3, so the remaining work is a construction of weight exactly three.

Block 3. **Cue.** $7 = 111_2$ and $289 = 100100001_2$. **Why natural.** A weight-three multiple of 7 and a weight-three multiple of 289 are handed to you by the numbers themselves. Since $2^c \bmod m$ depends only on c modulo the order of 2, each witness is really a pair of residue classes for the exponents. **Link to the next block.** Two congruence systems on the exponents, with moduli 3 and 272 (or the true order 136), must be solved simultaneously.

Block 4. **Cue.** Congruences on a with coprime moduli, and the same for b. **Why natural.** The Chinese remainder theorem is precisely the statement that such systems are solvable; $a = 8$ already satisfies both of its conditions, and $b = 5 + 272 = 277$ satisfies both of its. Coprimality of 7 and 289 then upgrades the two divisibilities to divisibility by 2023. **Link to the next block.** Write n explicitly and box the answer.

Layer 3: WHY

The truth underneath is that the residues of 2^c modulo an odd m form a purely periodic sequence whose period is the multiplicative order of 2, so any statement about which sums of powers of two are divisible by m is a statement about a finite cyclic group. Weight one asks whether 0 is in the group (never). Weight two asks whether -1 is in it, and -1 is in a cyclic subgroup of $(\mathbb{Z}/p)^\times$ exactly when that subgroup has even order. Weight three asks whether some $2^a + 2^b \equiv -1$, and once the modulus is a product of coprime pieces this decouples into independent alignment problems on the exponents.

The examiners chose the modulus with care. The factor 7 is there to block weight two (order 3, odd); the factor 17 alone would not block it, since $2^4 \equiv -1 \pmod{17}$. The square 17^2 is there so that you must handle a prime power, either by Euler's $\varphi(289) = 272$ or by lifting the order from 17 to 289. And both prime powers have weight three in binary, which is the guiding light for the construction.

The discernment mechanism: whenever a question mentions digits of a multiple of a fixed number, replace the digits by the residues of the base's powers, factor the modulus, and think in terms of orders. Train it by asking, for every odd m you meet, what the order of 2 is and whether it is even; that single bit of information decides the weight-two question forever.

Revolving concepts

Multiplicative order. For $\gcd(g, m) = 1$, $\mathrm{ord}_m(g)$ is the least $d \geq 1$ with $g^d \equiv 1 \pmod{m}$; then $g^s \equiv g^t$ if and only if $s \equiv t \pmod{d}$, and $d \mid \varphi(m)$ (Euler). Here $\mathrm{ord}_7(2) = 3$, $\mathrm{ord}_{17}(2) = 8$, $\mathrm{ord}_{289}(2) = 136$. See Book 1, field: Number theory, topic: multiplicative order.

The recursion, generalised. The residues $r_c = 2^c \bmod m$ satisfy $r_{c+1} \equiv 2r_c$, an invertible map on a finite set, so the sequence is purely periodic from $c = 0$ with period $\mathrm{ord}_m(2)$. Any sequence defined by iterating an invertible map on a finite set is purely periodic; if the map is not invertible (for instance $2^c \bmod m$ with m even) the sequence is eventually periodic after a pre-period given by the largest power of 2 dividing m. That is the general solution method for "powers modulo m": strip the part of m sharing factors with the base, then work in the unit group.

The even-order lemma. For an odd prime p and $p \nmid g$: -1 is a power of g modulo p if and only if $\mathrm{ord}_p(g)$ is even. Proof: the subgroup generated by g is cyclic of order d and contains an element of order 2 if and only if $2 \mid d$; and -1 is the unique element of order 2 in $(\mathbb{Z}/p)^\times$ because $x^2 \equiv 1$ has only the roots ± 1 modulo a prime. The same holds modulo odd prime powers.

Orders modulo prime powers. If $d = \mathrm{ord}_p(g)$ for an odd prime p, then $\mathrm{ord}_{p^e}(g) = d \cdot p^{\max(0, e - v_p(g^d - 1))}$; this is the lifting-the-exponent lemma in disguise. For $g = 2$, $p = 17$: $2^8 - 1 = 255 = 15 \cdot 17$ has $v_{17} = 1$, so $\mathrm{ord}_{289}(2) = 8 \cdot 17 = 136$, and then $2^{68} \equiv -1 \pmod{289}$. That is where the smaller witness $2^{77} + 2^{16} + 1$ comes from. See Book 1, field: Number theory, topic: lifting the exponent.

Chinese remainder theorem, used twice. Once on the modulus ($2023 \mid x$ if and only if $7 \mid x$ and $289 \mid x$), once on the exponents (a mod 3 and a mod 272 can be prescribed independently). See Book 1, field: Number theory, topic: Chinese remainder theorem.

Binary repunits. Every odd m divides $2^{\mathrm{ord}_m(2)} - 1$, a string of ones, so the minimum weight of a multiple of m is at most $\mathrm{ord}_m(2)$; the exact minimum has no closed form and is a genuine research-level question in general, which is why the examiners chose a modulus where it is 3.

Circuit Bank

- **WHEN** you see the number of ones in the binary expansion of a multiple of m **DO** write "weight k multiple exists iff $2^{a_1} + \cdots + 2^{a_k} \equiv 0 \pmod{m}$ for distinct a_i", then factor m and reduce the exponents modulo $\mathrm{ord}_m(2)$ **BECAUSE** the residues of powers of two are periodic with period equal to the order
- **WHEN** you must decide whether weight two is possible **DO** check, for one odd prime factor p, whether $\mathrm{ord}_p(2)$ is even **BECAUSE** $2^c \equiv -1 \pmod{p}$ holds for some c exactly when the order is even, and one odd-order prime factor is a complete obstruction
- **WHEN** you must construct a multiple of a composite m with a prescribed shape **DO** solve modulo each prime power with the smallest witness visible (often the binary expansion of the prime power itself), then align the exponent classes by CRT **BECAUSE** exponent conditions modulo coprime periods are simultaneously solvable
- **WHEN** you see a prime power modulus p^e with p odd **DO** use $\varphi(p^e)$ for a safe period, or lift the order from p with $v_p(g^d - 1)$ when you need the exact one **BECAUSE** Euler's theorem is always available; the exact order is only needed for sharper witnesses

Backtesting

1. First useful cue: $2023 = 7 \cdot 17^2$, with 7 and 289 both of binary weight three.
2. Method the expression suggested: an order obstruction modulo 7 for the lower bound and CRT on exponents for the upper bound.
3. Where you would hesitate, and the thought to reinforce: whether 17 must be checked in the lower bound. It must not; a lower bound needs one obstruction, and 7 provides it. Reinforce: prove impossibility with the smallest modulus that works.
4. The trap the question sets: hunting for the smallest n. The minimum weight is 3 regardless of which witness you find, and the smallest witness has 78 binary digits.
5. What makes you recognise it faster next time: "few ones in a multiple" is a question about the subgroup generated by 2, and "weight two" is the single bit "is the order even".
6. The reusable decision rule: lower bounds from odd orders, upper bounds from CRT alignment of exponents, and never touch the multiples themselves.

Mirror drills

1. (drill) What is the minimum number of ones in the binary expansion of a positive multiple of 23? *Answer:* 3; $\mathrm{ord}_{23}(2) = 11$ is odd so weight two is impossible, and $2^6 + 2^2 + 1 = 69 = 3 \cdot 23$.
2. (drill) For which odd primes p does some multiple of p have exactly two ones in binary? *Answer:* exactly those with $\mathrm{ord}_p(2)$ even, since $2^a + 2^b \equiv 0$ forces $2^{a-b} \equiv -1$; so $3, 5, 17$ qualify and $7, 23, 31, 89$ do not ($\mathrm{ord}_{89}(2) = 11$ because $2^{11} - 1 = 2047 = 23 \cdot 89$).

3. (drill) Given $2^4 \equiv -1 \pmod{17}$, show $2^{68} \equiv -1 \pmod{289}$. *Answer:* $2^{68} = (17-1)^{17} = \sum_k \binom{17}{k} 17^k (-1)^{17-k}$; the $k = 0$ term is -1, the $k = 1$ term is $17 \cdot 17$, and every later term has a factor 17^2.

3.9 B3. Expected length of the longest zigzag subsequence

A sequence $y_1, y_2, \ldots, y_k$ of real numbers is called zigzag *if $k = 1$, or if $y_2 - y_1, y_3 - y_2, \ldots, y_k - y_{k-1}$ are nonzero and alternate in sign. Let $X_1, X_2, \ldots, X_n$ be chosen independently from the uniform distribution on $[0, 1]$. Let $a(X_1, X_2, \ldots, X_n)$ be the largest value of k for which there exists an increasing sequence of integers $i_1, i_2, \ldots, i_k$ such that $X_{i_1}, X_{i_2}, \ldots, X_{i_k}$ is zigzag. Find the expected value of $a(X_1, X_2, \ldots, X_n)$ for $n \geq 2$.*

Read it out loud

This is an expectation of a combinatorial statistic of a random ordering. Say that sentence and half the question dies. The words "uniform distribution on $[0, 1]$" are dressing: continuous independent samples are almost surely distinct, and every one of the $n!$ orderings of $X_1, \ldots, X_n$ is equally likely by symmetry. The quantity a depends only on which of the X_i is bigger than which, so a is a function of a uniformly random permutation. No integral over $[0, 1]^n$ will ever be needed, and if you find yourself setting one up you have misread the question.

The second structural feature: a is defined as a maximum over an exponentially large family of subsequences, and expectation does not commute with maxima. So the only viable plan is to destroy the maximum, that is, to prove a structure theorem replacing the global optimisation by a count of local features. Once a is a sum of indicators, linearity of expectation finishes the job in one line, and linearity of expectation does not care that the indicators are wildly dependent. That is the whole design of the problem: it is a linearity of expectation problem hiding behind an extremal problem.

The fields wired in: order statistics only to the extent of exchangeability, extremal combinatorics for the structure theorem, and elementary probability for the last step. What is really being tested is whether you convert "longest alternating subsequence" into "number of direction changes", which is the standard structure theorem for this statistic.

Within thirty seconds you should have written two things: "only the ordering matters, so this is a random permutation", and "a must be a local count, guess a window of three consecutive terms". The second guess is cheap to justify in advance: the answer is asked for general $n \geq 2$, so it is almost certainly affine in n, and an affine answer with a non-integer slope smells of a probability like $2/3$ attached to $n - 2$ interior positions.

Approaches, ranked worst to best

1. **Integrate over the cube.** *Cue:* "uniform distribution on $[0, 1]$". *Circuit:* when you see a continuous density, integrate. You would need to partition $[0, 1]^n$ into regions where a is constant, which is exactly the combinatorial problem you were trying to avoid, now carrying a Jacobian. It costs enormous ink and answers nothing. The cue is a decoy planted by the setter.

2. **Compute the full distribution of a.** *Cue:* "find the expected value". *Circuit:* expectation equals sum over values times probability. You count permutations of $\{1, \ldots, n\}$ with longest zigzag exactly k, then sum k times that count. Those counts are the Eulerian-flavoured numbers for alternating runs; they exist, but you are computing a whole distribution to extract one moment. Ranked low because it is strictly more work than the question asks for.

3. **Markov chain on the last direction.** *Cue:* sequential structure, one new sample at a time. *Circuit:* model the process, condition on the newest term. Track the state "direction of the current run" and update as each X_i arrives. This works and gives $2/3$ per step, but you must justify that the new term extends or breaks the run with the right probabilities, and the state space is unnecessary once you see the local characterisation.

4. **Increment recursion on n.** *Cue:* "for $n \geq 2$", answer requested as a function of n. *Circuit:* define a_n, study $a_n - a_{n-1}$. Prove $a_n - a_{n-1} \in \{0, 1\}$ and equals 0 exactly when the last three terms are monotone, then induct. This is genuinely good and it is essentially the official route. It still needs the structure theorem to know what a_n counts, and it spends an induction where none is needed.

5. **Runs, then local extrema, then linearity.** *Cue:* a maximum over subsequences that you cannot integrate. *Circuit:* convert a global extremal statistic into a count of local features, then apply linearity of expectation. Prove $a = N + 1$ where N is the number of maximal monotone runs, note that $N - 1$ counts interior local extrema, and finish with $\Pr(\text{local extremum}) = 2/3$. Best: one structure theorem, no induction, no conditioning, and the same two moves work on every "expected length of the longest something" problem.

The ideal solution

The answer is $\boxed{\dfrac{2n+2}{3}}$.

Since the X_i are independent with a continuous distribution, they are almost surely pairwise distinct, and all $n!$ relative orderings are equally likely. Work on that event throughout, so that every difference $X_{i+1} - X_i$ is nonzero.

For a finite sequence $x = (x_1, \ldots, x_k)$ of distinct reals with $k \geq 2$, let $\varepsilon(x) = (\varepsilon_1, \ldots, \varepsilon_{k-1})$ be its sign word, $\varepsilon_i = \operatorname{sgn}(x_{i+1} - x_i) \in \{+, -\}$, and let $N(x)$ be one plus the number of indices i with $\varepsilon_i \neq \varepsilon_{i+1}$. Thus $N(x)$ is the number of maximal monotone runs of x, and x is zigzag exactly when $N(x) = k - 1$, that is, when its sign word alternates at every position.

Lemma 3.9.1. *If y is obtained from x by deleting one term, then $N(y) \leq N(x)$.*

Proof. Deleting x_1 or x_k deletes an end letter of the sign word, and deleting an end letter cannot increase the number of adjacent unequal pairs. Now delete an interior term x_i with $1 < i < k$. The two differences $d_{i-1} = x_i - x_{i-1}$ and $d_i = x_{i+1} - x_i$ are replaced by the single difference $d_{i-1} + d_i = x_{i+1} - x_{i-1}$, which is nonzero and whose sign is the sign of d_{i-1} or of d_i (if the two agree, the sum has that common sign; if they disagree, the sum has the sign of whichever is larger in magnitude). So the new sign word is obtained from $\varepsilon(x)$ by deleting one letter. Deleting the letter in position p of a word w replaces the two adjacent pairs (w_{p-1}, w_p) and (w_p, w_{p+1}) by the single pair (w_{p-1}, w_{p+1}), and

$$[\, w_{p-1} \neq w_{p+1} \,] \leq [\, w_{p-1} \neq w_p \,] + [\, w_p \neq w_{p+1} \,],$$

so the number of adjacent unequal pairs cannot increase. Hence $N(y) \leq N(x)$. □

Proposition 3.9.2. *For distinct reals $x_1, \ldots, x_n$ with $n \geq 2$, the longest zigzag subsequence has length $a(x) = N(x) + 1$.*

Proof. Lower bound. Write the maximal runs of x as index blocks $[p_0, p_1], [p_1, p_2], \ldots, [p_{N-1}, p_N]$ with $p_0 = 1$ and $p_N = n$, consecutive blocks having opposite directions. The subsequence $x_{p_0}, x_{p_1}, \ldots, x_{p_N}$ has consecutive differences that are nonzero (the terms are distinct) and alternate in sign (each block is monotone in the direction opposite to its neighbours), so it is a zigzag of length $N + 1$, giving $a(x) \geq N + 1$.

Upper bound. Let $x_{i_1}, \ldots, x_{i_m}$ be any zigzag subsequence with $m \geq 2$. It is obtained from x by deleting terms one at a time, so the lemma gives $N(x_{i_1}, \ldots, x_{i_m}) \leq N(x)$. But a zigzag of length m has $N = m - 1$. Hence $m - 1 \leq N(x)$, that is, $m \leq N(x) + 1$. □

Now localise. By definition $N(x) = 1 + \#\{i : 2 \leq i \leq n - 1,\ \varepsilon_{i-1} \neq \varepsilon_i\}$, and $\varepsilon_{i-1} \neq \varepsilon_i$ says exactly that x_{i-1}, x_i, x_{i+1} is not monotone, that is, x_i is a strict local extremum. Therefore, writing Z_i

for the indicator that X_{i-1}, X_i, X_{i+1} is not monotone,

$$a(X_1, \dots, X_n) = N + 1 = 2 + \sum_{i=2}^{n-1} Z_i.$$

For each fixed i the three variables X_{i-1}, X_i, X_{i+1} are independent and identically distributed with a continuous law, so their six relative orderings are equally likely, and exactly two of them (increasing and decreasing) are monotone. Hence $\Pr(Z_i = 1) = 1 - \frac{2}{6} = \frac{2}{3}$.

The Z_i are certainly not independent, but expectation is linear regardless, so

$$\mathbf{E}\big[a(X_1, \dots, X_n)\big] = 2 + \sum_{i=2}^{n-1} \mathbf{E}[Z_i] = 2 + (n-2) \cdot \frac{2}{3} = \frac{2n+2}{3}.$$

For $n = 2$ the sum is empty and the formula returns 2, which is correct since any two distinct reals form a zigzag. ■

Layer 1: WHAT

Block 1. Only the relative order of the X_i matters, and all orderings are equally likely: replace the sample by a uniformly random permutation and record that the values are almost surely distinct.

Block 2. Guess the structure theorem: the longest zigzag is forced by the run decomposition of the whole sequence, $a = N + 1$.

Block 3. Prove the easy half by exhibiting the run endpoints as an explicit zigzag.

Block 4. Prove the hard half by the deletion lemma: passing to a subsequence never increases the number of runs, and a zigzag of length m has exactly $m - 1$ runs.

Block 5. Localise: a new run starts exactly at an interior local extremum, so $a = 2 + \sum_{i=2}^{n-1} Z_i$.

Block 6. Apply linearity of expectation with $\Pr(Z_i = 1) = 2/3$.

Juxtaposition. The nominal route reads "uniform on $[0, 1]$" and "expected value" and starts building the distribution of a, either by integration over the cube or by counting permutations with a given longest zigzag. It will produce, after a great deal of ink, the same $2/3$ slope, because that slope was never about the distribution of a. The bifurcation point is the sentence "expectation does not commute with a maximum, so I must remove the maximum before I take any expectation". Everything upstream of that sentence is shared; everything downstream is a different problem.

Layer 2: HOW

Block 1. **Cue.** Independent continuous samples and a statistic defined purely by comparisons. **Why natural.** a is invariant under any strictly increasing map applied to all the X_i, and the uniform distribution has no other role than to make ties impossible and orderings symmetric. A statistic that only sees comparisons must be studied on the symmetric group, not on the cube. **Link to the next block.** Once the object is a random permutation, the only tools left are combinatorial, which forces you to look for combinatorial structure in a.

Block 2. **Cue.** A maximum over all increasing index sequences, which is far too large to enumerate. **Why natural.** Whenever a quantity is defined as the largest sub-object with a local alternating property, the extremal object is built greedily from the ambient object, and the greedy construction always uses turning points. So conjecture that a is determined by the turning points of the whole sequence and verify on small cases: $(1, 2, 3)$ gives $N = 1$, $a = 2$; $(2, 1, 3)$ gives $N = 2$, $a = 3$. **Link to the next block.** The conjecture $a = N + 1$ splits cleanly into a construction and a bound, which is the standard shape of an extremal proof.

Block 3. **Cue.** "There exists an increasing sequence of integers" is an existence claim. **Why natural.** Existence claims are discharged by exhibiting one object. The run endpoints are the canonical candidate because consecutive runs point in opposite directions by maximality, which is precisely the alternation required. **Link to the next block.** The construction gives $a \geq N + 1$; the burden shifts to proving nothing longer exists.

Block 4. **Cue.** An upper bound on all subsequences at once. **Why natural.** You cannot inspect all subsequences, so find a quantity that is monotone under the operation "delete one term". The number of runs is such a quantity, and the proof is at the level of the sign word: deleting a term merges two adjacent differences into their sum, and the sum inherits one of the two signs, so the sign word loses exactly one letter. Deleting a letter from a word cannot create new adjacent disagreements, by the triangle inequality for the indicator of inequality. **Link to the next block.** With N monotone under deletion and a zigzag of length m having $N = m - 1$, the bound $m \leq N + 1$ is immediate, closing the structure theorem.

Block 5. **Cue.** N is defined by a condition on consecutive pairs of differences, that is, on windows of three consecutive terms. **Why natural.** A count of positions where a local pattern occurs is by definition a sum of indicators, one per position. This is the move that converts a global statistic into something linearity of expectation can eat. **Link to the next block.** Now every Z_i has the same law by exchangeability, so one three-variable computation settles all of them.

Block 6. **Cue.** A sum of indicators inside an expectation. **Why natural.** Linearity of expectation holds with no independence assumption whatsoever, which is exactly why the localisation was worth so much. The individual probability is a symmetry count: six orderings of three distinct exchangeable values, two of them monotone. **Link to the next block.** Sum and simplify: $2 + \frac{2}{3}(n - 2)$.

Layer 3: WHY

The mathematical truth underneath is that the longest alternating subsequence is not really an optimisation at all. The map "sequence to sign word" turns the problem into a question about words in a two letter alphabet, and under that map taking a subsequence becomes deleting letters and merging, an operation that can only reduce the number of alternations. So the maximum is achieved by the most obvious candidate, and the extremal value is a linear functional of the ambient word. Any statistic with that shape, a maximum which is secretly a count, is one localisation away from linearity of expectation.

The discernment mechanism has two parts. First, ask what the statistic actually depends on. Here it depends only on comparisons, so the continuous distribution is decoration; recognising decoration early is worth more marks than any calculation. Second, ask whether the maximum is attained by a canonical object. If yes, the maximum is not an optimisation but a formula, and formulas can be expanded into sums of local terms. If you cannot see the canonical object, test $n = 3$ and $n = 4$ by hand: two minutes of small cases would have suggested $a = N + 1$ to anyone.

To train it, take the three standard "expected value of a maximum" traps and force each one through the same pipeline: identify the invariance, find the structure theorem, localise, then use linearity. The reason this pipeline beats conditioning arguments is that dependence between the local indicators is free. You should feel physically reluctant to compute a distribution when only a mean is requested.

What the setters nudged with: they wrote "for $n \geq 2$", telling you the answer is a formula in n, and they defined zigzag through consecutive differences rather than through any global property, pointing you at differences and therefore at windows of three. They also chose the

uniform distribution on $[0,1]$ rather than saying "a uniformly random permutation", because the disguise is the only difficulty they had left to add.

Revolving concepts

Zigzag (alternating) sequence. A sequence whose consecutive differences are nonzero and alternate in sign. Equivalent formulation: its sign word is $+-+-\cdots$ or $-+-+\cdots$. The longest alternating subsequence of a permutation of $\{1,\dots,n\}$ has mean $(2n+2)/3$ and is concentrated around it; the enumeration of permutations by number of alternating runs is classical and connects to the tangent and secant numbers through Andre's theorem on alternating permutations, where the count of alternating permutations of length n has exponential generating function $\tan x + \sec x$. See Book 1, field: Combinatorics, topic: alternating permutations.

Runs and descents. For a sequence of distinct reals, a descent is an index i with $x_i > x_{i+1}$, and the maximal monotone runs are the blocks cut out by the sign word. The number of runs is 1 plus the number of sign changes; the number of descents is the number of minus letters. Descent statistics on permutations are governed by Eulerian numbers. See Book 1, field: Combinatorics, topic: descents and Eulerian numbers.

Linearity of expectation. For any random variables $Y_1,\dots,Y_m$ with finite means, $\mathbf{E}[\sum Y_i] = \sum \mathbf{E}[Y_i]$, with no independence required. This is the single highest-yield probabilistic tool on the Putnam, because it converts a hard joint distribution into m easy marginals. See Book 1, field: Probability, topic: linearity and indicators.

Exchangeability and the six orderings. If $Y_1,\dots,Y_k$ are independent and identically distributed with a continuous law, then all $k!$ relative orderings have probability $1/k!$ and ties have probability zero. For $k = 3$ this gives the constant $2/3$ used above. See Book 1, field: Probability, topic: exchangeability and order statistics.

Recursion, generalised. The increment form of this problem is a first order linear recursion with constant increment: $\mathbf{E}[a_n] = \mathbf{E}[a_{n-1}] + c$ for $n \geq 3$ with $c = 2/3$ and initial value $\mathbf{E}[a_2] = 2$, whose solution is $\mathbf{E}[a_n] = \mathbf{E}[a_2] + c(n-2)$. The general statement worth installing: if a statistic A_n of the first n terms of an exchangeable sequence satisfies $A_n - A_{n-1} = \mathbf{1}[\text{a local pattern occupies the last } w \text{ terms}]$, then $\mathbf{E}[A_n] = \mathbf{E}[A_{w-1}] + p\,(n-w+1)$ where p is the probability of that pattern in a window of length w. Every "expected number of records", "expected number of local maxima", "expected number of ascents" question is this recursion with a different p (records give the non-constant increment $1/n$, which telescopes to a harmonic number instead). See Book 1, field: Sequences and recurrences, topic: first order recursions with constant increment.

Circuit Bank

- **WHEN** you see independent samples from a continuous distribution and a statistic defined only by comparisons **DO** replace the sample by a uniformly random permutation and record that ties have probability zero **BECAUSE** the statistic is invariant under increasing reparametrisation, so the density is decoration and only the ordering carries information
- **WHEN** you see an expectation of a maximum over exponentially many objects **DO** first prove a structure theorem that turns the maximum into a count of local features, and only then take expectations **BECAUSE** expectation does not commute with maxima, but it does commute with sums, so a maximum must be converted into a sum before any probability is used
- **WHEN** you see a count of positions at which a local pattern occurs **DO** write it as $\sum_i Z_i$ with Z_i an indicator, and use linearity of expectation **BECAUSE** linearity needs no independence, so heavy dependence between the Z_i costs you nothing

- **WHEN** you need an upper bound valid for all subsequences at once **DO** find a quantity that is monotone under deleting a single term, and induct on deletions **BECAUSE** a subsequence is a finite sequence of single deletions, so a one step monotonicity statement upgrades automatically to all subsequences

Backtesting

1. First useful cue: the statistic depends only on comparisons between the X_i, so the uniform distribution reduces to a uniformly random permutation with no ties.
2. Method the expression suggested: alternation is defined through consecutive differences, which points at sign words and therefore at maximal monotone runs.
3. Where you would hesitate, and the thought to reinforce: proving that no subsequence beats the run endpoints. Reinforce this: to bound every subsequence at once, find a quantity monotone under one deletion, because every subsequence is a chain of single deletions.
4. The trap the question sets: dressing a permutation statistic as a continuous probability question so that you set up an integral, or asking for a mean so that you compute a whole distribution.
5. What makes you recognise it faster next time: the phrase "the largest value of k for which there exists" inside an expectation is the signature of "structure theorem, then linearity of expectation".
6. The reusable decision rule: convert extremal statistics into local counts before touching probability, then let linearity absorb all the dependence.

Mirror drills

1. Let $X_1, \ldots, X_n$ be independent and uniform on $[0, 1]$. Find the expected number of indices i with $2 \leq i \leq n-1$ such that X_i is a strict local maximum. *Answer:* each interior index is a local maximum with probability $1/3$, so the expected number is $(n-2)/3$.
2. Find the expected number of ascents (indices i with $X_i < X_{i+1}$) among $X_1, \ldots, X_n$ independent and uniform on $[0, 1]$, and deduce the expected number of maximal monotone runs. *Answer:* $(n-1)/2$ ascents by symmetry, and the run count is 1 plus the number of sign changes, giving $1 + \frac{2}{3}(n-2)$, consistent with $\mathbf{E}[a] = (2n+2)/3$.
3. Show that a sequence of $n \geq 2$ distinct reals whose longest zigzag subsequence has length exactly 2 must be monotone. *Answer:* length 2 forces $N = 1$, that is, no sign change in the difference word, which is exactly monotonicity.

3.10 B4. Least time for a resetting piecewise parabola

For a nonnegative integer n and a strictly increasing sequence of real numbers $t_0, t_1, \ldots, t_n$, let $f(t)$ be the corresponding real-valued function defined for $t \geq t_0$ by the following properties:

(a) $f(t)$ is continuous for $t \geq t_0$, and is twice differentiable for all $t > t_0$ other than $t_1, \ldots, t_n$;

(b) $f(t_0) = 1/2$;

(c) $\lim_{t \to t_k^+} f'(t) = 0$ for $0 \leq k \leq n$;

(d) For $0 \leq k \leq n-1$, we have $f''(t) = k+1$ when $t_k < t < t_{k+1}$, and $f''(t) = n+1$ when $t > t_n$.

Considering all choices of n and $t_0, t_1, \ldots, t_n$ such that $t_k \geq t_{k-1} + 1$ for $1 \leq k \leq n$, what is the least possible value of T for which $f(t_0 + T) = 2023$?

Read it out loud

Read the conditions as mechanics and the problem becomes a gearbox. You are driving from height $1/2$ up to height 2023. Between consecutive marked times you are in gear k, meaning constant acceleration k. Condition (c) is the brutal one: at every marked time your speed is reset to zero. So each gear change buys you one extra unit of acceleration for the rest of the run, but it throws away all the speed you had built and, by the constraint $t_k \geq t_{k-1} + 1$, it costs you at least one unit of time. The question is how many times to change gear before making the final run.

Structurally, condition (c) makes the intervals independent. On (t_{k-1}, t_k) you start from rest with acceleration k, so the height gained is $\frac{k}{2}s_k^2$ where $s_k = t_k - t_{k-1}$ is the length of the interval, and nothing from the previous interval leaks in. Thus the total climb is a quadratic form in the gap vector, $\sum_k k s_k^2$, and the time is the linear form $\sum_k s_k$. This is now a clean finite dimensional optimisation problem with an integer parameter n: minimise a linear form subject to a quadratic form being fixed, with each of the first n variables bounded below by 1.

The fields: elementary integration to get the reduction, then convex optimisation to solve it, then a one variable inequality with an integer parameter to finish. What is really being tested is whether you understand that a convex function on a polytope is maximised at a vertex, and therefore that the optimum saturates all the gap constraints. Note that f is deliberately not differentiable at the t_k: the left derivative there is $ks_k > 0$ and the right derivative is 0. If you assume f' is continuous you are solving a different problem.

Within thirty seconds you should have written $s_k = t_k - t_{k-1}$, $f(t_k) - f(t_{k-1}) = \frac{k}{2}s_k^2$, and the reformulation: minimise $\sum_{k=1}^{n+1} s_k$ subject to $\sum_{k=1}^{n+1} k s_k^2 = 4045$ and $s_k \geq 1$ for $k \leq n$.

Approaches, ranked worst to best

1. **Global antiderivative.** *Cue:* "twice differentiable", "$f''(t) = k+1$". *Circuit:* integrate twice and impose the conditions. Writing a single formula for f on $[t_0, \infty)$ with all the constants of integration named, then imposing continuity and the derivative resets, is correct but drowns you in symbols before you ever reach the optimisation, and it hides the fact that the intervals do not interact. Ranked worst because it costs the most ink for the least structure.

2. **Lagrange multipliers, ignoring the floors.** *Cue:* "least possible value" plus a constraint. *Circuit:* constrained optimisation means stationary points of L. Minimising $\sum s_k$ subject to $\sum k s_k^2 = 4045$ gives $1 = 2\lambda k s_k$, so $s_k \propto 1/k$. This is a genuine stationary point, and it is the wrong one: on the positive part of the ellipsoid the linear form $\sum s_k$ is *maximised* there, not minimised. It also ignores $s_k \geq 1$ entirely. It gives a confidently wrong answer, which is worse than no answer.

3. **Numerical or heuristic search over n.** *Cue:* the integer parameter n is free. *Circuit:* when a parameter is a small integer, tabulate. Assume without proof that the optimum saturates $s_k = 1$, derive $T = n + \sqrt{\frac{4045}{n+1} - \frac{n}{2}}$, tabulate, and read off 29 at $n = 9$. This finds the answer fast and is exactly what you should do first for confidence, but the assumption is the entire content of the problem, so alone it scores near zero.

4. **Compactness plus exchange.** *Cue:* a minimum is asserted to exist. *Circuit:* show the minimum is attained, then perturb it. Prove the feasible set is compact so a minimiser exists, then show that if some $s_k > 1$ for $k \le n$ you can move that excess into the last gap and strictly increase the climb, contradicting minimality. This is correct and is the official route; it costs an existence argument and a monotonicity lemma about the intermediate value theorem before the perturbation can be run.

5. **Convexity: maximise the climb at fixed time.** *Cue:* a convex objective on a polytope of gap vectors. *Circuit:* a convex function on a compact convex set is maximised at an extreme point. Turn the question round: for a fixed total time T, how high can you get? The climb $\sum k s_k^2$ is convex, the gap vectors form a polytope, so the maximum sits at a vertex, and the vertices are exactly "all intermediate gaps at the floor 1". No existence argument, no perturbation, and the inequality that remains is a cubic that factorises. Best, because it explains why the optimum saturates instead of discovering it.

The ideal solution

The answer is $\boxed{29}$.

Reduction. Put $s_k = t_k - t_{k-1}$ for $1 \le k \le n$, so the hypothesis reads $s_k \ge 1$. Fix T and suppose $f(t_0 + T) = 2023$. First we may assume $t_0 + T \ge t_n$. Indeed, if not, let ν be the largest index with $t_\nu \le t_0 + T$ and discard the break points $t_{\nu+1}, \dots, t_n$: conditions (a) to (d) for the truncated data give exactly the same function on $[t_0, t_{\nu+1})$, since f'' equals $\nu + 1$ after t_ν in both descriptions, and the spacing hypothesis survives. So relabel and assume $t_0 + T \ge t_n$, and set $s_{n+1} = T - \sum_{k=1}^{n} s_k \ge 0$.

On (t_{k-1}, t_k) we have $f'' = k$ by (d), and $\lim_{t \to t_{k-1}^+} f'(t) = 0$ by (c), so $f'(t) = k(t - t_{k-1})$ there; by continuity of f,

$$f(t_k) - f(t_{k-1}) = \int_{t_{k-1}}^{t_k} k(t - t_{k-1})\, dt = \frac{k}{2} s_k^2 \qquad (1 \le k \le n),$$

and identically $f(t_0 + T) - f(t_n) = \frac{n+1}{2} s_{n+1}^2$. Summing and using $f(t_0) = 1/2$,

$$f(t_0 + T) = \frac{1}{2} + \frac{1}{2} \sum_{k=1}^{n+1} k\, s_k^2.$$

Hence $f(t_0 + T) = 2023$ is equivalent to

$$g(s) := \sum_{k=1}^{n+1} k\, s_k^2 = 4045, \qquad s_k \ge 1 \ (1 \le k \le n), \quad s_{n+1} \ge 0, \quad \sum_{k=1}^{n+1} s_k = T.$$

The concentration lemma. For an integer $\nu \ge 0$ and a real $T \ge \nu$ write

$$G(\nu, T) = \frac{\nu(\nu + 1)}{2} + (\nu + 1)(T - \nu)^2,$$

the value of g at the gap vector with ν gaps equal to 1 followed by a last gap of length $T - \nu$.

Lemma 3.10.1. *Let $n \ge 0$ be an integer and $T \ge n$. Over the set P of all $(s_1, \dots, s_{n+1})$ with $s_k \ge 1$ for $k \le n$, $s_{n+1} \ge 0$ and $\sum_{k=1}^{n+1} s_k = T$, the maximum of g is $G(0, T) = T^2$ if $n = 0$, and $\max\{G(n, T),\ G(n-1, T)\}$ if $n \ge 1$.*

Proof. P is a nonempty compact convex polytope and g is convex, being a nonnegative combination of the convex functions $s \mapsto s_k^2$. Every point of a polytope is a convex combination of its finitely many vertices, so by Jensen's inequality g attains its maximum over P at a vertex. Since P lies in the hyperplane $\sum s_k = T$ and its only other constraints are the $n+1$ coordinate bounds, a vertex is a feasible point at which at least n of those bounds are active. There are two possibilities.

If all n bounds $s_k \geq 1$ ($k \leq n$) are active, then $s_k = 1$ for $k \leq n$ and $s_{n+1} = T - n$, giving $g = \frac{n(n+1)}{2} + (n+1)(T-n)^2 = G(n, T)$.

Otherwise $s_{n+1} = 0$ and $s_k = 1$ for every $k \leq n$ except one index j, which forces $s_j = T - n + 1$. Then

$$g = \frac{n(n+1)}{2} - j + j(T-n+1)^2 = \frac{n(n+1)}{2} + j\big((T-n+1)^2 - 1\big).$$

Since $T \geq n$ we have $(T-n+1)^2 \geq 1$, so the bracket is nonnegative and the expression is largest at $j = n$, where it equals $\frac{n(n-1)}{2} + n(T-n+1)^2 = G(n-1, T)$. □

The bound. Suppose some admissible choice achieves $f(t_0 + T) = 2023$. By the reduction and the lemma there is an integer $\nu \geq 0$ with $\nu \leq T$ and

$$G(\nu, T) \geq 4045.$$

Assume for contradiction that $T < 29$. Then $\nu \leq T < 29$, so $0 \leq \nu \leq 28$, and $0 \leq T - \nu < 29 - \nu$, whence

$$\frac{\nu(\nu+1)}{2} + (\nu+1)(29-\nu)^2 > G(\nu, T) \geq 4045.$$

Multiplying by 2 and moving everything to one side, this says $R(\nu) > 0$ where

$$R(\nu) := \nu(\nu+1) + 2(\nu+1)(29-\nu)^2 - 8090 = 2\nu^3 - 113\nu^2 + 1567\nu - 6408.$$

But $R(9) = 1458 - 9153 + 14103 - 6408 = 0$, so $\nu - 9$ divides R, and

$$R(\nu) = (\nu - 9)\, Q(\nu), \qquad Q(\nu) = 2\nu^2 - 95\nu + 712.$$

The roots of Q are $\frac{95 \pm \sqrt{3329}}{4}$, and $57^2 = 3249 < 3329 < 3364 = 58^2$ gives $57 < \sqrt{3329} < 58$, so the smaller root lies in $(9.25,\ 9.5)$ and the larger in $(38,\ 38.25)$. Hence $Q(\nu) > 0$ for integers $\nu \leq 9$ and $Q(\nu) < 0$ for integers $10 \leq \nu \leq 38$. Therefore $R(\nu) \leq 0$ for every integer $0 \leq \nu \leq 28$: for $\nu \leq 9$ the first factor is nonpositive and the second positive, and for $10 \leq \nu \leq 28$ the first is positive and the second negative. This contradicts $R(\nu) > 0$. Hence $T \geq 29$.

The construction. Take $n = 9$ and $t_k = t_0 + k$ for $1 \leq k \leq 9$, so $s_1 = \cdots = s_9 = 1$, and run for a further $s_{10} = 20$. All spacing constraints hold with equality, and

$$g = \sum_{k=1}^{9} k \cdot 1 + 10 \cdot 20^2 = 45 + 4000 = 4045,$$

so $f(t_0 + 29) = \frac{1}{2} + \frac{1}{2} \cdot 4045 = 2023$ with $T = 9 + 20 = 29$. Therefore the least possible value of T is 29. ■

Layer 1: WHAT

Block 1. Read (c) as a reset: every interval starts from rest, so the intervals do not interact and the climb splits as a sum.

Block 2. Integrate once on each interval: $f' = k(t - t_{k-1})$, hence the climb on interval k is $\frac{k}{2}s_k^2$.

Block 3. Reformulate: minimise $T = \sum s_k$ subject to $\sum k s_k^2 = 4045$ with $s_k \geq 1$ for $k \leq n$, allowing n to vary.

Block 4. Invert the question: for a fixed time budget T, maximise the climb. The objective is convex on a polytope, so the maximum is at a vertex.

Block 5. Enumerate the vertices: all intermediate gaps at the floor, and one free gap. The best vertex is "all nine gaps at 1, everything else at the end", which is $G(\nu, T)$.

Block 6. Convert $G(\nu, T) \geq 4045$ with $T < 29$ into a cubic inequality in ν, factor out $\nu - 9$, and contradict.

Block 7. Exhibit $n = 9$, gaps $1, 1, \ldots, 1, 20$ and check the arithmetic exactly.

Juxtaposition. The nominal route writes a global formula for f, then optimises with Lagrange multipliers and gets $s_k \propto 1/k$, spreading the time over many small gaps. That is precisely backwards: it is the configuration that wastes the most time for a given climb. The bifurcation point is the sentence "the climb is a convex function of the gaps, so at fixed total time it is maximised at an extreme point, not at a stationary point". Convexity means the interior stationary point is a minimum of the climb, so any route that seeks a stationary point of the natural Lagrangian is walking away from the answer.

Layer 2: HOW

Block 1. **Cue.** Condition (c) prescribes the right hand limit of f' at every break point, not the value of f'. **Why natural.** A prescribed initial derivative on each interval means each interval is an independent initial value problem. The setter would have written "f is differentiable" if he had wanted the speed carried over; he wrote a one sided limit instead, which is the definition of a reset. Structural independence is the single most valuable thing you can extract from a piecewise definition, so hunt for it first. **Link to the next block.** Independent intervals means the total climb is a sum of per interval climbs, each depending on one gap only.

Block 2. **Cue.** Constant f'' on an interval with zero initial velocity. **Why natural.** This is the atomic case of "distance from rest under constant acceleration". Integrating twice gives $\frac{k}{2}s_k^2$ with no cross terms, and the absence of cross terms is exactly what makes the later optimisation tractable. **Link to the next block.** The climb is now a quadratic form and the time a linear form in the same variables, which is a standard optimisation shape.

Block 3. **Cue.** "Considering all choices of n and $t_0, \ldots, t_n$", and "least possible value". **Why natural.** The data $t_0, \ldots, t_n$ enters only through the gaps, and t_0 does not enter at all, so translate away t_0 and change variables to the gaps. Always reparametrise by the quantities the answer actually depends on; carrying redundant variables is how sign errors are born. **Link to the next block.** You now have a clean programme: minimise a linear form subject to a quadratic equality and box constraints, with a free integer n.

Block 4. **Cue.** A linear objective and a quadratic constraint, both with positive coefficients. **Why natural.** Minimising time at fixed climb and maximising climb at fixed time are the same problem read in two directions, and the second direction is the tractable one because the feasible set at fixed time is a polytope. That reversal is the whole move. A convex function on a compact convex set attains its maximum at an extreme point, since every point is a convex combination of vertices and Jensen pushes the value up to the largest vertex value. **Link to the next block.** Maximising over a polytope reduces an infinite problem to a finite list of candidate vertices.

Block 5. **Cue.** The constraints are exactly $n + 1$ coordinate bounds plus one hyperplane. **Why natural.** Vertices are where enough constraints are active to pin the point down, so at least n of the bounds must be tight. Reading off the two families costs one line each, and comparing them is a monotone comparison in j rather than an optimisation. **Link to the next block.** The second family collapses into the first with n replaced by $n - 1$, so the entire search is over the single family $G(\nu, T)$.

Block 6. **Cue.** A candidate optimum $T = 29$ obtained from a quick tabulation, and a family of inequalities indexed by an integer ν. **Why natural.** To prove a minimum you assume a better value and contradict. Substituting $T < 29$ makes the inequality polynomial in ν, and since $T = 29$ is achieved at $\nu = 9$, the polynomial must vanish at $\nu = 9$. Knowing the root in advance is what makes the factorisation instant rather than lucky. **Link to the next block.** The remaining quadratic has both roots outside the range where they could hurt, so the sign analysis is two lines.

Block 7. **Cue.** The problem asks for a value, so the bound is only half a solution. **Why natural.** An extremal answer is unproven until an example attains it. Exhibit the configuration and verify the arithmetic exactly, in integers: $45 + 4000 = 4045$. **Link to the next block.** Bound plus construction closes the problem.

Layer 3: WHY

Three mathematical truths make this route ideal.

First, resets decouple. A piecewise definition where each piece is given its own initial condition is a direct sum, not a coupled system, and the total is a sum of independent contributions. Recognising a direct sum turns a differential equation problem into an algebra problem.

Second, convexity decides where optima sit. The climb $\sum k s_k^2$ is convex in the gaps, so along any segment of feasible points it is maximised at an endpoint. This is why the optimum is "extreme": all intermediate gaps sit at the floor 1 and the entire remaining budget is dumped into the last gap, where the coefficient $n + 1$ is largest. The correct instinct, and it generalises far beyond this problem, is that maximising a convex objective means going to the corners, while minimising a convex objective means going to the interior stationary point. Half the marks on optimisation problems are decided by knowing which of those two situations you are in.

Third, the trade off is quantitative and small. Each extra gear costs one time unit and buys one extra unit of acceleration for the final run. With a fixed climb of 4045, more gears help until the fixed cost outweighs the gain, and the balance point is at $\nu = 9$. You can see it by differentiating $\nu \mapsto \nu + \sqrt{\frac{4045}{\nu+1} - \frac{\nu}{2}}$ informally, but you should not: since $T \geq \nu$ always, any $\nu \geq 30$ is immediately worse than 29, so only thirty integers survive and a finite check is both cheaper and safer.

The discernment mechanism: when you see a linear objective against a quadratic constraint with box constraints, ask which of the two is convex in the direction you are optimising, then decide corner versus interior before writing anything. If you cannot answer that, test the two extreme configurations by hand. Here, comparing "all time in one gap" with "time spread evenly" takes fifteen seconds and instantly reveals that concentration wins.

How the setter nudged: the constraint $t_k \geq t_{k-1} + 1$ exists purely to stop you from taking infinitely many gears with infinitesimal gaps, which tells you the floor is active at the optimum. The value 2023 was chosen so that $2 \cdot 2023 - 1 = 4045$ produces the exact integer answer $T = 29$ with $s_{10} = 20$; a clean answer is evidence that you have found the intended configuration.

Revolving concepts

Piecewise defined solutions of $f'' = c$. On an interval where f'' is the constant c and f' starts at 0 from the left endpoint t_{k-1}, we have $f'(t) = c(t - t_{k-1})$ and $f(t) - f(t_{k-1}) = \frac{c}{2}(t - t_{k-1})^2$. Continuity of f glues the pieces; f' may jump. See Book 1, field: Calculus and integrals, topic: piecewise initial value problems.

Convex functions and extreme points. A function g is convex if $g(\lambda x + (1-\lambda)y) \leq \lambda g(x) + (1-\lambda)g(y)$. A nonnegative combination of convex functions is convex, and $s \mapsto s^2$ is convex. Every point of a compact polytope P is a convex combination of its vertices, hence $\max_P g$ is attained at a vertex (Jensen). This is the finite dimensional case of Bauer's maximum principle. See Book 1, field: Inequalities, topic: convexity and extreme points.

Vertices of a polytope cut out by box constraints and one hyperplane. If $P = \{s \in \mathbb{R}^N : \sum s_k = T,\ s_k \geq \ell_k\}$, then the vertices of P are the points where at least $N-1$ of the bounds are active, that is, all coordinates but one sit at their floors. This is the standard "extreme points saturate all but one constraint" statement and is worth installing as a reflex. See Book 1, field: Inequalities, topic: linear programming vertices.

Lagrange multipliers and their failure mode. Stationary points of $\sum s_k - \lambda(\sum k s_k^2 - C)$ satisfy $s_k \propto 1/k$; this is a genuine critical point but it is the maximum of the linear form on the positive part of the ellipsoid, and it ignores inequality constraints entirely. Multipliers locate stationary points, never optima; the correct general statement needs the Karush, Kuhn and Tucker conditions with complementary slackness. See Book 1, field: Inequalities, topic: constrained optimisation and its traps.

Optimisation over an integer parameter. When a continuous optimisation is nested inside a choice of integer ν, solve the continuous problem for each ν to get $T(\nu)$, then bound the range of ν crudely (here $T \geq \nu$ kills $\nu \geq 30$ at once) and finish with a finite check or a factorised polynomial inequality. See Book 1, field: Inequalities, topic: mixed integer optimisation.

Recursion, generalised. The cumulative height is the telescoping sum $f(t_m) = f(t_0) + \sum_{k=1}^{m} \frac{k}{2} s_k^2$, that is, the first order recursion $h_m = h_{m-1} + \frac{m}{2} s_m^2$ with $h_0 = 1/2$. The general pattern to install: a recursion of the form $h_m = h_{m-1} + \varphi(m, s_m)$, in which the increment depends only on the current index and the current step size, always telescopes to $h_0 + \sum \varphi(k, s_k)$, and any optimisation of h subject to a constraint on $\sum s_k$ becomes a separable programme. Separability, in turn, is what allows a vertex argument, since separable convex objectives are convex. The same generalisation applies to every "greedy schedule with per stage cost" problem. See Book 1, field: Sequences and recurrences, topic: telescoping sums and separable objectives.

Circuit Bank

- **WHEN** you see a piecewise definition in which each piece is given its own initial condition **DO** treat the pieces as independent and write the total as a sum over pieces **BECAUSE** a prescribed one sided derivative at each break point is a reset, so no information crosses a break point and the problem is a direct sum

- **WHEN** you must minimise a linear form subject to a fixed quadratic form **DO** reverse the question and maximise the quadratic form at fixed linear form **BECAUSE** the level sets of the linear form are hyperplanes, so the reversed problem has a polytope as feasible set, and polytopes have finitely many vertices

- **WHEN** you are maximising a convex function on a compact convex set **DO** evaluate at the extreme points only **BECAUSE** every point is a convex combination of extreme points, so Jensen shows the maximum is attained at one of them; convex objectives are maximised at corners and minimised in the interior

- **WHEN** you must prove a specific number is the minimum of a family indexed by an integer **DO** assume a strictly better value, clear denominators to a polynomial inequality in the integer, and factor out the root you already know from the extremal configuration **BECAUSE** the extremal configuration tells you a root of the polynomial in advance, which turns an opaque factorisation into a division

- **WHEN** you have proved a sharp lower bound **DO** immediately construct the configuration that attains it and verify the arithmetic in integers **BECAUSE** an extremal problem is only half solved by a bound, and graders award the two halves separately

Backtesting

1. First useful cue: condition (c) prescribes the right hand limit of f' at each t_k, which is a reset and decouples the intervals.
2. Method the expression suggested: change variables to the gaps s_k, so that the climb is $\frac{1}{2}\sum k s_k^2$ and the time is $\sum s_k$.
3. Where you would hesitate, and the thought to reinforce: whether the optimum spreads the gaps or saturates them. Reinforce this: the climb is convex in the gaps, and convex objectives are maximised at corners, so the intermediate gaps sit at the floor 1 and the whole surplus goes into the last, highest geared interval.
4. The trap the question sets: applying Lagrange multipliers to get $s_k \propto 1/k$, which is the stationary point in the wrong direction and ignores the constraint $s_k \geq 1$; a second trap is assuming f' is continuous at the break points.
5. What makes you recognise it faster next time: a per interval acceleration that increases with the index, combined with a minimum spacing constraint, is a gearbox, and gearboxes are solved by choosing the number of gear changes and then running flat out in top gear.
6. The reusable decision rule: decouple at the resets, reparametrise by the gaps, reverse the optimisation so the feasible set is a polytope, then evaluate at vertices only.

Mirror drills

1. Maximise $s_1^2 + 2s_2^2 + 3s_3^2$ subject to $s_1 + s_2 + s_3 = 5$, $s_1 \geq 1$, $s_2 \geq 1$, $s_3 \geq 0$. *Answer:* check the vertices $(1,1,3)$, $(3,1,0)$, $(1,3,0)$, giving 30, 11, 19; the maximum is 30, attained by dumping the surplus into the highest coefficient.
2. In the problem above, replace 2023 by 5, so that the climb constraint becomes $\sum k s_k^2 = 9$. Find the least T. *Answer:* $T = 3$, attained both by $n = 0$ with a single gap $\sqrt{9} = 3$ and by $n = 1$ with gaps $(1, 2)$, since $1 + 2 \cdot 2^2 = 9$.
3. Show that without the constraints $s_k \geq 1$, the infimum of $\sum_{k=1}^{n+1} s_k$ subject to $\sum_{k=1}^{n+1} k s_k^2 = C$ tends to 0 as n grows. *Answer:* put everything in the last gap, $s_{n+1} = \sqrt{C/(n+1)}$ and all other gaps 0; the total is $\sqrt{C/(n+1)} \to 0$, which is exactly why the spacing hypothesis is present.

3.11 B5. When multiplication mod n is a square permutation

Determine which positive integers n have the following property: For all integers m that are relatively prime to n, there exists a permutation $\pi\colon \{1,2,\dots,n\} \to \{1,2,\dots,n\}$ such that $\pi(\pi(k)) \equiv mk \pmod{n}$ for all $k \in \{1,2,\dots,n\}$.

Read it out loud

Identify $\{1,2,\dots,n\}$ with $\mathbb{Z}/n\mathbb{Z}$ by sending n to 0. Since $\gcd(m,n) = 1$, the map $\sigma_m(x) = mx$ is a bijection of $\mathbb{Z}/n\mathbb{Z}$, and the displayed condition says precisely that $\pi \circ \pi = \sigma_m$ as permutations. So read the question out loud as: for which n is the multiplication by m permutation a square in the symmetric group, for every unit m?

That reading is the entire problem. "Square root of a permutation" is a solved question with a clean answer in terms of cycle type, so the problem splits into two independent halves: install the cycle type criterion, then compute the cycle type of multiplication by m on $\mathbb{Z}/n\mathbb{Z}$. The second half is number theory: the Chinese remainder theorem splits the ring, and on each prime power factor the cycle lengths are governed by the multiplicative order of m.

The trap is planted in plain sight. Nothing says π has to be multiplication by anything. If you look for $\pi(k) = ck$ with $c^2 \equiv m$, you are asking for every unit modulo n to be a square, which fails for all $n > 2$, and you will confidently produce the wrong answer set $\{1,2\}$. The permutation π is allowed to be arbitrary, and that freedom is what makes $n \equiv 2 \pmod 4$ work.

Fields: symmetric group theory (disjoint cycle decomposition, squares), and elementary number theory (Chinese remainder theorem, the structure of the unit group modulo a prime power). Within thirty seconds you should have written $\pi^2 = \sigma_m$, the word "cycle type", and the two cheap experiments: $m = -1$, which is always a unit, and m a primitive root modulo an odd prime factor.

Approaches, ranked worst to best

1. **Look for a linear π.** *Cue:* $\pi(\pi(k)) \equiv mk$ looks like composing multiplications. *Circuit:* when a functional equation is linear, guess a linear solution. Setting $\pi(k) = ck$ requires $c^2 \equiv m$ for every unit m, that is, every unit is a quadratic residue. Since $(\mathbb{Z}/n\mathbb{Z})^\times$ has even order for every $n > 2$, this fails for all $n > 2$ and yields the wrong answer. Worst, because it is fast, plausible and wrong.

2. **Brute force small n and conjecture.** *Cue:* "determine which positive integers". *Circuit:* when the answer is a set of integers, compute the first twenty. You find $1, 2, 6, 10, 14, 18$ and conjecture $n = 1$ or $n \equiv 2 \pmod 4$. Essential for confidence and for choosing what to prove, worth no marks alone, and dangerous if the pattern only starts later.

3. **Sign of the permutation.** *Cue:* a square must have sign $+1$. *Circuit:* the sign map is a homomorphism to $\{\pm 1\}$, so squares land in the kernel. Cheap and genuinely useful: for $4 \mid n$ and $m = -1$ the permutation is a product of $n/2 - 1$ transpositions, an odd number, so its sign is -1 and no square root exists. But the sign is too coarse in general: for $n = 9$, $m = 2$ the cycle type is a 6-cycle, a 2-cycle and a fixed point, sign $+1$, yet there is still no square root. Ranked here because it solves one case and silently fails on another.

4. **Ad hoc constructions case by case.** *Cue:* "there exists a permutation". *Circuit:* existence claims are met by construction. Building π by hand for each $n \equiv 2 \pmod 4$ without a general criterion is possible but you will be reinventing the pairing of equal cycles every time, and you will have no obstruction argument for the other n.

5. **Cycle type criterion plus the Chinese remainder theorem.** *Cue:* $\pi \circ \pi = \sigma_m$, a square in S_n; and $\gcd(m,n) = 1$, a unit. *Circuit:* squares in the symmetric group are characterised by cycle type; multiplication maps on $\mathbb{Z}/n\mathbb{Z}$ are analysed by splitting the ring. Prove once that a

permutation is a square exactly when every even cycle length occurs an even number of times, then read off the cycle type of σ_m in each of the three cases n odd, $n \equiv 2 \pmod 4$, $4 \mid n$. Best: one lemma, three short computations, and every case closed for the same structural reason.

The ideal solution

The property holds exactly for $\boxed{n = 1 \text{ and } n \equiv 2 \pmod 4}$.

Throughout, identify $\{1, \dots, n\}$ with $\mathbb{Z}/n\mathbb{Z}$ by sending n to 0, so that the required condition on π becomes $\pi \circ \pi = \sigma_m$, where $\sigma_m(x) = mx$ is a permutation of $\mathbb{Z}/n\mathbb{Z}$ because $\gcd(m, n) = 1$.

Lemma 3.11.1. *A permutation σ of a finite set is the square of some permutation if and only if, for every even $\ell \geq 2$, the number of ℓ-cycles in the disjoint cycle decomposition of σ is even.*

Proof. Suppose $\sigma = \tau^2$ and decompose τ into disjoint cycles. If c is a cycle of τ of odd length ℓ, then c^2 is again a single ℓ-cycle on the same support, of odd length. If c has even length ℓ, then c^2 splits into exactly two disjoint cycles of length $\ell/2$ on that support. Since disjoint cycles square independently, every cycle of σ of even length M arises from a cycle of τ of length $2M$, and each such cycle of τ contributes exactly two of them. Hence the number of M-cycles of σ is even.

Conversely, suppose the condition holds. Partition the cycles of σ into single cycles of odd length and pairs of cycles of equal even length; this is possible precisely because each even length occurs an even number of times. For an odd cycle c of length ℓ, put $\tau_c = c^{(\ell+1)/2}$, so that $\tau_c^2 = c^{\ell+1} = c$. For a pair of cycles $(a_1 a_2 \cdots a_\ell)$ and $(b_1 b_2 \cdots b_\ell)$ of the same length, put $\tau = (a_1 b_1 a_2 b_2 \cdots a_\ell b_\ell)$, whose square is $(a_1 a_2 \cdots a_\ell)(b_1 b_2 \cdots b_\ell)$. These pieces have pairwise disjoint supports, so their product τ satisfies $\tau^2 = \sigma$. □

Case $n = 1$. The only permutation is the identity, which is its own square. The property holds.

Case $n > 1$ odd. Write $n = p^e k$ with p an odd prime, $e \geq 1$ and $\gcd(p, k) = 1$; note k is odd. By the Chinese remainder theorem there is a ring isomorphism $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/p^e\mathbb{Z} \times \mathbb{Z}/k\mathbb{Z}$ under which multiplication by m becomes multiplication by the pair of residues of m. The group $(\mathbb{Z}/p^e\mathbb{Z})^\times$ is cyclic of order $L := p^{e-1}(p-1)$; choose g generating it and choose m with $m \equiv g \pmod{p^e}$ and $m \equiv 1 \pmod k$, which is a unit modulo n. Then $\sigma_m(u, v) = (gu, v)$.

Fix (u, v) and write $u = p^j u'$ with u' a unit and $0 \leq j \leq e$ (interpreting $j = e$ as $u = 0$). The cycle length of (u, v) is the least $t \geq 1$ with $g^t u \equiv u \pmod{p^e}$, that is, with $g^t \equiv 1 \pmod{p^{e-j}}$. Since the reduction $(\mathbb{Z}/p^e\mathbb{Z})^\times \to (\mathbb{Z}/p^{e-j}\mathbb{Z})^\times$ is surjective, g reduces to a generator, and the cycle length is $p^{e-j-1}(p-1)$ for $j < e$, and 1 for $j = e$. This equals L exactly when $j = 0$. So the elements lying in cycles of length L are precisely the $\varphi(p^e)k = Lk$ pairs with u a unit, and they form exactly k cycles. Now L is even because p is odd, and k is odd because n is odd. By the lemma, σ_m has no square root, so the property fails.

Case $n \equiv 2 \pmod 4$. Write $n = 2k$ with k odd, so $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/k\mathbb{Z}$. Any m coprime to n is odd, so $m \equiv 1 \pmod 2$ and $\sigma_m(\varepsilon, v) = (\varepsilon, mv)$. Hence both sets $\{0\} \times \mathbb{Z}/k\mathbb{Z}$ and $\{1\} \times \mathbb{Z}/k\mathbb{Z}$ are invariant, and the bijection $(0, v) \mapsto (1, v)$ conjugates the restriction of σ_m to the first onto its restriction to the second. Two conjugate permutations have the same cycle type, so every cycle length of σ_m occurs exactly twice as often as it does in the restriction, in particular an even number of times. By the lemma σ_m has a square root, for every such m, so the property holds.

Case $4 \mid n$. Take $m = n - 1$, which is coprime to n and satisfies $m \equiv -1 \pmod n$. The fixed points of σ_{-1} are the x with $2x \equiv 0 \pmod n$, namely $x = 0$ and $x = n/2$, and every other element is swapped with its negative, giving $(n-2)/2 = n/2 - 1$ transpositions. Since $4 \mid n$, the number $n/2$ is even, so $n/2 - 1$ is odd: there is an odd number of cycles of length 2. By the lemma σ_{-1} has no square root, so the property fails.

Every $n > 1$ is odd, congruent to 2 modulo 4, or divisible by 4, so the three cases together with $n = 1$ prove that the property holds exactly for $n = 1$ and $n \equiv 2 \pmod 4$. ■

Layer 1: WHAT

Block 1. Translate: $\pi(\pi(k)) \equiv mk$ means $\pi^2 = \sigma_m$ in the symmetric group of $\mathbb{Z}/n\mathbb{Z}$.

Block 2. Ask the right general question: which permutations are squares? Answer by cycle type, and prove it in both directions.

Block 3. Realise that the cycle type of σ_m is a number theory computation, and that the Chinese remainder theorem splits it into prime power pieces.

Block 4. Odd $n > 1$: choose m to be a primitive root on one odd prime power factor and trivial elsewhere, producing an odd number of cycles of even length $\varphi(p^e)$.

Block 5. $n \equiv 2 \pmod 4$: the factor $\mathbb{Z}/2\mathbb{Z}$ is inert because every unit is odd, so the whole permutation is two identical copies and every cycle length is doubled.

Block 6. $4 \mid n$: test $m = -1$, which gives two fixed points and an odd number of transpositions.

Juxtaposition. The nominal route treats $\pi(\pi(k)) \equiv mk$ as an equation to solve for a formula, tries $\pi(k) = ck$, discovers it needs every unit to be a quadratic residue, and answers $\{1, 2\}$. The bifurcation point is the sentence "π is only required to be a permutation, so this is a question about cycle type, not about arithmetic in $\mathbb{Z}/n\mathbb{Z}$". Once you refuse to look for a formula for π, the problem becomes routine.

Layer 2: HOW

Block 1. **Cue.** A permutation composed with itself, set equal to something explicit. **Why natural.** Composition of a permutation with itself is squaring in a group. As soon as an equation lives in a group you should stop treating it as a formula and start treating it as a question about that group's structure. The symmetric group's structure is its cycle type, so the equation $\pi^2 = \sigma$ must be a cycle type condition. **Link to the next block.** This turns "find π" into "is σ_m in the image of the squaring map", which is a classification question.

Block 2. **Cue.** Squaring a single cycle. **Why natural.** Everything about a permutation reduces to what happens on one cycle, because disjoint cycles commute and square independently. Squaring an ℓ-cycle is a step of size 2 around a cycle of length ℓ: if ℓ is odd the step generates the whole cycle, if ℓ is even it splits the cycle into the two parity classes. That single observation is the whole lemma, in both directions. **Link to the next block.** The forward direction gives an obstruction; the reverse direction gives a construction by interleaving, so the criterion is exact.

Block 3. **Cue.** $\gcd(m, n) = 1$ and multiplication modulo n. **Why natural.** Multiplication by a unit is an action of the cyclic group generated by m inside $(\mathbb{Z}/n\mathbb{Z})^\times$, and its cycles are the orbits. Orbits of a multiplicative action are computed factor by factor, so invoke the Chinese remainder theorem as a ring isomorphism, not merely as a congruence solver. **Link to the next block.** Now you can choose m freely on each factor, which is exactly the freedom needed to build a bad m.

Block 4. **Cue.** "For all integers m", so to disprove the property you need only one m. **Why natural.** Universally quantified hypotheses are attacked by choosing the extreme instance. The extreme instance of a multiplicative map is a generator, because it maximises orbit lengths and minimises the orbit count. Making m trivial on the complementary factor k keeps the count of long cycles equal to k, and k is odd for free. **Link to the next block.** An odd number of cycles of even length $\varphi(p^e)$ is exactly what the lemma forbids.

Block 5. **Cue.** $n \equiv 2 \pmod 4$ means $n = 2k$ with k odd, and every unit modulo 2 is 1. **Why natural.** The factor $\mathbb{Z}/2\mathbb{Z}$ cannot move, because the only unit there is the identity. So σ_m is literally the same permutation performed twice on two disjoint copies, which doubles every cycle count. Doubling is the cheapest possible way to satisfy the "even number of even cycles" condition, and it

works for every m at once, which is what the universal quantifier demands. **Link to the next block.** This is the positive case, and it needs no knowledge of the cycle structure on $\mathbb{Z}/k\mathbb{Z}$ whatsoever.

Block 6. **Cue.** $4 \mid n$ and the always available unit $m = -1$. **Why natural.** -1 is a unit modulo every n, and it is an involution, so its cycle type is trivial to compute: fixed points are the solutions of $2x \equiv 0$, of which there are exactly $\gcd(2, n)$, and everything else pairs up. Counting transpositions is then one subtraction, and the parity of that count is decided by n modulo 4. **Link to the next block.** The three cases are now exhaustive and the classification is complete.

Layer 3: WHY

The mathematical truth beneath the lemma is that squaring acts on cycles like a two step walk. On a cycle of odd length the step of size 2 is a generator of $\mathbb{Z}/\ell\mathbb{Z}$, so the cycle survives intact; on a cycle of even length the step of size 2 preserves parity, so the cycle splits into two halves of equal length. Therefore the squaring map on the symmetric group merges pairs of equal cycles and preserves odd cycles, and the image is exactly the set of permutations whose even cycle lengths pair up. Every question of the form "is this permutation a d-th power" is answered by the same analysis with d in place of 2, using $\gcd(d, \ell)$.

The second truth is that multiplication by a unit on $\mathbb{Z}/n\mathbb{Z}$ is a group action whose orbits you already know how to count: the orbit of x has size equal to the multiplicative order of m modulo $n/\gcd(x, n)$. The Chinese remainder theorem then makes the whole cycle type a product computation. That is why the odd case is beaten by a primitive root: cyclicity of $(\mathbb{Z}/p^e\mathbb{Z})^\times$ for odd p guarantees an element of maximal even order $\varphi(p^e)$, and the number of maximal orbits is then forced to be odd because everything else in sight is odd.

The discernment mechanism is a two part question you should ask on sight of any condition like $\pi \circ \pi =$ something. First: does the problem constrain π beyond being a permutation? Here it does not, so the answer depends only on conjugacy invariants, that is, on cycle type. Second: is the object whose cycle type I need a group action in disguise? Multiplication by m is, so its cycle type is a statement about orders, and orders factor through the Chinese remainder theorem. Train this by taking any permutation that is presented by a formula (multiplication, addition, a polynomial map, a shift) and computing its cycle type before doing anything else.

How the setter nudged: the phrase "for all integers m that are relatively prime to n" is a universal quantifier over the unit group, which is an invitation to pick your favourite unit, and the two favourites are -1, which exists always, and a primitive root, which exists on odd prime powers. The answer $n \equiv 2 \pmod 4$ then reads as "the factor of 2 must be present to duplicate everything, but not two factors of 2, since -1 then has an odd number of transpositions".

Revolving concepts

Disjoint cycle decomposition and conjugacy. Every permutation of a finite set is a product of disjoint cycles, uniquely up to order, and two permutations are conjugate in the symmetric group exactly when they have the same multiset of cycle lengths (the cycle type). Anything invariant under conjugation, such as being a square, must be a condition on cycle type. See Book 1, field: Abstract algebra, topic: symmetric groups and cycle type.

Squares in the symmetric group. σ is a square if and only if every even cycle length occurs an even number of times. More generally, σ is a d-th power if and only if for every ℓ, the cycles of length ℓ can be grouped so that each group of size g with $g \mid d$ and $\gcd(d/g, \ell) = 1$ merges into a cycle of length $g\ell$; the practical version to remember is the $d = 2$ case and the mechanism (a d step walk around a cycle of length ℓ splits it into $\gcd(d, \ell)$ cycles of length $\ell/\gcd(d, \ell)$). See Book 1, field: Abstract algebra, topic: powers and roots of permutations.

Chinese remainder theorem as a ring isomorphism. If $n = ab$ with $\gcd(a, b) = 1$ then $\mathbb{Z}/n\mathbb{Z} \cong \mathbb{Z}/a\mathbb{Z} \times \mathbb{Z}/b\mathbb{Z}$ as rings, hence $(\mathbb{Z}/n\mathbb{Z})^\times \cong (\mathbb{Z}/a\mathbb{Z})^\times \times (\mathbb{Z}/b\mathbb{Z})^\times$ as groups. Use it to prescribe m independently on each factor. See Book 1, field: Number theory, topic: Chinese remainder theorem.

Structure of the unit group modulo a prime power. For odd p, $(\mathbb{Z}/p^e\mathbb{Z})^\times$ is cyclic of order $p^{e-1}(p-1)$, so primitive roots exist; for $p = 2$ the group is trivial for $e = 1$, cyclic of order 2 for $e = 2$, and $\mathbb{Z}/2\mathbb{Z} \times \mathbb{Z}/2^{e-2}\mathbb{Z}$ for $e \geq 3$. The evenness of $p - 1$ for odd p is what supplies the even cycle length in the odd case. See Book 1, field: Number theory, topic: primitive roots.

Orbits of a multiplicative action. Under $x \mapsto mx$ on $\mathbb{Z}/n\mathbb{Z}$ with $\gcd(m, n) = 1$, the orbit of x has length equal to the multiplicative order of m modulo $n/\gcd(x, n)$. The set of x with a fixed gcd is a union of orbits, which is how the cycle count is obtained without listing anything. See Book 1, field: Number theory, topic: multiplicative order.

Recursion, generalised. The cycle lengths in the odd case follow the descending recursion $L_j = p^{e-j-1}(p-1)$, that is, $L_{j+1} = L_j/p$ for $0 \leq j \leq e-2$, terminating at $L_{e-1} = p - 1$ and then at a fixed point. This is the standard filtration recursion for a p-adic valuation: an element of valuation j behaves like an element of the smaller ring $\mathbb{Z}/p^{e-j}\mathbb{Z}$, so any orbit or order computation on $\mathbb{Z}/p^e\mathbb{Z}$ reduces by induction on e, with the answer for valuation j read off the answer for exponent $e - j$. The same reduction solves counting problems for solutions of congruences modulo prime powers, where lifting from p^e to p^{e+1} is Hensel's lemma. See Book 1, field: Number theory, topic: filtration by p-adic valuation.

Circuit Bank

- **WHEN** you see $\pi \circ \pi = \sigma$ with π unconstrained apart from being a permutation **DO** forget formulas and compute cycle types, then apply the criterion that every even cycle length must occur an even number of times **BECAUSE** being a square is invariant under conjugation, and cycle type is the complete conjugation invariant of a permutation
- **WHEN** you see a map given by multiplication by a unit modulo n **DO** split $\mathbb{Z}/n\mathbb{Z}$ by the Chinese remainder theorem and compute orbit lengths as multiplicative orders on each prime power factor **BECAUSE** multiplication by a unit is a group action, and orders and orbits are multiplicative across coprime factors
- **WHEN** you must disprove a statement quantified over all units m **DO** test $m = -1$ first, then a primitive root on one odd prime power factor **BECAUSE** -1 is a unit for every n and is an involution whose cycle type is immediate, while a primitive root maximises orbit length and therefore minimises and pins down the orbit count
- **WHEN** you need a permutation to have an even number of cycles of each even length **DO** look for an involution of the underlying set that commutes with the permutation and has no fixed points, so the set splits into two conjugate halves **BECAUSE** conjugate permutations have equal cycle types, so a free involution commuting with σ doubles every cycle count automatically

Backtesting

1. First useful cue: $\pi(\pi(k)) \equiv mk$ is the group equation $\pi^2 = \sigma_m$, and π carries no other constraint.
2. Method the expression suggested: classify squares in the symmetric group by cycle type, then compute the cycle type of multiplication by m using the Chinese remainder theorem.
3. Where you would hesitate, and the thought to reinforce: choosing the bad m in the odd case. Reinforce this: to break a universal statement about units, take a generator on one prime

power factor and the identity on the rest, so that the count of long cycles equals the size of the untouched factor.

4. The trap the question sets: guessing that π must itself be multiplication by some c, which reduces the problem to quadratic residues and gives the wrong answer set.
5. What makes you recognise it faster next time: any condition of the form "there is a permutation whose square is the given map" is a cycle type question, never a formula question.
6. The reusable decision rule: for equations inside a group, translate to structure theory of that group first; for permutations that structure is the cycle type.

Mirror drills

1. Show that the 4-cycle $(1\,2\,3\,4)$ is not the square of any permutation of $\{1,2,3,4\}$. *Answer:* it has exactly one cycle of the even length 4, and an odd count of even cycles is forbidden.
2. Show that every permutation of odd order is a square. *Answer:* all its cycle lengths are odd, so the criterion is vacuously satisfied; explicitly, if $\sigma^N = \mathrm{id}$ with N odd then $\sigma = (\sigma^{(N+1)/2})^2$.
3. Find an m coprime to 7 for which no permutation π of $\{1,\dots,7\}$ satisfies $\pi(\pi(k)) \equiv mk \pmod 7$. *Answer:* $m = 3$, a primitive root modulo 7, gives a single 6-cycle on the units together with the fixed point 0; one cycle of even length 6 is an odd count.

3.12 B6. Determinant of a matrix of representation counts

Let n be a positive integer. For i and j in $\{1, 2, \ldots, n\}$, let $s(i, j)$ be the number of pairs (a, b) of nonnegative integers satisfying $ai + bj = n$. Let S be the n-by-n matrix whose (i, j) entry is $s(i, j)$. For example, when $n = 5$, we have

$$S = \begin{bmatrix} 6 & 3 & 2 & 2 & 2 \\ 3 & 0 & 1 & 0 & 1 \\ 2 & 1 & 0 & 0 & 1 \\ 2 & 0 & 0 & 0 & 1 \\ 2 & 1 & 1 & 1 & 2 \end{bmatrix}.$$

Compute the determinant of S.

Read it out loud

Say the entry out loud before you say the word determinant: $s(i, j)$ counts the ways of writing n as a multiple of i plus a multiple of j. That is a convolution. Write $k = ai$ for the first part, so $n - k = bj$ for the second, and the count becomes

$$s(i, j) = \sum_{k=0}^{n} [\, i \mid k \,]\, [\, j \mid n - k \,],$$

where $[\, v \mid u \,]$ is 1 when v divides u and 0 otherwise, and every positive integer divides 0. A sum over an index k of a product of a term depending on (i, k) and a term depending on (k, j) is a matrix product. So S is not a mysterious table of counts; it is AB, where A and B are divisor incidence matrices.

Now the structural surprise that makes this a B6 and not a B2: the index k runs over $n + 1$ values, from 0 to n, while S is n-by-n. So A is n-by-$(n + 1)$ and B is $(n + 1)$-by-n, and there is no factorisation of $\det S$ as a product of two determinants. The tool for the determinant of a product of rectangular matrices is the Cauchy-Binet formula, which expresses $\det(AB)$ as a sum over maximal minors. Here that sum has exactly $n + 1$ terms, one for each column of A that you delete. The problem was built around this: one surplus column is the smallest interesting case of Cauchy-Binet.

The fields wired in: combinatorial counting to see the convolution, linear algebra (multilinearity of the determinant, Cauchy-Binet, sign of a column reversal), and multiplicative number theory (Möbius inversion and the partial sums of the Möbius function). What is really being tested is whether you can see a matrix as an operator built from a bilinear count, rather than as a table of numbers.

Knowledge gaps, stated plainly. Nothing in A-level Mathematics or Further Mathematics gives you Cauchy-Binet, the Möbius function, or Möbius inversion; those three must be installed before this problem is even approachable. If you do not have them, your only route is a long sequence of row and column operations on S itself, which is genuinely doable but is a bookkeeping marathon with two parity cases and a sign to track at every step. Install the three tools instead.

Within thirty seconds you should have written $s(i, j) = \sum_k [i \mid k][j \mid n - k]$, the sizes n by $n + 1$ and $n + 1$ by n, and the words "Cauchy-Binet, one surplus column".

Approaches, ranked worst to best

1. **Expand the determinant directly.** *Cue:* "compute the determinant". *Circuit:* a determinant is a signed sum over permutations. With entries as irregular as these, cofactor expansion produces nothing but noise. There is no pattern along any single row until you have already understood the structure, so this route never starts.

2. **Compute small cases and conjecture.** *Cue:* the answer must be a formula in n. *Circuit:* when the answer is a function of n, tabulate the first few values. You get $2, 2, -4, -4, 6, 6, -8, -8, \ldots$ for $n = 1, 2, 3, \ldots$, hence the conjecture $(-1)^{\lceil n/2 \rceil - 1} 2 \lceil n/2 \rceil$. Compulsory as a sanity check and as a target, worth close to nothing on its own, since "compute" on a B6 means "prove".

3. **Peel off rank one pieces.** *Cue:* the extreme terms $k = 0$ and $k = n$ of the convolution are special. *Circuit:* split a sum into a bulk term and boundary terms, then use the matrix determinant lemma. The $k = 0$ term contributes the rank one matrix with (i, j) entry $[\, j \mid n \,]$ and the $k = n$ term the rank one matrix with entry $[\, i \mid n \,]$, leaving a bulk product of an n-by-$(n-1)$ with an $(n-1)$-by-n matrix, which is singular. So S is a rank two update of a singular matrix and the determinant can be extracted. It works, but you end up computing the same minors that Cauchy-Binet would have handed you uniformly, with extra case analysis.

4. **Row and column reduction on S itself.** *Cue:* the entries for $i > n/2$ are visibly 0 or 1. *Circuit:* clear a matrix using its own structural rows. For $i > n/2$ the equation $ai + bj = n$ forces $a \in \{0, 1\}$, so those rows are almost indicator vectors; subtracting multiples of row $n-1$ and column $n-1$ makes most of the matrix an incidence pattern that can be cleared entry by entry, each clearance costing a sign. This is complete and elementary, and it is the official first solution, but it needs a separate endgame for n odd and n even and many sign decisions. Under time pressure it is where marks go missing.

5. **Factor as a product of divisor matrices, then Cauchy-Binet and Möbius inversion.** *Cue:* a bilinear count over a summation index, with $n + 1$ values of that index. *Circuit:* a convolution is a matrix product; a rectangular product needs Cauchy-Binet; divisor matrices are inverted by Möbius. Every step is forced by the previous one, each minor is a one line computation once you know the divisor matrix is unitriangular, and the final sum collapses using $\sum_{k \le n} C_k = 1$. Best, and it is the only route that explains why the answer depends on n only through $\lceil n/2 \rceil$.

The ideal solution

The answer is $\boxed{(-1)^{\lceil n/2 \rceil - 1}\, 2 \left\lceil \dfrac{n}{2} \right\rceil}$.

Step 1: factor S. For integers $u \ge 0$ and $v \ge 1$ write $[\, v \mid u \,] = 1$ if v divides u and 0 otherwise, so that $[\, v \mid 0 \,] = 1$ always. Let A be the n-by-$(n+1)$ matrix with rows indexed by $i \in \{1, \ldots, n\}$ and columns by $k \in \{0, 1, \ldots, n\}$, and B the $(n+1)$-by-n matrix with rows indexed by k and columns by j, defined by

$$A_{ik} = [\, i \mid k \,], \qquad B_{kj} = [\, j \mid n - k \,].$$

A pair (a, b) of nonnegative integers with $ai + bj = n$ is determined by $k := ai$, which satisfies $0 \le k \le n$, $i \mid k$ and $j \mid n - k$; conversely any such k gives the pair $(k/i, (n-k)/j)$. Hence

$$(AB)_{ij} = \sum_{k=0}^{n} [\, i \mid k \,][\, j \mid n - k \,] = s(i, j), \qquad \text{that is} \qquad S = AB.$$

Step 2: Cauchy-Binet. For an n-by-$(n+1)$ matrix A and an $(n+1)$-by-n matrix B, the Cauchy-Binet formula gives $\det(AB) = \sum_K \det(A_K) \det(B_K)$, the sum over all n-element subsets K of the common index set $\{0, \ldots, n\}$, where A_K keeps the columns in K and B_K the rows in K. Each such K omits exactly one index k, so writing $A^{(k)}$ for A with column k deleted and $B^{(k)}$ for B with row k deleted,

$$\det S = \sum_{k=0}^{n} \det\bigl(A^{(k)}\bigr) \det\bigl(B^{(k)}\bigr).$$

Step 3: express the minors of B through those of A. Since $(B^T)_{jk} = [\, j \mid n - k \,] = A_{j, n-k}$, the matrix B^T is A with its columns reversed. Deleting row k of B is deleting column k of B^T, which is

column $n-k$ of A; so $\left(B^{(k)}\right)^T$ is $A^{(n-k)}$ with its n columns in reversed order. Reversing n columns is a permutation of sign $(-1)^{n(n-1)/2}$, and transposing does not change a determinant, so

$$\det\left(B^{(k)}\right) = (-1)^{n(n-1)/2}\det\left(A^{(n-k)}\right).$$

Writing $D_k := \det(A^{(k)})$ we therefore have

$$\det S = (-1)^{n(n-1)/2}\sum_{k=0}^{n} D_k\, D_{n-k}.$$

Step 4: compute the maximal minors D_k. Let $v_0, v_1, \dots, v_n \in \mathbb{Z}^n$ be the columns of A and $e_1, \dots, e_n$ the standard basis. By definition $v_k = \sum_{d\mid k} e_d$ for $1 \le k \le n$, and $v_0 = e_1 + \cdots + e_n$.

First, $D_0 = \det(v_1, \dots, v_n)$ is the determinant of $([\, i \mid k\,])_{i,k=1}^{n}$, which is upper triangular because $i \mid k$ with $1 \le i, k \le n$ forces $i \le k$, and its diagonal entries are $[\, k \mid k\,] = 1$. Hence $D_0 = 1$ and $v_1, \dots, v_n$ is a basis of $\mathbb{Q}^n$.

Next, Möbius inversion applied to $v_k = \sum_{d\mid k} e_d$ gives $e_k = \sum_{d\mid k} \mu(d)\, v_{k/d}$ for $1 \le k \le n$. Summing over k and substituting $k = md$,

$$v_0 = \sum_{k=1}^{n} e_k = \sum_{k=1}^{n}\sum_{d\mid k}\mu(d)v_{k/d} = \sum_{m=1}^{n} C_m v_m, \qquad C_m := \sum_{1\le d\le n/m}\mu(d).$$

Fix k with $1 \le k \le n$. The matrix $A^{(k)}$ has columns $v_0, v_1, \dots, v_{k-1}, v_{k+1}, \dots, v_n$ in that order. Expanding v_0 by the displayed identity and using multilinearity, every term with $m \ne k$ repeats a column and dies, leaving

$$D_k = C_k \det(v_k, v_1, \dots, v_{k-1}, v_{k+1}, \dots, v_n) = (-1)^{k-1} C_k \det(v_1, \dots, v_n) = (-1)^{k-1} C_k,$$

the sign coming from moving v_k past the $k-1$ columns $v_1, \dots, v_{k-1}$. If we define $C_0 := -1$, then $D_0 = 1 = (-1)^{0-1} C_0$ as well, so

$$D_k = (-1)^{k-1} C_k \qquad (0 \le k \le n).$$

Step 5: evaluate the convolution. Since $(-1)^{k-1}(-1)^{n-k-1} = (-1)^n$,

$$\det S = (-1)^{n(n-1)/2}(-1)^n\,\Sigma, \qquad \Sigma := \sum_{k=0}^{n} C_k C_{n-k}.$$

Note that $C_m = \sum_{d\le\lfloor n/m\rfloor}\mu(d)$ depends only on $\lfloor n/m\rfloor$. In particular $C_m = \mu(1) = 1$ whenever $m > n/2$, and $C_m = \mu(1) + \mu(2) = 0$ whenever $n/3 < m \le n/2$; the second fact will be used only in the form $C_{n/2} = 0$ for even n.

The two extreme terms of Σ are $C_0 C_n$ and $C_n C_0$, each equal to -1, contributing -2. For $1 \le k \le n-1$ with $k \ne n/2$, exactly one of k and $n-k$ exceeds $n/2$, so $C_k C_{n-k} = C_{\min(k,n-k)}$; as k runs over that range each value $m < n/2$ occurs exactly twice, and if n is even the leftover term $k = n/2$ contributes $C_{n/2}^2 = 0$, which is also what extending the doubled sum up to $\lfloor n/2\rfloor$ contributes. Hence

$$\Sigma = -2 + 2\sum_{m=1}^{\lfloor n/2\rfloor} C_m.$$

That last sum is a standard collapse:

$$\sum_{m=1}^{n} C_m = \sum_{m=1}^{n}\sum_{1\le d\le n/m}\mu(d) = \sum_{\substack{d,m\ge 1\\ dm\le n}}\mu(d) = \sum_{f=1}^{n}\sum_{d\mid f}\mu(d) = \sum_{f=1}^{n}[\, f = 1\,] = 1,$$

using the defining property of the Möbius function. Since $C_m = 1$ for each of the $n - \lfloor n/2 \rfloor = \lceil n/2 \rceil$ values $m > \lfloor n/2 \rfloor$, we get $\sum_{m \le \lfloor n/2 \rfloor} C_m = 1 - \lceil n/2 \rceil$ and therefore

$$\Sigma = -2 + 2\left(1 - \left\lceil \frac{n}{2} \right\rceil\right) = -2\left\lceil \frac{n}{2} \right\rceil.$$

Step 6: the sign. Combining,

$$\det S = (-1)^{n(n-1)/2}(-1)^n \left(-2\left\lceil \frac{n}{2} \right\rceil\right) = (-1)^{n(n+1)/2+1}\, 2\left\lceil \frac{n}{2} \right\rceil.$$

Finally, $n(n+1)/2 + 1$ and $\lceil n/2 \rceil - 1$ have the same parity for every n. Both are periodic in n modulo 4: replacing n by $n+4$ changes $n(n+1)/2$ by $4n+10$, which is even, and changes $\lceil n/2 \rceil$ by 2. So it suffices to check $n = 0, 1, 2, 3$, where the two quantities are $(1, 2, 4, 7)$ and $(-1, 0, 0, 1)$, matching in parity in each case. Therefore

$$\det S = (-1)^{\lceil n/2 \rceil - 1}\, 2\left\lceil \frac{n}{2} \right\rceil.$$

As a check, $n = 1$ gives $S = [2]$ and the formula gives 2; $n = 5$ gives 6, which is the determinant of the matrix displayed in the statement. ■

Layer 1: WHAT

Block 1. Rewrite the entry as a convolution: $s(i, j) = \sum_{k=0}^{n} [i \mid k][j \mid n-k]$.

Block 2. Read the convolution as a matrix product $S = AB$ with divisor incidence matrices, and notice the index k has one more value than the matrix has rows.

Block 3. Apply Cauchy-Binet: $\det S$ is the sum over the $n+1$ ways of deleting one index.

Block 4. Identify B^T as A with reversed columns, so all the minors of B are minors of A up to one global sign.

Block 5. Compute the minors of A: the square block on columns 1 to n is unitriangular, and Möbius inversion writes the extra column v_0 in that basis with Mertens coefficients C_m.

Block 6. Evaluate $\sum_k C_k C_{n-k}$ using $C_m = 1$ for $m > n/2$ and the collapse $\sum_{m \le n} C_m = 1$.

Block 7. Reconcile the two signs into $(-1)^{\lceil n/2 \rceil - 1}$ by a periodicity check modulo 4.

Juxtaposition. The nominal route stares at the sample matrix for $n = 5$, notices that the last row and the first row look structured, and starts eliminating. It can be pushed through, and it is a legitimate solution, but every step is local and nothing tells you where the answer is heading until it arrives. The bifurcation point is the sentence "the entry is a sum over an intermediate index, so this matrix is a product". Once written, the rest of the solution is determined: rectangular product forces Cauchy-Binet, divisor matrix forces Möbius, and the sum over deletions forces a convolution of Mertens sums.

Layer 2: HOW

Block 1. Cue. The entry counts solutions of $ai + bj = n$ in nonnegative integers, a two part decomposition of n. **Why natural.** A count of ways to split a fixed total into a part from one family and a part from another family is by definition a convolution. Introducing the intermediate quantity $k = ai$ is not a substitution, it is naming the split point, which is the only structural feature the count has. **Link to the next block.** Convolutions are matrix products, so a matrix of convolution counts is a product of two matrices.

Block 2. Cue. The summation index runs over $\{0, 1, \ldots, n\}$, which has $n+1$ elements, while the matrix is n by n. **Why natural.** You must count the index set before choosing a tool. If k had run over n values you would get $\det S = \det A \det B$ and be finished in three lines; the surplus value $k = 0$, which exists because $a = 0$ is allowed, is precisely what makes the problem hard. Notice also that $k = 0$ and $k = n$ are the terms where one of the two parts is empty. **Link to the next block.** A rectangular factorisation is exactly the hypothesis of Cauchy-Binet.

Block 3. **Cue.** A product of an n-by-N and an N-by-n matrix, with $N = n + 1$. **Why natural.** Cauchy-Binet is the statement that the determinant of such a product is the sum over all n-element column selections of the product of the corresponding minors. With $N = n + 1$ there are only $n + 1$ selections, indexed by the deleted column, so the formula is a short sum rather than a combinatorial explosion. This is the only general tool that computes a determinant through a rectangular factorisation, so the cue determines the tool uniquely. **Link to the next block.** The problem is now to compute the $n + 1$ maximal minors of each factor.

Block 4. **Cue.** B_{kj} depends on $n - k$ where A_{jk} depends on k. **Why natural.** The two factors were built from the same divisor pattern, one indexed forwards and one backwards, because the equation $ai + bj = n$ is symmetric under swapping the roles of the two parts. Recognising B^T as A with reversed columns halves the work: you compute one family of minors, not two. The price is a single global sign, the sign of the reversal permutation. **Link to the next block.** Everything reduces to the numbers $D_k = \det A^{(k)}$.

Block 5. **Cue.** The divisor matrix restricted to columns 1 to n is triangular, and each column v_k is the sum of the basis vectors indexed by the divisors of k. **Why natural.** $i \mid k$ forces $i \leq k$, which is triangularity for free, so the columns $v_1, \ldots, v_n$ are a basis with determinant 1. The remaining column v_0 is the all ones vector, and the standard way to express a sum of basis vectors in terms of divisor sums is Möbius inversion, which is exactly the inverse of the operation "sum over divisors". Once $v_0 = \sum_m C_m v_m$, multilinearity kills every term but one in each minor, because all the other columns v_m are already present. **Link to the next block.** The minors are the Mertens sums up to sign, and the determinant becomes a convolution of those sums.

Block 6. **Cue.** A convolution $\sum_k C_k C_{n-k}$ where C_m depends only on $\lfloor n/m \rfloor$. **Why natural.** For m beyond the half way point $\lfloor n/m \rfloor = 1$ and $C_m = 1$, so in each pair $\{k, n - k\}$ one factor is 1 and the convolution degenerates into twice a partial sum. Partial sums of the Möbius function look hostile, but $\sum_{m \leq n} C_m$ is a double sum over $dm \leq n$ that reorders into $\sum_{f \leq n} \sum_{d \mid f} \mu(d)$, and the inner sum is the definition of μ as the inverse of the constant function 1. Never estimate a Mertens sum; reorder it. **Link to the next block.** The full sum is $-2\lceil n/2 \rceil$, and only the sign remains.

Block 7. **Cue.** Two sign factors, $(-1)^{n(n-1)/2}$ and $(-1)^n$, and a target expressed with a ceiling. **Why natural.** Both candidate exponents are periodic modulo 4 in n, so checking four values is a proof, not a spot check. Stating the periodicity first is what makes the four checks legitimate. **Link to the next block.** The two forms agree, so the answer can be written in the ceiling form.

Layer 3: WHY

The truth underneath is that this matrix is not a matrix of numbers but the Gram-like matrix of a bilinear pairing. The pairing is: given the divisor structure of the integers up to n, pair the indicator of multiples of i against the indicator of multiples of $n - k$ for j. The determinant of a Gram-like matrix built from two families of vectors in a space of dimension $n + 1$ is, by Cauchy-Binet, a sum over the ways of dropping one dimension. So the answer is a sum over the $n + 1$ possible "missing" split points, and the reason the final answer counts roughly $n/2$ is that only the split points beyond the half way mark contribute the trivial Mertens value 1.

The second truth is that the divisor matrix and the Möbius matrix are inverse to each other. The map "sum over divisors" is invertible on the integers up to n, and its inverse is the Möbius transform. That is why the minors are computable at all: any incidence matrix of a locally finite partial order with ones on the diagonal is unitriangular in a linear extension of the order, hence invertible over the integers, and its inverse is the Möbius function of the poset. Divisibility is the

poset here, and the classical Möbius function of number theory is precisely the Möbius function of that poset.

The discernment mechanism has three questions, and you should ask them in order on any determinant of a combinatorially defined matrix. First: is the entry a sum over an intermediate index? If yes, the matrix is a product, and you should factor before you compute. Second: do the factors have the same size as the matrix? If not, Cauchy-Binet, and count how many surplus columns there are, because that is the length of the resulting sum. Third: are the factors incidence matrices of an order? If yes, they are unitriangular, their determinants are 1, and their inverses are given by a Möbius function. Any one of those three questions answered wrongly turns this from a fifteen minute problem into a two hour one.

Train the discernment by refusing to compute any structured determinant before attempting a factorisation. Take the matrices $\gcd(i,j)$, $\lfloor n/\operatorname{lcm}(i,j)\rfloor$, the number of common divisors of i and j, and Smith's determinant; all of them factor through the divisor matrix, all of them are solved by the same three questions, and none of them should be attacked by elimination.

What the setter nudged with: the example matrix for $n = 5$ is printed in the statement, which is an invitation to compute small cases and conjecture, and the answer depends on n only through $\lceil n/2\rceil$, which is a strong hint that the split point $n/2$ matters. The count $s(i,j)$ was defined with nonnegative rather than positive integers, which is exactly what creates the surplus index values $k = 0$ and $k = n$ and therefore the need for Cauchy-Binet.

Revolving concepts

Cauchy-Binet formula. If A is m-by-N and B is N-by-m with $N \geq m$, then $\det(AB) = \sum_K \det(A_{\cdot,K})\det(B_{K,\cdot})$, summed over the m-element subsets K of $\{1,\dots,N\}$. For $N = m$ it reduces to $\det(AB) = \det A \det B$, and for $N < m$ it gives $\det(AB) = 0$. The special case $B = A^T$ gives $\det(AA^T) = \sum_K \det(A_{\cdot,K})^2$, the Gram determinant as a sum of squared minors. See Book 1, field: Matrices and determinants, topic: Cauchy-Binet.

Multilinearity and alternation of the determinant. The determinant is linear in each column and vanishes when two columns coincide; consequently, substituting a linear combination for one column leaves only the terms whose vector is not already among the other columns, and reordering columns multiplies by the sign of the permutation. Reversing m columns has sign $(-1)^{m(m-1)/2}$. See Book 1, field: Matrices and determinants, topic: multilinearity and column operations.

Möbius function and Möbius inversion. Define $\mu(1) = 1$, $\mu(d) = (-1)^r$ if d is a product of r distinct primes, and $\mu(d) = 0$ if d is divisible by a square. The defining property is $\sum_{d\mid f}\mu(d) = [f = 1]$. Inversion: if $g(k) = \sum_{d\mid k} h(d)$ for all k, then $h(k) = \sum_{d\mid k}\mu(d)g(k/d)$. In matrix language the divisor incidence matrix $Z = ([i \mid j])$ and the matrix $M = (\mu(j/i)[i \mid j])$ satisfy $MZ = ZM = I$. See Book 1, field: Number theory, topic: Möbius inversion.

Mertens sums. $C_m = \sum_{d\leq n/m}\mu(d)$ is a partial sum of the Möbius function, that is, the Mertens function evaluated at $\lfloor n/m\rfloor$. Individually these are erratic (the growth of the Mertens function is a deep open problem tied to the Riemann hypothesis), but the aggregate $\sum_{m\leq n} C_m$ collapses to 1 by reordering the double sum over $dm \leq n$. The moral: never evaluate a Mertens sum, only sum it. See Book 1, field: Number theory, topic: Dirichlet hyperbola reordering.

Dirichlet convolution and incidence algebras. The identity $s(i,j) = \sum_k [i \mid k][j \mid n-k]$ is an additive convolution, while $\sum_{d\mid f}\mu(d)$ is a multiplicative (Dirichlet) convolution; the problem sits at the crossing of the two, which is why it needs both a matrix product and a Möbius inversion. In general, the incidence matrix of any locally finite poset with ones on the diagonal is unitriangular in a linear extension, is invertible over the integers, and its inverse encodes the poset's Möbius function. See Book 1, field: Discrete structures and graphs, topic: incidence algebras and poset Möbius functions.

Recursion, generalised. Two recursions sit inside this solution. The first is the divisor recursion $v_k = \sum_{d\mid k} e_d$, whose inversion $e_k = \sum_{d\mid k}\mu(d)v_{k/d}$ is the general solution method for any relation

of the form $g = Zh$ with Z the divisor incidence operator: the solution is $h = Mg$ with M the Möbius operator, and the same statement holds verbatim for any locally finite poset with μ replaced by the poset's Möbius function. The second is the triangular recursion behind $\det D_0 = 1$: for a unitriangular system $x_k = \sum_{i<k} c_{ik} x_i + y_k$ the solution is obtained by forward substitution and the determinant is 1. The general lesson to install: whenever a family of vectors is indexed by a partial order and each vector is the sum of the basis vectors below it, that family is a basis, its transition matrix has determinant 1, and the inverse transition is the Möbius function of the order. See Book 1, field: Number theory, topic: divisor sum inversion.

Circuit Bank

- **WHEN** you see a matrix whose entry is a sum over an intermediate index of a product of two factors **DO** write the matrix as a product of two explicit matrices before attempting any determinant computation **BECAUSE** a determinant of a product is far more tractable than a determinant of a table, and factorisation exposes the structure that elimination only stumbles upon
- **WHEN** you have factored an n by n matrix as a product of an n by N and an N by n matrix with $N > n$ **DO** apply Cauchy-Binet and count the surplus, since the sum has one term per n-element subset of the N indices **BECAUSE** the determinant of a rectangular product is the sum of products of maximal minors, and with one surplus column the sum has only $n + 1$ terms
- **WHEN** you meet a matrix of divisibility indicators **DO** note it is unitriangular in the natural order, so its determinant is 1, and invert it with Möbius inversion **BECAUSE** $i \mid k$ forces $i \leq k$, so the divisibility incidence matrix is triangular with unit diagonal, and its inverse is the Möbius matrix
- **WHEN** you must evaluate a sum of partial sums of the Möbius function **DO** reorder the double sum by the product of the two indices, not by either index separately **BECAUSE** $\sum_{dm \leq n} \mu(d)$ groups by $f = dm$ into $\sum_{f \leq n} \sum_{d \mid f} \mu(d)$, which is 1; individual Mertens values are erratic but their aggregate telescopes
- **WHEN** your answer carries a sign of the form $(-1)^{q(n)}$ with q a quadratic in n **DO** prove the exponent is periodic modulo 4 and check four values **BECAUSE** a quadratic exponent changes by an even amount when n increases by 4, so four checks constitute a proof rather than a sample

Backtesting

1. First useful cue: the entry $s(i, j)$ counts splittings of n at an intermediate point $k = ai$, which is a convolution and therefore a matrix product.
2. Method the expression suggested: factor $S = AB$ with divisor incidence matrices of sizes n by $n + 1$ and $n + 1$ by n, then Cauchy-Binet.
3. Where you would hesitate, and the thought to reinforce: computing the maximal minors of A. Reinforce this: the columns indexed 1 to n are unitriangular hence a basis, so the only real work is expressing the surplus column in that basis, which is exactly what Möbius inversion does.
4. The trap the question sets: the surplus index. Allowing $a = 0$ and $b = 0$ makes the summation range $n + 1$ long, so the tempting factorisation into two square matrices is unavailable, and any solution that quietly assumes $\det S = \det A \det B$ is wrong.

5. What makes you recognise it faster next time: a combinatorially defined matrix whose entry counts a two part decomposition of a fixed total is always a product of two incidence matrices, and the length of the summation index decides whether you need Cauchy-Binet.
6. The reusable decision rule: factor first, count the index set second, invert incidence matrices with Möbius third, and reorder Mertens sums rather than evaluating them.

Mirror drills

1. Show that the n-by-n matrix with entries $[\,i \mid j\,]$ has determinant 1. *Answer:* a nonzero entry needs $i \mid j$ hence $i \le j$, so the matrix is upper triangular, and every diagonal entry is $[\,i \mid i\,] = 1$.
2. With $Z_{ij} = [\,i \mid j\,]$ and $M_{ij} = \mu(j/i)[\,i \mid j\,]$ for $1 \le i, j \le n$, prove $MZ = I$. *Answer:* the (i,k) entry is $\sum_{i \mid j \mid k} \mu(j/i)$, which after writing $j = id$ and $k = if$ becomes $\sum_{d \mid f} \mu(d) = [\,f = 1\,] = [\,i = k\,]$.
3. Let $A = \begin{pmatrix} 1 & 1 & 1 \\ 0 & 1 & 2 \end{pmatrix}$. Use Cauchy-Binet to compute $\det(AA^T)$ and check it directly. *Answer:* the three maximal minors are $1, 2, 1$, so the sum of their squares is 6; directly $AA^T = \begin{pmatrix} 3 & 3 \\ 3 & 5 \end{pmatrix}$ with determinant 6.

Chapter 4

Putnam 2022

4.1 A1. Lines meeting a logarithmic curve exactly once

Determine all ordered pairs of real numbers (a, b) such that the line $y = ax + b$ intersects the curve $y = \ln(1 + x^2)$ in exactly one point.

Read it out loud

Read the structure aloud before you write anything: this is a root-counting question for the equation $\ln(1 + x^2) = ax + b$, dressed up as a geometry question about a line and a curve. The moment you hear yourself say "count the solutions of $F(x) = 0$ where F depends on two parameters", the whole thing becomes single-variable calculus: monotonicity, turning points, and limits at $\pm\infty$. The fields wired into this problem are elementary real analysis (intermediate value theorem, monotonicity from the sign of the derivative), curve sketching, and a small amount of quadratic algebra to locate the turning points.

Now say the second sentence out loud: the parameter b appears only as an additive constant. That is the entire problem. If you write $g(x) = \ln(1 + x^2) - ax$, then the equation is $g(x) = b$, and counting intersections is counting preimages of a horizontal level under a fixed function g. For a one-variable function the number of preimages of a level is completely controlled by the shape of g: how many monotone pieces it has and what the values at the turning points are. So the two parameters are not symmetric in difficulty. You sweep b for free once you know the shape of g, and the only real work is classifying the shape of g as a varies.

The third thing to say aloud is that $\ln(1 + x^2)$ is even and its derivative $2x/(1 + x^2)$ is a bounded function, bounded by 1 in absolute value, with equality only at $x = \pm 1$. That bound is exactly the sort of guiding light the setters plant: the number 1 appearing as the extreme slope of the curve is why $|a| = 1$ is the threshold in the answer. A curve whose slope never exceeds 1 can be crossed at most once by a line of slope 1 or steeper.

Within thirty seconds you should have written down three things: $g(x) = \ln(1 + x^2) - ax$, the count is the number of solutions of $g(x) = b$, and $g'(x) = \dfrac{2x}{1 + x^2} - a$, whose sign is the sign of $-(ax^2 - 2x + a)$. A quadratic in the numerator means at most two turning points, which means at most three monotone pieces, which means the answer will be a case split on how many turning points there are: none, one degenerate one, or two.

Approaches, ranked worst to best

1. **Tangency algebra.** *Cue:* "exactly one point" sounds like a repeated root. *Circuit:* discriminant circuit from quadratics. This is the zombie route imported from GCSE and A-level: exactly one intersection means tangency, so impose $\ln(1 + x^2) = ax + b$ and $2x/(1 + x^2) = a$ simultaneously and eliminate. It ranks last because the premise is false here. Tangency is neither necessary nor sufficient for a single intersection with a transcendental curve: the line of slope 2 through the origin meets the curve once and is nowhere tangent to it, while the tangent line at $x = 1$ with $a = 1/2$ touches once and crosses twice more. The circuit "one solution equals repeated root" is a polynomial circuit and you must refuse to fire it outside polynomials.

2. **Two-parameter sweep of $F(x) = \ln(1 + x^2) - ax - b$.** *Cue:* two parameters in one function. *Circuit:* sign analysis of a general function. Differentiate F, find its stationary points, evaluate F there, and demand that exactly one of the values crosses zero. Correct but wasteful: you are carrying b through every derivative computation even though F' does not contain b. Every line of algebra costs ink and invites a skipped-calculation error. It ranks above the tangency fantasy because it at least ends in the right place.

3. **Convexity or graph-shape narration.** *Cue:* a logarithm of a quadratic, so an S-shaped graph. *Circuit:* inflection and asymptotic slope. Note that $\ln(1 + x^2)$ is concave outside $[-1, 1]$ and

convex inside, sketch, and read off the answer. This is fast and gives you the shape instantly, but a sketch is not a proof and the boundary case $|a| = 1$ is exactly where a sketch lies to you. Use it to know the answer in ten seconds, then upgrade it.

4. **Level sets of $g(x) = \ln(1 + x^2) - ax$.** *Cue:* b appears only additively. *Circuit:* separate the parameters, sweep the free one for free. Rewrite the equation as $g(x) = b$ and classify the monotone pieces of g. The count of solutions of $g(x) = b$ is then read off a one-line shape statement. This is the ideal route: it kills one parameter at zero cost, reduces the problem to a single quadratic $ax^2 - 2x + a = 0$, and makes the case split (0, 1, or 2 turning points) forced rather than invented. Add the symmetry reduction $a \mapsto -a$ from the evenness of $\ln(1 + x^2)$ and you halve the work again.

The ideal solution

Write $f(x) = \ln(1 + x^2)$ and, for a fixed real a,

$$g(x) = f(x) - ax, \qquad g'(x) = \frac{2x}{1 + x^2} - a = -\frac{ax^2 - 2x + a}{1 + x^2}.$$

The line $y = ax + b$ meets $y = f(x)$ in exactly one point if and only if the equation $g(x) = b$ has exactly one real solution.

Symmetry reduction. Since f is even, (x_0, y_0) lies on both $y = f(x)$ and $y = ax + b$ if and only if $(-x_0, y_0)$ lies on both $y = f(x)$ and $y = -ax + b$. Hence (a, b) is a solution pair if and only if $(-a, b)$ is. It therefore suffices to treat $a \geq 0$ and then reflect.

Asymptotics, used twice below. As $x \to +\infty$ we have $f(x)/x \to 0$, so for $a > 0$, $g(x) = x(f(x)/x - a) \to -\infty$ and $g(-x) = x(f(x)/x + a) \to +\infty$. For $a = 0$, $g = f \geq 0$ with $g(x) \to +\infty$ at both ends.

Case $a = 0$. Here $g = f$ is strictly decreasing on $(-\infty, 0]$ and strictly increasing on $[0, \infty)$, with minimum $f(0) = 0$. So $g(x) = b$ has no solution for $b < 0$, exactly one solution ($x = 0$) for $b = 0$, and exactly two for $b > 0$. The only admissible pair is $(a, b) = (0, 0)$.

Case $a \geq 1$. For all real x, the arithmetic mean geometric mean inequality gives $1 + x^2 \geq 2|x|$, hence $\left|\frac{2x}{1+x^2}\right| \leq 1$, with equality only at $x = \pm 1$. Therefore $g'(x) \leq 1 - a \leq 0$, and g' vanishes at no more than the single point $x = 1$ (which happens only when $a = 1$). A continuous function with $g' \leq 0$ everywhere and $g' = 0$ at finitely many points is strictly decreasing; combined with $g(-\infty) = +\infty$ and $g(+\infty) = -\infty$ and continuity, g is a strictly decreasing bijection from $\mathbb{R}$ onto $\mathbb{R}$. Hence $g(x) = b$ has exactly one solution for every real b. So every pair with $a \geq 1$ works, and by symmetry every pair with $|a| \geq 1$ works.

Case $0 < a < 1$. The stationary points of g are the roots of $ax^2 - 2x + a = 0$, namely

$$r_\pm = \frac{1 \pm \sqrt{1 - a^2}}{a}, \qquad 0 < r_- < 1 < r_+, \qquad r_- r_+ = 1.$$

Both roots are real and distinct because $0 < a < 1$, and both are positive because their sum $2/a$ and product 1 are positive. Since $ax^2 - 2x + a$ is an upward parabola, $g'(x) = -(ax^2 - 2x + a)/(1 + x^2)$ is negative on $(-\infty, r_-)$, positive on (r_-, r_+), and negative on (r_+, ∞). Thus g strictly decreases to a local minimum $g(r_-)$, strictly increases to a local maximum $g(r_+)$, then strictly decreases; and $g(-\infty) = +\infty$, $g(+\infty) = -\infty$.

On each of the three maximal monotone intervals, g is injective, and the intermediate value theorem identifies the exact range attained on each: $(g(r_-), +\infty)$, $[g(r_-), g(r_+)]$ and $(-\infty, g(r_+))$ respectively. Counting how many of these ranges contain a given level b, and noting $g(r_-) < g(r_+)$:

$$\#\{x : g(x) = b\} = \begin{cases} 1, & b < g(r_-) \text{ or } b > g(r_+), \\ 2, & b = g(r_-) \text{ or } b = g(r_+), \\ 3, & g(r_-) < b < g(r_+). \end{cases}$$

So for $0 < a < 1$ the admissible b are exactly $b < g(r_-)$ and $b > g(r_+)$. Finally, at a stationary point r we have $a(1+r^2) = 2r$, so $1 + r^2 = 2r/a$ and the two thresholds simplify:

$$g(r_\pm) = \ln\left(\frac{2r_\pm}{a}\right) - ar_\pm.$$

Answer. The line $y = ax + b$ meets $y = \ln(1+x^2)$ in exactly one point precisely when

$$\boxed{|a| \geq 1 \text{ (any } b), \quad \text{or} \quad (a,b) = (0,0), \quad \text{or} \quad 0 < |a| < 1 \text{ with } b < g(r_-) \text{ or } b > g(r_+),}$$

where, writing $\alpha = |a|$ and $r_\pm = \dfrac{1 \pm \sqrt{1-\alpha^2}}{\alpha}$, the thresholds are $g(r_\pm) = \ln(1 + r_\pm^2) - \alpha r_\pm = \ln(2r_\pm/\alpha) - \alpha r_\pm$.

Consistency check. As $\alpha \to 1^-$ we get $r_\pm \to 1$ and both thresholds tend to $\ln 2 - 1$, so the forbidden band closes up and the answer matches the $|a| = 1$ case continuously. As $\alpha \to 0^+$ we get $g(r_-) \to 0$ and $g(r_+) \to +\infty$, so the admissible set shrinks to $b < 0$, and in the limit $a = 0$ only $b = 0$ survives, matching the separate treatment of $a = 0$.

Layer 1: WHAT

Block 1. Rewrite the intersection condition as a root count for $\ln(1+x^2) - ax - b = 0$.

Block 2. Move b to the other side: the condition is $g(x) = b$ with $g(x) = \ln(1+x^2) - ax$. One parameter is now free.

Block 3. Use evenness of $\ln(1+x^2)$ to reduce to $a \geq 0$.

Block 4. Differentiate: $g'(x) = \frac{2x}{1+x^2} - a$, so stationary points solve the quadratic $ax^2 - 2x + a = 0$; the discriminant $4 - 4a^2$ splits the problem at $|a| = 1$.

Block 5. For $a \geq 1$, use the bound $|2x/(1+x^2)| \leq 1$ to conclude g is strictly decreasing, hence a bijection onto $\mathbb{R}$: exactly one solution for every b.

Block 6. For $0 < a < 1$, get two stationary points $r_\pm$ with $r_- r_+ = 1$, read off the down, up, down shape, and count preimages of b by monotone pieces: 1, 2 or 3.

Block 7. Handle $a = 0$ separately since the quadratic degenerates, and simplify the thresholds using $a(1+r^2) = 2r$.

Juxtapose this with the nominal route. The zombie writes $\ln(1+x^2) = ax + b$, cannot solve it, panics, and reaches for tangency: differentiate both sides, set $2x/(1+x^2) = a$, substitute back, and try to force a "repeated root" condition on a transcendental equation. That route produces the curve of tangency lines and then declares victory, missing every steep line and every line that clears the hump entirely, which is most of the answer set.

The bifurcation point is the single line where you decide what "exactly one point" means. The zombie translates it as "tangency". You translate it as "the level set $g^{-1}(b)$ is a singleton". Everything downstream is forced by that one translation.

Layer 2: HOW

Block 1. **Cue.** Two curves, "intersects in exactly one point", with one of the curves a line. **Why natural.** An intersection of $y = F(x)$ and $y = G(x)$ is by definition a solution of $F(x) = G(x)$. Nothing geometric is being asked; the picture is scaffolding. Convert to an equation immediately, because equations have derivatives and pictures do not. **Link to the next block.** Once you have an equation with two parameters, ask which parameter is cheap.

Block 2. **Cue.** The constant b appears additively and nowhere else. **Why natural.** A parameter that appears only as an additive constant is a level, and levels are swept for free: the solution count of $g(x) = b$ as b varies is completely determined by the shape of the single function g. Refusing to

separate them means differentiating a function that carries a redundant symbol through every line. **Link to the next block.** With b demoted to a level, the only remaining question is the shape of g, which is a derivative question.

Block 3. **Cue.** $\ln(1+x^2)$ is even. **Why natural.** Reflecting in the y-axis fixes the curve and maps the line $y = ax + b$ to $y = -ax + b$. A symmetry of the configuration that fixes one object and acts on the parameters is always worth using before the work starts, not after, because it halves the case analysis rather than merely checking it. **Link to the next block.** You may now assume $a \geq 0$, which makes the sign analysis of the quadratic clean.

Block 4. **Cue.** $g'(x) = \frac{2x}{1+x^2} - a$: a rational function set to zero. **Why natural.** Putting the derivative over the common denominator $1 + x^2 > 0$ makes the sign of g' the sign of a quadratic, and a quadratic has a discriminant. The discriminant $4 - 4a^2$ is where the case split comes from; you do not invent the threshold $|a| = 1$, the algebra hands it to you. **Link to the next block.** The discriminant sign gives three regimes: $|a| > 1$ (no stationary points), $|a| = 1$ (one, degenerate), $|a| < 1$ (two).

Block 5. **Cue.** $2x/(1+x^2)$ with $1 + x^2 \geq 2|x|$ staring at you. **Why natural.** The derivative of the curve is bounded by 1 in modulus, and that is the geometric content of the whole problem: a curve whose steepest tangent has slope 1 cannot be crossed twice by a line of slope at least 1, because between two crossings the mean value theorem would demand a tangent of slope exactly $a \geq 1$ at an interior point, and slope 1 is attained only at $x = 1$. Writing it as "$g' \leq 0$ with equality at finitely many points, hence g strictly decreasing" is the rigorous version and costs one clause. **Link to the next block.** Strict monotonicity plus the two infinite limits gives a bijection, which settles every b at once.

Block 6. **Cue.** Two stationary points with product $r_- r_+ = 1$ from $ax^2 - 2x + a = 0$. **Why natural.** Vieta reads off the product as $a/a = 1$ without solving, which tells you one root is inside $(0, 1)$ and one outside, matching the fact that the curve's steepest point is at $x = 1$. The down, up, down shape then forces the level count, and the count is a statement about monotone pieces, not about tangency. **Link to the next block.** Turning point values become the thresholds, so evaluate g at $r_\pm$ and simplify.

Block 7. **Cue.** $r_\pm$ satisfies $ar^2 - 2r + a = 0$, so $1 + r^2$ is expressible. **Why natural.** Never substitute a messy root into a logarithm if the defining equation of that root simplifies the argument first. From $a(1 + r^2) = 2r$ you get $\ln(1 + r^2) = \ln(2r/a)$, which removes a nested square root from inside a logarithm. This is the same laziness principle you use for symmetric-function problems: use the relation, not the root. **Link to the next block.** The final answer is now a clean inequality on b for each a.

Layer 3: WHY

The mathematical truth beneath this route is that the number of solutions of $g(x) = b$ equals the number of maximal monotone intervals of g whose range contains b. That statement, which is just the intermediate value theorem plus injectivity on monotone pieces, is the whole engine. It converts a two-parameter counting problem into: find the turning points, evaluate, compare with b. Nothing else is needed, and in particular no notion of tangency ever appears, because tangency is a statement about multiplicity of contact and the question asked about cardinality of contact. Those two agree for polynomials and disagree in general; A-level trains the polynomial reflex hard, and this problem is built to punish it.

The second truth is the boundedness of $x \mapsto 2x/(1+x^2)$ by 1. This is the derivative of a logarithm of a quadratic, and it is bounded because $\ln(1+x^2)$ grows like $2\ln|x|$, which is sublinear. Sublinear growth is exactly the reason no line of positive slope can be beaten forever, which is what forces at least one intersection for every $a \neq 0$. So the two ends of the analysis, "at least one" and "at most one", come from the same source: growth slower than linear.

The discernment mechanism, the thing to install so that you choose this route inside sixty seconds in the arena: count the parameters, and ask of each whether it is a level, a scale, or a genuine shape parameter. A level (additive constant) is free: quotient it out. A scale (multiplicative constant) is often removable by rescaling. Only a shape parameter, here a, deserves case analysis. Most multi-parameter Putnam problems in analysis collapse the moment you sort the parameters into those three bins. Train it by taking any past question with two parameters and asking, before writing anything, which parameter you can eliminate by translation, scaling, or symmetry.

What the examiners nudged with: they chose $\ln(1+x^2)$ rather than any old curve because its derivative has a clean extremal value of exactly 1, so the threshold $|a| = 1$ is memorable and the case split is forced; and they wrote "exactly one point" rather than "is tangent to" precisely to see whether you would fire the tangency circuit unbidden. The word "determine all ordered pairs" is also a command term: it tells you the answer is a description of a region in the (a, b) plane, so a full-marks script must present a classification with every case, not a single family.

Revolving concepts

Level sets and preimage counting. For continuous $g : \mathbb{R} \to \mathbb{R}$ with finitely many stationary points, split $\mathbb{R}$ into maximal intervals on which g is strictly monotone. On each, g is injective and its image is an interval determined by the endpoint limits. The count $\#g^{-1}(b)$ is then the number of those image intervals containing b, with boundary levels shared by adjacent pieces. See Book 1, field: Real analysis, topic: monotonicity and the intermediate value theorem.

Strict monotonicity from a nonstrict derivative bound. If g is differentiable with $g' \leq 0$ on an interval and g' vanishes only on a set with empty interior (in particular a finite set), then g is strictly decreasing. Proof: for $u < v$, $g(v) - g(u) = \int_u^v g'$, and a nonpositive integrand that is not identically zero on (u, v) has negative integral. Knowing this saves you from the common error of asserting only "nonincreasing" at $a = 1$ and losing the boundary case. See Book 1, field: Real analysis, topic: mean value theorem and monotonicity.

The bound $\frac{2|x|}{1+x^2} \leq 1$. Equivalent to $(|x|-1)^2 \geq 0$; equality only at $|x| = 1$. This is the arithmetic mean geometric mean inequality in its two-term form, and it recurs constantly: the same expression is sin of the tangent half-angle substitution, since $x = \tan(\theta/2)$ gives $2x/(1+x^2) = \sin\theta$. That substitution also parametrises this problem's answer: writing $a = \sin\theta$ with $\theta \in (0, \pi/2)$ gives $r_- = \tan(\theta/2)$ and $r_+ = \cot(\theta/2)$. See Book 1, field: Inequalities, topic: two-term AM GM; and field: Complex numbers, topic: Weierstrass substitution.

Vieta on the stationary quadratic. For $ax^2 - 2x + a = 0$ the roots satisfy $r_- + r_+ = 2/a$ and $r_- r_+ = 1$. Reciprocal pairs of roots are the signature of a palindromic quadratic, and any palindromic equation $ax^2 + cx + a = 0$ has roots $r, 1/r$. Wherever you meet a palindromic polynomial, divide by $x^{n/2}$ and substitute $t = x + 1/x$: that is the general reduction, and it halves the degree. See Book 1, field: Polynomials, topic: palindromic polynomials and reciprocal roots.

Recursion note. There is no recursion, induction, or telescoping in this problem, and it is worth saying so explicitly: the case analysis is a finite partition of parameter space, not an inductive structure. Do not manufacture induction where a partition suffices.

Asymptotic comparison $\ln(1+x^2)/x \to 0$. The general fact is that $(\ln x)^\alpha = o(x^\beta)$ for all $\alpha > 0$, $\beta > 0$. It is what guarantees that every line of nonzero slope eventually escapes the curve, hence at least one crossing. See Book 1, field: Real analysis, topic: growth hierarchies.

Circuit Bank

- **WHEN** you see a parameter appearing only as an additive constant in an equation $F(x, b) = 0$ **DO** isolate it as a level, $g(x) = b$, and study the shape of g once **BECAUSE** the solution count as the level varies is read off the monotone pieces of g, so one shape analysis answers infinitely many parameter values at zero extra cost
- **WHEN** you see "exactly one intersection point" for a non-polynomial curve **DO** count preimages under a level map, never impose tangency **BECAUSE** tangency is multiplicity of contact and the question asks for cardinality of contact; the two coincide only for polynomials, and assuming otherwise discards the entire steep-line family
- **WHEN** you see a stationary condition that is a rational function set to zero **DO** clear the positive denominator and read the sign of the resulting polynomial **BECAUSE** the discriminant of that polynomial hands you the case split and the threshold value of the parameter without you having to guess where the behaviour changes
- **WHEN** you must evaluate a function at an ugly root r of a known quadratic **DO** use the defining relation to rewrite the argument before substituting **BECAUSE** the relation $ar^2 - 2r + a = 0$ turns $\ln(1 + r^2)$ into $\ln(2r/a)$, killing a nested surd; substituting the explicit root first costs three lines and invites a skipped-calculation error
- **WHEN** you see an even (or odd) function paired with a parameterised line **DO** use the reflection to halve the parameter range before any computation **BECAUSE** a symmetry acting on parameters is a free quotient of the case analysis, and applying it after the algebra means you paid for both halves

Backtesting

1. First useful cue: b occurs only as an additive constant, so the problem is a level-set count for $g(x) = \ln(1 + x^2) - ax$, not a two-parameter search.
2. Method the expression suggested: differentiate g, clear the denominator $1 + x^2 > 0$, and let the discriminant of $ax^2 - 2x + a$ dictate the case split at $|a| = 1$.
3. Where you would hesitate, and the thought to reinforce: at $a = 1$, where $g' \leq 0$ with equality at $x = 1$. Reinforce that a nonpositive derivative vanishing at finitely many points still gives strict monotonicity; do not weaken the conclusion to "nonincreasing" and then fail to conclude injectivity.
4. The trap the question sets: the phrase "exactly one point" baiting the tangency circuit. The tangency answer set is a one-dimensional curve in the (a, b) plane; the true answer set has nonempty interior. If your answer is a curve rather than a region, you have fired the wrong circuit.
5. What makes you recognise it faster next time: the shape of the answer. "Determine all ordered pairs" means a region, and regions come from inequalities on levels, which come from turning point values.
6. The reusable decision rule: separate parameters into levels, scales and shape parameters; eliminate levels by making them the right-hand side, eliminate scales by rescaling, and case-split only on shape parameters.

Mirror drills

1. For which real c does $y = cx$ meet $y = \ln x$ ($x > 0$) in exactly one point? *Answer:* Set $h(x) = \ln x / x$, whose maximum is $1/e$ at $x = e$; the line meets the curve twice for $0 < c < 1/e$,

once for $c = 1/e$ and once for $c \le 0$, and never for $c > 1/e$. The circuit is the same: isolate the parameter, here as a level of $\ln x / x$ rather than an additive level.

2. Determine the number of real solutions of $e^x = ax + b$ in terms of (a, b). *Answer:* With $g(x) = e^x - ax$, $g' = e^x - a$ has one zero when $a > 0$ and none when $a \le 0$; so for $a \le 0$ there is exactly one solution for every b, and for $a > 0$ there are 0, 1 or 2 solutions according as $b < g(\ln a)$, $b = g(\ln a)$ or $b > g(\ln a)$, where $g(\ln a) = a - a \ln a$.

3. Show that $\left|\frac{2x}{1+x^2}\right| \le 1$ for all real x, with equality only at $x = \pm 1$, and deduce that no line of slope 2 is tangent to $y = \ln(1 + x^2)$. *Answer:* The inequality is $(|x| - 1)^2 \ge 0$; a tangent of slope 2 would need $2x/(1 + x^2) = 2$, impossible, so the slope-2 lines all cross transversally and each crosses exactly once.

4.2 A2. Maximum negative coefficients of a squared polynomial

Let n be an integer with $n \geq 2$. Over all real polynomials $p(x)$ of degree n, what is the largest possible number of negative coefficients of $p(x)^2$?

Read it out loud

Say the structure aloud: you are asked for an extremal count, over all degree n real polynomials, of how many entries of the coefficient vector of $p(x)^2$ can be negative. So there are two halves and you must write both: an upper bound valid for every p, and one explicit p attaining it. Any answer that supplies only one half scores about half the marks, so budget your time in two blocks from the start.

The second thing to say is what a square does to a coefficient sequence. If $p(x) = \sum_{i=0}^{n} a_i x^i$ then $p(x)^2 = \sum_{k=0}^{2n} b_k x^k$ with $b_k = \sum_{i+j=k} a_i a_j$: the coefficient sequence of the square is the convolution of the coefficient sequence with itself. Convolution is the only structure in the problem, so every idea must be a statement about convolutions.

The third observation is free and it is the whole upper bound in embryo: the two end coefficients of a convolution square are $b_0 = a_0^2$ and $b_{2n} = a_n^2$. Both are squares of real numbers, so neither is negative, and $b_{2n} > 0$ because $\deg p = n$. There are $2n+1$ coefficients in total, so at most $2n-1$ of them can be negative, and the answer is $2n-1$, $2n-2$, or lower. Within thirty seconds you should have written $b_k = \sum_{i+j=k} a_i a_j$, noted $b_0 = a_0^2 \geq 0$ and $b_{2n} = a_n^2 > 0$, and asked the only remaining question: can $b_1, \ldots, b_{2n-1}$ all be negative?

The fields wired in are polynomial algebra (convolution of coefficient sequences), sign propagation arguments of the sort that also appear in formal power series square roots, and an extremal construction. There is no analysis here and no need for roots; everything happens in the coefficient ring.

Approaches, ranked worst to best

1. **Search for a counterexample at small n by hand.** *Cue:* "largest possible number" invites experiment. *Circuit:* small cases first. Setting $n = 2$ and grinding through $(a_2x^2 + a_1x + a_0)^2$ will convince you that 2 negatives is achievable and 3 is not, but for the general n it proves nothing and eats twenty minutes. Experiment to find the answer, never to write the answer. Ranked last because it produces no proof of either half.

2. **Evaluate at special points.** *Cue:* a squared polynomial is nonnegative everywhere. *Circuit:* substitute $x = 1$, $x = -1$, integrate against a kernel. From $p(1)^2 = \sum_k b_k \geq 0$ you get $b_0 + b_{2n} \geq -\sum_{k=1}^{2n-1} b_k$, which is a true statement and a useless one: b_0 and b_{2n} are unbounded, so no contradiction can come out. Global evaluations cannot see the sign pattern of individual coefficients; they only see weighted totals. Ranked here because it is the first reflex of anyone trained on "squares are nonnegative" and it dies immediately.

3. **Descartes' rule of signs on $p(x)^2$.** *Cue:* the sign pattern $+, -, \ldots, -, +$ has exactly two sign changes. *Circuit:* sign changes bound positive roots. Two sign changes bound the number of positive roots of p^2 by two, and every real root of p^2 has even multiplicity. This is a genuine structural constraint, but it constrains roots, not coefficients, and a polynomial with that sign pattern and no positive root is not immediately absurd. You would still have to go back to the coefficients. Ranked above evaluation because it at least engages with the sign pattern.

4. **Sign propagation through the convolution, plus a palindromic construction.** *Cue:* b_k is a convolution and you are told about signs of b. *Circuit:* invert the convolution one coefficient at a time and track signs. Assume all of $b_1, \ldots, b_{2n-1}$ are negative and read the convolution as a recursion for the a_k: each b_k determines a_k from $a_0, \ldots, a_{k-1}$. Signs propagate, all the a_k with $k \geq 1$ are forced negative, and then $b_{2n-1} = 2a_{n-1}a_n$ is forced positive, which is the

contradiction. For the construction, take a palindromic p so that p^2 is palindromic and you only have to check half the coefficients. This is the ideal route: it uses the only structure present, it is four lines long, and the palindromy halves the verification.

The ideal solution

The answer is $\boxed{2n-2}$.

Write $p(x) = \sum_{i=0}^{n} a_i x^i$ with $a_n \neq 0$, and $p(x)^2 = \sum_{k=0}^{2n} b_k x^k$, so that

$$b_k = \sum_{\substack{i+j=k \\ 0 \le i,j \le n}} a_i a_j \qquad (0 \le k \le 2n).$$

Upper bound. We have $b_0 = a_0^2 \geq 0$ and $b_{2n} = a_n^2 > 0$, so at most the $2n-1$ coefficients $b_1, \ldots, b_{2n-1}$ can be negative. Suppose, for contradiction, that all of them are negative.

Since $b_1 = 2a_0a_1 < 0$, both a_0 and a_1 are nonzero and of opposite signs. Replacing p by $-p$ changes neither p^2 nor the degree, so we may assume $a_0 > 0$; then $a_1 < 0$.

Claim 4.2.1. *$a_k < 0$ for every k with $1 \le k \le n$.*

Proof. If not, let k be the least index in $\{1, \ldots, n\}$ with $a_k \geq 0$; by the previous paragraph $k \geq 2$, and $a_i < 0$ for $1 \le i \le k-1$. Since $k \le n$, every index i with $0 \le i \le k$ satisfies $i \le n$ and $k - i \le n$, so

$$b_k = 2a_0a_k + \sum_{i=1}^{k-1} a_i a_{k-i}.$$

The first term is ≥ 0 because $a_0 > 0$ and $a_k \geq 0$. Every summand $a_i a_{k-i}$ with $1 \le i \le k-1$ is a product of two negative numbers, hence positive. Therefore $b_k \geq 0$, contradicting $b_k < 0$. □

Now use $n \geq 2$, so that $n - 1 \geq 1$ and the claim applies to both a_{n-1} and a_n. Then

$$b_{2n-1} = 2a_{n-1}a_n > 0,$$

contradicting $b_{2n-1} < 0$. Hence at most $2n-2$ of the coefficients of $p(x)^2$ are negative.

Construction attaining $2n-2$. Put

$$p(x) = n(x^n + 1) - 2Q(x), \qquad Q(x) = x + x^2 + \cdots + x^{n-1},$$

which has degree n since its leading coefficient is $n \neq 0$ (for $n = 2$, read $Q(x) = x$). Then

$$p(x)^2 = n^2(x^n+1)^2 - 4n(x^n+1)Q(x) + 4Q(x)^2.$$

Read off coefficients. The first term contributes n^2 to each of b_0, b_{2n} and $2n^2$ to b_n, and nothing else. The second term contributes $-4n$ to each of $b_1, \ldots, b_{n-1}$ (from Q) and to each of $b_{n+1}, \ldots, b_{2n-1}$ (from x^nQ), and nothing to b_n. For the third term, for $1 \le i \le n$ the coefficient of x^i in $Q(x)^2$ is the number of pairs (j,k) with $1 \le j,k \le n-1$ and $j + k = i$, which is $i - 1$. Hence

$$b_i = -4n + 4(i-1) \qquad (1 \le i \le n-1),$$

and each of these is at most $-4n + 4(n-2) = -8 < 0$. Also $b_n = 2n^2 + 4(n-1) > 0$ and $b_0 = b_{2n} = n^2 > 0$.

Finally, p is palindromic: its coefficient list is $n, -2, -2, \ldots, -2, n$, unchanged on reversal, that is $x^n p(1/x) = p(x)$. Therefore $x^{2n}p(1/x)^2 = p(x)^2$, so $b_k = b_{2n-k}$ for all k, and the negativity of $b_1, \ldots, b_{n-1}$ gives the negativity of $b_{n+1}, \ldots, b_{2n-1}$ at once. So exactly $2(n-1) = 2n-2$ coefficients of $p(x)^2$ are negative, and the bound is attained.

Check at $n = 2$. $p(x) = 2x^2 - 2x + 2$, $p(x)^2 = 4x^4 - 8x^3 + 12x^2 - 8x + 4$, with exactly $2 = 2n-2$ negative coefficients. At $n = 3$, $p(x) = 3x^3 - 2x^2 - 2x + 3$ and $p(x)^2 = 9x^6 - 12x^5 - 8x^4 + 26x^3 - 8x^2 - 12x + 9$, with exactly 4 negative coefficients.

Layer 1: WHAT

Block 1. Name the object: b_k is the self-convolution $\sum_{i+j=k} a_i a_j$.

Block 2. Free bound: $b_0 = a_0^2 \geq 0$ and $b_{2n} = a_n^2 > 0$, so the count is at most $2n-1$, and the entire question is whether $2n-1$ is achievable.

Block 3. Assume it is, and normalise: $b_1 = 2a_0a_1 < 0$ forces $a_0 \neq 0$, and $p \mapsto -p$ lets you take $a_0 > 0$.

Block 4. Read the convolution as a rule that determines the sign of a_k from the signs of $a_0, \ldots, a_{k-1}$ and of b_k, and take a least counterexample: all a_k with $1 \leq k \leq n$ are negative.

Block 5. Look at the other end of the sequence, where the same negativity is fatal: $b_{2n-1} = 2a_{n-1}a_n > 0$. Contradiction, so the count is at most $2n-2$.

Block 6. Build an example. Force the sign pattern by making the two end coefficients large and positive and all interior coefficients equal and negative, so that every middle coefficient of the square is dominated by the cross term $-4n$.

Block 7. Use palindromy to check only half the coefficients, then count: $2n-2$.

Juxtapose with the nominal route. The zombie multiplies out $\left(\sum a_i x^i\right)^2$ in full generality, stares at $2n+1$ simultaneous sign conditions in $n+1$ unknowns, and starts casework on $n = 2, 3, 4$. He never notices that the conditions are not independent: they are a triangular system, and a triangular system in which each new equation determines one new unknown should be run as a recursion, not solved as a system.

The bifurcation point is the line where you decide what the hypothesis "$b_k < 0$ for all $1 \leq k \leq 2n-1$" is for. The zombie treats it as $2n-1$ constraints to satisfy. You treat it as $2n-1$ sign inputs to a recursion that outputs the signs of the a_k, and then you check the output against the top end of the sequence. Same data, opposite direction of travel.

Layer 2: HOW

Block 1. **Cue.** "Coefficients of $p(x)^2$". **Why natural.** Squaring a polynomial is convolving its coefficient sequence with itself. Writing $b_k = \sum_{i+j=k} a_i a_j$ is not a step, it is a translation into the only language in which the question has content. The moment it is written, the extreme coefficients b_0 and b_{2n} become visibly special because their sums have a single term. **Link to the next block.** A one-term sum is a square, and squares have a sign.

Block 2. **Cue.** $b_0 = a_0^2$, $b_{2n} = a_n^2$. **Why natural.** These two coefficients cannot be negative, so the answer is at most $2n-1$ before any work. The habit here is: in any extremal counting problem, spend the first thirty seconds finding the free bound, because it tells you the size of the gap you must still close. Here the gap is exactly one, which is a strong hint that the truth is $2n-2$ and that a single parity-like or endpoint obstruction is waiting. **Link to the next block.** To close a gap of one, assume the extreme case and hunt for a contradiction.

Block 3. **Cue.** $b_1 = 2a_0a_1$, the second simplest convolution. **Why natural.** Assuming $b_1 < 0$ immediately gives $a_0 \neq 0$, which is what makes the recursion well posed: a_0 is the pivot you divide by. And since p and $-p$ have the same square, fixing the sign of a_0 is a free normalisation, not an assumption. **Link to the next block.** With $a_0 > 0$ fixed, each b_k pins down the sign of a_k.

Block 4. **Cue.** The identity $b_k = 2a_0a_k + \sum_{i=1}^{k-1} a_i a_{k-i}$. **Why natural.** Isolating the two terms that involve a_k is the standard move for a triangular system: it turns the convolution into a formula for the new unknown in terms of the old ones. If all previous a_i ($1 \leq i \leq k-1$) are negative, the residual sum is a sum of products of negatives, hence positive, so $2a_0a_k$ must be even more negative than b_k. The negativity of a_k is therefore forced, not chosen. Phrasing it as a least counterexample rather than as an induction saves a sentence and removes the base case fuss. **Link to the next block.**

Once you know the sign of every a_k, look for a coefficient of p^2 whose sign is now determined and compare with the hypothesis.

Block 5. **Cue.** b_{2n-1} is the mirror image of b_1. **Why natural.** The convolution is symmetric, so the near-top coefficients have exactly as few terms as the near-bottom ones: $b_{2n-1} = 2a_{n-1}a_n$. The hypothesis pushed all the a_k negative, and a product of two negatives is positive. The contradiction lands at the far end of the sequence because that is the only other place where the convolution has a single term. Notice that this is precisely where the condition $n \geq 2$ is spent: at $n = 1$ the indices collide, $b_{2n-1} = b_1 = 2a_0a_1$, and there is no contradiction, which is why the problem excludes $n = 1$ (and indeed for $n = 1$ the answer is 1, not 0). **Link to the next block.** The bound is $2n - 2$; now produce a polynomial attaining it.

Block 6. **Cue.** You need $2n - 2$ negatives, that is every coefficient except b_0, b_n, b_{2n}. **Why natural.** The proof of the bound tells you the shape of the example: the obstruction was the pair of end coefficients of p, so make those large and positive and all interior coefficients negative. Then the cross terms $2a_0a_i$ and $2a_na_{n+i}$ dominate. Choosing all interior coefficients equal to -2 and both ends equal to n makes the dominant cross term $-4n$ and the competing positive term at most $4(n-2)$, so the margin is 8 and no fine tuning is needed. Choosing round constants that leave slack is a design decision: an example with slack is an example you can verify in two lines. **Link to the next block.** Verify the coefficients, but only half of them.

Block 7. **Cue.** The coefficient list of p reads the same backwards. **Why natural.** A palindromic polynomial satisfies $x^{\deg}p(1/x) = p(x)$, and this property is multiplicative, so p^2 is palindromic too and $b_k = b_{2n-k}$. That halves the verification and it also explains why b_n is the natural exceptional coefficient: it is the fixed point of the reflection $k \mapsto 2n - k$. **Link to the next block.** Count: $n - 1$ negatives below the middle, $n - 1$ above, total $2n - 2$.

Layer 3: WHY

The mathematical truth underneath is that the map $a \mapsto b$, $b_k = \sum_{i+j=k} a_ia_j$, is triangular with respect to the index: b_k involves a_k only through the term $2a_0a_k$, plus terms with strictly smaller indices. Whenever $a_0 \neq 0$ this is invertible, and the inverse is the formal square root recursion

$$a_0 = \sqrt{b_0}, \qquad a_k = \frac{1}{2a_0}\left(b_k - \sum_{i=1}^{k-1} a_ia_{k-i}\right) \quad (k \geq 1),$$

which computes the coefficients of $\sqrt{b(x)}$ as a formal power series. The problem is a sign question about this recursion: if the input has one positive coefficient followed by all negative ones, the output has one positive coefficient followed by all negative ones. The prototype is $\sqrt{1-4x} = 1 - 2\sum_{m\geq 1} C_{m-1}x^m$ with C_j the Catalan numbers, where every coefficient after the constant is negative. That is exactly the phenomenon the problem forbids you from realising with a polynomial, because a polynomial square root must stop, and stopping is what creates the positive coefficient $b_{2n-1} = 2a_{n-1}a_n$ at the top.

So the real content is: an infinite series can have the sign pattern $+, -, -, -, \ldots$ and be a perfect square, but a polynomial of degree $2n$ cannot, because a polynomial has two ends and the argument that forces all interior signs at the bottom end contradicts what the top end forces. Finiteness is the obstruction. That is the single sentence to remember.

The discernment mechanism: when a problem gives you sign information about a convolution, decide which direction the information flows. Sign data on the product plus the triangular structure means the data flows to the factors, one index at a time; and once you know the factors, all remaining coefficients of the product are determined in sign, so you should immediately test the coefficients

you have not used yet. Here the unused coefficient is b_{2n-1}, and it is the one that kills the hypothesis. Train the reflex by asking, at the end of any sign propagation: which hypotheses have I not yet spent? The one you have not spent is usually the contradiction.

The examiners' nudges: they said "degree n" rather than "at most degree n" so that $a_n \neq 0$ and $b_{2n} > 0$ is guaranteed, and they imposed $n \geq 2$ precisely because the argument at the top end needs two distinct indices $n-1$ and n. The phrase "largest possible number" is the command term telling you that a bound alone is worth part marks and a construction alone is worth part marks, so write both.

Revolving concepts

Convolution of coefficient sequences. For $p(x) = \sum a_i x^i$ and $q(x) = \sum c_j x^j$, $[x^k](pq) = \sum_{i+j=k} a_i c_j$. The extreme coefficients of a product are single products: $[x^0](pq) = a_0 c_0$ and $[x^{\deg p + \deg q}](pq) = a_{\deg p} c_{\deg q}$. See Book 1, field: Polynomials, topic: coefficient convolution.

The formal square root recursion (generalised recursion). Given a formal power series $b(x) = b_0 + b_1 x + \cdots$ with $b_0 > 0$, there is a unique power series $a(x)$ with $a(0) > 0$ and $a^2 = b$, given by $a_0 = \sqrt{b_0}$ and, for $k \geq 1$,

$$a_k = \frac{1}{2a_0}\left(b_k - \sum_{i=1}^{k-1} a_i a_{k-i}\right).$$

This is the general solution method for any quadratic convolution recursion: isolate the two terms containing the highest unknown, and read the rest as a known history. The same pattern solves Catalan-type recursions $c_k = \sum_{i=0}^{k-1} c_i c_{k-1-i}$ (whose generating function satisfies $C = 1 + xC^2$, giving $C = (1 - \sqrt{1-4x})/(2x)$), and more generally any equation $F(x, y) = 0$ solved for y as a power series by the implicit function theorem for formal series. See Book 1, field: Generating functions, topic: quadratic functional equations and Catalan numbers.

Palindromic (self-reciprocal) polynomials. p of degree d is palindromic if $x^d p(1/x) = p(x)$, equivalently its coefficient list is a palindrome. Products of palindromic polynomials are palindromic; the roots come in pairs $r, 1/r$; and substituting $t = x + 1/x$ after dividing by $x^{d/2}$ halves the degree in the even case. Here it is used purely as a labour-saving symmetry on the coefficient list. See Book 1, field: Polynomials, topic: palindromic polynomials.

Descartes' rule of signs. The number of positive real roots of a real polynomial, counted with multiplicity, is at most the number of sign changes in its coefficient list, and has the same parity. Useful to know here so that you can rule the approach out fast rather than wonder about it. See Book 1, field: Polynomials, topic: sign rules and root location.

Extremal problems with a two-part answer. Any "largest possible number" question is answered by a pair (bound, construction). The bound should use only structural facts; the construction should be chosen with slack so that verification is cheap. See Book 1, field: Combinatorics, topic: extremal constructions.

Circuit Bank

- **WHEN** you see coefficients of a product or a square **DO** write the convolution $b_k = \sum_{i+j=k} a_i c_j$ and immediately evaluate the two extreme coefficients **BECAUSE** extreme coefficients are single products, so they carry unconditional sign information that the middle coefficients do not, and that is usually the free half of an extremal bound
- **WHEN** you are given sign or size data on every coefficient of a product **DO** isolate the highest unknown, $b_k = 2a_0 a_k + \sum_{i=1}^{k-1} a_i a_{k-i}$, and propagate signs from the bottom index upwards **BECAUSE** the convolution is triangular in the index, so the data determines the factors one coefficient at a time; solving it as a simultaneous system instead is the zombie route

- **WHEN** a sign propagation has determined all the unknowns **DO** test the hypotheses you have not yet spent, starting from the opposite end of the index range **BECAUSE** the contradiction in a finite problem almost always lives where the structure stops, because the propagation argument used only one end
- **WHEN** you need an explicit extremal example **DO** copy the shape of the obstruction from your own proof of the bound, and choose constants with a wide margin **BECAUSE** the proof of the upper bound tells you which coefficients must be positive, so the example is the sign pattern of the proof; wide margins turn verification into one inequality instead of a case check
- **WHEN** you are checking many coefficients of a product **DO** make the factor palindromic **BECAUSE** palindromy is multiplicative, so $b_k = b_{2n-k}$ and you verify half the list; the middle index is the fixed point of the reflection and is the natural exception

Backtesting

1. First useful cue: coefficients of a square means self-convolution, and the extreme coefficients $b_0 = a_0^2$, $b_{2n} = a_n^2$ are unconditionally nonnegative, giving the free bound $2n - 1$.
2. Method the expression suggested: assume the free bound is attained, isolate a_k in $b_k = 2a_0a_k + \sum_{i=1}^{k-1} a_i a_{k-i}$, and propagate signs upward from $k = 1$.
3. Where you would hesitate, and the thought to reinforce: at the normalisation $a_0 > 0$. Reinforce that p and $-p$ have the same square, so fixing that sign costs nothing; hesitating here and carrying two cases doubles the work for no reason.
4. The trap the question sets: stopping at the free bound $2n - 1$ and asserting it is attained, or producing a construction with $2n - 2$ negatives and not proving that $2n - 1$ is impossible. The gap is exactly one, and it is where all the marks are.
5. What makes you recognise it faster next time: a convolution with sign data is a triangular system; triangular systems are recursions; recursions with sign data are sign propagations; sign propagations end by testing an unused coefficient at the far end.
6. The reusable decision rule: in extremal coefficient problems, get the free bound from the extreme coefficients, then attack the single remaining case by assuming it and propagating.

Mirror drills

1. Show that if p is a real polynomial of degree $n \geq 1$ with $p(0) \neq 0$, then $p(x)^2$ cannot have all of $b_1, \ldots, b_{2n-1}$ negative and $b_0 < 0$ simultaneously. *Answer:* $b_0 = a_0^2 > 0$ always, so the hypothesis $b_0 < 0$ is already impossible; the point of the drill is to see that the end coefficients of a square are the only ones with a forced sign.
2. Prove that the coefficients of $\sqrt{1 - 4x}$, after the constant term 1, are all negative. *Answer:* Write $\sqrt{1 - 4x} = 1 - 2xC(x)$ where C is the Catalan generating function; squaring gives $C = 1 + xC^2$, and induction on the recursion $c_k = \sum_{i=0}^{k-1} c_i c_{k-1-i}$ with $c_0 = 1$ gives $c_k > 0$ for all k, hence all coefficients of $\sqrt{1 - 4x}$ beyond the constant are negative. This is the infinite model that the finite problem forbids.
3. What is the largest possible number of negative coefficients of $p(x)^3$ for p real of degree n, when $n = 1$? *Answer:* $p^3 = (a_1x + a_0)^3$ has coefficients $a_0^3, 3a_0^2a_1, 3a_0a_1^2, a_1^3$; taking $a_0 > 0 > a_1$ makes the middle two negative and the ends of opposite sign, so with $a_0 > 0$ exactly two are negative, and taking $a_0 < 0 < a_1$ gives two as well; the maximum is 2. The drill is to notice that odd powers do not have the "end coefficients are squares" protection, so the whole structure of the argument changes.

4.3 A3. Counting Lyness sequences modulo a prime

Let p be a prime number greater than 5. Let $f(p)$ denote the number of infinite sequences $a_1, a_2, a_3, \ldots$ such that $a_n \in \{1, 2, \ldots, p-1\}$ and

$$a_n a_{n+2} \equiv 1 + a_{n+1} \pmod{p}$$

for all $n \geq 1$. Prove that $f(p)$ is congruent to 0 or 2 $\pmod 5$.

Read it out loud

Read the structure: the values live in $\{1, \ldots, p-1\}$, which is exactly $\mathbb{F}_p^\times$, the nonzero residues, so every a_n is invertible. The relation $a_n a_{n+2} = 1 + a_{n+1}$ can therefore be divided through:

$$a_{n+2} = \frac{1 + a_{n+1}}{a_n}.$$

That is a second order nonlinear recurrence over a field, and a second order recurrence is a first order dynamical system on pairs: the state is (a_n, a_{n+1}) and the map is $T(x, y) = \left(y, \frac{1+y}{x}\right)$. So the problem is: count the points of $\mathbb{F}_p^\times \times \mathbb{F}_p^\times$ whose entire forward orbit under T stays inside $\mathbb{F}_p^\times \times \mathbb{F}_p^\times$, and prove that count is 0 or 2 modulo 5.

Now the loudest cue in the whole problem: the modulus you are asked about is 5, and 5 has nothing to do with p. When a counting problem asks for an answer modulo a prime q that is unrelated to the data, there are only two mechanisms in the whole of mathematics that produce such a statement. Either the count factorises into an expression whose residue you can compute directly, or there is a group of order q acting on the set being counted, and the count is congruent to the number of fixed points modulo q. The second mechanism is far more common and far more explanatory. So before doing anything else, ask: where is the action of $\mathbb{Z}/5$? The only candidate group action on a set of sequences is the shift, and a shift generates $\mathbb{Z}/5$ exactly when every admissible sequence has period 5.

That is the thought you should have inside the first minute, and it tells you what to compute: the first several terms, hoping the sequence closes up after five. Within thirty seconds you should have written $a_{n+2} = (1 + a_{n+1})/a_n$ and started computing a_3, a_4, a_5, a_6 in terms of a_1 and a_2. The fields wired in are finite field arithmetic, discrete dynamical systems (a birational map of the plane), group actions and orbit counting, and the theory of quadratic residues.

Approaches, ranked worst to best

1. **Compute $f(p)$ for small primes and guess.** *Cue:* "prove $f(p) \equiv 0$ or 2" invites data. *Circuit:* tabulate and pattern-match. Running $p = 7, 11, 13$ by hand gives $20, 72, 110$ and the residues $0, 2, 0$. Useful reconnaissance, worthless as a proof, and expensive by hand. Ranked last, but do run one small case mentally as a sanity check on your algebra.

2. **Treat it as a linear recurrence in disguise.** *Cue:* recurrences suggest characteristic equations. *Circuit:* substitute $a_n = c\lambda^n$ or take logarithms. The relation is multiplicative on the left and additive on the right, so no substitution linearises it globally. Chasing a closed form for a_n in terms of a_1, a_2 is exactly the zombie route: it is not that the algebra is hard, it is that there is nothing at the end of it. Ranked here because it burns the most time for the least return.

3. **Direct count: prove period 5, characterise admissible (a_1, a_2), count, reduce mod 5.** *Cue:* a periodic sequence is determined by finitely many exclusions. *Circuit:* count the complement of a union of lines. This works and it is the natural second thought. Periodicity reduces admissibility to the three conditions $a_1 \neq -1$, $a_2 \neq -1$, $a_1 + a_2 \neq -1$, giving $f(p) = (p-2)(p-3)$, a product of two consecutive integers; you then check the residues of p modulo

5. It is complete and it also delivers the exact value, which the problem did not ask for. It ranks second only because it explains the "0 or 2" as an accident of arithmetic rather than as structure.

4. **Orbit counting under the shift.** *Cue:* the answer is asked modulo 5 and the sequences have period 5. *Circuit:* a group of prime order q acting on a finite set makes $|S| \equiv \#\mathrm{Fix} \pmod q$. Prove period 5, note the shift generates a $\mathbb{Z}/5$ action on the set of admissible sequences, and conclude that $f(p)$ is congruent modulo 5 to the number of constant admissible sequences, which is the number of roots of $x^2 - x - 1$ in $\mathbb{F}_p$, that is 0 or 2 because the discriminant 5 is nonzero. This is the ideal route: it is five lines after the periodicity computation, it never counts anything, and it explains exactly why 2 and not some other residue appears, namely because a quadratic with nonzero discriminant has 0 or 2 roots.

The ideal solution

Throughout, identify $\{1, 2, \dots, p-1\}$ with $\mathbb{F}_p^\times$ and work in $\mathbb{F}_p$. Call a sequence *admissible* if it satisfies the two conditions in the statement, and let S be the set of admissible sequences, so $f(p) = |S|$.

Step 1: the recurrence is a map on pairs. Each a_n is a unit, so the relation $a_n a_{n+2} = 1 + a_{n+1}$ is equivalent to $a_{n+2} = (1 + a_{n+1})a_n^{-1}$. Hence an admissible sequence is determined by (a_1, a_2), and S is finite (there are at most $(p-1)^2$ candidates).

Step 2: every admissible sequence has period 5. Let $\{a_n\} \in S$ and write $x = a_1$, $y = a_2$. Using the recurrence repeatedly, and dividing only by terms of the sequence, which are units:

$$a_3 = \frac{1+y}{x}, \qquad a_4 = \frac{1+a_3}{y} = \frac{1+x+y}{xy}, \qquad a_5 = \frac{1+a_4}{a_3} = \frac{\frac{(1+x)(1+y)}{xy}}{\frac{1+y}{x}} = \frac{1+x}{y},$$

$$a_6 = \frac{1+a_5}{a_4} = \frac{\frac{1+x+y}{y}}{\frac{1+x+y}{xy}} = x, \qquad a_7 = \frac{1+a_6}{a_5} = \frac{1+x}{\frac{1+x}{y}} = y.$$

In the computation of a_5 we used $1 + a_4 = 1 + \frac{1+x+y}{xy} = \frac{xy+x+y+1}{xy} = \frac{(1+x)(1+y)}{xy}$, and in that of a_6 we used $1 + a_5 = \frac{y+1+x}{y}$. Since $(a_6, a_7) = (a_1, a_2)$ and the sequence is determined by consecutive pairs, $a_{n+5} = a_n$ for all $n \geq 1$.

Step 3: the shift acts. Define $\sigma : S \to S$ by $\sigma(a_1, a_2, a_3, \dots) = (a_2, a_3, a_4, \dots)$. This lands in S: the shifted sequence takes values in $\mathbb{F}_p^\times$, and its recurrence at index n is the original recurrence at index $n+1$, which holds. By Step 2, $\sigma^5 = \mathrm{id}_S$, so σ generates an action of the cyclic group $\mathbb{Z}/5\mathbb{Z}$ on the finite set S.

Step 4: orbits and fixed points. Every orbit of this action has size dividing 5, hence size 1 or 5. An orbit has size 1 exactly when σ fixes the sequence, that is when $a_{n+1} = a_n$ for all n, that is when the sequence is constant. Partitioning S into orbits,

$$f(p) = |S| = \#\{\text{constant admissible sequences}\} + 5 \cdot \#\{\text{orbits of size } 5\},$$

so

$$f(p) \equiv \#\{\text{constant admissible sequences}\} \pmod 5.$$

Step 5: count the constant admissible sequences. The constant sequence $a_n = a$ with $a \in \mathbb{F}_p^\times$ is admissible if and only if $a \cdot a = 1 + a$, that is

$$a^2 - a - 1 = 0.$$

Since $p > 5$ we have p odd and $5 \neq 0$ in $\mathbb{F}_p$, so we may complete the square: $a^2 - a - 1 = 0 \iff (2a-1)^2 = 5$. The map $a \mapsto 2a - 1$ is a bijection of $\mathbb{F}_p$, so the number of solutions a equals the

number of square roots of 5 in $\mathbb{F}_p$, which is 2 if 5 is a nonzero quadratic residue modulo p and 0 otherwise (it cannot be 1, since $5 \neq 0$ means the two square roots $\pm\sqrt{5}$ are distinct). Note also that no root of $a^2 - a - 1$ is zero, since $0^2 - 0 - 1 = -1 \neq 0$, so every root does lie in $\mathbb{F}_p^\times$ and gives a genuine admissible sequence.

Combining Steps 4 and 5, $f(p) \equiv 0$ or $2 \pmod 5$, as required. ■

Remark: the exact value, as an independent check. By Step 2 an admissible sequence exists for a pair $(x, y) \in \mathbb{F}_p^\times \times \mathbb{F}_p^\times$ precisely when the five values $a_1, \ldots, a_5$ are all nonzero, and from the formulas above those five values are

$$x, \quad y, \quad \frac{1+y}{x}, \quad \frac{1+x+y}{xy}, \quad \frac{1+x}{y},$$

so the conditions are $x \neq -1$, $y \neq -1$ and $x + y \neq -1$. There are $p - 2$ choices of x in $\mathbb{F}_p^\times$ with $x \neq -1$, likewise for y, and for each such x exactly one forbidden value $y = -1 - x$, which is itself neither 0 nor -1 (it is 0 only if $x = -1$, and -1 only if $x = 0$). Hence

$$f(p) = (p-2)^2 - (p-2) = (p-2)(p-3).$$

Since $p > 5$ is prime, $p \not\equiv 0 \pmod 5$, and $(p-2)(p-3)$ modulo 5 for $p \equiv 1, 2, 3, 4$ is $2, 0, 0, 2$ respectively. This agrees with Step 5: $p \equiv \pm 1 \pmod 5$ is exactly the condition for 5 to be a quadratic residue modulo p, by quadratic reciprocity applied to $5 \equiv 1 \pmod 4$.

Remark: the invariant. The three conditions $x \neq -1, y \neq -1, x + y \neq -1$ are exactly the statement that

$$I(x, y) = \frac{(1+x)(1+y)(1+x+y)}{xy}$$

is nonzero, and a direct computation shows I is invariant under the map $T(x, y) = \left(y, \frac{1+y}{x}\right)$. So admissibility says precisely that the orbit lies on a nonzero level set of I.

Layer 1: WHAT

Block 1. Recognise $\{1, \ldots, p-1\} = \mathbb{F}_p^\times$ and divide: $a_{n+2} = (1 + a_{n+1})/a_n$.

Block 2. Recognise the second order recurrence as a map on pairs; the sequence is determined by (a_1, a_2).

Block 3. Read the target modulus 5 as an instruction to find a $\mathbb{Z}/5$ action; the only candidate is the shift, so compute terms and hope for period 5.

Block 4. Compute a_3, a_4, a_5, a_6, a_7 and find $(a_6, a_7) = (a_1, a_2)$: period 5.

Block 5. The shift is a $\mathbb{Z}/5$ action on the finite set S; orbits have size 1 or 5, so $|S| \equiv \#\mathrm{Fix} \pmod 5$.

Block 6. Fixed points are constant sequences, so they solve $a^2 = 1 + a$; a quadratic with nonzero discriminant over $\mathbb{F}_p$ (p odd) has 0 or 2 roots. Done.

Juxtapose with the nominal route. The zombie sees a recurrence and tries to solve it: he attempts a substitution to linearise, or he computes $a_3, a_4, a_5, \ldots$ and keeps going hoping for a closed form, or he sets up the count of (a_1, a_2) and then tries to justify the residues by arithmetic on $(p-2)(p-3)$ without ever asking why the number 5 should appear. The direct count is a perfectly good solution, but the zombie version of it stops at "and now check four cases", learning nothing.

The bifurcation point is the sentence in which you interpret the number 5 in "modulo 5". The zombie reads it as the modulus he must eventually reduce to. You read it as the order of a group that must be acting, and that reading tells you what to compute before you compute it.

Layer 2: HOW

Block 1. Cue. $a_n \in \{1, \ldots, p-1\}$, that is, never 0 modulo p. **Why natural.** The value set was specified as the nonzero residues, not as all residues, and the only reason to specify that is

invertibility. Dividing by a_n is therefore licensed by the hypothesis, and the recurrence becomes explicit rather than implicit. Whenever a problem restricts values to units, the intended move is division. **Link to the next block.** An explicit recurrence of order two is a dynamical system on pairs.

Block 2. Cue. Two previous terms determine the next. **Why natural.** The state of a second order recurrence is the pair (a_n, a_{n+1}), and the recurrence is then a single map T applied repeatedly. This reframing matters because the natural objects (orbits, periods, invariants, fixed points) all belong to the language of maps, not of sequences. **Link to the next block.** A map with finite order gives a group action, and group actions give congruences.

Block 3. Cue. The answer is requested modulo 5, and 5 is not related to p. **Why natural.** A congruence modulo a prime q for the size of a set almost always comes from a $\mathbb{Z}/q$ action, because for a group of prime order every orbit has size 1 or q. This is the same mechanism behind the necklace proof of Fermat's little theorem and behind Cauchy's theorem. Since S is a set of sequences, the only natural $\mathbb{Z}/5$ candidate is the shift, and for the shift to have order 5 you need period 5. **Link to the next block.** So the periodicity is not a lucky observation; it is the hypothesis your strategy requires, and you go and test it.

Block 4. Cue. A short nonlinear recurrence with small coefficients. **Why natural.** You test periodicity by computing, and the computation is short if you keep everything as a single fraction in x and y and factorise numerators as you go. The key simplification is that $1 + a_4 = \frac{(1+x)(1+y)}{xy}$ factorises, which is what makes a_5 collapse to $\frac{1+x}{y}$ and produces the symmetry between a_3 and a_5 under swapping x and y. That symmetry is a signal that the orbit is closing. **Link to the next block.** Period 5 confirmed, the group action exists.

Block 5. Cue. A group of prime order acting on a finite set. **Why natural.** The orbit-stabiliser theorem says orbit sizes divide the group order, and 5 is prime, so orbits have size 1 or 5. Summing sizes gives $|S| \equiv \#\{\text{orbits of size } 1\} \pmod 5$. Note what you have avoided: you never computed $|S|$. **Link to the next block.** Fixed points of a shift are constant sequences, which are governed by one equation.

Block 6. Cue. A constant sequence substituted into $a_n a_{n+2} = 1 + a_{n+1}$. **Why natural.** Substituting the simplest possible object into the defining relation is always the first probe, and here it produces $a^2 - a - 1 = 0$, whose discriminant is 5. That is the second appearance of 5, and it is the real reason the answer is "0 or 2": a quadratic over a field of odd characteristic has 0, 1 or 2 roots, and it has exactly 1 only when the discriminant vanishes, which is excluded by $p > 5$. The hypothesis $p > 5$ is therefore doing precise work, not decorating the statement. **Link to the next block.** Two possible fixed point counts, 0 and 2, give the two possible residues.

Layer 3: WHY

Three mathematical truths make this route ideal. First, the counting principle: if a group G of prime order q acts on a finite set S, then $|S| \equiv |S^G| \pmod q$, because non-fixed orbits all have size exactly q. This is the cheapest congruence-producing machine in mathematics, and it converts a hard counting problem into an easy fixed point problem. Second, the structural fact that the Lyness recurrence $x_{n+2} = (x_{n+1} + 1)/x_n$ is periodic with period 5 over any field where the terms are defined. Third, the fact that over a field of odd characteristic a quadratic has 0 or 2 roots unless its discriminant vanishes.

The discernment mechanism for choosing this route is a single question: what is the modulus for? Whenever a problem asks you to prove a congruence modulo a small prime for the cardinality

of a set of structured objects (sequences, colourings, tuples, subsets), look for the group of that order acting by symmetry of the structure. Shifts and rotations are the usual suspects for sequences and necklaces; conjugation and coordinate permutation for tuples. Then the whole problem reduces to identifying the fixed objects, which are always the "degenerate" or "constant" ones, and those are governed by a single equation.

To train it: take the classical proofs of Fermat's little theorem by necklaces, of Wilson's theorem by pairing x with x^{-1}, and of Fermat's two squares theorem by an involution, and in each case name the group, the set, and the fixed points. Once those three slots are automatic you will fill them without prompting.

What the examiners nudged with: the value set $\{1, \dots, p-1\}$ nudges you to divide; the condition $p > 5$ nudges you to notice that 5 is invertible and that the discriminant is nonzero; and the modulus 5 nudges you to look for the pentagonal symmetry. Three separate appearances of the number 5 in one problem is not a coincidence, it is a signature: the sequence has period 5, the group has order 5, and the discriminant is 5. All three come from the same source, the pentagon recurrence.

Revolving concepts

The Lyness recurrence and its generalisation. The recurrence

$$x_{n+2} = \frac{x_{n+1} + \alpha}{x_n}$$

is the Lyness map. For $\alpha = 1$ it is periodic with period 5 for all initial data where it is defined; for general α it is not periodic but preserves the invariant

$$I_\alpha(x, y) = \frac{(x+\alpha)(y+\alpha)(x+y+\alpha)}{xy},$$

whose level sets are elliptic curves, and the map acts as translation by a point on each. The $\alpha = 1$ case is degenerate: the translation is by a 5-torsion point. This is the general recursion behind the problem, and its general solution method is: look for a rational invariant, then interpret the map as translation on the level curve. See Book 1, field: Sequences and recurrences, topic: Lyness recurrences and integrable maps.

The pentagon recurrence and cluster algebras. The relation $a_n a_{n+2} = 1 + a_{n+1}$ is the exchange relation of the rank two cluster algebra of type A_2, and its periodicity with period $5 = h + 2$ (where $h = 3$ is the Coxeter number of A_2) is the smallest instance of Zamolodchikov periodicity. The same framework gives the Laurent phenomenon: every a_n is a Laurent polynomial in a_1, a_2 with integer coefficients, which is exactly what you observe when computing a_3, a_4, a_5. Flag this as a knowledge gap: the theory of cluster algebras is well beyond A-level and beyond most undergraduate courses, but the single fact "this recurrence has period 5" is what you need installed. See Book 1, field: Sequences and recurrences, topic: Laurent phenomenon and Zamolodchikov periodicity.

Group actions and congruences. If a finite group G acts on a finite set S, then S is partitioned into orbits and $|\mathrm{Orb}(s)| = [G : \mathrm{Stab}(s)]$ divides $|G|$. When $|G| = q$ is prime, orbits have size 1 or q, so $|S| \equiv |S^G| \pmod q$. Standard consequences: Fermat's little theorem via $\mathbb{Z}/q$ acting on strings of length q by rotation, and Cauchy's theorem. See Book 1, field: Abstract algebra, topic: group actions, orbit-stabiliser, and fixed point congruences.

Quadratic equations over $\mathbb{F}_p$ and quadratic residues. For odd p, $ax^2 + bx + c = 0$ with $a \neq 0$ has 2, 1 or 0 roots in $\mathbb{F}_p$ according as the discriminant $b^2 - 4ac$ is a nonzero square, zero, or a non-square. Exactly half of $\mathbb{F}_p^\times$ are squares. Quadratic reciprocity for 5 (which is 1 modulo 4) gives $\left(\frac{5}{p}\right) = \left(\frac{p}{5}\right)$, so 5 is a residue precisely when $p \equiv \pm 1 \pmod 5$. The roots of $x^2 - x - 1$ are the golden ratio and its conjugate, and their existence in $\mathbb{F}_p$ governs the period of the Fibonacci sequence modulo p. See Book 1, field: Number theory, topic: quadratic residues and reciprocity.

Counting a complement of a union. The count $(p-2)^2-(p-2)$ is inclusion and exclusion in its simplest form: from the grid of allowed (x,y) remove the line $x+y=-1$, checking that the removed line meets the grid in exactly $p-2$ points. See Book 1, field: Combinatorics, topic: inclusion and exclusion.

Circuit Bank

- **WHEN** you are asked to prove a congruence modulo a small prime q for the number of structured objects **DO** look for a group of order q acting on the set, then count only the fixed objects **BECAUSE** orbits of a prime order group have size 1 or q, so the cardinality is congruent to the number of fixed points; this converts counting into solving one equation
- **WHEN** you see a recurrence whose values are restricted to the units of a ring **DO** divide immediately to make it explicit, then treat it as a map on the state vector of previous terms **BECAUSE** the restriction to units exists precisely to license the division, and an explicit map brings orbits, periods and invariants into play
- **WHEN** you suspect a nonlinear recurrence is periodic **DO** compute successive terms as single fractions and factorise every numerator as you go **BECAUSE** periodicity shows up as unexpected factorisations, and a symmetric expression appearing (here a_3 and a_5 swapping under $x \leftrightarrow y$) is the sign that the orbit is about to close
- **WHEN** you have a map T and want to know which orbits are defined for all time **DO** look for a rational invariant I with $I \circ T = I$, and read the degeneracy conditions off the zeros and poles of I **BECAUSE** an invariant turns an infinite condition (all iterates defined) into a finite one (a nonzero level set), which is exactly what the three exclusions $x, y, x+y \neq -1$ encode
- **WHEN** a fixed point count is needed for a shift action on sequences **DO** substitute a constant sequence into the defining relation **BECAUSE** shift-invariant means constant, so the fixed points are always the solutions of the one-variable version of the relation, and over a field that is usually a quadratic with a computable discriminant

Backtesting

1. First useful cue: the modulus 5 is unrelated to p, so a $\mathbb{Z}/5$ action must exist; the only candidate on sequences is the shift, so test for period 5.
2. Method the expression suggested: divide by a_n to get $a_{n+2} = (1+a_{n+1})/a_n$, iterate five times keeping everything as one fraction, and confirm $(a_6, a_7) = (a_1, a_2)$.
3. Where you would hesitate, and the thought to reinforce: at the claim that the shift maps admissible sequences to admissible sequences. Reinforce that the shifted sequence satisfies the same recurrence indexed one lower, and that its values are the same set of units; one sentence, but a grader will look for it.
4. The trap the question sets: computing $f(p) = (p-2)(p-3)$ and then checking four residue cases without noticing that the "0 or 2" is really the number of roots of a quadratic. The arithmetic route is correct and scores, but it hides the mechanism, and the mechanism is what transfers to the next problem.
5. What makes you recognise it faster next time: three appearances of the number 5 (the modulus, the period, the discriminant) all trace back to the same pentagonal structure. When a number repeats across unrelated-looking parts of a problem, it is one structure seen three times.
6. The reusable decision rule: congruence for a cardinality means a group action of that order; find the group, then count fixed points only.

Mirror drills

1. Let q be prime and let S be the set of strings of length q over an alphabet of size m that are not constant. Show $q \mid |S|$ and deduce Fermat's little theorem. *Answer:* Cyclic rotation gives a $\mathbb{Z}/q$ action whose fixed points are exactly the constant strings, so $m^q \equiv m \pmod{q}$; the same three-slot argument as A3.

2. Show that the recurrence $x_{n+2} = (x_{n+1} + 1)/x_n$ over $\mathbb{Q}$ with $x_1 = 2$, $x_2 = 3$ is periodic, and state the invariant. *Answer:* The orbit is $2, 3, 2, 1, 1, 2, 3, \ldots$ of period 5, and $I(x, y) = \frac{(1+x)(1+y)(1+x+y)}{xy}$ takes the value $\frac{3 \cdot 4 \cdot 6}{6} = 12$ at every consecutive pair.

3. For which primes p does $x^2 \equiv x + 1 \pmod{p}$ have a solution? *Answer:* Exactly when 5 is a quadratic residue modulo p, that is when $p = 5$ or $p \equiv \pm 1 \pmod{5}$; by reciprocity, since $5 \equiv 1 \pmod{4}$, this is the condition $\left(\frac{p}{5}\right) = 1$. These are the primes for which the golden ratio exists in $\mathbb{F}_p$.

4.4 A4. Expected weighted sum before the first ascent

Suppose that $X_1, X_2, \ldots$ are real numbers between 0 and 1 that are chosen independently and uniformly at random. Let $S = \sum_{i=1}^{k} X_i/2^i$, where k is the least positive integer such that $X_k < X_{k+1}$, or $k = \infty$ if there is no such integer. Find the expected value of S.

Read it out loud

Say the structure: this is a sum with a random number of terms, where the random number is a stopping time defined by "the first time something happens". The summands carry deterministic geometric weights 2^{-i}, and the randomness sits in two places at once: in the values X_i and in how many of them survive. That double dependence is the entire difficulty, and it is also the entire hint, because there is exactly one standard device for a sum whose length is random: write the sum as an infinite sum of indicator-weighted terms and use linearity of expectation.

So write

$$S = \sum_{i=1}^{\infty} \frac{X_i}{2^i} \mathbf{1}[\, i \le k \,],$$

and the problem becomes: for each fixed i, compute $E[X_i \mathbf{1}[i \le k]]$. The point of this rewriting is that the sum is now over a fixed index set, and every difficulty has been pushed into a single expectation for each i.

Now read the stopping rule again, carefully, because the second cue is hidden there. The variable k is the least index with $X_k < X_{k+1}$. The event $\{k = i\}$ involves X_{i+1} and is two-sided: a run of descents followed by one ascent. But the event $\{k \ge i\}$ involves only $X_1, \ldots, X_i$ and is one-sided:

$$\{k \ge i\} = \{X_1 \ge X_2 \ge \cdots \ge X_i\},$$

because $k \ge i$ says precisely that none of the first $i-1$ comparisons was an ascent. Almost surely there are no ties, so this is a strictly decreasing run. That is the whole problem: the indicator you need is a tail event, and the tail event of a "first time" stopping rule is always the clean one.

The fields wired in are: elementary probability (linearity, independence), exchangeability of independent identically distributed continuous random variables, order statistics and simplex volumes, the Beta integral, and the exponential series. Within thirty seconds you should have written S as an indicator sum and identified $\{i \le k\}$ as a descending run.

Approaches, ranked worst to best

1. **Simulate or compute the first few cases.** *Cue:* a strange stopping rule. *Circuit:* small cases first. Computing $E[S \mid k = 1]$, $E[S \mid k = 2]$ by hand will produce numbers but no pattern that you can sum, because the conditional distribution of $X_1, \ldots, X_i$ given $k = i$ is an awkward object. Ranked last: it is work without structure.

2. **Condition on the value of k.** *Cue:* "k is the least positive integer such that". *Circuit:* law of total expectation. Write $E[S] = \sum_i P(k = i) E[S \mid k = i]$. This is the reflex most people have, and it is bad here. Conditioning on $\{k = i\}$ imposes both a descending run on $X_1, \ldots, X_i$ and an ascent at the end, so the conditional law of X_i is distorted by information about X_{i+1}, and you must carry a two-sided constraint through every integral. It ranks above blind experiment because it does terminate, but it costs several times the ink.

3. **Full joint density of the descending run.** *Cue:* the event is a region in $[0,1]^i$. *Circuit:* integrate over an explicit simplex. Write $E[X_i \mathbf{1}[i \le k]]$ as an i-fold integral over $\{1 > x_1 > \cdots > x_i > 0\}$ of x_i, and evaluate by iterated integration. Correct, and only slightly more expensive than the ideal route; the iterated integral rebuilds the factorial $(i-1)!$ from scratch each time. Ranked third because it works but recomputes a standard fact.

4. **Indicator sum plus conditioning on the last variable.** *Cue:* a random number of terms with a "first time" rule. *Circuit:* linearity of expectation over a tail event, then exchangeability. Write S as an indicator sum, replace $\{i \le k\}$ by the descending run event, condition on $X_i = t$, and split that event into two independent pieces: "the earlier variables all exceed t", of probability $(1-t)^{i-1}$, and "the earlier variables are in decreasing order", of probability $1/(i-1)!$ by exchangeability. One Beta integral finishes it. This is the ideal route: it never conditions on k, it never writes a multiple integral, and the factorial appears as a symmetry count rather than as the output of a computation.

The ideal solution

The answer is $\boxed{2e^{1/2} - 3}$, approximately 0.29744.

Setup. Since the X_i are independent with a continuous distribution, almost surely no two are equal, and we work on that full-measure event throughout. Extend the sum defining S by zero summands beyond k, so that

$$S = \sum_{i=1}^{\infty} \frac{X_i}{2^i} \mathbf{1}[\, i \le k\,].$$

Every summand is nonnegative and $0 \le S \le \sum_{i \ge 1} 2^{-i} = 1$, so S is a bounded nonnegative random variable and, by Tonelli's theorem (equivalently, monotone convergence for a series of nonnegative terms), expectation may be taken term by term:

$$E[S] = \sum_{i=1}^{\infty} \frac{1}{2^i} E[X_i \mathbf{1}[\, i \le k\,]]\,.$$

Step 1: the indicator is a descending run. By definition k is the least positive integer j with $X_j < X_{j+1}$. Hence $k \ge i$ if and only if no $j \in \{1, \dots, i-1\}$ satisfies $X_j < X_{j+1}$, that is, if and only if $X_1 > X_2 > \cdots > X_i$ (almost surely, using that ties have probability zero). In particular this event depends only on $X_1, \dots, X_i$.

Step 2: the expectation of one summand. Fix $i \ge 1$ and condition on $X_i = t$. The event $\{X_1 > X_2 > \cdots > X_i\}$ is the intersection of

$$A = \{X_1, \dots, X_{i-1} \text{ all lie in } (t, 1)\}, \qquad B = \{X_1 > X_2 > \cdots > X_{i-1}\},$$

and these two events are independent, because A depends only on the unordered set of values $\{X_1, \dots, X_{i-1}\}$ while B depends only on which of the $(i-1)!$ orderings of those values occurs, and for independent identically distributed continuous variables the ordering is uniform on the $(i-1)!$ possibilities and independent of the value set. Thus, given $X_i = t$,

$$P(A) = (1-t)^{i-1}, \qquad P(B) = \frac{1}{(i-1)!}, \qquad P(X_1 > \cdots > X_i \mid X_i = t) = \frac{(1-t)^{i-1}}{(i-1)!}.$$

(For $i = 1$ both factors are empty products equal to 1, and the formula correctly gives probability 1.) Since X_i is uniform on $(0, 1)$,

$$E[X_i \mathbf{1}[i \le k]] = \int_0^1 t\, \frac{(1-t)^{i-1}}{(i-1)!}\, dt = \frac{1}{(i-1)!} \cdot \frac{1!\,(i-1)!}{(i+1)!} = \frac{1}{(i+1)!},$$

where the middle equality is the Beta integral $\int_0^1 t^{a-1}(1-t)^{b-1} dt = \frac{(a-1)!(b-1)!}{(a+b-1)!}$ with $a = 2$, $b = i$; if you prefer to avoid quoting it, substitute $u = 1 - t$ and expand:

$$\int_0^1 t(1-t)^{i-1} dt = \int_0^1 (1-u)u^{i-1} du = \frac{1}{i} - \frac{1}{i+1} = \frac{1}{i(i+1)},$$

and $\frac{1}{(i-1)!\, i(i+1)} = \frac{1}{(i+1)!}$.

Step 3: sum the series. Therefore

$$E[S] = \sum_{i=1}^{\infty} \frac{1}{2^i (i+1)!} = 2 \sum_{i=1}^{\infty} \frac{(1/2)^{i+1}}{(i+1)!} = 2 \sum_{j=2}^{\infty} \frac{(1/2)^j}{j!} = 2 \left(e^{1/2} - 1 - \frac{1}{2} \right) = 2e^{1/2} - 3.$$

The case $k = \infty$ is null. From Step 1, $P(k \geq i) = P(X_1 > \cdots > X_i) = 1/i!$, since by exchangeability all $i!$ orderings of $X_1, \ldots, X_i$ are equally likely. Hence $P(k = \infty) = \lim_{i \to \infty} 1/i! = 0$, so S is defined almost surely, and however S is defined on the null event $\{k = \infty\}$ (the full series converges there anyway) the expectation is unaffected.

Sanity checks. First, $P(k = i) = \frac{1}{i!} - \frac{1}{(i+1)!} = \frac{i}{(i+1)!}$, and $\sum_{i \geq 1} \frac{i}{(i+1)!} = \sum_{i \geq 1} \left(\frac{1}{i!} - \frac{1}{(i+1)!} \right) = 1$ by telescoping, confirming that k is almost surely finite. Second, $E[k] = \sum_{i \geq 1} P(k \geq i) = \sum_{i \geq 1} \frac{1}{i!} = e - 1 \approx 1.718$, so typically only one or two terms survive, and indeed $E[S] \approx 0.297$ sits just below $E[X_1/2] = 1/4$ plus a small contribution from the later terms. Third, the crude bound $E[S] \leq E[X_1]/2 + \sum_{i \geq 2} 2^{-i}/i! = 1/4 + \ldots$ is comfortably consistent.

Layer 1: WHAT

Block 1. Rewrite the random-length sum as an infinite indicator sum $S = \sum_i 2^{-i} X_i \mathbf{1}[i \leq k]$.

Block 2. Justify term-by-term expectation: nonnegative summands, Tonelli.

Block 3. Convert the indicator: $\{i \leq k\}$ is the descending run $\{X_1 > \cdots > X_i\}$, an event in the first i variables only.

Block 4. Condition on $X_i = t$ and factor the run event into "all earlier values exceed t" times "the earlier values are in decreasing order", independent, with probabilities $(1-t)^{i-1}$ and $1/(i-1)!$.

Block 5. Integrate: $E[X_i \mathbf{1}] = \frac{1}{(i-1)!} \int_0^1 t(1-t)^{i-1} dt = \frac{1}{(i+1)!}$.

Block 6. Sum $\sum_i 2^{-i}/(i+1)!$ by shifting the index into the exponential series: $2e^{1/2} - 3$.

Block 7. Check that $k < \infty$ almost surely, so nothing was lost.

Juxtapose with the nominal route. The zombie writes $E[S] = \sum_i P(k = i) E[S \mid k = i]$, computes $P(k = i) = \frac{i}{(i+1)!}$, and then has to find the conditional expectation of a whole descending run given that it ends with an ascent. That requires the joint density of $(X_1, \ldots, X_i)$ restricted to a region cut out by $i - 1$ descent conditions and one ascent condition involving an extra variable, and it produces nested integrals that do not simplify until the very end. The answer is the same; the ink is four times as much and the error probability rises with it.

The bifurcation point is the choice of which event to use: $\{k = i\}$ or $\{k \geq i\}$. Everything else follows. Say it as a rule: for a stopping time defined as a first passage, always work with the tail event, because the tail event is the complement of a conjunction of "nothing happened yet" conditions and therefore involves no extra variable.

Layer 2: HOW

Block 1. **Cue.** A sum whose upper limit is a random variable. **Why natural.** You cannot integrate over a random index set, but you can integrate over a fixed index set with an indicator. Rewriting $\sum_{i=1}^{k}$ as $\sum_{i=1}^{\infty} \mathbf{1}[i \leq k]$ separates the two sources of randomness, and linearity of expectation does not care that the indicators and the summands are dependent. That last point is the one people forget: linearity needs no independence at all. **Link to the next block.** Now each term is a single expectation, and the only question is which event the indicator describes.

Block 2. **Cue.** An infinite sum of nonnegative random variables. **Why natural.** Interchanging an infinite sum with an expectation is an interchange of two limits and needs a licence. For nonnegative terms the licence is free: Tonelli, or equivalently monotone convergence. Writing one clause naming it is the difference between a full-marks script and a script with a hole in it. Do not write "clearly we may interchange". **Link to the next block.** With the interchange justified, compute one term.

Block 3. **Cue.** "k is the least positive integer such that $X_k < X_{k+1}$". **Why natural.** "Least index with a property" always means: the tail event $\{k \geq i\}$ is "the property failed at every earlier index". Here the property is an ascent, so its failure at every earlier index is a descending run. This is the definitional reading of a first passage time, and it is also the reason the event does not touch X_{i+1}: the constraint at index i has not yet been tested. **Link to the next block.** The indicator now depends on exactly the variables $X_1, \ldots, X_i$, so condition on the one you also need, namely X_i.

Block 4. **Cue.** You need $E[X_i \cdot \mathbf{1}(\text{run})]$, a product of a value and an ordering condition. **Why natural.** Condition on the variable that appears in both roles, $X_i = t$. The remaining condition on $X_1, \ldots, X_{i-1}$ splits into a magnitude condition (all bigger than t) and an ordering condition (decreasing), and these are independent for independent identically distributed continuous variables: the unordered set of values and the ordering pattern are independent, with the ordering uniform over $(i-1)!$ possibilities. That independence is exchangeability, and it is the reason a factorial appears without any integration. **Link to the next block.** One clean one-dimensional integral remains.

Block 5. **Cue.** $\int_0^1 t(1-t)^{i-1}dt$. **Why natural.** This is a Beta integral, and even without quoting the Beta function the substitution $u = 1 - t$ turns it into $\int_0^1 (u^{i-1} - u^i)du = \frac{1}{i} - \frac{1}{i+1}$. The resulting collapse $\frac{1}{(i-1)!\,i(i+1)} = \frac{1}{(i+1)!}$ is the signal that you have taken the right route: an answer of the form $\frac{1}{(i+1)!}$ is exactly what can be summed against 2^{-i}. **Link to the next block.** A reciprocal factorial against a geometric weight is the exponential series.

Block 6. **Cue.** $\sum_i \frac{1}{2^i(i+1)!}$, a factorial with a shifted index. **Why natural.** To use $e^x = \sum_j x^j/j!$ the exponent and the factorial must carry the same index. Multiplying and dividing by 2 to convert 2^{-i} into $2^{-(i+1)}$ aligns them, after which the sum is the exponential series with its first two terms removed. Index alignment before summation is a mechanical habit; never try to sum a series whose exponent and factorial disagree. **Link to the next block.** Evaluate and subtract the missing terms.

Block 7. **Cue.** The problem itself flagged the possibility $k = \infty$. **Why natural.** The statement went out of its way to define $k = \infty$ when no ascent occurs, which is an instruction to say something about that case. $P(k \geq i) = 1/i!$ makes it a null event in one line. A grader looks for this, and it costs one sentence. **Link to the next block.** Nothing is lost; the computed expectation is the answer.

Layer 3: WHY

The truths beneath this route are three. First, linearity of expectation holds for arbitrarily dependent summands, which is why the indicator decomposition is legitimate and why it is the universal tool for random-length sums; the analogous identity $E[N] = \sum_{i \geq 1} P(N \geq i)$ for nonnegative integer N is the same idea with $X_i \equiv 1$. Second, exchangeability: for independent identically distributed continuous random variables, the vector of ranks is uniform on the symmetric group and independent of the vector of order statistics. That single fact is what makes $P(X_1 > \cdots > X_i) = 1/i!$ a symmetry count rather than an integral, and it is what lets you factor the descending run event into an independent magnitude part and ordering part. Third, the volume of the simplex $\{t < x_{i-1} < \cdots < x_1 < 1\}$ is $(1-t)^{i-1}/(i-1)!$, which is the same factorial appearing geometrically.

The discernment mechanism: when a probability problem contains a stopping rule, immediately classify it as a first passage and write down its tail event. Then ask whether the quantity you want is a sum, and if so, decompose it by index and use linearity. The combination "first passage plus linearity of expectation over indicators" solves a large fraction of all Putnam probability problems, because it is the only method that does not require knowing the joint law of the stopped sequence.

To train it, take any problem containing the words "the first time", rewrite the target as a sum of indicator-weighted terms, and force yourself to express every indicator using only tail events

$\{k \geq i\}$. If an event you have written mentions a variable with index greater than i, you have used $\{k = i\}$ somewhere and should back up.

What the examiners nudged with: the weights 2^{-i} are there so that the final series is summable against a reciprocal factorial and produces a clean closed form with $e^{1/2}$, and the explicit mention of $k = \infty$ is a deliberate invitation to check that the event is null. The choice of the uniform distribution on $[0, 1]$ is what makes $P(A) = (1-t)^{i-1}$ so clean; the same method works for any continuous distribution, with $(1 - F(t))^{i-1}$ in place of $(1-t)^{i-1}$, which is a good way to see that only continuity and independence were used.

Revolving concepts

Linearity of expectation and the indicator decomposition. For any random variables $Y_i \geq 0$ and any event structure, $E\left[\sum_i Y_i\right] = \sum_i E[Y_i]$ by Tonelli, with no independence required. The special case $E[N] = \sum_{i \geq 1} P(N \geq i)$ for a nonnegative integer valued N is the tail-sum formula, and the continuous analogue is $E[Y] = \int_0^\infty P(Y > y)\, dy$. See Book 1, field: Probability, topic: linearity and tail sums.

Exchangeability and ranks. If $X_1, \ldots, X_n$ are independent and identically distributed with a continuous distribution, then almost surely they are distinct, all $n!$ orderings are equally likely, and the rank vector is independent of the order statistics. Consequences: $P(X_1 > \cdots > X_n) = 1/n!$; the probability that the maximum occurs at a specified index is $1/n$; the expected number of records among n such variables is $\sum_{j=1}^{n} 1/j$. See Book 1, field: Probability, topic: exchangeability, ranks and records.

Iterated integration and the simplex, with the recursion generalised. Define $F_i(t) = P(X_1 > X_2 > \cdots > X_{i-1} > t)$ for uniform variables on $[0, 1]$. Then $F_1(t) = 1$ and

$$F_{i+1}(t) = \int_t^1 F_i(u)\, du,$$

because the smallest of the run must exceed t and the rest form a run above it. This recursion is solved by $F_i(t) = \frac{(1-t)^{i-1}}{(i-1)!}$, proved by induction: repeated integration of a power raises the exponent and divides by the new exponent, generating the factorial. The general statement is that the m-fold iterated integral of the constant 1 over a nested region is the volume of a simplex, $\frac{(1-t)^m}{m!}$, and more generally the m-fold iterated integral of g equals $\frac{1}{(m-1)!}\int_t^1 (u-t)^{m-1} g(u) du$ (Cauchy's formula for repeated integration). This is the general recursion behind every factorial that appears in an ordering problem. See Book 1, field: Probability, topic: order statistics; and field: Calculus and integrals, topic: repeated integration.

The Beta integral. $\int_0^1 t^{a-1}(1-t)^{b-1} dt = \frac{\Gamma(a)\Gamma(b)}{\Gamma(a+b)}$, which for positive integers reads $\frac{(a-1)!(b-1)!}{(a+b-1)!}$. In this problem it is the case $a = 2$, and it can be avoided by one substitution, but knowing it converts a two-line computation into a one-line one. See Book 1, field: Calculus and integrals, topic: Beta and Gamma functions.

Exponential series with shifted index. $\sum_{j \geq 0} x^j/j! = e^x$; to evaluate $\sum_{i \geq 1} x^i/(i+1)!$ multiply by x and shift: $x \sum_{i \geq 1} \frac{x^i}{(i+1)!} = \sum_{j \geq 2} \frac{x^j}{j!} = e^x - 1 - x$. Here $x = 1/2$ and the prefactor is $1/x = 2$. See Book 1, field: Sequences and recurrences, topic: manipulating standard series by index shifting.

Circuit Bank

- **WHEN** you see a sum with a random number of terms **DO** rewrite it as $\sum_{i=1}^{\infty} (\text{term}_i) \mathbf{1}[i \leq k]$ and apply linearity of expectation **BECAUSE** linearity needs no independence, so the dependence between the indicator and the term is harmless, and the index set becomes fixed
- **WHEN** you see a stopping time defined as "the least index such that P holds" **DO** work with the tail event $\{k \geq i\} = \{\text{P failed at indices } 1, \ldots, i-1\}$, never with $\{k = i\}$ **BECAUSE** the

tail event involves only the variables up to index i, whereas $\{k = i\}$ drags in the next variable and distorts every conditional law

- **WHEN** you need the expectation of a value times an ordering indicator **DO** condition on the variable that appears in both, then split the remaining event into a magnitude part and an ordering part **BECAUSE** for independent identically distributed continuous variables these two parts are independent, so the factorial appears as a symmetry count and no multiple integral is needed
- **WHEN** you must interchange an infinite sum and an expectation **DO** check nonnegativity and name Tonelli or monotone convergence in one clause **BECAUSE** the interchange is a limit exchange and is not automatic; for nonnegative terms it is free, and naming it costs one line and secures the marks
- **WHEN** a series has x^i against $(i + c)!$ **DO** multiply by x^c to align the exponent with the factorial, then subtract the missing initial terms of the exponential series **BECAUSE** closed forms only appear once the index of the power and the index of the factorial agree, so alignment is a mechanical prerequisite, not an insight

Backtesting

1. First useful cue: the sum has a random upper limit, so write it as an indicator sum over a fixed index set.
2. Method the expression suggested: identify $\{i \leq k\}$ as the descending run $\{X_1 > \cdots > X_i\}$, condition on $X_i = t$, factor into magnitude and ordering, integrate, sum.
3. Where you would hesitate, and the thought to reinforce: at the independence of the magnitude event and the ordering event. Reinforce that for independent identically distributed continuous variables, the set of values and the ordering pattern are independent, with all $(i - 1)!$ orders equally likely; this is exchangeability, and it is the single most useful fact in Putnam probability.
4. The trap the question sets: conditioning on $\{k = i\}$, which entangles X_{i+1} and turns a one-line computation into nested integrals. The second trap is ignoring the $k = \infty$ clause, which the problem deliberately inserted.
5. What makes you recognise it faster next time: the phrase "the least positive integer such that" is the signature of a first passage time, and a first passage time is always handled through its tail event.
6. The reusable decision rule: random-length sum means indicator decomposition; first passage means tail event; ordering condition means exchangeability.

Mirror drills

1. With $X_1, X_2, \ldots$ independent uniform on $[0, 1]$ and k as in the problem, find $E[k]$. *Answer:* $E[k] = \sum_{i \geq 1} P(k \geq i) = \sum_{i \geq 1} \frac{1}{i!} = e - 1$. Same tail event, tail-sum formula instead of a weighted sum.
2. Let N be the least n such that $X_1 + \cdots + X_n > 1$, for independent uniforms on $[0, 1]$. Show $E[N] = e$. *Answer:* $P(N > n) = P(X_1 + \cdots + X_n \leq 1) = \frac{1}{n!}$, the volume of the standard simplex, so $E[N] = \sum_{n \geq 0} P(N > n) = \sum_{n \geq 0} \frac{1}{n!} = e$. The same simplex factorial, arrived at by volume rather than by ordering.
3. Replace the uniform distribution by an arbitrary continuous distribution with cumulative distribution function F, and compute $E\left[\sum_{i=1}^{k} X_i / 2^i\right]$ in terms of F. *Answer:* The identical

argument gives $\sum_{i\geq 1} \frac{1}{2^i(i-1)!} \int_{-\infty}^{\infty} t\,(1-F(t))^{i-1} dF(t)$; only continuity and independence were used, and the uniform case is $F(t) = t$ on $[0,1]$.

4.5 A5. Domino placing game on 2022 squares

Alice and Bob play a game on a board consisting of one row of 2022 consecutive squares. They take turns placing tiles that cover two adjacent squares, with Alice going first. By rule, a tile must not cover a square that is already covered by another tile. The game ends when no tile can be placed according to this rule. Alice's goal is to maximize the number of uncovered squares when the game ends; Bob's goal is to minimize it. What is the greatest number of uncovered squares that Alice can ensure at the end of the game, no matter how Bob plays?

Read it out loud

Read the position, not the board. At any moment the uncovered squares form a disjoint union of maximal runs, and a run of length ℓ is the only thing that matters about it: a tile placed in that run splits it into two runs of lengths i and $\ell - 2 - i$ for some $0 \le i \le \ell - 2$, and runs never interact. So a position is a finite multiset of nonnegative integers, a move replaces one entry $\ell \ge 2$ by two entries summing to $\ell - 2$, and the game ends when every entry is 0 or 1. The final score is the number of entries equal to 1.

Say the second thing aloud: this is a *scoring* game on a disjunctive sum, not a normal-play game. Nobody is trying to make the last move; the players are trying to control a number. That single sentence rules out the entire Sprague-Grundy apparatus, which computes who wins a last-move game and says nothing about scores. It also warns you that positions do not simply add: the value of a two-component position is not the sum of the component values, because the players choose where to move.

Third, two reformulations that cost nothing and buy a lot. Each tile covers exactly two squares, so if d tiles are placed the score is $2022 - 2d$: Alice is trying to make the game *short* and Bob to make it *long*. And the final score equals the number of odd-length runs at the end, because a terminal run has length 0 or 1; moreover splitting ℓ into i and $\ell - 2 - i$ preserves parity of the total, so an odd run always produces exactly one odd run and an even run produces either two odd runs or none. Hence the number of odd runs never decreases, it increases by exactly 2 when a player splits an even run into two odd pieces, and the final score is the initial number of odd runs plus twice the number of such splits. Alice wants those splits, Bob wants to avoid them, and Bob can always avoid them: from any run of length $\ell \ge 2$ he may play $(0, \ell - 2)$, which preserves every parity.

The fields wired in are combinatorial game theory (scoring games, disjunctive sums, the notion of a switch and its mean value and temperature), induction on a well-founded order, and elementary number theory modulo 7. Within thirty seconds you should have written the position as a multiset, the move as $\ell \mapsto (i, \ell - 2 - i)$, and the two value functions $A(P)$ (Alice to move) and $B(P)$ (Bob to move). Then you should compute small cases, because the periodicity you are going to need is not something you can guess.

Approaches, ranked worst to best

1. **Symmetry or mirroring strategy.** *Cue:* 2022 is even and the row is symmetric. *Circuit:* pairing strategy for the second player. Mirroring works in normal-play games where the goal is to move last. Here it controls nothing: Bob mirroring Alice's split produces exactly the odd runs Alice wanted. Ranked last because it is the reflex that costs the most time before revealing that it answers a different question.

2. **Sprague-Grundy values of the octal game.** *Cue:* splitting a row by placing dominoes is the game Dawson's Kayles. *Circuit:* nim-value tables and periodicity. The Grundy theory of this move rule is genuinely classical and genuinely periodic, and it is completely useless here: Grundy values decide who makes the last move, and the question asks for a score. Ranked

here because it is a real theory misapplied, which is more dangerous than an obviously wrong idea.

3. **Single-interval recursion with a following strategy.** *Cue:* the game starts as one interval. *Circuit:* define $a(n)$, $b(n)$ for one interval, then bound them. Define $a(n)$ and $b(n)$ as the values of a single interval of length n with Alice, respectively Bob, to move, compute them for $n \leq 6$, and prove $a(n) = a(n-7)+1$, $b(n) = b(n-7)+1$ by having the second player answer in whichever component the first player just used. This is the standard route and it does reach 290. Its cost is that the following strategy breaks down whenever a component runs out, so you must allow a player to pass, and then you must prove that a pass is worthless, which is the statement that having the move never hurts. That lemma is not free.

4. **Closed form for every position via mean value and temperature.** *Cue:* positions are multisets and values must be computed for all of them anyway. *Circuit:* write each component as a switch: a fixed mean plus or minus one, and let the parity of the number of hot components decide. Assign to each run length n a mean $\pi(n)$ and a temperature $h(n) \in \{0,1\}$, and prove that the value of any position is $\sum \pi(n_i)$ shifted by $+1$ or -1 according to whether the number of hot components is odd, with the sign favouring whoever is to move. The whole induction then rests on one inequality about π, which has a two-line proof. This is the ideal route: it solves every position rather than every interval, it needs no pass convention, and it makes the answer 290 a one-line evaluation.

The ideal solution

The answer is $\boxed{290}$.

Formalisation. A position is a finite multiset $P = \{n_1, \ldots, n_r\}$ of nonnegative integers (the lengths of the maximal uncovered runs; entries equal to 0 may be kept or discarded, it makes no difference below). A move replaces some entry $n \geq 2$ by two entries i and j with $i, j \geq 0$ and $i + j = n - 2$. A position with all entries at most 1 is terminal, and its score is the number of entries equal to 1. Each move decreases $\Sigma(P) := \sum_i n_i$ by exactly 2, so the game terminates and we may define, by induction on $\Sigma(P)$, the optimal scores

$$A(P) = \max_{\mu} B(P^{\mu}), \qquad B(P) = \min_{\mu} A(P^{\mu}),$$

over all moves μ, with $A(P) = B(P) = \text{score}(P)$ at terminal positions. Here A is the value with Alice (the maximiser) to move and B with Bob (the minimiser) to move. The answer to the problem is $A(\{2022\})$.

The two statistics. For an integer $n \geq 0$ write $n = 7q + r$ with $0 \leq r \leq 6$ and define

$$\pi(n) = q + \pi_r, \qquad h(n) = h_r,$$

where the residue tables are

r	0	1	2	3	4	5	6
π_r	0	1	0	1	1	1	1
h_r	0	0	0	0	1	0	1

Call n *hot* if $h(n) = 1$, that is if $n \equiv 4$ or $6 \pmod 7$. For a position put

$$\Pi(P) = \sum_i \pi(n_i), \qquad t(P) = \#\{i : n_i \text{ hot}\}, \qquad \epsilon(P) = t(P) \bmod 2 \in \{0,1\}.$$

Theorem 4.5.1. *For every position P,*

$$A(P) = \Pi(P) + \epsilon(P), \qquad B(P) = \Pi(P) - \epsilon(P).$$

Before the proof, note the answer it gives: $2022 = 7 \cdot 288 + 6$, so $\pi(2022) = 288 + 1 = 289$ and 2022 is hot, whence $t = 1$, $\epsilon = 1$ and

$$A(\{2022\}) = 289 + 1 = 290.$$

We need two lemmas about the single function π. Write $c_r := 7\pi_r - r$, so that

$$(c_0, c_1, c_2, c_3, c_4, c_5, c_6) = (0,\ 6,\ -2,\ 4,\ 3,\ 2,\ 1), \qquad \pi(n) = \frac{n + c_r}{7} \quad (r = n \bmod 7),$$

and observe that every c_r lies in the interval $[-2, 6]$.

Lemma 4.5.2 (Almost additivity)**.** *If $i, j \geq 0$ and $i + j = n - 2$, then $\Delta\Pi := \pi(i) + \pi(j) - \pi(n) \in \{-1, 0, 1\}$.*

Proof. Let s, u, r be the residues of i, j, n modulo 7. Then

$$7\,\Delta\Pi = (i + c_s) + (j + c_u) - (n + c_r) = (i + j - n) + c_s + c_u - c_r = c_s + c_u - c_r - 2.$$

Since each c lies in $[-2, 6]$ we get $c_s + c_u \in [-4, 12]$ and $-c_r \in [-6, 2]$, so $7\Delta\Pi \in [-12, 12]$. A multiple of 7 in that range is -7, 0 or 7, so $\Delta\Pi \in \{-1, 0, 1\}$. □

Lemma 4.5.3 (Parity)**.** *For every $n \geq 0$, $\pi(n) + h(n) \equiv n \pmod 2$. Consequently, if $i + j = n - 2$ and δ denotes the parity of $h(i) + h(j) + h(n)$, then $\Delta\Pi \equiv \delta \pmod 2$, and therefore, by the previous lemma,*

$$|\Delta\Pi| \leq \delta \in \{0, 1\}.$$

Proof. From the tables, $\pi_r + h_r$ takes the values $0, 1, 0, 1, 2, 1, 2$ for $r = 0, \ldots, 6$, and these are congruent to r modulo 2 in every case. Hence $\pi(n) + h(n) = q + \pi_r + h_r \equiv q + r \equiv 7q + r = n \pmod 2$. For the consequence, add the three instances:

$$\Delta\Pi = \big(\pi(i) + \pi(j) - \pi(n)\big) \equiv \big(i + h(i)\big) + \big(j + h(j)\big) - \big(n + h(n)\big) \equiv (i + j - n) + \delta \equiv \delta \pmod 2,$$

using $i + j - n = -2$ and $-h(n) \equiv h(n)$. Finally, if $\delta = 0$ then $\Delta\Pi$ is an even element of $\{-1, 0, 1\}$, hence 0; if $\delta = 1$ then $|\Delta\Pi| \leq 1 = \delta$ by the previous lemma. □

Lemma 4.5.4 (Attainment)**.** *Let $n \geq 2$. Then:*

(a) the split $(1, n - 3)$ if $n \geq 3$, and the split $(0, 0)$ if $n = 2$, has $\Delta\Pi \geq 0$;

(b) the split $(0, n - 2)$ has $\Delta\Pi \leq 0$;

(c) if n is hot then the split $(1, n - 3)$ has $\Delta\Pi = 1$;

(d) if n is hot then the split $(2, n - 4)$ has $\Delta\Pi = -1$.

Proof. All four are evaluations of $7\Delta\Pi = c_s + c_u - c_r - 2$ from the proof of the first lemma, where $r = n \bmod 7$.

For (b), $s = 0$ and $u \equiv r - 2$, so $7\Delta\Pi = c_0 + c_{r-2} - c_r - 2$, which for $r = 0, 1, \ldots, 6$ equals $0, -7, 0, 0, -7, 0, 0$ respectively; hence $\Delta\Pi \in \{0, -1\}$.

For (a) and (c) with $n \geq 3$, $s = 1$ and $u \equiv r - 3$, so $7\Delta\Pi = c_1 + c_{r-3} - c_r - 2$, which for $r = 0, 1, \ldots, 6$ equals $7, 0, 7, 0, 7, 0, 7$; hence $\Delta\Pi = 1$ when r is even and 0 when r is odd, and in particular $\Delta\Pi \geq 0$ always and $\Delta\Pi = 1$ when $r \in \{4, 6\}$. For $n = 2$ the only split is $(0, 0)$, giving $7\Delta\Pi = c_0 + c_0 - c_2 - 2 = 0$.

For (d), a hot n has $r \in \{4, 6\}$ and hence $n \geq 4$, so the split $(2, n - 4)$ exists; here $s = 2$ and $u \equiv r - 4$, so $7\Delta\Pi = c_2 + c_{r-4} - c_r - 2$, which equals $-2 + c_0 - c_4 - 2 = -7$ for $r = 4$ and $-2 + c_2 - c_6 - 2 = -7$ for $r = 6$; hence $\Delta\Pi = -1$. □

Proof of the Theorem. Induct on $\Sigma(P)$. If P is terminal then every entry is 0 or 1, and $\pi(0) = h(0) = 0$, $\pi(1) = 1$, $h(1) = 0$; so $\Pi(P)$ is the number of entries equal to 1, which is the score, while $t(P) = 0$ and $\epsilon(P) = 0$. Hence $A(P) = B(P) = \Pi(P)$, as claimed.

Now let P be non-terminal, so some entry $n \geq 2$ exists. A move μ replaces such an n by i, j with $i + j = n - 2$; write $\Delta\Pi$ and δ as above for that move. Then $\Pi(P^\mu) = \Pi(P) + \Delta\Pi$ and $t(P^\mu) = t(P) + h(i) + h(j) - h(n)$, so $\epsilon(P^\mu) = (\epsilon(P) + \delta) \bmod 2$. Write $\Pi = \Pi(P)$ and $\epsilon = \epsilon(P)$. By the inductive hypothesis,

$$B(P^\mu) = \Pi + \Delta\Pi - \epsilon(P^\mu), \qquad A(P^\mu) = \Pi + \Delta\Pi + \epsilon(P^\mu).$$

Alice to move, $\epsilon = 0$. Here $\epsilon(P^\mu) = \delta$, so $B(P^\mu) = \Pi + \Delta\Pi - \delta \leq \Pi$ by the Parity Lemma, with equality exactly when $\Delta\Pi = \delta$, that is (again by the Parity Lemma) exactly when $\Delta\Pi \geq 0$. Part (a) of the Attainment Lemma supplies such a move. Hence $A(P) = \Pi = \Pi + \epsilon$.

Alice to move, $\epsilon = 1$. Here $\epsilon(P^\mu) = 1 - \delta$, so $B(P^\mu) = \Pi + \Delta\Pi + \delta - 1 \leq \Pi + 2\delta - 1 \leq \Pi + 1$, using $\Delta\Pi \leq \delta \leq 1$. Equality forces $\delta = 1$ and $\Delta\Pi = 1$. Since $t(P)$ is odd, P contains a hot entry n, and part (c) gives a split of it with $\Delta\Pi = 1$; then $\delta \equiv \Delta\Pi \equiv 1$ automatically. Hence $A(P) = \Pi + 1 = \Pi + \epsilon$.

Bob to move, $\epsilon = 0$. Here $A(P^\mu) = \Pi + \Delta\Pi + \delta \geq \Pi$, with equality exactly when $\Delta\Pi \leq 0$; part (b) supplies such a move, so $B(P) = \Pi = \Pi - \epsilon$.

Bob to move, $\epsilon = 1$. Here $A(P^\mu) = \Pi + \Delta\Pi + 1 - \delta \geq \Pi + 1 - 2\delta \geq \Pi - 1$, with equality forcing $\delta = 1$ and $\Delta\Pi = -1$; since $t(P)$ is odd there is a hot entry, and part (d) supplies such a split. Hence $B(P) = \Pi - 1 = \Pi - \epsilon$.

This completes the induction. □

Reading the answer off. For the original board, $P = \{2022\}$, and $A(\{2022\}) = \pi(2022) + 1 = 289 + 1 = 290$. Alice's strategy is explicit: she always splits a hot run n as $(1, n - 3)$, and after any Bob move she does the same in whichever run is then hot; Bob's optimal replies are $(0, n - 2)$ from any run when the number of hot runs is even, and $(2, n - 4)$ from a hot run otherwise.

Consistency checks. For a single interval the theorem gives $A(\{n\}) = \pi(n) + h(n)$ and $B(\{n\}) = \pi(n) - h(n)$, that is, in terms of $r = n \bmod 7$,

$$A(\{n\}) = \left\lfloor \tfrac{n}{7} \right\rfloor + (0, 1, 0, 1, 2, 1, 2)_r, \qquad B(\{n\}) = \left\lfloor \tfrac{n}{7} \right\rfloor + (0, 1, 0, 1, 0, 1, 0)_r,$$

which one can verify directly for $n \leq 6$: for example $A(\{4\}) = 2$ because Alice splits 4 as $(1, 1)$, while $B(\{4\}) = 0$ because Bob splits 4 as $(0, 2)$ and the remaining 2 is forced. Also $A(\{n\}) \equiv B(\{n\}) \equiv n \pmod 2$, as the Parity Lemma requires and as the covering argument demands, since every tile removes two squares.

Layer 1: WHAT

Block 1. Abstract the board into a multiset of run lengths; a move is $n \mapsto (i, n - 2 - i)$; terminal score is the number of 1s.

Block 2. Note the game is a scoring game on a disjunctive sum, so component values do not simply add and Grundy theory does not apply.

Block 3. Compute small single-interval values and observe the period 7 pattern, then guess the correct general statistics: a mean π and a hot indicator h.

Block 4. State the closed form $A = \Pi + \epsilon$, $B = \Pi - \epsilon$ for every position, where ϵ is the parity of the number of hot components.

Block 5. Prove almost additivity of π, by writing $\pi(n) = (n + c_{n \bmod 7})/7$ with all $c_r \in [-2, 6]$: the difference is a multiple of 7 trapped in $[-12, 12]$.

Block 6. Prove the parity lemma $\pi + h \equiv n$, which upgrades $|\Delta\Pi| \leq 1$ to $|\Delta\Pi| \leq \delta$.

Block 7. Exhibit the four attainment splits $(1, n - 3)$, $(0, n - 2)$, $(2, n - 4)$ and run the induction, then evaluate at $n = 2022$.

Juxtapose with the nominal route. The zombie fixes on the single interval, defines $a(n)$ and $b(n)$ only for intervals, and then discovers that after the first move he has two intervals and no theory for them. He patches this with a following strategy, discovers it breaks when one interval is exhausted, patches that with a pass convention, and then owes a proof that passing is worthless. Every patch is forced by the original decision to define values only for intervals.

The bifurcation point is the moment you decide what object to compute a value for. If you compute values for intervals, you are permanently fighting the disjunctive sum. If you compute values for all positions, the sum structure works for you, because the value of a position turns out to be a sum of per-component numbers plus a one-bit correction.

Layer 2: HOW

Block 1. **Cue.** "A tile must not cover a square already covered" and the board is a row. **Why natural.** Covered squares partition the uncovered ones into runs, and a tile lives inside exactly one run. Nothing about a run matters except its length, and nothing about the arrangement of runs matters at all. Reducing to a multiset is the definition-level move: two positions with the same multiset have identical game trees. **Link to the next block.** A move on a multiset is a split, so the game is a splitting game.

Block 2. **Cue.** The goal is a number, not the last move. **Why natural.** Combinatorial game theory splits sharply into normal play (last move wins, Sprague-Grundy applies) and scoring play (Grundy values are meaningless). Naming which one you are in prevents twenty minutes spent computing nim-values for Dawson's Kayles, which is a real and famous table that answers a different question. **Link to the next block.** In scoring play the right tools are per-component mean values and temperatures.

Block 3. **Cue.** You cannot guess the answer for 2022 without data. **Why natural.** Small case computation here is not fishing, it is the only way to discover the modulus 7. Compute $A(\{n\})$ and $B(\{n\})$ for $n \leq 15$ by hand or by systematic recursion, notice $A(\{n\}) = A(\{n-7\}) + 1$ and the residue patterns $(0,1,0,1,2,1,2)$ and $(0,1,0,1,0,1,0)$, and then read them as mean plus or minus temperature: $\pi = (A+B)/2$, $h = (A-B)/2$. That last step is the important one, because π and h are the objects that behave well under disjunctive sums, whereas A and B are not. **Link to the next block.** A mean and a $0/1$ temperature per component is exactly the data of a switch, and sums of switches are governed by parity.

Block 4. **Cue.** Each component is worth $\pi \pm 1$ depending on who takes it. **Why natural.** A component with $h = 1$ behaves like a two-sided option: whoever moves in it first swings the value by one in his own favour. In a sum of such switches the players alternately claim them, so the maximiser claims $\lceil t/2 \rceil$ of the t hot components and the minimiser claims the rest, and the net effect is exactly ± 1 according to the parity of t. Guessing the closed form $A = \Pi + \epsilon$ is therefore not a leap; it is the standard behaviour of a sum of switches of equal temperature. **Link to the next block.** A guessed closed form must now be proved by induction, and the induction needs to know how π and h change under a split.

Block 5. **Cue.** π is a floor function plus a bounded residue correction. **Why natural.** Any function of the form $\lfloor n/m \rfloor + (\text{table in } n \bmod m)$ can be written as $(n + c_{n \bmod m})/m$ with c bounded, and then differences of π along a relation like $i + j = n - 2$ are multiples of $1/m$ with an explicit bound. The trap of doing casework over the 49 residue pairs is avoided entirely: a multiple of 7 in $[-12, 12]$ has only three possible values. This is the laziness principle at its purest, and it is the technical heart of the solution. **Link to the next block.** Almost additivity gives $|\Delta\Pi| \leq 1$, which is not yet enough; you need to know when $\Delta\Pi$ can be nonzero.

Block 6. **Cue.** The score is always congruent to n modulo 2. **Why natural.** Every tile removes two squares, so the parity of the number of uncovered squares is invariant. In the language of the statistics this says $\pi + h \equiv n$, and combining it with $|\Delta\Pi| \leq 1$ gives $|\Delta\Pi| \leq \delta$: a split that does not change the parity of the hot count cannot change Π at all. That is the sharp statement the induction consumes. **Link to the next block.** Bounds in hand, produce the moves that attain them.

Block 7. **Cue.** Bounds without witnesses prove only one inequality. **Why natural.** Every optimal value statement needs both a bound over all moves and a move attaining it, and the attaining moves should be simple and uniform: split off one square, split off nothing, split off a pair. Checking these three splits is seven evaluations each of $c_s + c_u - c_r - 2$, which is arithmetic, not casework on the game. **Link to the next block.** The induction closes, and $2022 \equiv 6 \pmod 7$ gives 290.

Layer 3: WHY

The deep reason the answer has the form it does is that this game is, in the language of combinatorial game theory, a sum of *switches* all of temperature 1. A component of length n is worth $\pi(n)$ to both players if it is cold ($h = 0$), and worth $\pi(n) + 1$ to whoever moves in it first if it is hot ($h = 1$). In a sum of switches of equal temperature the correct play is greedy: take a hot component, because the value of doing so is exactly the temperature and nothing later can be worth more. The players alternate taking hot components, so the maximiser gains one net unit precisely when the number of hot components is odd. Everything about the 7 is contained in the two residue tables; everything about the game dynamics is contained in the word "switch".

Why 7? Honestly: because the small-case computation says so. There is no shortcut that produces 7 from the rules without computing; what the theory provides is the certainty that some such periodic pattern must exist (the value function of an interval grows linearly with a bounded correction, so it is eventually periodic in n modulo something) and the method to verify a guessed pattern in a few lines once you have it. Do not pretend the 7 was deduced. Name it as a computed input, verify it rigorously, and move on. That is what a knowledge gap looks like when handled honestly.

The discernment mechanism: in any scoring game on a disjunctive sum, do not compute values of positions; compute two numbers per component, the mean and the temperature, and then ask how the temperatures combine. If all temperatures are equal, parity decides. If they differ, the greedy rule "move in the hottest component" is the standard heuristic and is provably correct for sums of switches. Train it by taking any splitting or removal game with a numerical payoff, computing A and B for small components, forming $\pi = (A + B)/2$ and $h = (A - B)/2$, and looking for a pattern in π and h separately. Patterns in A and B separately are usually invisible; patterns in π and h usually are not.

What the examiners nudged with: they chose 2022, and $2022 = 7 \cdot 288 + 6$, so the residue is 6, one of the two hot residues, which guarantees that the answer is not simply $\lfloor n/7 \rfloor$ and that a solver who found the period but not the residue table will be off by two. They also phrased the goals asymmetrically ("Alice's goal is to maximize, Bob's to minimize") to make clear that this is a scoring game and to warn you off last-move theory.

Revolving concepts

Scoring games versus normal play. In normal play the last player to move wins, and every position has a Sprague-Grundy value; the value of a disjunctive sum is the exclusive-or of the values. In scoring play the payoff is a number and Grundy theory does not apply; the relevant invariants are the pair of optimal values (A, B) with each player to move, from which one forms the mean $(A + B)/2$ and the temperature $(A - B)/2$. See Book 1, field: Games and invariants, topic: scoring games, means and temperatures.

Switches and sums of switches. A switch is a position worth $m + \tau$ to the player who moves there first and $m - \tau$ to the other, with mean m and temperature τ. In a sum of switches with equal temperature τ, the value with the maximiser to move is $\sum m_i + \tau \cdot (t \bmod 2)$ where t is the number of switches with $\tau > 0$. That is precisely the theorem proved above, with $\tau = 1$. See Book 1, field: Games and invariants, topic: switches and the greedy hot-move principle.

Splitting games and their move rule. A move $n \mapsto (i, n-2-i)$ with $i + j = n - 2$ is the octal game move rule of Dawson's Kayles; in normal play its Grundy sequence is eventually periodic with period 34. Knowing this and knowing it is irrelevant here is worth more than not knowing it. See Book 1, field: Games and invariants, topic: octal games and Grundy periodicity.

Parity invariants. Every tile covers two squares, so the number of uncovered squares is invariant modulo 2 throughout; equivalently the number of odd runs never decreases and changes only by $+2$. The general principle: find a quantity that every legal move changes in a restricted way, and it will bound the outcome. Here it gives the free bound "the final score is at least the current number of odd runs". See Book 1, field: Games and invariants, topic: monovariants and parity.

Writing $\lfloor n/m \rfloor + f(n \bmod m)$ as $(n + c_{n \bmod m})/m$, and the recursion it encodes. Any function of this shape satisfies the recursion $\pi(n) = \pi(n-m) + 1$, whose general solution is exactly $\lfloor n/m \rfloor$ plus an arbitrary function of the residue; the constants $c_r = mf(r) - r$ are bounded, and differences along any linear relation among the arguments are then multiples of $1/m$ with an explicit range. The general lesson: a quantity satisfying $\pi(n) = \pi(n-m) + 1$ should be linearised in exactly this way before you compare it at several arguments, because it converts m^2 residue cases into one interval bound. See Book 1, field: Number theory, topic: floor functions and periodic corrections.

Circuit Bank

- **WHEN** you see a combinatorial game whose payoff is a number rather than who moves last **DO** compute the pair of values (A, B) for small components and immediately form the mean $(A + B)/2$ and the temperature $(A - B)/2$ **BECAUSE** means add over disjunctive sums and temperatures decide by parity, whereas A and B individually obey no addition rule at all, which is why raw value tables look patternless
- **WHEN** a game position decomposes into non-interacting components **DO** define the value function on multisets of components, never only on the starting shape **BECAUSE** a theory defined only on single components forces following strategies, pass conventions and extra lemmas, all of which are patches for a bad choice of state space
- **WHEN** you must compare $\lfloor n/m \rfloor + f(n \bmod m)$ at several related arguments **DO** write it as $(n + c_{n \bmod m})/m$ with $c_r = mf(r) - r$ and bound the c_r **BECAUSE** the comparison becomes a multiple of m trapped in a small interval, which replaces m^2 residue cases by one line
- **WHEN** you have proved a bound on every move **DO** produce explicit attaining moves, uniform in the parameter, before declaring the value **BECAUSE** an optimal-value claim is two inequalities; bounds without witnesses prove only that the value is at most (or at least) the claimed one, which is half the marks
- **WHEN** every move changes some count by a fixed amount **DO** name the invariant and read off the free bound it gives **BECAUSE** here each tile removes two squares, so the score parity is fixed and the number of odd runs is a monovariant, which immediately gives the lower bound that Bob can never beat

Backtesting

1. First useful cue: the board decomposes into independent runs, so the state is a multiset and the game is a disjunctive sum with a numerical payoff.

2. Method the expression suggested: compute (A, B) for small single runs, convert to mean and temperature, guess the closed form for all positions, prove it by induction on the total length.
3. Where you would hesitate, and the thought to reinforce: at the temptation to define values only for a single interval. Reinforce that in a splitting game the very first move creates two components, so any theory not defined on multisets will need repairs; choose the larger state space at the start.
4. The trap the question sets: assuming the value of a sum is the sum of the values. It is not: with three runs of length 4, the value with Bob to move is 2, not the sum of the individual Bob values. Only the means add; the temperatures interact through parity.
5. What makes you recognise it faster next time: a scoring game with independent components is a sum of switches; look for the temperature first and the mean second.
6. The reusable decision rule: for scoring games, tabulate (A, B) for small components, take half the sum and half the difference, and look for a pattern in those two sequences rather than in A and B.

Mirror drills

1. Compute $A(\{4\})$ and $B(\{4\})$ directly and check them against the formula. *Answer:* Alice splits 4 as $(1, 1)$ and scores 2; Bob splits 4 as $(0, 2)$, the forced reply exhausts the 2, and the score is 0. The formula gives $\pi(4) = 1$, $h(4) = 1$, so $A = 2$ and $B = 0$.
2. In the same game, what is the value with Bob to move on a position of three runs of length 4? *Answer:* $\Pi = 3$, $t = 3$ so $\epsilon = 1$, giving $B = 3 - 1 = 2$. Note this is not $b(4) + b(4) + b(4) = 0$, which is the whole point of the drill.
3. For which n does Alice moving first on a single run of length n do strictly better than Bob moving first? *Answer:* Exactly when n is hot, that is $n \equiv 4$ or $6 \pmod 7$; then $A - B = 2$, and otherwise $A = B$. Equivalently, exactly when the run is a genuine switch rather than a number.

4.6 A6. Odd power sums and interval lengths

Let n be a positive integer. Determine, in terms of n, the largest integer m with the following property: There exist real numbers $x_1, \ldots, x_{2n}$ with $-1 < x_1 < x_2 < \cdots < x_{2n} < 1$ such that the sum of the lengths of the n intervals

$$[x_1^{2k-1}, x_2^{2k-1}], [x_3^{2k-1}, x_4^{2k-1}], \ldots, [x_{2n-1}^{2k-1}, x_{2n}^{2k-1}]$$

is equal to 1 for all integers k with $1 \leq k \leq m$.

Read it out loud

Strip the geometry off first. The map $t \mapsto t^{2k-1}$ is strictly increasing on $\mathbb{R}$ for every $k \geq 1$, so each interval $[x_{2i-1}^{2k-1}, x_{2i}^{2k-1}]$ is genuinely an interval and its length is $x_{2i}^{2k-1} - x_{2i-1}^{2k-1}$. The condition is therefore

$$\sum_{i=1}^{n} \left(x_{2i}^{2k-1} - x_{2i-1}^{2k-1} \right) = 1 \qquad (1 \leq k \leq m),$$

that is, a system of m equations on the alternating odd power sums of the $2n$ numbers. So the intervals are decoration: this is a symmetric function problem in disguise, and it is asking how many odd power sums you can prescribe.

Now the move that turns it into a clean statement. The exponents are odd, so a minus sign in front of x^{2k-1} can be absorbed into the variable: $-x^{2k-1} = (-x)^{2k-1}$. That converts the alternating sum into an honest power sum of a modified list. And the right hand side is the constant 1, which is itself 1^{2k-1}, so it too can be absorbed as one more element of the list, with a sign. Define the multiset of $2n+1$ real numbers

$$R = \{ x_1,\ x_3,\ \ldots,\ x_{2n-1} \} \cup \{ -x_2,\ -x_4,\ \ldots,\ -x_{2n} \} \cup \{ 1 \}.$$

Then for every k the $(2k-1)$st power sum of R is

$$s_{2k-1}(R) = \sum_{i=1}^{n} x_{2i-1}^{2k-1} - \sum_{i=1}^{n} x_{2i}^{2k-1} + 1 = -1 + 1 = 0$$

precisely when the kth condition holds. So the problem says: how many consecutive odd power sums of a $(2n+1)$-element multiset of reals can vanish, when $2n$ of the elements lie in $(-1, 1)$ and one of them is 1?

That reformulation carries its own answer in embryo, because odd power sums measure the failure of a multiset to be symmetric under $x \mapsto -x$: if $R = -R$ then every odd power sum vanishes, and conversely you expect that enough vanishing odd power sums forces $R = -R$. But R contains 1 and every other element has modulus strictly less than 1, so R cannot be symmetric. The whole upper bound is the quantitative form of "enough".

The fields wired in are symmetric function theory (power sums, elementary symmetric functions, the Newton identities), polynomial parity, and, for the construction, roots of unity and the cosine parametrisation of Chebyshev nodes. Within thirty seconds you should have written the alternating sum, absorbed the signs into the variables, adjoined the element 1, and written "all odd power sums up to s_{2m-1} vanish".

Approaches, ranked worst to best

1. **Solve the system for small n numerically and guess.** *Cue:* "determine in terms of n". *Circuit:* small cases first. For $n = 1$ the system is $x_2 - x_1 = 1$ and $x_2^3 - x_1^3 = 1$; substituting $x_2 = x_1 + 1$ into the second gives $3x_1(x_1 + 1) = 0$, so $x_1 \in \{0, -1\}$ and hence $x_2 \in \{1, 0\}$, both excluded by the strict inequalities. So $m = 1$ when $n = 1$. This is worth doing, because it reveals that the obstruction is exactly that the extremal configuration wants to touch the endpoint 1. As a proof strategy it is hopeless for general n. Ranked last as a method, first as reconnaissance.

2. **Count equations against unknowns.** *Cue:* m equations in $2n$ unknowns. *Circuit:* dimension count. With $2n$ free variables one expects to satisfy $2n$ conditions, suggesting $m = 2n$. The answer is n, so the count is not merely imprecise, it is off by a factor of two. The reason is that a naive count ignores the constraint $|x_j| < 1$: the system does have solutions with $2n$ conditions, but they force some $|x_j| \geq 1$, exactly as in the case $n = 1$. A dimension count can never see a boundary. Ranked here as a warning: it is the argument that makes people confidently write the wrong answer.

3. **Generating function of odd power sums plus residues.** *Cue:* power sums suggest a generating function. *Circuit:* recover a multiset from the poles of a rational function. Set $f(z) = \sum_{x \in R} \frac{xz}{1-x^2z^2} = \sum_{k \geq 1} s_{2k-1}(R) z^{2k-1}$. Vanishing of enough coefficients forces $f \equiv 0$ if the pole count is controlled, and the residue of f at $z = 1/x$ recovers both x and its multiplicity. This works, and it is the route that generalises to a comparison of two arbitrary multisets, but it needs a preliminary reduction (cancel pairs $\{x, -x\}$, which contribute nothing to odd power sums, until none remain) and careful bookkeeping of pole orders against the order of vanishing at $z = 0$. Ranked third only because that bookkeeping is where a script under time pressure springs a leak.

4. **Newton identities on the monic polynomial with root multiset R.** *Cue:* power sums with a bound on how many are prescribed, and a fixed number of elements. *Circuit:* convert power sums to elementary symmetric functions, then read off polynomial parity. Vanishing of $s_1, s_3, \ldots, s_{2m-1}$ propagates through the Newton identities to $e_1 = e_3 = \cdots = e_{2m-1} = 0$, and for $m \geq n+1$ that kills every odd elementary symmetric function of the $2n+1$ roots, which says exactly that the polynomial is odd. An odd polynomial with root 1 has root -1, and $-1 \notin R$. This is the ideal route: it is fully elementary, it needs no reduction step, and the final contradiction is a single evaluation.

The ideal solution

The answer is $\boxed{m = n}$.

Reformulation. Since $t \mapsto t^{2k-1}$ is strictly increasing, the sum of the lengths of the n intervals equals $\sum_{i=1}^{n} \left(x_{2i}^{2k-1} - x_{2i-1}^{2k-1} \right)$, so the requirement is

$$\sum_{i=1}^{n} x_{2i}^{2k-1} - \sum_{i=1}^{n} x_{2i-1}^{2k-1} = 1 \qquad (1 \leq k \leq m). \tag{$*$}$$

Part 1: $m \leq n$.

Suppose real numbers $x_1, \ldots, x_{2n}$ with $|x_j| < 1$ satisfy $(*)$ for $k = 1, \ldots, m$ (the ordering is not needed for this half). Define the monic polynomial of degree $N := 2n+1$

$$p(x) = (x-1) \prod_{i=1}^{n} (x - x_{2i-1})\,(x + x_{2i}),$$

whose roots, listed with multiplicity, form the multiset

$$R = \{1\} \cup \{x_1, x_3, \ldots, x_{2n-1}\} \cup \{-x_2, -x_4, \ldots, -x_{2n}\}.$$

Let s_j denote the jth power sum of R and e_j its jth elementary symmetric function, so that

$$p(x) = \sum_{j=0}^{N} (-1)^j e_j\, x^{N-j}, \qquad e_0 = 1.$$

Because $2k-1$ is odd we have $(-x_{2i})^{2k-1} = -x_{2i}^{2k-1}$, so for $1 \le k \le m$,

$$s_{2k-1} = 1 + \sum_{i=1}^{n} x_{2i-1}^{2k-1} - \sum_{i=1}^{n} x_{2i}^{2k-1} = 1 - 1 = 0,$$

using $(*)$.

Lemma 4.6.1. *If $s_j = 0$ for every odd j with $1 \le j \le J$, then $e_j = 0$ for every odd j with $1 \le j \le \min(J, N)$.*

Proof. The Newton identities for a monic polynomial with N roots state that for $1 \le j \le N$,

$$j\, e_j = \sum_{i=1}^{j} (-1)^{i-1} e_{j-i}\, s_i.$$

Argue by strong induction on odd $j \le \min(J, N)$. In the sum on the right, split by the parity of i. If i is odd then $i \le j \le J$, so $s_i = 0$ and the term vanishes. If i is even then $j - i$ is odd and $1 \le j - i < j$, so $e_{j-i} = 0$ by the inductive hypothesis and the term vanishes. Hence $j\, e_j = 0$, and since $j \ge 1$ we get $e_j = 0$. □

Now suppose, for contradiction, that $m \ge n+1$. Then $s_j = 0$ for all odd $j \le 2m-1$, and $2m - 1 \ge 2n + 1 = N$, so the lemma gives $e_j = 0$ for *every* odd j with $1 \le j \le N$. In the expansion $p(x) = \sum_{j=0}^{N} (-1)^j e_j x^{N-j}$, the term of index j has degree $N - j = 2n + 1 - j$, which is even exactly when j is odd. Therefore every even-degree coefficient of p vanishes, that is, p is an odd polynomial:

$$p(-x) = -p(x) \qquad \text{for all } x.$$

But $p(1) = 0$ by construction, so $p(-1) = -p(1) = 0$ and -1 is a root of p. The roots of p are 1 and the $2n$ numbers $x_1, x_3, \ldots, x_{2n-1}, -x_2, -x_4, \ldots, -x_{2n}$, all of which have modulus strictly less than 1. So -1 is not a root, a contradiction. Hence $m \le n$.

Part 2: $m = n$ is attained.

Take

$$x_j = \cos \frac{(2n+1-j)\pi}{2n+1} \qquad (j = 1, \ldots, 2n).$$

As j runs from 1 to $2n$ the angle $\frac{(2n+1-j)\pi}{2n+1}$ decreases strictly from $\frac{2n\pi}{2n+1}$ to $\frac{\pi}{2n+1}$, staying in the open interval $(0, \pi)$ on which the cosine is strictly decreasing and takes values in $(-1, 1)$. Hence $-1 < x_1 < x_2 < \cdots < x_{2n} < 1$, as required.

Fix k with $1 \le k \le n$. Since $2k-1$ is odd, $(-1)^{j(2k-1)} = (-1)^j$, so the sum of the lengths is

$$L_k := \sum_{i=1}^{n} \left(x_{2i}^{2k-1} - x_{2i-1}^{2k-1} \right) = \sum_{j=1}^{2n} (-1)^j x_j^{2k-1} = \sum_{j=1}^{2n} \left((-1)^j x_j \right)^{2k-1}.$$

Substitute $\ell = 2n + 1 - j$, so that ℓ runs over $1, \ldots, 2n$ and $(-1)^j = (-1)^{2n+1-\ell} = -(-1)^\ell$. Using $\cos(\ell\pi + \theta) = (-1)^\ell \cos\theta$,

$$(-1)^j x_j = -(-1)^\ell \cos \frac{\ell\pi}{2n+1} = -\cos\left(\ell\pi + \frac{\ell\pi}{2n+1} \right) = -\cos \frac{2\pi(n+1)\ell}{2n+1}.$$

Let $\zeta = e^{2\pi i (n+1)/(2n+1)}$. Since $2(n+1) - (2n+1) = 1$, we have $\gcd(n+1, 2n+1) = 1$, so ζ is a primitive $(2n+1)$st root of unity, and $\cos \frac{2\pi(n+1)\ell}{2n+1} = \frac{\zeta^\ell + \zeta^{-\ell}}{2}$. Therefore, using that $2k-1$ is odd to pull out the minus sign,

$$L_k = -\sum_{\ell=1}^{2n} \left(\frac{\zeta^\ell + \zeta^{-\ell}}{2} \right)^{2k-1} = -\frac{1}{2^{2k-1}} \sum_{t=0}^{2k-1} \binom{2k-1}{t} \sum_{\ell=1}^{2n} \zeta^{\ell(2k-1-2t)},$$

after expanding by the binomial theorem and exchanging the two finite sums.

For each t, put $u = 2k-1-2t$. Then u is odd, so $u \neq 0$, and $|u| \leq 2k-1 \leq 2n-1 < 2n+1$; hence $u \not\equiv 0 \pmod{2n+1}$ and ζ^u is a root of unity different from 1. Consequently

$$\sum_{\ell=1}^{2n} \zeta^{\ell u} = \left(\sum_{\ell=0}^{2n} \zeta^{\ell u}\right) - 1 = 0 - 1 = -1.$$

Substituting,

$$L_k = -\frac{1}{2^{2k-1}} \sum_{t=0}^{2k-1} \binom{2k-1}{t} \cdot (-1) = \frac{1}{2^{2k-1}} \cdot 2^{2k-1} = 1.$$

So the configuration satisfies $(*)$ for every k with $1 \leq k \leq n$, and $m = n$ is attained. Combining the two parts, the largest such m is n. ■

Where the construction stops, and why. The only step that used $k \leq n$ was the bound $|u| \leq 2k-1 < 2n+1$. At $k = n+1$ the extreme terms $t = 0$ and $t = 2k-1$ give $u = \pm(2n+1) \equiv 0$, so those two inner sums equal $2n$ instead of -1, and $L_{n+1} = 1 - \frac{(2n+1)\cdot 2}{2^{2n+1}} \neq 1$. For example at $n = 1$ the configuration is $x_1 = -\frac{1}{2}$, $x_2 = \frac{1}{2}$, with $L_1 = 1$ and $L_2 = \frac{1}{4}$; at $n = 2$ one gets $L_1 = L_2 = 1$ and $L_3 = \frac{11}{16}$. The construction is tight, not merely sufficient.

Remark: what the upper bound really says. A multiset R satisfies $R = -R$ if and only if all of its odd power sums vanish. Part 1 is the statement that for a multiset of size $2n+1$, vanishing of the first $n+1$ odd power sums already forces this symmetry, which then contradicts $1 \in R$ and $|x| < 1$ for all other elements. The number $n+1$ is exactly what is needed because the polynomial has $n+1$ odd-index elementary symmetric functions among $e_1, \ldots, e_{2n+1}$.

Layer 1: WHAT

Block 1. Strip the geometry: interval lengths are differences of odd powers, so the condition is an alternating odd power sum equal to 1.

Block 2. Absorb the alternating signs into the variables, legitimate because the exponents are odd, and absorb the constant 1 as an extra element 1 of the list.

Block 3. Restate: the odd power sums $s_1, s_3, \ldots, s_{2m-1}$ of a $(2n+1)$-element multiset R all vanish.

Block 4. Convert power sums to elementary symmetric functions with the Newton identities: all odd-index e_j vanish up to index $2m-1$.

Block 5. If $m \geq n+1$ this covers every odd index up to the degree, so p has only odd-degree terms: p is an odd polynomial.

Block 6. An odd polynomial with root 1 has root -1; but every element of R other than 1 has modulus below 1. Contradiction, so $m \leq n$.

Block 7. For the construction, take the $2n$ interior Chebyshev-type nodes $\cos\frac{(2n+1-j)\pi}{2n+1}$, convert the alternating sum to a sum over $(2n+1)$st roots of unity, and use $\sum_{\ell=1}^{2n} \zeta^{\ell u} = -1$ for $u \not\equiv 0$.

Juxtapose with the nominal route. The zombie counts m equations against $2n$ unknowns, concludes $m = 2n$, and either stops or wastes the session trying to build a configuration. If he does get as far as suspecting $m = n$, he tries to prove it by manipulating the m equations directly, subtracting them from one another and hoping for a factorisation, which never comes because the equations are not polynomial identities in any small number of quantities: they are m separate statements about $2n$ variables that only become tractable in the language of symmetric functions.

The bifurcation point is the sentence where you decide what a "sum of odd powers with signs" is. The zombie sees m constraints. You see a power sum of one multiset, because odd exponents absorb signs, and once it is a power sum the entire theory of symmetric functions is available.

Layer 2: HOW

Block 1. Cue. Intervals with endpoints x^{2k-1}, and the exponent $2k-1$ is odd. **Why natural.** $t \mapsto t^{2k-1}$ is strictly increasing on all of $\mathbb{R}$, which is the only reason $[x_{2i-1}^{2k-1}, x_{2i}^{2k-1}]$ is a legitimate

interval; the odd exponent was chosen by the setters for exactly that purpose and for the sign absorption to come. So write the length as a difference immediately and forget the geometry. **Link to the next block.** A signed sum of odd powers is asking to be turned into an unsigned one.

Block 2. **Cue.** $-x^{2k-1} = (-x)^{2k-1}$, valid for every odd exponent simultaneously. **Why natural.** This is the definition-level fact that makes the whole problem uniform in k. Any device that works for one k but not all is useless here; sign absorption works for all odd exponents at once, so it does not disturb the family of conditions. Likewise $1 = 1^{2k-1}$ for all k, so the constant can be moved into the list as an element rather than sitting on the right hand side. **Link to the next block.** You now have a single multiset whose odd power sums all vanish.

Block 3. **Cue.** "All odd power sums vanish" is a symmetry statement. **Why natural.** If a multiset is invariant under negation then odd power sums cancel in pairs. The converse is the substance. Recognising the target as "prove the multiset is negation-symmetric" tells you exactly what contradiction to look for at the end, namely an element that has no negative partner. The element 1 is that partner-less element, and it is partner-less precisely because the problem confined all the x_j to $(-1, 1)$. **Link to the next block.** To get from power sums to a structural statement about the multiset, pass to the polynomial with those roots.

Block 4. **Cue.** Finitely many power sums of finitely many numbers. **Why natural.** Power sums and elementary symmetric functions are two coordinate systems on the same object, and the Newton identities are the change of coordinates. The identities are triangular, so vanishing propagates: the parity split in the sum is not a computation, it is the observation that odd plus even is odd. Notice the shape of the induction: it needs only the terms with i odd (killed by hypothesis) and those with i even (killed by the inductive hypothesis on a smaller odd index). **Link to the next block.** Vanishing odd elementary symmetric functions is a statement about which coefficients of p are zero.

Block 5. **Cue.** $p(x) = \sum_j (-1)^j e_j x^{N-j}$ with $N = 2n + 1$ odd. **Why natural.** The index j and the degree $N - j$ have opposite parities when N is odd, so "all odd-index e_j vanish" is literally "all even-degree coefficients vanish", which is literally "p is an odd function". Choosing N odd is therefore not an accident of the construction: including the extra element 1 made the degree $2n + 1$, and that is what turns a symmetric function statement into a parity statement. **Link to the next block.** An odd polynomial has a negation-symmetric root multiset.

Block 6. **Cue.** p odd and $p(1) = 0$. **Why natural.** $p(-1) = -p(1) = 0$ needs no theory. This is why the extra element was chosen to be 1 and not some other constant: the right hand side of the problem is 1, and 1 sits exactly on the boundary of the forbidden region for the x_j. The strict inequality $x_{2n} < 1$ in the statement is therefore load-bearing, not cosmetic, and the $n = 1$ computation shows what happens when it is relaxed: the solution degenerates to $\{0, 1\}$. **Link to the next block.** The bound is n; now construct.

Block 7. **Cue.** You need $2n$ numbers in $(-1, 1)$ with n prescribed alternating odd power sums. **Why natural.** Prescribing many power sums at once is what roots of unity do best, because $\sum_{\ell=1}^{N-1} \zeta^{\ell u} = -1$ for every $u \not\equiv 0 \pmod N$, uniformly in u. So look for the x_j as cosines of equally spaced angles, whose odd powers expand by the binomial theorem into exactly such sums; then the binomial coefficients sum to 2^{2k-1} and cancel the normalising power of 2. The alternating signs and the reversal $\ell = 2n + 1 - j$ are there to convert $(-1)^\ell \cos \frac{\ell\pi}{2n+1}$ into a single cosine of a multiple angle, using $\cos(\ell\pi + \theta) = (-1)^\ell \cos\theta$. **Link to the next block.** The uniformity in u holds precisely while $|u| \le 2k - 1 < 2n + 1$, which is the constraint $k \le n$, and it fails at $k = n + 1$, matching the upper bound exactly.

Layer 3: WHY

Three truths make this route ideal. First, odd power sums are exactly the coordinates that detect the failure of a finite multiset of reals to be symmetric about the origin; a multiset is negation-symmetric if and only if all its odd power sums vanish, and finitely many of them suffice once you have as many as there are odd-index elementary symmetric functions. Second, the Newton identities are a triangular change of coordinates between power sums and elementary symmetric functions, so vanishing conditions transfer between the two systems with an induction that respects parity. Third, for a monic polynomial of odd degree, "all odd-index elementary symmetric functions vanish" and "the polynomial is an odd function" are the same statement, which converts algebra into a one-line evaluation at $x = -1$.

The discernment mechanism to install: whenever a problem prescribes several power sums, or several moments, or several integrals of powers, do not manipulate the equations. Encode the data as a polynomial (or a generating function) whose roots are the objects, and translate the prescriptions into vanishing coefficients. Then the question "how many conditions can hold" becomes "how many coefficients can vanish before the polynomial degenerates", and degeneration is always visible. The same reflex solves problems about moments determining a measure, about Vandermonde systems, and about equal power sums of two sets (the Prouhet, Tarry and Escott problem, whose whole content is that equal power sums up to a high enough order force equality of the multisets).

Train it by taking any statement of the form "the first r power sums of these numbers are $c_1, \dots, c_r$" and immediately writing the first r elementary symmetric functions via the Newton identities, then asking what the resulting polynomial looks like. If a pattern of zeros appears, name the corresponding structural property (odd, even, palindromic, divisible by x^d), because that property is the real content.

What the examiners nudged with. Every constant in the statement is doing work: the odd exponents $2k-1$ permit sign absorption; the target value 1 matches the excluded endpoint of the interval $(-1, 1)$ so that the contradiction is "-1 is a root"; the count $2n$ of variables makes the auxiliary polynomial have odd degree $2n+1$; and the strictness of $-1 < x_1$ and $x_{2n} < 1$ is what makes the extremal configuration unattainable at $m = n+1$. A statement in which every numeral is load-bearing is a statement engineered around a single mechanism, and finding that mechanism is the whole task.

Knowledge gaps, stated plainly. This problem is not approachable at all without three things installed. (i) The elementary symmetric functions e_j and power sums s_j of a multiset, and the relation $p(x) = \sum_j (-1)^j e_j x^{N-j}$ between e_j and the coefficients. (ii) The Newton identities $je_j = \sum_{i=1}^{j} (-1)^{i-1} e_{j-i} s_i$ for $1 \le j \le N$; without them there is no bridge from the given data (power sums) to the usable data (coefficients). (iii) The geometric series of roots of unity, $\sum_{\ell=0}^{N-1} \zeta^{\ell u} = 0$ for ζ a primitive Nth root of unity and $u \not\equiv 0 \pmod N$; without it the construction cannot be verified. None of these is in A-level Further Mathematics. A fourth item, partial fractions and residues of rational functions, is needed only for the alternative route through the generating function $\sum \frac{xz}{1-x^2z^2}$, and can be skipped if you take the Newton route. Do not attempt this problem in the arena until (i), (ii) and (iii) are automatic.

Revolving concepts

Power sums, elementary symmetric functions, and the Newton identities. For a multiset $\{r_1, \dots, r_N\}$ with e_j the jth elementary symmetric function and $s_j = \sum_i r_i^j$, the identities

$$j\,e_j = \sum_{i=1}^{j} (-1)^{i-1} e_{j-i} s_i \quad (1 \le j \le N), \qquad s_j = \sum_{i=1}^{N} (-1)^{i-1} e_i\, s_{j-i} \quad (j > N)$$

let you pass freely between the two coordinate systems. The second form is itself a linear recursion: once $e_1, \dots, e_N$ are known, the power sums satisfy a linear recurrence of order N with those

coefficients, whose general solution is $s_j = \sum_i c_i r_i^j$. Generalise the recursion this way: any sequence defined by a linear recurrence with characteristic polynomial p is a combination of jth powers of the roots of p, and conversely power sums of roots satisfy the recurrence read off from the coefficients. See Book 1, field: Polynomials, topic: symmetric functions and the Newton identities.

Negation symmetry and odd power sums. A finite multiset $R \subset \mathbb{R}$ satisfies $R = -R$ if and only if $s_j(R) = 0$ for every odd j, if and only if the monic polynomial with root multiset R is an even or an odd function according to the parity of $|R|$. The quantitative version, that finitely many odd power sums suffice, is what Part 1 proves. See Book 1, field: Polynomials, topic: parity of polynomials and root symmetry.

Equal power sums force equal multisets. If two multisets of complex numbers of sizes at most N have equal jth power sums for $j = 1, \ldots, N$, they are equal. This is the Prouhet, Tarry and Escott circle of ideas and is the same statement as invertibility of the Vandermonde matrix (r_i^j) for distinct r_i, whose determinant is $\prod_{i<j}(r_j - r_i)$. See Book 1, field: Polynomials, topic: Newton identities and Vandermonde.

Roots of unity sums. For ζ a primitive Nth root of unity, $\sum_{\ell=0}^{N-1} \zeta^{\ell u}$ equals N if $N \mid u$ and 0 otherwise. This is the discrete orthogonality relation and the engine of every explicit construction of this type. Combined with $\cos\theta = \frac{e^{i\theta}+e^{-i\theta}}{2}$ and the binomial theorem, it evaluates $\sum_\ell (\cos \alpha\ell)^d$ in closed form. See Book 1, field: Complex numbers, topic: roots of unity and orthogonality.

Chebyshev nodes. The numbers $\cos \frac{j\pi}{N}$ are the extremal nodes of the Chebyshev polynomial T_N, and $\cos \frac{(2j-1)\pi}{2N}$ are its roots. Configurations of equally spaced angles are the default guess whenever many symmetric conditions must hold at once, because their power sums collapse via roots of unity. See Book 1, field: Polynomials, topic: Chebyshev polynomials and node sets.

Generating function of odd power sums. $\sum_{k\geq 1} s_{2k-1} z^{2k-1} = \sum_{x\in R} \frac{xz}{1-x^2z^2}$, a rational function whose poles at $z = \pm 1/x$ have residues determining x and its multiplicity. This is the alternative route and the one that generalises to comparing two multisets directly. See Book 1, field: Generating functions, topic: power sums and partial fractions.

Circuit Bank

- **WHEN** you see a signed sum of odd powers of several variables **DO** absorb every sign into its variable, since $-x^{2k-1} = (-x)^{2k-1}$, and absorb any constant right hand side as an extra element of the list **BECAUSE** the family of conditions then becomes 'these power sums of one multiset vanish', which puts the whole of symmetric function theory at your disposal, uniformly in the exponent
- **WHEN** a problem prescribes several power sums of an unknown finite set **DO** form the monic polynomial with those roots and translate the prescriptions into its coefficients via the Newton identities **BECAUSE** power sums and coefficients are two coordinate systems linked by a triangular change of variables, and structural facts (parity, divisibility, degeneration) are visible only in the coefficient system
- **WHEN** all odd-index elementary symmetric functions of a monic polynomial of odd degree vanish **DO** conclude that the polynomial is an odd function and evaluate at a known root and its negative **BECAUSE** for odd degree N, the index j and the degree $N - j$ have opposite parity, so vanishing odd-index e_j is exactly vanishing even-degree coefficients, and an odd polynomial has a negation-symmetric root multiset
- **WHEN** you must construct a configuration satisfying many power sum conditions at once **DO** try equally spaced angles, so that the powers expand into sums $\sum_\ell \zeta^{\ell u}$ over roots of unity **BECAUSE** those sums take the same value -1 for every $u \not\equiv 0$, so the binomial expansion collapses uniformly in k, and the range of k for which it works is read off from the condition $|u| < N$

- **WHEN** a dimension count suggests more conditions can be met than the answer allows **DO** suspect a boundary constraint and test the smallest case explicitly **BECAUSE** counting equations against unknowns is blind to inequalities; here the extremal solution for $n = 1$ is $(0, 1)$, which sits exactly on the forbidden endpoint, and that is the whole obstruction

Backtesting

1. First useful cue: the exponents are odd, so signs absorb into variables and the alternating sum of lengths is a single power sum of a modified multiset; the constant 1 joins the multiset as the element 1.
2. Method the expression suggested: form the monic polynomial with that root multiset, push the vanishing power sums through the Newton identities to vanishing odd-index elementary symmetric functions, and read the result as polynomial parity.
3. Where you would hesitate, and the thought to reinforce: at the parity bookkeeping in the Newton induction. Reinforce the split by parity of the summation index: odd i kills s_i by hypothesis, even i leaves $j - i$ odd and smaller, killed by induction. Written that way it is two clauses, not a computation.
4. The trap the question sets: the equation count. With $2n$ unknowns it is natural to expect $2n$ conditions to be satisfiable, and the true answer is half of that. The gap is caused entirely by the open interval $(-1, 1)$, which a dimension count cannot see.
5. What makes you recognise it faster next time: every numeral in the statement is load-bearing. Odd exponents permit sign absorption, the value 1 equals the excluded endpoint, and $2n$ variables give an auxiliary polynomial of odd degree. When a statement is that tightly engineered, look for a single mechanism.
6. The reusable decision rule: prescribed power sums become prescribed coefficients; then structural degeneration of the polynomial is the obstruction, and the number of conditions you can impose is the number of coefficients you can kill before it degenerates.

Mirror drills

1. Show that if the multiset $\{r_1, \ldots, r_N\}$ of complex numbers has $s_j = 0$ for $j = 1, \ldots, N$, then all $r_i = 0$. *Answer:* By the Newton identities and induction, $e_j = 0$ for $j = 1, \ldots, N$, so the monic polynomial is x^N and every root is 0. Same mechanism, with no parity refinement.
2. Solve the case $n = 1$ of the problem by hand: find all $-1 < x_1 < x_2 < 1$ with $x_2 - x_1 = 1$ and $x_2^3 - x_1^3 = 1$. *Answer:* Substituting $x_2 = x_1 + 1$ into the second equation gives $3x_1^2 + 3x_1 + 1 = 1$, so $x_1 \in \{0, -1\}$ and $x_2 \in \{1, 0\}$; both violate the strict inequalities, so $m = 1$ for $n = 1$. The degenerate solution $\{0, 1\}$ shows that the boundary is the obstruction.
3. Evaluate $\sum_{\ell=1}^{2n} \cos \frac{2\pi\ell}{2n+1}$ and $\sum_{\ell=1}^{2n} \cos^3 \frac{2\pi\ell}{2n+1}$ for $n \geq 2$. *Answer:* Both equal -1: the first is $\sum_{\ell=1}^{2n} \frac{\zeta^\ell + \zeta^{-\ell}}{2} = \frac{-1-1}{2} = -1$, and expanding the cube gives $\frac{1}{8}\left(\binom{3}{0} + \binom{3}{1} + \binom{3}{2} + \binom{3}{3}\right)(-1) = -1$, using that each exponent $\pm 3, \pm 1$ is nonzero modulo $2n + 1$ when $2n + 1 > 3$. This is the construction's computation in miniature.

4.7 B1. Parity of exponential series coefficients

Suppose that $P(x) = a_1x + a_2x^2 + \cdots + a_nx^n$ is a polynomial with integer coefficients, with a_1 odd. Suppose that $e^{P(x)} = b_0 + b_1x + b_2x^2 + \cdots$ for all x. Prove that b_k is nonzero for all $k \geq 0$.

Read it out loud

An exponential of an integer polynomial with no constant term, expanded as a power series; prove that no coefficient dies. Three fields are wired in: generating functions (the object $e^{P(x)}$ is an exponential generating function in disguise), elementary number theory (the hypotheses are integrality and the parity of a_1), and one line of calculus (the exponential is the function whose logarithmic derivative is P').

What is really being tested is whether you look at the right object. The coefficients b_k are rational numbers with denominators dividing $k!$, and a rational number is a poor thing to prove nonzero: there is no modulus to work in. The hypotheses are the guiding lights. "Integer coefficients" and "a_1 odd" are parity hypotheses, and a parity hypothesis is planted so that you prove something is odd. An odd integer is nonzero. Thus the statement you should actually prove is stronger and discrete: $k!\,b_k$ is an odd integer for every k.

Within thirty seconds you should have written $c_k = k!\,b_k$, checked $c_0 = e^{P(0)} = 1$ and $c_1 = a_1$, and asked how c_{k+1} is built from $c_0, \ldots, c_k$. The cleanest engine for "how are the coefficients of e^P built" is not the series $\sum_m P^m/m!$; it is the differential equation $y' = P'(x)\,y$, because that equation is what e^P is, definitionally, and a first-order linear equation with polynomial coefficient turns into a finite linear recursion.

Approaches, ranked worst to best

1. **Multinomial grind.** *Cue:* the explicit series $e^P = \sum_m P(x)^m/m!$. *Circuit:* expand and track coefficients. (zombie route: b_k becomes a sum over compositions of k of multinomial coefficients divided by $m!$; you cannot see the denominators cancel, let alone the parity, without first normalising by $k!$, and once you have normalised you should be using a better engine anyway)

2. **Exponential formula and involutions.** *Cue:* e^P as the generating function of set partitions with coloured blocks. *Circuit:* count, then cancel in pairs with a sign-reversing involution. (correct: $k!\,b_k$ is a signed count of set partitions of $\{1, \ldots, k\}$ whose blocks of size j carry one of $|a_j|$ colours, and a fixed-point-free involution on the non-trivial partitions leaves only the all-singletons partition, which is odd; but it needs the exponential formula plus a carefully checked involution, the most theory for the least ink saved)

3. **Multiplicative decomposition.** *Cue:* $e^P = e^{a_1x}\,e^{a_2x^2}\cdots e^{a_nx^n}$. *Circuit:* show each factor other than e^x lies in a multiplicatively closed class of series with even normalised coefficients. (works: series $\sum c_kx^k/k!$ with $c_0 = 1$ multiply by $e_k = \sum_m \binom{k}{m} c_m d_{k-m}$, so "all c_k even for $k \geq 1$" is preserved; you then need $2^k/k!$ to have even numerator, which is Legendre's formula $v_2(k!) < k$, and $(mj)!/m!$ even for $j \geq 2$; three small lemmas where one would do)

4. **Truncated series at the top degree.** *Cue:* $k!\,b_k = [x^k]\sum_{m\leq k} \frac{k!}{m!}P^m$. *Circuit:* the $m = k$ term contributes a_1^k, every $\frac{k!}{m!}$ with $m \leq k-2$ is even. (the official first solution: two lines, but with a boundary case, k odd and $m = k-1$, where $k!/m! = k$ is odd and you must compute the x^k coefficient of $k\,P^{k-1}$ by hand as $k(k-1)a_1^{k-2}a_2$; a place to drop a point under time pressure)

5. **Differential equation to recursion.** *Cue:* e^P is the solution of $y' = P'(x)\,y$ with $y(0) = 1$. *Circuit:* convert the equation into a recursion for $c_k = k!\,b_k$ and reduce mod 2. (ideal: the

recursion has integer coefficients that are falling factorials, every term except $a_1 c_k$ is visibly even, one induction, no boundary cases)

The ideal solution

Write $y(x) = e^{P(x)}$. Since P is a polynomial, y is entire, so its Taylor series $\sum_{k\geq 0} b_k x^k$ converges for every real x and may be differentiated term by term. The chain rule gives

$$y'(x) = P'(x)\, y(x), \qquad P'(x) = \sum_{j=1}^{n} j a_j x^{j-1}.$$

Comparing the coefficient of x^k on both sides, with the convention $b_m = 0$ for $m < 0$,

$$(k+1)\, b_{k+1} = \sum_{j=1}^{n} j a_j\, b_{k+1-j} \qquad (k \geq 0).$$

Set $c_k = k!\, b_k$. Multiply the display by $k!$ and use $k!/(k+1-j)! = k(k-1)\cdots(k+2-j)$, a product of $j-1$ consecutive integers (empty for $j = 1$):

$$c_{k+1} = a_1 c_k + \sum_{j=2}^{n} j a_j\, k(k-1)\cdots(k+2-j)\; c_{k+1-j},$$

where terms with $k+1-j < 0$ vanish.

Claim 4.7.1. *For every $k \geq 0$, c_k is an odd integer.*

Proof. Induction on k. First $c_0 = b_0 = e^{P(0)} = 1$. Suppose $c_0, \ldots, c_k$ are odd integers. Every term of the recursion for c_{k+1} is then an integer. The term $a_1 c_k$ is odd. For $j = 2$ the term is $2a_2 k\, c_{k-1}$, even. For $j \geq 3$ the factor $k(k-1)\cdots(k+2-j)$ is a product of $j-1 \geq 2$ consecutive integers, one of which is even, so the term is even. Hence $c_{k+1} \equiv a_1 c_k \equiv 1 \pmod 2$. □

Therefore $b_k = c_k/k!$ is a nonzero rational number for every $k \geq 0$. In fact $k!\, b_k$ is an odd integer, which is the statement worth remembering.

Layer 1: WHAT

Block 1: strengthen the target. "$b_k \neq 0$" becomes "$k!\, b_k$ is an odd integer", a discrete statement with a modulus to work in. Block 2: choose the engine. e^P is characterised by $y' = P'y$, $y(0) = 1$; write that, not the series $\sum P^m/m!$. Block 3: read off the recursion $(k+1)b_{k+1} = \sum_j j a_j b_{k+1-j}$ and normalise to $c_k = k!\, b_k$, which clears every denominator and produces falling factorials as coefficients. Block 4: reduce mod 2. Every term except $a_1 c_k$ carries an even factor: the explicit 2 when $j = 2$, or a product of at least two consecutive integers when $j \geq 3$. Block 5: close by induction from $c_0 = 1$.

Juxtaposition. The nominal route expands $e^P = \sum_m P^m/m!$ and tries to compute b_k; it is not wrong, and pushed correctly it becomes approach 4, but it forces a boundary case and it never tells you why the parity behaves. The bifurcation point is Block 1 combined with Block 2: the moment you decide to study $k!\, b_k$ rather than b_k, and to let the differential equation, not the series, generate the coefficients. After that the proof writes itself.

Layer 2: HOW

Block 1. **Cue.** "Integer coefficients", "a_1 odd", "prove nonzero". Integrality plus a parity hypothesis is a request for a parity conclusion. **Why natural.** A rational number has no arithmetic to exploit; an integer does. Denominators of b_k divide $k!$ because e^P is a sum of $P^m/m!$, so $k!\, b_k$ is the natural integer candidate. Odd implies nonzero, so proving the stronger discrete statement is not extra work, it is the only route that has a chance. **Link to the next block.** Now you need a description of $k!\, b_k$ that shows integrality and parity together; that means a recursion.

Block 2. **Cue.** The function is $e^{P(x)}$, and P is a polynomial. **Why natural.** The exponential of a polynomial is the unique solution of $y' = P'(x)\,y$ with $y(0) = 1$. A linear differential equation with polynomial coefficient is a finite linear recursion on Taylor coefficients: that is what Further Maths series solutions of differential equations are for. The series $\sum P^m/m!$ hides the same recursion behind multinomial sums. **Link to the next block.** Differentiate termwise (legitimate: y is entire) and compare coefficients.

Block 3. **Cue.** The recursion $(k+1)b_{k+1} = \sum_j j a_j b_{k+1-j}$ has a factor $k+1$ on the left. **Why natural.** Multiplying by $k!$ converts $(k+1)b_{k+1}$ into c_{k+1} and each b_{k+1-j} into c_{k+1-j} times $k!/(k+1-j)!$, a falling factorial with $j-1$ factors. Falling factorials are integers, so $c_k \in \mathbb{Z}$ falls out for free. **Link to the next block.** Look at each coefficient mod 2.

Block 4. **Cue.** Products of consecutive integers. **Why natural.** A product of m consecutive integers equals $m!$ times a binomial coefficient, hence is divisible by $m!$; for $m \geq 2$ it is even. The $j = 2$ term has no such product (only one factor, k), but it carries the explicit factor 2 from $j a_j = 2a_2$. So the only odd term is $a_1 c_k$. **Link to the next block.** Induction from $c_0 = 1$ needs nothing else.

Block 5. **Cue.** A recursion with all lower terms even except one. **Why natural.** $c_{k+1} \equiv a_1 c_k \pmod 2$ and a_1 is odd, so parity is inherited. This is also why a_1 odd is exactly the right hypothesis: with a_1 even, c_k is even for $k \geq 1$ and the conclusion can fail (e^{x^2} has every odd coefficient zero). **Link to the next block.** Translate back: $b_k = c_k/k! \neq 0$.

Layer 3: WHY

The mathematical truth underneath is that e^P for an integer polynomial P is an exponential generating function with integer coefficients, and those coefficients satisfy a linear recursion with polynomial (falling factorial) coefficients because e^P satisfies a first-order linear differential equation. Any series satisfying such an equation has this property; the exponential is only the simplest case. Once the recursion is written, parity is a matter of looking at which coefficients are even, and falling factorials of length at least two are always even.

The discernment mechanism is the question "what modulus?". When a problem asks you to prove that a quantity is nonzero, and the hypotheses are arithmetic (integers, odd, prime), the examiners are telling you to find an integer-valued normalisation and a modulus in which the quantity is a unit. Here the normalisation is $k!$ and the modulus is 2. To train this, take every "prove nonzero" statement you meet and ask: nonzero in which ring, and what is the cheapest congruence that implies it?

The nudge was explicit. They did not write $P(x) = x + a_2x^2 + \cdots$; they wrote "$a_1$ odd", and they wrote "$e^{P(x)} = b_0 + b_1 x + \cdots$ for all x" to make you think analytically about something that is arithmetic. Seeing through the analytic costume to the parity statement is the whole problem.

Revolving concepts

Exponential generating functions. For a sequence (c_k), the series $\sum c_k x^k/k!$ is its exponential generating function; products of such series correspond to binomial convolutions $e_k = \sum_m \binom{k}{m} c_m d_{k-m}$, which is why integrality and parity classes are preserved under multiplication. The series $e^{P(x)}$ with P a polynomial without constant term is the exponential generating function of "structures built from connected pieces", the exponential formula: $k!\,[x^k]\,e^P$ counts set partitions of a k-set with each block of size j weighted by a_j. See Book 1, field: Generating functions, topic: exponential generating functions and the exponential formula.

Differential equations to recursions (the generalised recursion). If $y = \sum b_k x^k$ satisfies $y' = Q(x)\,y$ with $Q(x) = \sum_{i \geq 0} q_i x^i$ a polynomial, then $(k+1)b_{k+1} = \sum_i q_i b_{k-i}$, and with $c_k = k!\,b_k$

this becomes $c_{k+1} = \sum_i q_i\, k(k-1)\cdots(k-i+1)\, c_{k-i}$. Special cases you should recognise on sight: $Q = 1 + x$ gives $c_{k+1} = c_k + kc_{k-1}$, the involution numbers $1, 1, 2, 4, 10, 26, \ldots$; $Q = e^x$ gives $c_{k+1} = \sum_m \binom{k}{m} c_m$, the Bell numbers; $Q = 2x$ gives $c_{k+1} = 2kc_{k-1}$, the coefficients of e^{x^2}. The general solution method is always the same: the recursion is linear with polynomial coefficients, so any arithmetic property (integrality, congruences, growth rate) is proved by induction along it. See Book 1, field: Sequences and recurrences, topic: linear recursions with polynomial coefficients.

Falling factorials and divisibility. $k(k-1)\cdots(k-m+1) = m!\binom{k}{m}$ is divisible by $m!$; in particular any product of two consecutive integers is even. The 2-adic valuation of $k!$ is $v_2(k!) = k - s_2(k)$, where $s_2(k)$ is the number of ones in the binary expansion of k (Legendre's formula); this is the lemma behind approach 3. See Book 1, field: Number theory, topic: valuations of factorials.

Circuit Bank

- **WHEN** you see "prove nonzero" with integer and parity hypotheses **DO** write an integer normalisation of the quantity and prove it is odd, or a unit mod a well-chosen prime **BECAUSE** because a rational number has no modulus to work in and an odd integer is automatically nonzero
- **WHEN** you see the coefficients of $e^{P(x)}$, or of any y with $y' = Q(x)\, y$ **DO** write the recursion $(k+1)b_{k+1} = \sum_i q_i b_{k-i}$, then pass to $c_k = k!\, b_k$ **BECAUSE** because a first-order linear differential equation with polynomial coefficient is a finite linear recursion, and the $k!$ normalisation turns its coefficients into integer falling factorials
- **WHEN** you see a product of consecutive integers inside a recursion **DO** mark it as divisible by the factorial of its length **BECAUSE** because it equals that factorial times a binomial coefficient, so it is even as soon as it has two factors

Backtesting

1. First useful cue: "integer coefficients, a_1 odd, prove $b_k \neq 0$". Parity in, parity out.
2. Method the expression suggested: e^P satisfies $y' = P'y$, so read the coefficients from the differential equation and normalise by $k!$.
3. Where you would hesitate, and the thought to reinforce: whether $k!\, b_k$ is even an integer. Do not stop to prove it separately; the recursion proves integrality and parity in the same induction.
4. The trap the question sets: the analytic phrasing "for all x" invites you to think about convergence or about e^P as a function; nothing about the function matters beyond the one differentiation.
5. What makes you recognise it faster next time: any exponential of a polynomial, or any e^f with f having integer Taylor coefficients, is an exponential generating function of integers, and its coefficients obey a recursion you can write in one line.
6. The reusable decision rule: to prove a Taylor coefficient is nonzero, multiply by $k!$ and prove a congruence.

Mirror drills

1. (drill) The involution numbers $I_k = k!\,[x^k]\, e^{x+x^2/2}$ satisfy $I_{k+1} = I_k + kI_{k-1}$. Explain why the parity argument of B1 breaks and why I_k is even for all $k \geq 2$. *Answer:* $a_2 = 1/2$ is not an integer, so the $j = 2$ term is $k\, c_{k-1}$ with no factor 2; then $I_2 = 2$ and $I_{k+1} \equiv I_k + kI_{k-1}$ gives even for ever by induction on pairs.

2. (drill) Let P have integer coefficients with $P(0) = 0$ and a_1 even. Show $k!\,b_k$ is even for all $k \geq 1$, and give a P with some $b_k = 0$. *Answer:* $c_{k+1} \equiv a_1 c_k \equiv 0$; take $P = x^2$, where every odd coefficient vanishes.

3. (drill) Let p be an odd prime and $P(x) = x + x^p$. Prove $k!\,b_k \equiv 1 \pmod{p}$ for all $k \geq 0$. *Answer:* the recursion is $c_{k+1} = c_k + p\,k(k-1)\cdots(k+2-p)\,c_{k+1-p}$, and the second term carries the factor $ja_j = p$, so $c_{k+1} \equiv c_k \pmod{p}$ and $c_0 = 1$; in fact $c_k = 1$ exactly for $k < p$.

4.8 B2. Finite sets closed under the cross product

Let $\times$ represent the cross product in $\mathbb{R}^3$. For what positive integers n does there exist a set $S \subset \mathbb{R}^3$ with exactly n elements such that

$$S = \{v \times w : v, w \in S\}?$$

Read it out loud

A finite set of vectors in three dimensions that reproduces itself exactly under all pairwise cross products. The fields wired in are vector algebra (the cross product and its two defining properties), a scrap of linear algebra (orthogonal sets and spanning), and a finiteness argument of the kind that appears in group theory: a set closed under an operation cannot contain anything whose "size" grows under that operation.

What is really being tested is whether you know what the cross product controls. It is not associative, it has no identity, and it is not a group operation, so the algebraic reflexes you would use for a multiplicative closure problem do not transfer. What it does have is exactly three facts: $v \times v = 0$, the output is orthogonal to both inputs, and $\|v \times w\| = \|v\|\|w\| \sin\theta$. That last one is the only scalar law available, and a scalar law plus finiteness is a rigidity statement waiting to happen.

Within thirty seconds you should have written two lines. First, S is nonempty, so pick any $v \in S$ and note $0 = v \times v \in S$: the zero vector is forced. Second, the norm identity, with the observation that if u and z are perpendicular then $\|u \times z\| = \|u\|\|z\|$ exactly, with no sine factor to spoil it. Multiplication by a fixed constant, applied over and over inside a finite set, is a contradiction unless the constant is 1. That sentence is the whole solution.

The answer, once you see the shape, is the unit octahedron plus its centre: $\{0, \pm v_1, \pm v_2, \pm v_3\}$ for an orthonormal basis, giving $n = 7$, together with the degenerate $S = \{0\}$, giving $n = 1$.

Approaches, ranked worst to best

1. **Coordinate brute force.** *Cue:* the operation is given by an explicit bilinear formula in coordinates. *Circuit:* write everything in components and solve. (zombie route: $3n$ unknowns and n^2 vector equations with no structure imposed; you cannot even write down the condition $S = \{v \times w\}$ without choosing which product equals which element, so the case count explodes before any algebra happens)
2. **Small case hunting.** *Cue:* the phrase "for what positive integers n". *Circuit:* compute $n = 1, 2, 3, \ldots$ by hand and guess. (useful reconnaissance only: you quickly find $\{0\}$ and the seven element octahedral set, but the entire difficulty is proving nothing else exists, and enumeration gives you no leverage on that)
3. **Identify the algebra.** *Cue:* a bilinear antisymmetric product on $\mathbb{R}^3$. *Circuit:* recognise the Lie algebra $\mathfrak{so}(3)$, or the imaginary quaternions. (correct and illuminating: the answer set is the image of the quaternion group $\{\pm 1, \pm i, \pm j, \pm k\}$, and the Jacobi identity explains why the structure is rigid; but no amount of structure theory removes the need for the finiteness argument, so this is theory bought at full price and used at a discount)
4. **Orthogonality first.** *Cue:* $v \times w$ is perpendicular to both v and w. *Circuit:* force all pairs of elements to be perpendicular, then count. (the right second move made first: to convert "$v \times w \in S$" into "$v \perp w$" you need to know that $\|v \times w\| = 1$, which requires the norm classification you have not yet done, so starting here stalls immediately)
5. **Norms first.** *Cue:* "exactly n elements" (finiteness) sitting next to an operation with a multiplicative norm law. *Circuit:* iterate the operation to build a geometric progression of norms, and let finiteness kill it. (ideal: one induction shows every nonzero element is a unit vector, after which orthogonality is a one line consequence of $\sin\theta = 1$ and the count is forced)

The ideal solution

The answers are $n = 1$ and $n = 7$.

Throughout, write $\|\cdot\|$ for the Euclidean norm and use the two defining properties of the cross product: $v \times w$ is orthogonal to both v and w, and $\|v \times w\| = \|v\|\|w\| \sin\theta$ where $\theta \in [0, \pi]$ is the angle between v and w. In particular $v \times w = 0$ exactly when v and w are linearly dependent.

Suppose S has exactly $n \geq 1$ elements and $S = \{v \times w : v, w \in S\}$. Since S is nonempty, choose $v \in S$; then $0 = v \times v \in S$. So $0 \in S$ always.

Case 1: every pair of elements of S is linearly dependent. Then every cross product of elements of S vanishes, so $S = \{0\}$ and $n = 1$. Conversely $S = \{0\}$ does satisfy $S = \{v \times w\}$, so $n = 1$ occurs.

Case 2: S contains two linearly independent vectors.

Claim 4.8.1. *Every nonzero $u \in S$ satisfies $\|u\| = 1$.*

Proof. Suppose $u \in S$ is nonzero with $\|u\| \neq 1$. Since S contains two linearly independent vectors, at least one of them, call it z, is not a scalar multiple of u. Define

$$u_1 = u \times z, \qquad u_{k+1} = u \times u_k \quad (k \geq 1).$$

Every u_k lies in S, because S is closed under cross products of its own elements. We show by induction that $u_k \neq 0$ and $u_k \perp u$ for all $k \geq 1$. For $k = 1$: u and z are independent, so $u_1 \neq 0$, and $u_1 \perp u$ by construction. If $u_k \neq 0$ and $u_k \perp u$, then u and u_k are independent (a nonzero vector orthogonal to u is not a multiple of u), so $u_{k+1} = u \times u_k \neq 0$, and $u_{k+1} \perp u$ by construction.

Because $u \perp u_k$ for $k \geq 1$, the sine factor is 1 and the norms satisfy exactly

$$\|u_{k+1}\| = \|u\|\,\|u_k\|, \qquad \text{so} \qquad \|u_k\| = \|u\|^{k-1}\|u_1\|.$$

Since $\|u_1\| > 0$ and $\|u\| \neq 1$, the numbers $\|u_k\|$ are pairwise distinct, so the vectors u_k are pairwise distinct. Thus S contains infinitely many elements, contradicting $|S| = n$. □

Claim 4.8.2. *Any two nonzero elements of S are either equal up to sign or orthogonal.*

Proof. Let $u, z \in S$ be nonzero; by the previous claim both are unit vectors. If they are linearly dependent then $z = \pm u$. If they are independent then $u \times z$ is a nonzero element of S, hence a unit vector, so

$$1 = \|u \times z\| = \|u\|\|z\| \sin\theta = \sin\theta,$$

forcing $\theta = \pi/2$, that is $u \perp z$. □

Now pick $v_1, v_2 \in S$ linearly independent, and set $v_3 = v_1 \times v_2 \in S$. By the claims, v_1 and v_2 are orthogonal unit vectors, so v_3 is a unit vector orthogonal to both, and (v_1, v_2, v_3) is an orthonormal basis of $\mathbb{R}^3$ with

$$v_1 \times v_2 = v_3, \quad v_2 \times v_3 = v_1, \quad v_3 \times v_1 = v_2.$$

Reversing the order of each product changes the sign, so S contains all seven of

$$0, \quad v_1, \quad v_2, \quad v_3, \quad -v_1 = v_3 \times v_2, \quad -v_2 = v_1 \times v_3, \quad -v_3 = v_2 \times v_1,$$

and these are pairwise distinct. Conversely, let $u \in S$ be nonzero. By the second claim u is, for each i, either $\pm v_i$ or orthogonal to v_i. If u were orthogonal to all of v_1, v_2, v_3, then u would be orthogonal to a spanning set of $\mathbb{R}^3$, forcing $u = 0$; so $u = \pm v_i$ for some i. Therefore

$$S = \{0, \pm v_1, \pm v_2, \pm v_3\}, \qquad n = 7.$$

It remains to confirm that such an S really does satisfy the condition, since the problem demands equality and not merely closure. Take any orthonormal basis (v_1, v_2, v_3) with $v_3 = v_1 \times v_2$ and put $T = \{0, \pm v_1, \pm v_2, \pm v_3\}$. A product with 0, or of two parallel elements, gives $0 \in T$. A product

$(\varepsilon_1 v_i) \times (\varepsilon_2 v_j)$ with $i \neq j$ and $\varepsilon_1, \varepsilon_2 \in \{1, -1\}$ equals $\pm v_k$ where k is the third index, and as the signs and the order of i and j vary, both v_k and $-v_k$ are attained. Hence $\{v \times w : v, w \in T\} = T$ exactly.

The possible values are therefore

$$\boxed{n = 1 \text{ and } n = 7}.$$

Layer 1: WHAT

Block 1: S is nonempty, and $v \times v = 0$, so $0 \in S$ with no work at all.

Block 2: split on the only structural dichotomy available: either all of S lies on one line through the origin, in which case every product dies and $S = \{0\}$, or S contains two independent vectors.

Block 3: ask what quantity the cross product controls, and answer: the norm, multiplicatively, exactly when the arguments are perpendicular. Build the orbit $u_1 = u \times z$, $u_{k+1} = u \times u_k$, which stays perpendicular to u forever, so the norms form a geometric progression with ratio $\|u\|$.

Block 4: finiteness forbids a nonconstant geometric progression, so $\|u\| = 1$ for every nonzero $u \in S$.

Block 5: with all nonzero elements on the unit sphere, $u \times z \in S$ forces $\sin\theta = 1$, so distinct directions in S are perpendicular.

Block 6: perpendicular unit vectors in $\mathbb{R}^3$ come in at most three directions, each contributing two elements; adding 0 gives exactly seven, and there is no room for an eighth.

Block 7: verify both directions of the equality for the seven element set.

The nominal route goes: "finite set, explicit operation, let me parametrise the elements and solve". It writes $v_i = (x_i, y_i, z_i)$, expands determinants, and drowns. The bifurcation point is Block 3, at the question *which scalar function of a vector does this operation multiply?* The zombie answers "components"; components are not multiplied by anything. The ideal answer is "the norm", and it is multiplied exactly, with no error term, precisely in the situation the operation itself manufactures (perpendicular arguments). Once you have a multiplicative quantity inside a finite set, the problem is over.

Layer 2: HOW

Block 1. **Cue.** $S = \{v \times w : v, w \in S\}$ with v and w ranging independently, so in particular $w = v$ is allowed. **Why natural.** Closure conditions are always tested first on the diagonal, because $v * v$ is the cheapest element you can manufacture. Here $v \times v = 0$ by antisymmetry, so the zero vector is in S for free. This also tells you $n = 1$ is a live candidate rather than a formality. **Link to the next block.** Knowing $0 \in S$, the real question is what the nonzero elements can be.

Block 2. **Cue.** The cross product of dependent vectors vanishes. **Why natural.** The operation has a large "kernel": any set of pairwise parallel vectors produces only 0. That is a genuinely different regime from the general one, so it must be split off before you argue about norms; otherwise the phrase "choose z not parallel to u" is unjustified. **Link to the next block.** In the nondegenerate case you now have two independent elements to feed the machine.

Block 3. **Cue.** "Exactly n elements" sitting next to a norm law. **Why natural.** Finiteness is never decoration. It is there because some construction inside the problem generates an infinite family unless a parameter is pinned. The only parameter in sight is $\|u\|$, and the only law is $\|u \times z\| = \|u\|\|z\| \sin\theta$. The sine is a nuisance in general, so you engineer it away: cross with u repeatedly, and after the first step every vector in the orbit is perpendicular to u, so $\sin\theta = 1$ permanently. **Link to the next block.** The orbit has norms $\|u\|^{k-1}\|u_1\|$, which is a geometric progression.

Block 4. Cue. A geometric progression of distinct positive numbers inside a finite set. **Why natural.** Distinct norms mean distinct vectors, and infinitely many of them. The ratio must therefore be 1. This is the same mechanism that forces a finite multiplicative subgroup of the positive reals to be trivial, and it is why every element of a finite set closed under a norm multiplicative operation lies on the unit sphere or at the origin. **Link to the next block.** All nonzero elements are unit vectors; now the sine factor becomes readable.

Block 5. Cue. $\|u \times z\| = \sin\theta$ once u and z are unit vectors, and the left side must be 1 because $u \times z$ is a nonzero element of S. **Why natural.** An equation between a norm and a sine converts algebra into geometry with no intermediate steps. $\sin\theta = 1$ has the unique solution $\theta = \pi/2$ in $[0, \pi]$. **Link to the next block.** Distinct directions are now orthogonal, which is a counting constraint.

Block 6. Cue. Pairwise orthogonal nonzero vectors in $\mathbb{R}^3$. **Why natural.** They are linearly independent, so there are at most three of them; each direction contributes exactly the two unit vectors $\pm v_i$; and any further element would have to be orthogonal to a spanning set, hence zero. The count $1 + 6 = 7$ is forced, not chosen. **Link to the next block.** Only the verification remains.

Block 7. Cue. The problem says $S = \{v \times w\}$, an equality of sets. **Why natural.** Everything so far proves the inclusion "S is contained in the octahedral set" and that the seven elements are produced. A grader will want the reverse inclusion checked explicitly, that is, that every one of the seven elements arises as a product and that no product escapes the set. Both are one line each, and skipping them is how a full solution becomes a nine.

Layer 3: WHY

The mathematical truth underneath is that the cross product carries a *multiplicative valuation*. Whenever an operation $*$ on a set admits a function N into the positive reals with $N(a * b) = N(a)N(b)$ on the relevant part of the domain, any finite $*$ closed set must have all its values of N equal to 1, because $N(a), N(a)^2, N(a)^3, \ldots$ would otherwise be infinitely many distinct values. You have met this before in three costumes: a finite subgroup of $\mathbb{R}^{>0}$ is trivial; a finite set of complex numbers closed under multiplication lies in $\{0\} \cup$ (roots of unity); a finite set of monic integer polynomials closed under multiplication must consist of degree zero elements. The cross product is a fourth costume, with the twist that its valuation is multiplicative only under a side condition (perpendicularity), and the operation itself supplies that side condition after one application. That is the reason the orbit $u_{k+1} = u \times u_k$ is the natural object rather than an arbitrary product of elements.

The discernment mechanism is a question you should ask automatically whenever the words "finite" and "closed" appear together: *what does this operation do to size?* If the answer is "multiplies it", finiteness forces size 1. If the answer is "adds to it" (degree, length, number of prime factors), finiteness forces the additive quantity to be zero. If the answer is "does not change it", you learn nothing and must find another invariant. Training this is a matter of building a reflex list of size functions: modulus, norm, determinant, degree, trace, order, cardinality of support.

What the examiners nudged toward is visible in the phrasing. They wrote "with exactly n elements" rather than "finite", which puts the cardinality in your face; and they wrote the condition as an equality of sets rather than as closure, which is the hint that the answer will be a rigid, self reproducing configuration rather than anything you can pad. The rigidity itself is not an accident: $\{0, \pm v_1, \pm v_2, \pm v_3\}$ is the set of imaginary parts of the quaternion group $\{\pm 1, \pm i, \pm j, \pm k\}$, and the reason no other configuration survives is that $(\mathbb{R}^3, \times)$ is a simple Lie algebra whose only finite self reproducing subsets are the trivial one and the standard octahedron.

Revolving concepts

The cross product. For $v, w \in \mathbb{R}^3$, $v \times w$ is the unique vector with $v \times w \perp v, w$, with $\|v \times w\| = \|v\|\|w\| \sin\theta$, and with $(v, w, v \times w)$ right handed. It is bilinear and antisymmetric, so $v \times v = 0$; it is not associative, and instead satisfies the Jacobi identity $u \times (v \times w) + v \times (w \times u) + w \times (u \times v) = 0$ and the triple product expansion $u \times (v \times w) = v(u \cdot w) - w(u \cdot v)$. See Book 1, field: Geometry, topic: vector product.

The map $L_u(x) = u \times x$ and its recursion. Fix $u \neq 0$. Then L_u kills the line $\mathbb{R}u$ and maps $u^\perp$ to itself; on that plane it is the composition of a rotation by $\pi/2$ about u with a scaling by $\|u\|$. The recursion $u_{k+1} = L_u(u_k)$ is therefore a linear recursion with matrix $\|u\|R$, R a rotation by a quarter turn, so in general

$$u_k = \|u\|^{k-1} R^{k-1} u_1, \qquad \|u_k\| = \|u\|^{k-1}\|u_1\|.$$

Generalise the pattern: for any linear map A acting on a subspace as a similarity of ratio r, the orbit of a nonzero vector has norms $r^k\|u_1\|$, so a finite orbit forces $r = 1$, and when $r = 1$ the orbit is finite exactly when the rotation part has rational angle. Here $r = 1$ gives $L_u^4 = \mathrm{id}$ on $u^\perp$, period 4, which is precisely why the answer set has $\pm$ pairs. See Book 1, field: Linear algebra, topic: orthogonal maps and orbits.

Multiplicative valuations and finiteness. If $N : X \to \mathbb{R}^{>0}$ satisfies $N(a * b) = N(a)N(b)$ and X is finite and closed under $*$, then $N \equiv 1$. Standard companions: a finite subgroup of the multiplicative group of a field is cyclic; a finite set of complex numbers closed under multiplication is μ_m or $\mu_m \cup \{0\}$. See Book 1, field: Abstract algebra, topic: finite closure arguments.

Orthogonal sets. Nonzero pairwise orthogonal vectors are linearly independent, so at most dim many; a vector orthogonal to a spanning set is zero. See Book 1, field: Linear algebra, topic: orthonormal bases.

Quaternions. For purely imaginary quaternions u, v one has $uv = -(u \cdot v) + u \times v$, so the cross product is the imaginary part of quaternion multiplication. The seven element answer is $\{0\} \cup \{\pm i, \pm j, \pm k\}$ in disguise. See Book 1, field: Abstract algebra, topic: quaternions.

Circuit Bank

- **WHEN** a problem gives a finite set closed under an operation **DO** ask at once what scalar quantity the operation multiplies, and iterate the operation on one element **BECAUSE** finiteness forbids a nonconstant geometric progression, so that quantity is pinned to 1 and the problem collapses to a rigidity statement
- **WHEN** an operation satisfies $v * v = e$ for a distinguished element e, and a set equals its own image **DO** write e into the set immediately **BECAUSE** the diagonal $w = v$ is always legal in a closure condition and is the cheapest element you can manufacture
- **WHEN** a cross product is forced to be a unit vector **DO** read the angle straight off $\|v \times w\| = \|v\|\|w\| \sin\theta$ **BECAUSE** the norm identity converts an algebraic membership condition into the geometric statement $\theta = \pi/2$ in a single line
- **WHEN** a norm law carries an awkward factor such as $\sin\theta$ **DO** engineer the factor to equal 1 by feeding the operation its own output **BECAUSE** the cross product outputs vectors perpendicular to its inputs, so the second application onwards is automatically in the clean case
- **WHEN** the problem states $S = \{\text{image}\}$ rather than image $\subseteq S$ **DO** verify both inclusions explicitly at the end **BECAUSE** an equality of sets is two statements, and the reverse inclusion is where constructed examples are usually checked or lost

Backtesting

1. First useful cue: the words "exactly n elements" next to an operation with a multiplicative norm law; finiteness is the active hypothesis, not the question.
2. Method the expression suggested: iterate $u \mapsto u \times u_k$ to manufacture perpendicular arguments, turning $\|v \times w\| = \|v\|\|w\| \sin\theta$ into an exact geometric progression of norms.
3. Where you would hesitate, and the thought to reinforce: at the start, when the problem looks like a search over configurations. Reinforce this: you are not searching, you are proving that a scalar is pinned. Configuration searches in $\mathbb{R}^3$ are hopeless; scalar rigidity arguments are three lines.
4. The trap the question sets: forgetting that the condition is an equality, so that a set which is merely closed under the cross product (for example any set of parallel vectors together with 0) looks like a counterexample; and, inside the norm argument, forgetting to check $u_k \neq 0$ at every step, without which the "distinct norms" claim is empty.
5. What makes you recognise it faster next time: the pairing of a finiteness hypothesis with any operation that has a size law. That pairing has exactly one purpose in every problem that contains it.
6. The reusable decision rule: finite plus closed plus multiplicative size equals size 1. Then use the structural properties (here orthogonality) to count.

Mirror drills

1. Find all finite sets $S \subset \mathbb{C}$ with $S = \{zw : z, w \in S\}$. *Answer:* the modulus is multiplicative, so finiteness forces every nonzero element onto the unit circle; the nonzero part is then a finite nonempty subset of the circle group closed under multiplication, hence a subgroup, hence μ_m. The answers are $\{0\}$, μ_m and $\mu_m \cup \{0\}$ for $m \geq 1$.
2. Suppose only that $S \subset \mathbb{R}^3$ is finite and $\{v \times w : v, w \in S\} \subseteq S$. Classify S. *Answer:* the same two claims apply verbatim, so either all elements of S are parallel (and then S is any such set that contains 0 if it has at least two elements, since all products vanish), or $S = \{0, \pm v_1, \pm v_2, \pm v_3\}$ for an orthonormal basis. The equality hypothesis was only ever used to rule out padding by parallel vectors.
3. For a unit vector u and any $x \perp u$, compute $u \times (u \times (u \times (u \times x)))$. *Answer:* x. On $u^\perp$ the map L_u is a quarter turn, so L_u^4 is the identity; this period 4 is exactly why the seven element set closes up after $\pm$ pairs.

4.9 B3. Iterated recolouring of the positive reals

Assign to each positive real number a colour, either red or blue. Let D be the set of all distances $d > 0$ such that there are two points of the same colour at distance d apart. Recolour the positive reals so that the numbers in D are red and the numbers not in D are blue. If we iterate this recolouring process, will we always end up with all the numbers red after a finite number of steps?

Read it out loud

A two colouring of the positive reals is fed into a machine that returns a new two colouring: a number is red in the output exactly when it occurs as the gap between two points of the same colour in the input. The question is whether every orbit of this machine reaches the all red colouring in finitely many steps. Three fields are wired in: colouring combinatorics of the Ramsey type (monochromatic pairs at prescribed distances), the elementary arithmetic of the reals under translation, and the dynamics of an iterated map on the space of colourings.

Notice the double role played by the positive reals. They are simultaneously the points being coloured and the distances being classified. That is the whole design of the problem: a statement about the colours of the points $x/2$ and $3x/2$ becomes a statement about the colour of the number x, because $3x/2 - x/2 = x$. Any solution has to exploit that self reference; nothing else is available.

The decisive cue is that the question asks about the colour red, which is defined by an existential statement ("there exist two points of the same colour at distance d"), and existential statements are weak. So chase the other side. A number d is blue exactly when

$$\chi(t) \neq \chi(t+d) \quad \text{for every } t > 0,$$

a statement quantified over all t, and therefore enormously strong: it says the input colouring is a perfect alternation along every arithmetic progression of common difference d. Within thirty seconds you should have written that line down, and then the immediate consequence: two colours means "different" composed with "different" is "same", so $\chi(t+2d) = \chi(t)$ for all t, and hence $2d$ is red in the output. Blue numbers are already fighting each other.

The answer is yes, and in fact two steps always suffice; one step does not, so the constant 2 is sharp.

Approaches, ranked worst to best

1. **Hunt for a counterexample colouring.** *Cue:* the phrasing "will we always". *Circuit:* construct a colouring whose orbit never turns all red. (zombie route: you find, in a minute, that the colouring "red if $\lfloor t \rfloor$ is even" leaves the number 1 blue after one step, and you then spend an hour trying to iterate that idea into a genuine counterexample; the search cannot terminate because there is no counterexample, and it teaches you nothing about why)

2. **Measure theoretic difference sets.** *Cue: D* is a union of difference sets of the colour classes. *Circuit:* apply the Steinhaus theorem: a set of positive measure has a difference set containing an interval about 0. (fails as stated: the colouring is completely arbitrary, so the colour classes need not be measurable, and the hypothesis of the Steinhaus theorem is unavailable; even after patching, it would only give an interval of red numbers, not all of them)

3. **Ramsey and van der Waerden machinery.** *Cue:* monochromatic configurations at prescribed gaps. *Circuit:* quote a theorem guaranteeing monochromatic arithmetic progressions. (heavy artillery aimed at a target that a three point pigeonhole hits: the classical theorems are about colourings of the integers and would need translation and care, and they still would not organise the two step structure)

4. **Doubling only.** *Cue: d* blue forces $2d$ red. *Circuit:* iterate the doubling relation to fill up the line. (a genuine and correct first lemma, and it does half the work, but doubling alone cannot reach a number from below; you can never conclude that a given x is red from statements about $2x, 4x, 8x, \ldots$, so the argument stalls one case short)
5. **Alternation plus the three point pigeonhole.** *Cue:* blue is a for all statement; three points and two colours force a monochromatic pair. *Circuit:* convert blue into perfect alternation, then split any blue gap at its midpoint. (ideal: two short lemmas, then a four point chain $x, 2x, 3x, 4x$ splits into exactly two cases, one killed by doubling and one killed by halving)

The ideal solution

The answer is yes: for every initial colouring, all positive reals are red after two steps, and they stay red thereafter.

Write a colouring as a function χ from $\mathbb{R}^+$ to $\{\text{red}, \text{blue}\}$. For a colouring χ let

$$D(\chi) = \{d > 0 : \chi(t) = \chi(t+d) \text{ for some } t > 0\},$$

which is exactly the set D of the problem, and let $\Phi(\chi)$ be the colouring whose red set is $D(\chi)$. Let χ_0 be the initial colouring, $\chi_n = \Phi^n(\chi_0)$, and write R_n and B_n for the red and blue sets of χ_n, so that $R_{n+1} = D(\chi_n)$ and B_{n+1} is its complement in $\mathbb{R}^+$.

Unwinding the definition gives the only fact we shall use about blue numbers:

$$d \in B_{n+1} \iff \chi_n(t) \neq \chi_n(t+d) \text{ for every } t > 0. \tag{$*$}$$

Lemma 4.9.1 (Doubling)**.** *If $d \in B_{n+1}$ then $2d \in R_{n+1}$.*

Proof. By $(*)$, $\chi_n(t) \neq \chi_n(t+d)$ and $\chi_n(t+d) \neq \chi_n(t+2d)$ for every $t > 0$. There are only two colours, so $\chi_n(t) = \chi_n(t+2d)$. Thus t and $t+2d$ are two points of the same colour at distance $2d$, so $2d \in D(\chi_n) = R_{n+1}$. □

Lemma 4.9.2 (Splitting)**.** *If $d \in B_{n+1}$ then for every s with $0 < s < d$, at least one of s and $d - s$ lies in R_{n+1}. In particular, taking $s = d/2$, we get $d/2 \in R_{n+1}$.*

Proof. Fix any $t > 0$ and look at the three points t, $t+s$, $t+d$. By $(*)$ the outer two have different colours, so the middle one has the same colour as exactly one of them. If $\chi_n(t+s) = \chi_n(t)$ then $s \in D(\chi_n)$; if $\chi_n(t+s) = \chi_n(t+d)$ then $d - s \in D(\chi_n)$. The special case $s = d/2$ makes the two alternatives identical, so no case split survives. □

Theorem 4.9.3. $R_2 = \mathbb{R}^+$.

Proof. Suppose for contradiction that $x \in B_2$ for some $x > 0$. By $(*)$ applied at level $n = 1$, $\chi_1(t) \neq \chi_1(t+x)$ for all $t > 0$. Applying this at $t = x, 2x, 3x$ gives

$$\chi_1(x) \neq \chi_1(2x), \qquad \chi_1(2x) \neq \chi_1(3x), \qquad \chi_1(3x) \neq \chi_1(4x),$$

so with two colours the four numbers split as $\chi_1(x) = \chi_1(3x)$ and $\chi_1(2x) = \chi_1(4x)$, and these two common colours are different. Exactly one of the following two cases holds.

Case A: $2x$ and $4x$ are blue. Apply the Doubling Lemma to $2x \in B_1$: it gives $4x \in R_1$, contradicting $4x \in B_1$.

Case B: x and $3x$ are blue. Apply the Splitting Lemma at the midpoint to each of them: $x/2 \in R_1$ and $3x/2 \in R_1$. So $x/2$ and $3x/2$ are two points of the same colour under χ_1, and their distance is

$$\frac{3x}{2} - \frac{x}{2} = x.$$

Hence $x \in D(\chi_1) = R_2$, contradicting $x \in B_2$.

Both cases are impossible, so B_2 is empty and $R_2 = \mathbb{R}^+$. □

Finally, the all red colouring is a fixed point: if every positive real is red, then for every $d > 0$ the points t and $t + d$ are both red, so $d \in D$ and the recolouring returns all red again. Hence $R_n = \mathbb{R}^+$ for every $n \geq 2$, and the answer to the question is

$$\boxed{\text{yes, and two steps always suffice}}.$$

The constant 2 cannot be improved. Colour t red when $\lfloor t \rfloor$ is even and blue when $\lfloor t \rfloor$ is odd. Then $\lfloor t+1 \rfloor = \lfloor t \rfloor + 1$ for every t, so t and $t+1$ always receive different colours and $1 \notin D(\chi_0)$; that is, $1 \in B_1$ and one step is not enough.

Layer 1: WHAT

Block 1: name the machine and unwind what blue means: d is blue in the output exactly when the input colouring alternates with step d, a statement quantified over all t.

Block 2: two colours turn alternation into periodicity, $\chi(t + 2d) = \chi(t)$, so a blue d forces $2d$ to be red. Blue numbers cannot contain a number together with its double.

Block 3: three points and two colours force a monochromatic pair, so a blue d forces at least one of s and $d - s$ to be red for each $s \in (0, d)$; choosing $s = d/2$, the fixed point of $s \mapsto d - s$, removes the ambiguity and gives $d/2$ red.

Block 4: assume x is blue after two steps. Alternation with step x splits $\{x, 2x, 3x, 4x\}$ into the pair $\{x, 3x\}$ and the pair $\{2x, 4x\}$, monochromatic and of opposite colours in χ_1.

Block 5: if $\{2x, 4x\}$ is the blue pair, doubling applied to $2x$ makes $4x$ red. Contradiction.

Block 6: if $\{x, 3x\}$ is the blue pair, halving makes $x/2$ and $3x/2$ both red; their difference is exactly x, so x is red at the next step. Contradiction.

Block 7: conclude, note stability of the all red colouring, and exhibit the floor parity colouring to show one step is genuinely not enough.

The nominal route reads "will we always" as an invitation to search for a counterexample, and spends its time constructing colourings. The bifurcation point is Block 1: the decision to argue about blue rather than red. Red is an existential property, so knowing a number is red gives you one point pair somewhere on the line, chosen by an adversary. Blue is a universal property, so knowing a number is blue gives you information at every single point of the line. In a problem where you must derive strong conclusions, you must chase the strong hypothesis, which means arguing by contradiction from the complement.

Layer 2: HOW

Block 1. **Cue.** "D is the set of all distances d such that there are two points...", an existential definition, and a question that asks for everything to become red. **Why natural.** To prove a set is everything, assume some element is missing and derive a contradiction. The missing element is blue, and negating the existential quantifier turns the definition inside out into a for all statement about the previous colouring. Writing $(*)$ down is not a step, it is a translation, and translations of quantifiers are always free. **Link to the next block.** You now have a functional equation for the colouring, $\chi_n(t) \neq \chi_n(t + d)$ for all t, and functional equations are things you compose with themselves.

Block 2. **Cue.** "either red or blue": exactly two colours. **Why natural.** With two colours, "different" is an involution: apply the alternation twice and you return to the original colour. So a colouring that alternates with step d is periodic with period $2d$, and periodicity manufactures monochromatic pairs at distance $2d$ for free. This is the only place in the entire solution where the number of colours matters, and it matters absolutely: with three colours the lemma is false. **Link to the next block.** Doubling gives you an upward mechanism; you will also need a downward one.

Block 3. **Cue.** Three points, two colours. **Why natural.** The pigeonhole in its smallest form. If you place a third point anywhere strictly between t and $t+d$, it must agree with one of the ends, and the two possible gaps are s and $d-s$. That is an "or", and an "or" is a nuisance. The way to remove it is to make the two alternatives coincide, which happens exactly at the midpoint $s = d/2$, the fixed point of the involution $s \mapsto d-s$. This is why $3x/2$ appears in the standard write ups: it is the midpoint of x and $2x$, nothing more. **Link to the next block.** You now have both a doubling and a halving rule for blue numbers.

Block 4. **Cue.** The assumption $x \in B_2$ gives alternation with step x for the colouring χ_1, and you have two rules that convert blue numbers into red ones. **Why natural.** Alternation with step x means the colours of $x, 2x, 3x, 4x, \ldots$ alternate, so the odd multiples share a colour and the even multiples share the other. Testing four consecutive multiples is the smallest window that contains both a doubling pair ($2x$ and $4x$) and a halving pair (x and $3x$) whose halves differ by x. You take four because the target distance is x and you need $3x/2 - x/2$. **Link to the next block.** The case split is now forced by the two colours, and each case has a rule waiting for it.

Block 5. **Cue.** $2x$ and $4x$ both blue. **Why natural.** The Doubling Lemma applied to the smaller one immediately says the larger is red. This is why doubling exists: it forbids a blue number and its double from coexisting, and the even multiples in an alternating chain are exactly such a pair. **Link to the next block.** One case down.

Block 6. **Cue.** x and $3x$ both blue, and you are trying to prove something about the number x itself. **Why natural.** The Splitting Lemma halves them, and $3x/2 - x/2 = x$ is exactly the distance you need. This is the self referential heart of the problem: the colours of two points produce the colour of a number, and the number produced is the very one you assumed blue. **Link to the next block.** Contradiction in both cases, so the assumption was false.

Block 7. **Cue.** The question says "after a finite number of steps". **Why natural.** A finite bound invites you to state the exact one and to show it is best possible. The all red colouring is a fixed point (trivially), and the floor parity example shows one step can fail, so 2 is the answer and the solution is complete rather than merely correct.

Layer 3: WHY

The mathematical truth beneath this problem is that being blue is almost impossible. A blue number d certifies that the colouring is a proper 2 colouring of the graph on $\mathbb{R}^+$ whose edges join points at distance d. That graph is a disjoint union of two way infinite paths, one for each coset of $d\mathbb{Z}$, so such colourings exist in abundance, which is why the first step can fail. But a blue number at the second stage demands that the colouring χ_1, which is itself already defined by a difference condition, be one of these rigid alternating colourings, and the alternating structure then forces relations between the colours of $x/2, x, 3x/2, 2x, 3x, 4x$ that cannot all hold. The problem is a two step rigidity cascade: an arbitrary colouring is unconstrained, its recolouring is mildly constrained, and the recolouring of that is over constrained.

The discernment mechanism is a rule about quantifiers. When a problem defines a set by "the values that are realised", the useful hypothesis is always non membership, because non membership is universally quantified and gives you information at every point, whereas membership hands you a single witness chosen by an adversary. Train this by making it mechanical: before doing anything else, write both the definition and its negation in full quantifier form, and look at which of the two you would rather have as a hypothesis. Nine times out of ten it is the negation, and the proof will be by contradiction.

The second mechanism is smaller but reusable: when a pigeonhole produces an "or" between two quantities s and $d - s$ that you get to choose, choose the symmetric point where the two coincide. Symmetrisation to kill a disjunction is a standard reflex, and it is why the solution never has to consider four cases instead of two.

What the examiners nudged toward is visible in two places. They specified exactly two colours, which is what makes doubling work, and they asked for a finite number of steps rather than for the limit behaviour, which signals that the true answer is an explicit small integer. A problem that says "finite" almost always has answer 1, 2 or 3, and the write up should name the number.

Revolving concepts

Difference sets and distance sets. If A_{red} and A_{blue} partition $\mathbb{R}^+$, then $D = ((A_{\text{red}} - A_{\text{red}}) \cup (A_{\text{blue}} - A_{\text{blue}})) \cap \mathbb{R}^+$. Difference sets are the natural home of this problem. The classical theorem in the area is the Steinhaus difference theorem: if $A \subseteq \mathbb{R}$ has positive Lebesgue measure then $A - A$ contains an open interval around 0. It cannot be used here because the colouring is arbitrary and its classes need not be measurable, which is a knowledge gap worth naming rather than hiding. See Book 1, field: Real analysis, topic: difference sets and Steinhaus.

Distance graphs and proper colourings. For fixed $d > 0$, let G_d be the graph on $\mathbb{R}^+$ with $t \sim t + d$. Each connected component is a path, so G_d is bipartite and has proper 2 colourings; a number d is blue for χ exactly when χ is a proper colouring of G_d. Generalising to the graph with edges at distances in a set E, the chromatic number of the line with a forbidden distance set is a large classical topic (the Hadwiger and Nelson problem is its planar cousin). See Book 1, field: Discrete structures and graphs, topic: distance graphs.

Alternation as a recursion, and its generalisation. The condition "$\chi(t + d) \neq \chi(t)$ for all t" is the recursion $u_{k+1} = u_k + 1$ over $\mathbb{F}_2$ along each coset $t + d\mathbb{Z}$, with solution $u_k = u_0 + k$ and hence period 2. The general statement is: if a colouring by m colours satisfies $\chi(t + d) = \sigma(\chi(t))$ for a fixed permutation σ of the colours, then $\chi(t + kd) = \sigma^k(\chi(t))$, and χ is periodic with period $\text{ord}(\sigma) \cdot d$. With $m = 2$ the only fixed point free permutation is the transposition, of order 2, which is why $2d$ is forced red. With $m = 3$ the permutation can be a 3 cycle, the period is $3d$, and the doubling lemma is false: colour t by $\lfloor t \rfloor \bmod 3$ and both 1 and 2 fail to be realised as monochromatic distances. See Book 1, field: Combinatorics, topic: colourings and periodicity.

Iterated maps on colourings and fixed points. The recolouring Φ acts on the set of all 2 colourings. Our theorem says Φ^2 is the constant map onto the all red colouring. Consequently the all red colouring is the unique fixed point of Φ, and there are no periodic orbits of any length: if $\Phi(\chi) = \chi$ then $\chi = \Phi^2(\chi)$, which is all red. See Book 1, field: Logic and set theory, topic: iterated maps and fixed points.

Quantifier negation. $d \notin \{d : \exists t,\ \chi(t) = \chi(t + d)\}$ means $\forall t,\ \chi(t) \neq \chi(t + d)$. Elementary, and the single most valuable move in the problem. See Book 1, field: Logic and set theory, topic: quantifiers.

Circuit Bank

- **WHEN** a set is defined as "the values that are realised" and you must prove it is everything **DO** assume an element is missing and write the negation as a for all statement **BECAUSE** membership gives you one adversarial witness, non membership gives you information at every point, so the complement is always the stronger hypothesis
- **WHEN** a two colouring satisfies $\chi(t) \neq \chi(t + d)$ for all t **DO** write $\chi(t + 2d) = \chi(t)$ at once **BECAUSE** with exactly two colours "different" composed with "different" is "same", so alternation is periodicity of period $2d$ and monochromatic pairs at distance $2d$ appear for free
- **WHEN** a pigeonhole gives you an "or" between two quantities s and $d - s$ with s at your

disposal **DO** choose the symmetric value $s = d/2$ **BECAUSE** the midpoint is the fixed point of $s \mapsto d - s$, so the disjunction collapses and no case analysis survives

- **WHEN** a process is iterated and the question says "after a finite number of steps" **DO** find the unique fixed point, prove an explicit bound, and produce an example showing the bound is sharp **BECAUSE** a finiteness question of this shape almost always has a small explicit answer, and the sharpness example is what turns a correct proof into a complete one

Backtesting

1. First useful cue: D is defined by an existential quantifier and the target is "everything is in D", so the working hypothesis must be that some d is not, written out in full as a for all statement.
2. Method the expression suggested: convert blue into alternation, then extract two rules, one that pushes upward (doubling, from two colours) and one that pushes downward (halving, from the three point pigeonhole at the midpoint).
3. Where you would hesitate, and the thought to reinforce: at the choice of the window $x, 2x, 3x, 4x$. Reinforce this: you are trying to prove that x itself is red, so you need two same coloured points differing by exactly x; the halves $x/2$ and $3x/2$ are the cheapest such pair, and they come from the odd multiples x and $3x$, so the window must reach as far as $4x$ to close the case split.
4. The trap the question sets: the phrase "will we always" invites a counterexample hunt, and the floor parity colouring rewards that hunt with a false positive at step one. Recognise that a one step obstruction is not an obstruction.
5. What makes you recognise it faster next time: any problem in which the objects being coloured and the objects being classified live in the same set is self referential by design, and the solution will use an arithmetic identity connecting a pair of points to the number they define (here $3x/2 - x/2 = x$).
6. The reusable decision rule: negate the defining quantifier, exploit that two colours make alternation an involution, and symmetrise to kill disjunctions.

Mirror drills

1. Show that there is a two colouring of $\mathbb{R}^+$ with no two points of the same colour at distance 1, and prove that any such colouring satisfies $\chi(t+2) = \chi(t)$ for all t. *Answer:* colour t by the parity of $\lfloor t \rfloor$; the alternation condition composed with itself gives period 2.
2. Show that the Doubling Lemma fails for three colours, and identify the exact line of the proof that breaks. *Answer:* colour t by $\lfloor t \rfloor$ mod 3; neither 1 nor 2 is realised as a monochromatic distance, so a blue number and its double coexist. The broken line is "there are only two colours, so $\chi(t) = \chi(t+2d)$": with three colours the alternation can be a 3 cycle of period 3.
3. Prove that the all red colouring is the only colouring fixed by the recolouring map Φ. *Answer:* if $\Phi(\chi) = \chi$ then $\chi = \Phi^2(\chi)$, and Φ^2 is constantly all red by the theorem.

4.10 B4. Cyclic sequences whose triples are progressions

Find all integers n with $n \geq 4$ for which there exists a sequence of distinct real numbers $x_1, \ldots, x_n$ such that each of the sets

$$\{x_1, x_2, x_3\}, \{x_2, x_3, x_4\}, \ldots,$$
$$\{x_{n-2}, x_{n-1}, x_n\}, \{x_{n-1}, x_n, x_1\}, \text{ and } \{x_n, x_1, x_2\}$$

forms a 3-term arithmetic progression when arranged in increasing order.

Read it out loud

A cyclic sequence of pairwise distinct reals in which every window of three consecutive terms, read as an unordered set, is an arithmetic progression. The fields wired in are elementary number theory (a 2-adic scale and a reduction mod 3), the combinatorics of cyclic overlapping windows, and an explicit construction problem. As always with "find all n", the work splits into an obstruction and a construction, and the answer will be a congruence class with the small cases inspected by hand.

The key structural feature is the overlap. Consecutive windows share two of their three elements, so their common differences cannot be unrelated. In a 3-term arithmetic progression with common difference d, the distance between any two members is d or $2d$, and nothing else. Two consecutive windows must therefore agree about the gap $|x_{m+1} - x_{m+2}|$, which forces the ratio d_{m+1}/d_m into $\{1/2, 1, 2\}$. That single line is the engine: it says all the common differences sit on one multiplicative lattice $2^{\mathbb{Z}} d_1$, which is exactly the situation in which you normalise.

Within thirty seconds you should have written three things: that a set $\{a, b, c\}$ is an arithmetic progression exactly when one of the three is the average of the other two (equivalently, $a + b + c$ is three times the middle term); that consecutive common differences are related by a factor $1/2$, 1 or 2; and that scaling and translating the whole sequence changes nothing, so you may put the smallest common difference at 1 and start counting.

The answer is: exactly the multiples of 3 that are at least 9.

Approaches, ranked worst to best

1. **Case analysis on which term is the middle.** *Cue:* the arithmetic progression condition is a linear equation once you know the middle term. *Circuit:* write $x_m + x_{m+2} = 2x_{m+1}$ for each choice and solve the linear system. (zombie route: there are three choices of middle term in each of n windows, so 3^n linear systems; even for $n = 6$ that is 729 systems, and the exponential structure hides the fact that the choices are heavily constrained by the overlaps)

2. **Computer or hand search.** *Cue:* "find all integers n". *Circuit:* enumerate small cases and guess the pattern. (reconnaissance only: it tells you $n = 9$ works and $n = 6$ does not, which is most of the answer and none of the proof, and it will not produce the general construction)

3. **Transfer matrix on window patterns.** *Cue:* each window is labelled by which of its three positions holds the middle term, and consecutive labels are constrained. *Circuit:* encode admissible label words and count closed walks. (workable and genuinely close to a solution: the ratio law is exactly the transition rule of this automaton, and the cyclic closure condition is a closed walk; but the bookkeeping is heavy and you still need an arithmetic invariant to rule out $3 \nmid n$ and a construction for the survivors)

4. **Multiplicative closure of the ratios.** *Cue:* $d_{m+1}/d_m \in \{1/2, 1, 2\}$ and the cycle returns to its start. *Circuit:* take logarithms and note the exponents sum to zero around the cycle. (true, useful, and worth stating as a remark, but on its own it gives no congruence condition on n; it constrains the pattern of doublings, not the length)

5. **Normalise, then reduce modulo 3.** *Cue:* all common differences are powers of 2, hence units modulo 3. *Circuit:* scale the smallest difference to 1, shift to integers, and observe each window is a complete residue system mod 3. (ideal: three lines give $3 \mid n$, a short forced chase kills $n = 6$, and an explicit concatenation of three progressions covers every remaining case)

The ideal solution

The answer is: the multiples of 3 that are at least 9, that is $n \in \{9, 12, 15, \dots\}$. Throughout, "distinct" means pairwise distinct, indices are read modulo n, and T_m denotes the window $\{x_m, x_{m+1}, x_{m+2}\}$. Since the x_i are pairwise distinct, each T_m has three distinct elements, so it is an arithmetic progression with a common difference $d_m > 0$: sorted, $T_m = \{t_m, t_m + d_m, t_m + 2d_m\}$.

Lemma 4.10.1 (Ratio law)**.** *For every m, $d_{m+1} \in \{\frac{1}{2}d_m,\ d_m,\ 2d_m\}$.*

Proof. The distance between two distinct members of a 3-term progression with common difference d is either d or $2d$. The elements x_{m+1} and x_{m+2} lie in both T_m and T_{m+1}, so

$$|x_{m+1} - x_{m+2}| \in \{d_m, 2d_m\} \cap \{d_{m+1}, 2d_{m+1}\},$$

and the four possible matches give $d_{m+1} = d_m$, $d_{m+1} = 2d_m$, $d_{m+1} = \frac{1}{2}d_m$ or $d_{m+1} = d_m$ again. □

Lemma 4.10.2 (Normalisation)**.** *If a sequence exists for a given n, then one exists in which every d_m is a power of 2 with $\min_m d_m = 1$, every x_i is an integer, and the minimum is attained at $m = 1$.*

Proof. The hypothesis is invariant under $x_i \mapsto \lambda x_i + c$ for $\lambda \neq 0$, under cyclic rotation of the indices, and under reversal of the sequence, since all of these carry the family of cyclic windows to itself. By the ratio law every d_m equals $2^{k_m} d_1$ for some integer k_m, so the d_m take finitely many values on one multiplicative lattice and a smallest one exists; rotate so that it occurs at $m = 1$ and scale so that it equals 1. Then every d_m is a power of 2, in particular a positive integer. Consecutive terms lie in a common window, so $x_{m+1} - x_m \in \{\pm d_m, \pm 2d_m\} \subset \mathbb{Z}$, whence every x_i differs from x_1 by an integer; translate by $-x_1$. □

Proposition 4.10.3. *If such a sequence exists then $3 \mid n$.*

Proof. Normalise as above. Each d_m is a power of 2 and is therefore not divisible by 3, so in the progression $t_m, t_m + d_m, t_m + 2d_m$ the three residues modulo 3 are pairwise distinct: each window T_m is a complete residue system modulo 3. Now T_m and T_{m+1} share x_{m+1} and x_{m+2}, which occupy two of the three residue classes, so x_{m+3} and x_m must both occupy the remaining one:

$$x_{m+3} \equiv x_m \pmod 3 \qquad \text{for all } m.$$

(The same conclusion follows from the fact that the sum of a 3-term progression is three times its middle term, so $x_m + x_{m+1} + x_{m+2} \equiv 0 \pmod 3$ for every m, and subtracting consecutive instances gives $x_{m+3} \equiv x_m$.) Thus the map $m \mapsto x_m \bmod 3$ on $\mathbb{Z}/n$ is invariant under the shift by 3. If $3 \nmid n$ then 3 is a generator of $\mathbb{Z}/n$, so that map is constant, contradicting the fact that x_1, x_2, x_3 lie in three different residue classes. Hence $3 \mid n$. □

Proposition 4.10.4. *There is no such sequence with $n = 6$.*

Proof. Normalise as above, so $d_1 = 1 = \min_m d_m$, every d_m is a power of 2, the x_i are integers, and $T_1 = \{x_1, x_2, x_3\}$ is a set of three consecutive integers; translating, $T_1 = \{0, 1, 2\}$. The reversal $x_i \mapsto x_{4-i}$ (indices mod 6) preserves the hypothesis and exchanges x_1 with x_3, and the reflection $x \mapsto 2 - x$ preserves the hypothesis and fixes T_1. Under the group these generate, the six orderings of $\{0, 1, 2\}$ fall into the two classes represented by $(x_1, x_2, x_3) = (0, 1, 2)$ and $(x_1, x_2, x_3) = (0, 2, 1)$.

In both cases $\{x_2, x_3\} = \{1, 2\}$, a pair at distance 1. Since every $d_m \geq 1$, a gap of 1 inside T_2 must equal d_2 (it cannot equal $2d_2 \geq 2$), so $d_2 = 1$ and T_2 is a set of three consecutive integers containing 1 and 2 adjacently: either $\{0, 1, 2\}$, forcing $x_4 = 0 = x_1$, or $\{1, 2, 3\}$. Distinctness leaves $x_4 = 3$.

Case A: $(x_1, x_2, x_3) = (0, 1, 2)$. Then $x_3 = 2$ and $x_4 = 3$ are at distance 1, so as above $d_3 = 1$ and $T_3 \in \{\{1, 2, 3\}, \{2, 3, 4\}\}$, giving $x_5 \in \{1, 4\}$; but $1 = x_2$, so $x_5 = 4$. Repeating with $\{x_4, x_5\} = \{3, 4\}$ gives $d_4 = 1$ and $x_6 \in \{2, 5\}$; but $2 = x_3$, so $x_6 = 5$. The window $\{x_5, x_6, x_1\} = \{0, 4, 5\}$ has gaps 4 and 1, so it is not an arithmetic progression.

Case B: $(x_1, x_2, x_3) = (0, 2, 1)$. Then $x_3 = 1$ and $x_4 = 3$ are at distance 2, so $d_3 \in \{1, 2\}$. If $d_3 = 1$ then 1 and 3 are the two endpoints of T_3 and $x_5 = 2 = x_2$, excluded. So $d_3 = 2$ and $T_3 \in \{\{-1, 1, 3\}, \{1, 3, 5\}\}$, giving $x_5 \in \{-1, 5\}$.

If $x_5 = -1$, then x_4 and x_5 are at distance 4, so $d_4 \in \{2, 4\}$; $d_4 = 2$ makes -1 and 3 the endpoints and $x_6 = 1 = x_3$, excluded, so $d_4 = 4$ and $x_6 \in \{-5, 7\}$. The window $\{x_5, x_6, x_1\}$ is then $\{-5, -1, 0\}$ with gaps 4 and 1, or $\{-1, 0, 7\}$ with gaps 1 and 7; neither is an arithmetic progression.

If $x_5 = 5$, then x_4 and x_5 are at distance 2, so $d_4 \in \{1, 2\}$, giving $x_6 = 4$ (from $d_4 = 1$) or $x_6 \in \{1, 7\}$ (from $d_4 = 2$), of which $1 = x_3$ is excluded. The window $\{x_5, x_6, x_1\}$ is then $\{0, 4, 5\}$ with gaps 4 and 1, or $\{0, 5, 7\}$ with gaps 5 and 2; neither is an arithmetic progression.

Every branch fails, so no sequence exists for $n = 6$. □

Proposition 4.10.5. *For every integer $m \geq 2$ there is such a sequence with $n = 3m + 3$.*

Proof. Concatenate the three progressions

$$(1, 3, 5, \ldots, 4m - 1), \quad (4m - 2), \quad (4m, 4m - 4, \ldots, 4, 0), \quad (2),$$

that is, take the $2m$ odd numbers from 1 to $4m - 1$ in increasing order, then $4m - 2$, then the $m + 1$ multiples of 4 from $4m$ down to 0 in decreasing order, then 2. The length is $2m + 1 + (m + 1) + 1 = 3m + 3$.

The terms are distinct: the first block consists of odd numbers, the third block of multiples of 4, and the two remaining terms $4m - 2$ and 2 are congruent to 2 modulo 4 and are distinct because $m \geq 2$.

Every cyclic window is an arithmetic progression. Windows inside the first block are $\{4i - 3, 4i - 1, 4i + 1\}$, of common difference 2. Windows inside the third block are $\{4j - 4, 4j, 4j + 4\}$, of common difference 4. The seven remaining windows, in order, are

$$\{4m - 3, 4m - 1, 4m - 2\}, \quad \{4m - 1, 4m - 2, 4m\}, \quad \{4m - 2, 4m, 4m - 4\},$$
$$\{4m, 4m - 4, 4m - 8\}, \quad \{4, 0, 2\}, \quad \{0, 2, 1\}, \quad \{2, 1, 3\},$$

with common differences $1, 1, 2, 4, 2, 1, 1$ respectively (the fourth is already of the interior type when $m \geq 2$). All are arithmetic progressions. □

Combining, the values of $n \geq 4$ for which a sequence exists are

$$\boxed{n \equiv 0 \pmod 3 \text{ with } n \geq 9}.$$

For orientation, the case $m = 2$ reads $1, 3, 5, 7, 6, 8, 4, 0, 2$.

Layer 1: WHAT

Block 1: read the condition as a statement about unordered sets, so each window has a well defined positive common difference d_m.

Block 2: consecutive windows overlap in two terms, and in a progression every internal gap is d or $2d$; this forces $d_{m+1}/d_m \in \{1/2, 1, 2\}$.

Block 3: therefore all d_m lie on the lattice $2^{\mathbb{Z}} d_1$. Scale the smallest to 1 and rotate it to position 1; every difference becomes a power of 2 and every term becomes an integer after a shift.

Block 4: powers of 2 are units modulo 3, so each window occupies all three residue classes; hence $x_{m+3} \equiv x_m \pmod 3$, and on a cycle of length n this forces $3 \mid n$ unless the residues are constant, which they are not.

Block 5: for $n = 6$ the normalisation pins $\{x_1, x_2, x_3\} = \{0, 1, 2\}$, the symmetries cut six orderings to two, and the minimality of d_1 forces every subsequent term; five branches survive to x_6 and all five fail the closing window.

Block 6: for $n = 3m + 3$ with $m \geq 2$, concatenate an increasing progression of odds, a single term, a decreasing progression of multiples of 4, and one more term; check the seven junction windows and the distinctness by residues modulo 4.

The nominal route fixes attention on the terms x_i, of which there are n unknowns, and on the choice of middle term in each window, of which there are 3^n patterns. The bifurcation point is Block 2: look at the differences, not the terms. The differences are a single positive number together with a walk of exponents in $\{-1, 0, 1\}$, which is one unknown and a bounded walk. That is the reduction that makes everything else possible, and it is available only because the windows overlap in two places rather than one.

Layer 2: HOW

Block 1. **Cue.** "forms a 3-term arithmetic progression when arranged in increasing order", and "distinct real numbers". **Why natural.** Sorting is what makes the common difference well defined and positive; distinctness is what stops d_m from being 0. Naming d_m turns a geometric condition on a set into a numerical parameter you can chase around the cycle. **Link to the next block.** You now have n parameters $d_1, \ldots, d_n$ attached to overlapping windows, so the next question is how consecutive ones interact.

Block 2. **Cue.** Consecutive windows share two elements. **Why natural.** The only distances available inside a 3-term progression of difference d are d and $2d$. A shared pair therefore has to be describable by both windows, which is a two by two matching and gives a ratio in $\{1/2, 1, 2\}$. Whenever consecutive constraints overlap in all but one element, the correct first move is always to write the compatibility relation between their local parameters. **Link to the next block.** The differences live on a multiplicative lattice, which is the standard setup for a normalisation.

Block 3. **Cue.** A finite set of positive reals closed under ratios that are powers of 2, plus a problem invariant under scaling and translation. **Why natural.** Scaling is free, so put the smallest difference at 1; then every difference is a positive integer power of 2. Consecutive terms differ by $\pm d_m$ or $\pm 2d_m$, all integers, so after one translation the whole sequence is integral. Normalising is not cosmetic: it is what makes a congruence argument legal, since congruences require integers. **Link to the next block.** Integers plus differences that are powers of 2 is a request for a modulus in which 2 is a unit.

Block 4. **Cue.** Every common difference is a power of 2, so no d_m is divisible by 3. **Why natural.** In $\{t, t + d, t + 2d\}$ with $\gcd(d, 3) = 1$, the three residues mod 3 are distinct, so a window is a complete residue system. Two consecutive windows share two elements, so the third elements agree modulo 3, and the residue pattern is 3 periodic. On $\mathbb{Z}/n$ a shift by 3 generates everything unless $3 \mid n$, and a constant residue pattern is impossible. The choice of modulus is forced: you want the smallest modulus in which all possible differences are invertible and in which a 3-term progression is spread out, and 3 is it. **Link to the next block.** Only the boundary cases $n = 6$ and the construction remain.

Block 5. **Cue.** $n = 6$ is the one multiple of 3 in range that the construction will not reach, and the normalisation has already pinned the first window. **Why natural.** With $d_1 = 1$ minimal and $T_1 = \{0, 1, 2\}$, every later step is forced: a gap of 1 between two known consecutive terms cannot

be $2d$, so $d = 1$ and there are only two candidate windows, one of which repeats an earlier term. Distinctness is doing the pruning, which is why the chase terminates instead of branching. Reversal and the reflection $x \mapsto 2 - x$ cut the six starting orders to two before the chase begins, which is the difference between two cases and six. **Link to the next block.** The obstruction is complete; now produce examples.

Block 6. Cue. A construction is needed for every multiple of 3 from 9 upwards, so it must be a family with a free parameter. **Why natural.** Long stretches of a single progression give interior windows for free, so the design problem is only about the junctions. You need to leave one progression and enter another with two junction windows that are themselves progressions, and you need the cycle to close through two extra terms. Choosing the blocks to be odd numbers and multiples of 4, with the two connectors congruent to 2 modulo 4, makes distinctness a residue check rather than an inequality chase.

Layer 3: WHY

Two invariants run this problem, one multiplicative and one additive, and they are linked by the normalisation.

The multiplicative invariant is the ratio law. Writing $d_{m+1} = 2^{\varepsilon_m} d_m$ with $\varepsilon_m \in \{-1, 0, 1\}$ and going once around the cycle gives $\sum_{m=1}^{n} \varepsilon_m = 0$: the number of doublings equals the number of halvings. This is the general shape of every cyclic problem with local multiplicative constraints. Take logarithms, and a local constraint becomes a bounded walk that must return to its starting point.

The additive invariant is the residue pattern modulo 3. It is available only because the ratios are powers of 2, so that 2 is a unit in the chosen modulus and a progression cannot collapse into fewer than three residue classes. The general principle: to convert a statement about gaps into a statement about labels, reduce modulo a q in which every possible gap is invertible; then arithmetic progressions become rainbow patterns and periodicity arguments become available.

The discernment mechanism is a question you should ask of any cyclic or chain problem: what is the local compatibility relation between consecutive constraints, and what does it force after one full lap? Here the lap gives two things, the exponent sum zero and the residue period 3, and the second is the one that constrains n. Training this means always converting a global existence question into a one step recursion plus a closure condition, before attempting any construction.

What the examiners nudged toward is visible in the phrase "when arranged in increasing order". Had the windows been progressions in the given order, the sequence would have constant difference d and the cycle would force $nd = 0$, killing everything at once; the reordering is exactly what creates the three way choice, hence the ratio law, hence the whole problem. The lower bound $n \geq 4$ is also a nudge: it quietly removes $n = 3$, where the three windows are all the same set, so that the only interesting small case left for you to inspect is $n = 6$.

Revolving concepts

Three term progressions. A set $\{a, b, c\}$ of distinct reals is an arithmetic progression exactly when one element is the mean of the other two, equivalently when $a + b + c$ is three times one of them. Consequences used here: the internal distances are d and $2d$ only; the sum is three times the middle term, which is why sums of consecutive windows differ by $x_{m+3} - x_m$ and give the mod 3 periodicity in one line. See Book 1, field: Sequences and recurrences, topic: arithmetic progressions.

Overlapping window recursions and their generalisation. A constraint imposed on every window of length k in a cyclic sequence becomes a one step recursion on the local parameters, together with a closure condition after n steps. Here the recursion is $d_{m+1} = 2^{\varepsilon_m} d_m$ with $\varepsilon_m \in \{-1, 0, 1\}$, whose general solution is $d_m = 2^{\varepsilon_1 + \cdots + \varepsilon_{m-1}} d_1$, and the closure condition is $\sum_{m=1}^{n} \varepsilon_m = 0$.

The same pattern governs every cyclic problem with local multiplicative constraints: take logarithms, obtain a bounded walk, and impose that it returns to its start. See Book 1, field: Sequences and recurrences, topic: cyclic recursions and closure.

Normalisation by the symmetry group. The hypothesis is invariant under the affine maps $x \mapsto \lambda x + c$ with $\lambda \neq 0$, under cyclic rotation, and under reversal. Spending this group to place the extremal object (the smallest common difference) at a standard position is the standard extremal normalisation, and it is what reduces the $n = 6$ analysis from six starting orders to two. See Book 1, field: Combinatorics, topic: normalisation and extremal choice.

Reduction modulo a unit. If every difference appearing in a configuration is invertible modulo q, then progressions occupy distinct residues mod q, and metric statements become colouring statements. Here $q = 3$ and the differences are powers of 2. The general slogan: choose the modulus so that the step sizes become units. See Book 1, field: Number theory, topic: choosing a modulus.

Cyclic periodicity. A function on $\mathbb{Z}/n$ invariant under the shift by p is constant on the cosets of the subgroup generated by p, which has index $\gcd(p, n)$. So p periodicity plus non constancy forces $\gcd(p, n) > 1$, and here $p = 3$ gives $3 \mid n$. See Book 1, field: Number theory, topic: cyclic groups and periods.

Circuit Bank

- **WHEN** consecutive constraints overlap in all but one element **DO** write the compatibility relation between their local parameters immediately **BECAUSE** overlapping windows convert a global search over all terms into a one step recursion on a single parameter, which is the difference between 3^n cases and one lemma
- **WHEN** local parameters are related by bounded multiplicative factors and the problem is scale invariant **DO** take the smallest, scale it to 1, and check that everything becomes integral **BECAUSE** a multiplicative lattice with a minimum is exactly the setup for a reduction to integers, and congruence arguments are illegal until you have integers
- **WHEN** all the step sizes in a configuration are units modulo some small q **DO** reduce modulo q and look at which residues each configuration occupies **BECAUSE** a progression whose difference is invertible occupies distinct residues, so a metric condition becomes a colouring condition and periodicity arguments become available
- **WHEN** a cyclic object forces a pattern of period p **DO** conclude $\gcd(p, n) > 1$ or the pattern is constant **BECAUSE** the shift by p on $\mathbb{Z}/n$ generates a subgroup of index $\gcd(p, n)$, so a nonconstant p periodic pattern cannot exist when $\gcd(p, n) = 1$
- **WHEN** the answer is a congruence class with small exceptions **DO** handle the smallest case by a forced chase from the normalised extremal position, and build the general case as blocks with only junctions to check **BECAUSE** minimality plus distinctness prunes the chase to a handful of branches, and interior windows of a long block are free

Backtesting

1. First useful cue: consecutive windows share two elements, and in a 3-term progression only the gaps d and $2d$ occur; that pairing gives the ratio law before any algebra.
2. Method the expression suggested: pass from the terms to the common differences, normalise the smallest to 1, become integral, and then reduce modulo 3.
3. Where you would hesitate, and the thought to reinforce: at the choice of modulus. Reinforce this: the modulus is not guessed, it is determined by the requirement that every possible difference be a unit and that a progression have three distinct residues. Powers of 2 are units

modulo every odd number, and 3 is the smallest odd modulus in which a progression is a complete residue system.

4. The trap the question sets: assuming the answer is "all multiples of 3" and never testing $n = 6$; the construction needs two connector terms with the same residue mod 4, and at $m = 1$ they collide, which is exactly why 6 fails.
5. What makes you recognise it faster next time: cyclic sequences with a local condition on every window are always a one step recursion plus a closure condition; if the local condition is multiplicative, take logarithms, and if it is additive, sum around the cycle.
6. The reusable decision rule: differences before terms, normalise before congruences, and always inspect the smallest member of the surviving congruence class by hand.

Mirror drills

1. Show that in any such sequence, the number of indices m with $d_{m+1} = 2d_m$ equals the number with $d_{m+1} = \frac{1}{2}d_m$. *Answer:* write $d_{m+1} = 2^{\varepsilon_m} d_m$ and sum ε_m around the cycle; the sum telescopes to $\log_2(d_1/d_1) = 0$.
2. For which $n \geq 3$ is there a cyclic sequence of distinct reals with x_m, x_{m+1}, x_{m+2} an arithmetic progression in that order for every m? *Answer:* none. The condition says $x_{m+1} - x_m$ is a constant d, and summing around the cycle gives $nd = 0$, so $d = 0$ and the terms coincide. This is exactly what the phrase "when arranged in increasing order" is there to prevent.
3. Verify that $1, 3, 5, 7, 6, 8, 4, 0, 2$ works and identify the three blocks and the junction windows. *Answer:* the blocks are $(1, 3, 5, 7)$ of difference 2, $(8, 4, 0)$ of difference 4, and the connectors 6 and 2; the junction windows are $\{5, 7, 6\}, \{7, 6, 8\}, \{6, 8, 4\}, \{4, 0, 2\}, \{0, 2, 1\}, \{2, 1, 3\}$, with differences $1, 1, 2, 2, 1, 1$.

4.11 B5. When zero is the most likely weighted sum

For $0 \leq p \leq 1/2$, let $X_1, X_2, \ldots$ be independent random variables such that

$$X_i = \begin{cases} 1 & \text{with probability } p, \\ -1 & \text{with probability } p, \\ 0 & \text{with probability } 1-2p, \end{cases}$$

for all $i \geq 1$. Given a positive integer n and integers $b, a_1, \ldots, a_n$, let $P(b, a_1, \ldots, a_n)$ denote the probability that $a_1X_1 + \cdots + a_nX_n = b$. For which values of p is it the case that

$$P(0, a_1, \ldots, a_n) \geq P(b, a_1, \ldots, a_n)$$

for all positive integers n and all integers $b, a_1, \ldots, a_n$?

Read it out loud

A lazy symmetric three point random walk step, rescaled by arbitrary integer weights and summed, with the question: for which p is the origin always the most likely outcome, whatever the weights? The fields wired in are probability (independence and distributions of integer valued sums), generating functions (a probability of an exact integer value is a coefficient), and harmonic analysis on the circle (the coefficient is a Fourier coefficient, and "largest at 0" is a positivity statement).

What is really being tested is whether you can name the right closed class. The property "the distribution has its maximum at 0" is a statement about one coefficient among many, and coefficientwise statements do not survive convolution. So the reflex must be: transform. Write the distribution of $a_1X_1 + \cdots + a_nX_n$ as a product, because independence is multiplicative after transforming, and then ask what property of a single factor is both multiplicative and strong enough to force the maximum to sit at 0.

Within thirty seconds you should have written

$$E\left[x^S\right] = \prod_{k=1}^{n} f(x^{a_k}), \qquad f(x) = px + px^{-1} + (1-2p),$$

and then set $x = e^{i\theta}$ to get $f(e^{i\theta}) = 1 - 2p + 2p\cos\theta$, a real function whose minimum over θ is $1 - 4p$. The threshold in the answer is already visible: $p = 1/4$ is exactly where a single step stops having a nonnegative transform, equivalently where the palindromic quadratic $pt^2 + (1-2p)t + p$ stops having real roots, its discriminant being $(1-2p)^2 - 4p^2 = 1 - 4p$.

The answer is $0 \leq p \leq 1/4$.

Approaches, ranked worst to best

1. **Injection between representations.** *Cue:* two probabilities are being compared, and each is a sum over solutions of a linear equation. *Circuit:* build an injection from the solutions of $\sum a_k X_k = b$ into the solutions of $\sum a_k X_k = 0$. (zombie route: the terms carry unequal weights, since a solution with j nonzero coordinates has probability $p^j(1-2p)^{n-j}$, so an injection proves nothing unless it does not decrease the number of zero coordinates, and no such injection exists in general; the counterexample family for $p > 1/4$ is precisely a case where the target has more solutions than the source)

2. **Induction on n with a unimodality hypothesis.** *Cue:* sums of independent symmetric variables. *Circuit:* prove the distribution is symmetric and unimodal, hence maximal at 0. (fails at the first step: the weights a_k put the summands on different lattices, so unimodality is not available, and the class of distributions whose maximum sits at 0 is not visibly closed under convolution, so the induction has no invariant class to live in)

3. **Threshold hunting by examples.** *Cue:* "for which values of p". *Circuit:* test small configurations until the inequality breaks. (necessary reconnaissance, and it does deliver the exact threshold if you pick the right family, but it proves only one of the two halves)
4. **Laurent polynomial factorisation.** *Cue:* $f(x) = px + px^{-1} + (1-2p)$ is palindromic. *Circuit:* factor $f(x) = \frac{p}{\alpha}(x+\alpha)(x^{-1}+\alpha)$ and write the whole product as $g(x)g(x^{-1})$, then compare coefficients by Cauchy and Schwarz. (correct, fully elementary, and it exposes the discriminant $1-4p$ as the source of the threshold; it ranks just below the transform route only because it needs a separate word for $p=0$ and a normalisation from Laurent polynomials to genuine polynomials)
5. **Characteristic functions.** *Cue:* a probability of an exact integer value is a Fourier coefficient, and independence is multiplication. *Circuit:* prove the transform of one step is nonnegative and note that nonnegativity is multiplicative. (ideal: three standard facts, no edge cases, and the threshold appears as the minimum of $1-2p+2p\cos\theta$)

The ideal solution

The answer is

$$\boxed{0 \le p \le \tfrac{1}{4}}.$$

Write $S = a_1X_1 + \cdots + a_nX_n$, an integer valued random variable taking finitely many values.

Part 1: every p with $1/4 < p \le 1/2$ fails.

First dispose of $p = 1/2$: take $n = 1$, $a_1 = 1$, $b = 1$. Then $P(0,1) = 1-2p = 0 < 1/2 = P(1,1)$.

Now let $1/4 < p < 1/2$ and take $a_k = 2^{k-1}$ for $k = 1, \dots, n$, with $b = 1$.

The only $(X_1, \dots, X_n) \in \{0, \pm 1\}^n$ with $\sum_k 2^{k-1}X_k = 0$ is the zero vector. Indeed, if some coordinate is nonzero, let j be the largest index with $X_j \neq 0$; then $\left|\sum_{k<j} 2^{k-1}X_k\right| \le 2^{j-1} - 1$, so the total has absolute value at least $2^{j-1} - (2^{j-1}-1) = 1$. Hence

$$P(0, a_1, \dots, a_n) = (1-2p)^n.$$

Next, the solutions of $\sum_k 2^{k-1}X_k = 1$ are exactly the n vectors

$$(\underbrace{-1, \dots, -1}_{j-1}, 1, 0, \dots, 0), \qquad j = 1, \dots, n,$$

since $2^{j-1} - (2^{j-2} + \cdots + 1) = 1$. To see there are no others, let j again be the largest index with $X_j \neq 0$. If $X_j = -1$ the sum is at most $-2^{j-1} + (2^{j-1}-1) = -1 < 1$, so $X_j = 1$ and the remaining coordinates must sum to $1 - 2^{j-1} = -(2^{j-1}-1)$, the least value they can attain, which happens only when $X_k = -1$ for every $k < j$. The vector indexed by j has exactly j nonzero coordinates, so

$$P(1, a_1, \dots, a_n) = \sum_{j=1}^{n} p^j (1-2p)^{n-j}.$$

Dividing, and writing $r = \dfrac{p}{1-2p} > 0$ (legitimate since $p < 1/2$),

$$\frac{P(1, a_1, \dots, a_n)}{P(0, a_1, \dots, a_n)} = \sum_{j=1}^{n} r^j.$$

Now $r > 1/2$ if and only if $2p > 1-2p$, that is $p > 1/4$. If $r \ge 1$ the partial sums are at least n and exceed 1 as soon as $n \ge 2$; if $1/2 < r < 1$ the partial sums increase to $r/(1-r) > 1$, so they exceed 1 for all large n. Either way there is an n with $P(1, a_1, \dots, a_n) > P(0, a_1, \dots, a_n)$, so p fails.

Part 2: every p with $0 \le p \le 1/4$ works.

For an integer valued random variable X taking finitely many values, define its characteristic function

$$\varphi_X(\theta) = E\left[e^{iX\theta}\right] = \sum_{\ell\in\mathbb{Z}} \Pr(X=\ell)\, e^{i\ell\theta},$$

a finite trigonometric sum. We use three facts, each one line.

Inversion. For every integer b,

$$\frac{1}{2\pi}\int_0^{2\pi} e^{-ib\theta}\varphi_X(\theta)\,d\theta = \sum_{\ell}\Pr(X=\ell)\cdot\frac{1}{2\pi}\int_0^{2\pi} e^{i(\ell-b)\theta}\,d\theta = \Pr(X=b),$$

because the inner integral is 1 when $\ell = b$ and 0 otherwise.

Independence. If X and Y are independent then $\varphi_{X+Y} = \varphi_X\varphi_Y$, since $E[e^{i(X+Y)\theta}] = E[e^{iX\theta}]E[e^{iY\theta}]$.

Scaling. For an integer a, $\varphi_{aX}(\theta) = E[e^{iaX\theta}] = \varphi_X(a\theta)$.

For our steps,

$$\varphi_{X_k}(\theta) = (1-2p) + pe^{i\theta} + pe^{-i\theta} = 1 - 2p + 2p\cos\theta \geq 1 - 4p \geq 0$$

for all θ, using $p \leq 1/4$. Hence, by independence and scaling,

$$\varphi_S(\theta) = \prod_{k=1}^{n}\varphi_{X_k}(a_k\theta) = \prod_{k=1}^{n}\big(1 - 2p + 2p\cos(a_k\theta)\big) \geq 0.$$

Finally, for any integer b, inversion and $|e^{-ib\theta}| = 1$ give

$$P(b,a_1,\ldots,a_n) = \frac{1}{2\pi}\int_0^{2\pi} e^{-ib\theta}\varphi_S(\theta)\,d\theta \leq \frac{1}{2\pi}\int_0^{2\pi}|\varphi_S(\theta)|\,d\theta = \frac{1}{2\pi}\int_0^{2\pi}\varphi_S(\theta)\,d\theta = P(0,a_1,\ldots,a_n),$$

where the last two steps use $\varphi_S \geq 0$ and inversion at $b = 0$. This holds for every n and all integers $b, a_1, \ldots, a_n$, which completes the proof.

Layer 1: WHAT

Block 1: recognise $P(b,a_1,\ldots,a_n)$ as a coefficient. Independence says the distribution of S is a convolution, and convolutions become products only after a transform, so transform first and ask questions later.

Block 2: on the unit circle the transform of one step is $1 - 2p + 2p\cos\theta$, a real function, and the transform of S is the product of such factors at the arguments $a_k\theta$.

Block 3: identify the correct multiplicative property. "The coefficient at 0 is the largest" is not multiplicative; "the function is nonnegative" is, and it implies the former, because the coefficient at 0 is the average of the function and every other coefficient is bounded by the average of the modulus.

Block 4: the single step transform has minimum $1 - 4p$ at $\theta = \pi$, so nonnegativity holds exactly when $p \leq 1/4$. Sufficiency is finished.

Block 5: for necessity, design the weights so that 0 is starved: powers of 2 make 0 uniquely representable while $b = 1$ has a chain of n representations whose probabilities form a geometric sequence with ratio $r = p/(1-2p)$.

Block 6: the ratio of the two probabilities is $\sum_{j=1}^{n} r^j$, which can exceed 1 exactly when $r > 1/2$, that is when $p > 1/4$. The two halves meet at the same number by two different looking computations, and Layer 3 explains why they must.

The nominal route reads the problem as a comparison of two counts and starts building injections between solution sets. The bifurcation point is Block 1: the decision to stop looking at individual outcomes and to encode the whole distribution in one function. Everything after that is forced, because there is exactly one property of a single factor that both multiplies and controls coefficients.

Layer 2: HOW

Block 1. Cue. Independent summands, integer values, and a probability of an exact value. **Why natural.** Independence is a multiplicative hypothesis, but only after you leave the space of distributions and enter the space of transforms. A probability of an exact integer value is a coefficient, and coefficients of products are convolutions, so the dictionary runs in the direction you need. Writing $E[x^S] = \prod f(x^{a_k})$ costs one line and converts the entire hypothesis into a formula. **Link to the next block.** You now need to decide where to evaluate x.

Block 2. Cue. The Laurent polynomial $f(x) = px + px^{-1} + (1-2p)$ is palindromic, that is invariant under $x \mapsto x^{-1}$. **Why natural.** Palindromic means real on the unit circle, since $x^{-1} = \bar{x}$ there. That is the cue to set $x = e^{i\theta}$: the object becomes a real valued function of a real variable, and real valued functions have signs, and signs are things you can use. Off the circle the product is a complex mess with no order structure. **Link to the next block.** On the circle the coefficient extraction becomes an integral.

Block 3. Cue. You must prove one coefficient dominates all the others, uniformly over n and over the weights. **Why natural.** Uniform statements over an unbounded family need a property closed under the operation that builds the family, and the operation here is multiplication of transforms. So ask which properties are multiplicative. "Nonnegative on the circle" is multiplicative and immediately gives coefficient domination through the triangle inequality for integrals, since the zeroth coefficient is the mean of the function and every other coefficient is at most the mean of its modulus. That is the whole engine, and it is the standard notion of a positive definite sequence. **Link to the next block.** It remains to check nonnegativity for a single factor.

Block 4. Cue. $1 - 2p + 2p\cos\theta$ with $\cos\theta$ ranging over $[-1, 1]$. **Why natural.** The minimum is at $\theta = \pi$, where the two symmetric atoms at ± 1 interfere destructively with each other and constructively against the atom at 0. The value $1 - 4p$ is a discriminant in disguise, and the condition $p \le 1/4$ is exactly the condition under which the step distribution is a self convolution of a signed measure, that is, a square. **Link to the next block.** Sufficiency is complete; the threshold now has to be shown sharp.

Block 5. Cue. To break "the mode is at 0", starve the origin. **Why natural.** $P(0)$ is a sum over all representations of 0, so make that sum a single term. Weights $1, 2, 4, \ldots, 2^{n-1}$ do it, because binary expansions with digits $0, \pm 1$ represent 0 only trivially. Meanwhile $b = 1$ picks up the whole chain $1, -1+2, -1-2+4$, and so on: the borrow chain of binary arithmetic. Choosing the weights so that one target has a unique representation and a neighbouring target has a geometric family is the standard design for this kind of counterexample. **Link to the next block.** The two probabilities are now a single term against a geometric sum.

Block 6. Cue. A geometric series with ratio $r = p/(1-2p)$ compared to 1. **Why natural.** $\sum_{j \ge 1} r^j > 1$ exactly when $r > 1/2$, and $r > 1/2$ is $2p > 1 - 2p$, that is $p > 1/4$. No case analysis on p is needed and no asymptotics beyond the elementary geometric series, provided you divide by $P(0)$ first instead of expanding both sides.

Layer 3: WHY

The mathematical truth underneath is the theory of positive definite sequences. A sequence $(c_b)_{b \in \mathbb{Z}}$ is the Fourier coefficient sequence of a nonnegative function on the circle exactly when it is positive definite, and every positive definite sequence satisfies $|c_b| \le c_0$. Our problem asks for the values of p making the step distribution positive definite, because positive definiteness is preserved by

the two operations the problem applies: independent sums (products of transforms) and integer rescaling $X \mapsto aX$ (which replaces θ by $a\theta$). The property actually asked about, "the maximum is at 0", is a consequence of positive definiteness but is not itself closed under these operations, which is why every attempt to induct directly on n stalls. This is the general lesson: when a property must survive an operation, strengthen it to the property the transform makes multiplicative.

The two routes to sufficiency are the same proof in two alphabets. Since $p \le 1/4$, the palindromic quadratic $pt^2 + (1-2p)t + p$ has real roots α and $1/\alpha$ with $\alpha > 0$, so for $p > 0$

$$f(x) = \frac{p}{\alpha}\,(x+\alpha)(x^{-1}+\alpha), \qquad \prod_k f(x^{a_k}) = g(x)\,g(x^{-1}), \quad g(x) = \Big(\frac{p}{\alpha}\Big)^{n/2} \prod_k (x^{a_k}+\alpha),$$

after multiplying g by a power of x to clear negative exponents, which does not change $g(x)g(x^{-1})$. Writing $g(x) = \sum_j c_j x^j$, the coefficient of x^0 in $g(x)g(x^{-1})$ is $\sum_j c_j^2$ and the coefficient of x^b is $\sum_j c_j c_{j+b}$, and $\sum_j c_j c_{j+b} \le \sum_j c_j^2$ by Cauchy and Schwarz. Setting $x = e^{i\theta}$ turns $g(x)g(x^{-1})$ into $|g(e^{i\theta})|^2$, so the factorisation and the nonnegativity of the transform are literally the same statement; the general fact behind it is the Fejer and Riesz theorem, that every nonnegative trigonometric polynomial is $|g(e^{i\theta})|^2$ for some polynomial g. That is also why the threshold appears twice, once as a discriminant $1-4p$ and once as a geometric ratio $r \le 1/2$: both are the statement that the quadratic factors over the reals.

The discernment mechanism is short. When a problem quantifies over an unbounded family built by one operation, find the transform under which the operation becomes multiplication, then find the multiplicative property that implies what you want. To train it, keep a list of the transforms and the operations they linearise: generating functions for convolution of sequences, characteristic functions for sums of independent variables, logarithms for products, determinants for composition of maps.

What the examiners nudged toward is the number $1/2$ in the range $0 \le p \le 1/2$, which is there only to make $1-2p$ a probability, and the phrase "for all positive integers n", which tells you the answer is a uniform statement and hence that the counterexample must be a family growing with n, not a single configuration.

Revolving concepts

Characteristic functions of integer valued variables. $\varphi_X(\theta) = E[e^{iX\theta}]$ is a finite trigonometric sum; inversion recovers $\Pr(X = b)$ as the bth Fourier coefficient; independence gives $\varphi_{X+Y} = \varphi_X \varphi_Y$; and $\varphi_{aX}(\theta) = \varphi_X(a\theta)$. All four are one line proofs from the definition. See Book 1, field: Probability, topic: characteristic functions.

Positive definiteness. A sequence is positive definite exactly when it is the Fourier coefficient sequence of a nonnegative function (the Herglotz theorem), and then $|c_b| \le c_0$. The class of integer valued random variables with $\varphi \ge 0$ is closed under independent sums and integer scalings, which is what makes it the right invariant class here. See Book 1, field: Probability, topic: positive definite sequences.

Palindromic Laurent polynomials and Fejer and Riesz. A Laurent polynomial invariant under $x \mapsto x^{-1}$ is real on the unit circle; it is nonnegative there exactly when it equals $g(x)g(x^{-1})$ for a polynomial g. For the quadratic $px + px^{-1} + (1-2p)$ this is the condition that $pt^2 + (1-2p)t + p$ has real roots, that is $1 - 4p \ge 0$. See Book 1, field: Polynomials, topic: reciprocal polynomials.

Signed binary representations. With digits in $\{0, \pm 1\}$ and weights 2^{k-1}, the number 0 has only the trivial representation, and 1 has exactly the n representations given by the borrow chain $2^{j-1} - 2^{j-2} - \cdots - 1 = 1$. The general fact: weights that grow at least as fast as the partial sums of their predecessors make representations essentially unique, which is the standard way to design a target with a prescribed number of preimages. See Book 1, field: Number theory, topic: digit systems.

The recursion behind the geometric sum, generalised. Let $Q_n = P(1; 1, 2, \ldots, 2^{n-1}) = \sum_{j=1}^{n} p^j (1-2p)^{n-j}$. Splitting on whether the last weight is used gives the first order linear recursion

$$Q_n = (1-2p)\, Q_{n-1} + p^n, \qquad Q_0 = 0.$$

The general solution of $u_n = \alpha u_{n-1} + \beta^n$ with $u_0 = 0$ is $u_n = \beta\, \dfrac{\beta^n - \alpha^n}{\beta - \alpha}$ when $\alpha \neq \beta$, and $u_n = n\beta^n$ when $\alpha = \beta$. Here $\alpha = 1-2p$ and $\beta = p$ give $Q_n = p\dfrac{(1-2p)^n - p^n}{1-3p}$, and the degenerate case $\alpha = \beta$ is exactly $p = 1/3$. Dividing by $(1-2p)^n$ first, as in the ideal solution, replaces all of this by a geometric series and removes the degenerate case entirely. See Book 1, field: Sequences and recurrences, topic: first order linear recursions with geometric forcing.

Circuit Bank

- **WHEN** a probability that a sum of independent integer valued variables equals an exact value **DO** write it as a coefficient of a Laurent polynomial, or as a Fourier coefficient of a product **BECAUSE** independence becomes multiplication only after transforming, and exact value conditions become coefficient extractions
- **WHEN** you must show the coefficient at 0 dominates every other coefficient **DO** prove the generating function is nonnegative on the unit circle **BECAUSE** the zeroth coefficient is the mean of the function and every other coefficient is bounded by the mean of its modulus, so nonnegativity is exactly the right strengthening
- **WHEN** a property must survive an operation applied unboundedly often **DO** strengthen it to a property that the relevant transform makes multiplicative **BECAUSE** only multiplicative classes are closed under products, and a class that is not closed gives an induction with nothing to induct on
- **WHEN** you must break a claim that the mode sits at 0 **DO** choose weights making 0 uniquely representable while a neighbouring value collects a geometric family of representations **BECAUSE** the ratio of the two probabilities then becomes a geometric series, and the threshold is where its sum crosses 1
- **WHEN** a threshold of the form $1/4$ or $1-4p$ appears **DO** look for the discriminant of a palindromic quadratic **BECAUSE** thresholds in positivity problems are almost always the point where a factorisation stops being real

Backtesting

1. First useful cue: independence plus an exact integer target, which is the standing invitation to write the probability as a coefficient and then evaluate on the unit circle.
2. Method the expression suggested: transform, note that the single step transform is $1-2p+2p\cos\theta$, and read the threshold off its minimum.
3. Where you would hesitate, and the thought to reinforce: at the step "nonnegativity of the transform implies the maximum is at 0". Reinforce this: the zeroth Fourier coefficient is the average of the function, and averaging a nonnegative function against a unimodular factor can only shrink it. One line, no theory quoted.
4. The trap the question sets: believing the inequality can be proved combinatorially by comparing solution counts. The counts genuinely go the wrong way for $p > 1/4$, so any purely combinatorial argument must be false; the weights matter, not the counts.

5. What makes you recognise it faster next time: any question of the form "is the distribution of a weighted sum maximal at 0" is a positive definiteness question, and the answer is always governed by the sign of the single step transform.
6. The reusable decision rule: transform first, strengthen to the multiplicative property, and design counterexamples by starving the target you want to be small.

Mirror drills

1. Compute the characteristic function of one step and find exactly the p for which it is nonnegative for all θ. *Answer:* $\varphi(\theta) = 1 - 2p + 2p\cos\theta$, with minimum $1 - 4p$ at $\theta = \pi$; nonnegative for all θ exactly when $p \leq 1/4$.
2. With $p = 1/3$, so that each X_i is uniform on $\{-1, 0, 1\}$, compare $\Pr(X_1 + 2X_2 = 0)$ and $\Pr(X_1 + 2X_2 = 1)$. *Answer:* $1/9$ against $2/9$; the representations of 1 are $(1, 0)$ and $(-1, 1)$, while 0 has only $(0, 0)$.
3. Show that if X and Y are independent integer valued variables with $\varphi_X \geq 0$ and $\varphi_Y \geq 0$, then $\varphi_{X+Y} \geq 0$ and $\varphi_{aX} \geq 0$ for every integer a. *Answer:* $\varphi_{X+Y} = \varphi_X \varphi_Y$, a product of nonnegative functions, and $\varphi_{aX}(\theta) = \varphi_X(a\theta)$, a reparametrisation. This is precisely the closure that the naive property "maximum at 0" lacks.

4.12 B6. A continuous functional equation on the positive reals

Find all continuous functions $f : \mathbb{R}^+ \to \mathbb{R}^+$ such that

$$f(xf(y)) + f(yf(x)) = 1 + f(x+y)$$

for all $x, y > 0$. Here $\mathbb{R}^+$ denotes the set of positive real numbers, excluding 0.

Read it out loud

A symmetric functional equation in which the left side feeds products $xf(y)$ into f and the right side feeds a sum $x+y$ into f, with an additive constant 1 sitting between them. The fields wired in are functional equations, but the tools that actually do the work are all real analysis: limits superior and inferior, the intermediate value theorem applied to the image of an interval, compactness on closed bounded intervals, and the iteration of an affine contraction. No algebraic identity will finish this problem. It is an analysis problem wearing a functional equation costume.

Three structural features are worth naming before you compute anything. First, the equation is symmetric under swapping x and y, so the diagonal $y = x$ is special: it collapses the two left hand terms into one and produces

$$2f(xf(x)) = 1 + f(2x),$$

which is the single most used line of the whole solution. It says that the value of f along the curve $h(x) := xf(x)$ is controlled by the value of f at $2x$. The unknown of this problem is not f; it is h.

Second, the equation has a scaling symmetry: if f is a solution then so is $x \mapsto f(cx)$ for every $c > 0$, because writing the original equation at the pair (cx, cy) gives exactly the equation for the new function. So solutions come in one parameter families, the answer will contain a free positive constant, and any proof of uniqueness must spend that symmetry on a normalisation at some point.

Third, the mismatch between products on the left and a sum on the right is the real content. For $f(xf(y))$ and $f(x+y)$ to be comparable when x is large, you need $xf(y)$ and $x+y$ to be comparable, that is you need f to behave like 1 near 0 and like $1/x$ near infinity. That is the shape of the answer before you have proved anything: $f(x) = \dfrac{1}{1+cx}$ with $c \geq 0$, and $xf(x) \to 1/c$.

Within thirty seconds you should have written the diagonal identity, named $h(x) = xf(x)$, and verified the family $1/(1+cx)$. Those three lines are worth real marks on their own, and they are the skeleton of everything that follows.

Approaches, ranked worst to best

1. **Substitution hunting.** *Cue:* it is a functional equation, so try $x = y$, $x = 1$, $y = 1$, $f(1) = t$. *Circuit:* plug in special values until something cancels. (zombie route: the domain is $\mathbb{R}^+$ with no distinguished point and no value of f given, so nothing cancels; the scaling symmetry guarantees that no substitution can ever pin a value, because there is no canonical point to pin it at)

2. **Assume differentiability and differentiate.** *Cue:* the equation looks like it wants a derivative. *Circuit:* differentiate in x at fixed y and solve the resulting differential equation. (illegal and unrecoverable: only continuity is given, and continuous solutions of functional equations can behave very badly; a grader will delete the entire solution at the first unproved differentiation)

3. **Guess the family and claim rigidity.** *Cue:* the shape $1/(1+cx)$ is forced by the product versus sum mismatch. *Circuit:* verify the family and assert uniqueness. (worth two or three marks and no more: the verification is three lines and the uniqueness is the entire problem)

4. **Monotonicity first.** *Cue:* the equation compares $f(xf(y))$ with $f(x+y)$, so if f were injective you could cancel. *Circuit:* prove f is strictly decreasing, hence injective, then compare arguments. (a genuine and useful tool, but it is available only after you know $f < 1$ everywhere, which needs the asymptotics at both ends; used as an opening move it stalls, used at the end it gives a clean alternative finish)

5. **Pin the two boundary behaviours, then propagate inwards.** *Cue:* a functional equation on an open half line with only continuity has no interior structure to exploit, so the information must come from the ends. *Circuit:* prove $f(0^+) = 1$; split off the case $f \geq 1$; prove $f(\infty) = 0$ and $xf(x) \to 1$; normalise the scale; take limits inside the equation. (ideal: every step is a limit taken inside the equation, and the whole proof is the study of the single function $h(x) = xf(x)$)

The ideal solution

The solutions are exactly

$$\boxed{f(x) = \frac{1}{1+cx} \quad \text{for a constant } c \geq 0,}$$

the case $c = 0$ being the constant function 1. Throughout, $E(x,y)$ denotes the given equation and $h(x) = xf(x)$.

Step 0: these functions work. With $f(x) = 1/(1+cx)$ we have $xf(y) = x/(1+cy)$, so

$$f(xf(y)) = \frac{1}{1+\frac{cx}{1+cy}} = \frac{1+cy}{1+cx+cy}, \qquad f(yf(x)) = \frac{1+cx}{1+cx+cy},$$

and the sum is $\dfrac{2+cx+cy}{1+cx+cy} = 1 + \dfrac{1}{1+c(x+y)} = 1 + f(x+y)$.

Also record the scaling symmetry: if f solves E and $c > 0$, then $g(x) = f(cx)$ solves E, since $E(cx, cy)$ reads $f(cx\, f(cy)) + f(cy\, f(cx)) = 1 + f(c(x+y))$, which is $g(xg(y)) + g(yg(x)) = 1 + g(x+y)$.

Now let f be any continuous solution.

Step 1: $\lim_{x\to 0^+} f(x) = 1$**.**

First, f is bounded near 0. For $0 < x < 1$, $E(x,1)$ and $f(f(x)) > 0$ give $f(xf(1)) < 1 + f(x+1) \leq 1 + \max_{[1,2]} f$, the maximum existing by continuity on a compact interval. As x ranges over $(0,1)$, $xf(1)$ ranges over $(0, f(1))$, so f is bounded on $(0, f(1))$. Consequently $h(x) = xf(x) \to 0$ as $x \to 0^+$.

Write $L_+ = \limsup_{t\to 0^+} f(t)$ and $L_- = \liminf_{t\to 0^+} f(t)$; both are finite by the previous paragraph, and $L_- \geq 0$.

Lemma 4.12.1. $\limsup_{x\to 0^+} f(h(x)) = L_+$ *and* $\liminf_{x\to 0^+} f(h(x)) = L_-$.

Proof. Since $h(x) \to 0^+$ as $x \to 0^+$, the inequalities $\limsup_{x\to 0^+} f(h(x)) \leq L_+$ and $\liminf_{x\to 0^+} f(h(x)) \geq L_-$ are immediate. For the reverse inequalities, note that h is continuous and positive on $(0,\infty)$ with $h(x) \to 0$ as $x \to 0^+$, so for every $\delta > 0$ the image $h((0,\delta))$ contains the interval $(0, h(\delta/2))$: given t in that interval, choose $x_1 \in (0, \delta/2)$ with $h(x_1) < t$, and apply the intermediate value theorem to h on $[x_1, \delta/2]$. Now fix $\eta > 0$ and $\delta > 0$. By definition of L_+ there are points t arbitrarily close to 0 with $f(t) > L_+ - \eta$; choose such a t with $t < h(\delta/2)$ and write $t = h(x)$ with $x \in (0,\delta)$. Hence $\sup_{0<x<\delta} f(h(x)) > L_+ - \eta$ for every δ and η, giving $\limsup_{x\to 0^+} f(h(x)) \geq L_+$. The liminf statement is identical with the inequalities reversed. □

Now take $y = x$ in E:

$$2f(h(x)) = 1 + f(2x) \qquad \text{for all } x > 0. \tag{D}$$

Taking lim sup as $x \to 0^+$ in (D) gives $2L_+ = 1 + L_+$, so $L_+ = 1$; taking lim inf gives $2L_- = 1 + L_-$, so $L_- = 1$. Hence $f(x) \to 1$ as $x \to 0^+$.

Step 2: if $f(x) \geq 1$ for all $x > 0$, then $f \equiv 1$.

Suppose $f \geq 1$ everywhere and, for contradiction, that $\lambda := f(y) > 1$ for some $y > 0$. Since $f(yf(x)) \geq 1$, the equation $E(x, y)$ gives

$$f(\lambda x) \leq f(x + y) \qquad \text{for all } x > 0,$$

that is, writing $u = \lambda x$ and $T(u) = u/\lambda + y$,

$$f(u) \leq f(T(u)) \qquad \text{for all } u > 0.$$

The affine map T has slope $1/\lambda \in (0, 1)$ and fixed point $c_0 = \lambda y/(\lambda - 1) > 0$, and $T^k(u) - c_0 = (u - c_0)\lambda^{-k} \to 0$. Iterating the inequality and using continuity of f at c_0,

$$f(u) \leq f(T^k(u)) \xrightarrow[k\to\infty]{} f(c_0),$$

so $M := \sup_{x>0} f(x) \leq f(c_0) < \infty$: the function is bounded.

Next, $f \geq 1$ gives $h(x) = xf(x) \geq x \to \infty$, while $h(x) \to 0$ as $x \to 0^+$ by Step 1. Since h is continuous on the interval $(0, \infty)$, the intermediate value theorem shows h takes every positive value. So given $t > 0$, choose x with $h(x) = t$; then (D) gives

$$f(t) = \frac{1 + f(2x)}{2} \leq \frac{1 + M}{2}.$$

Taking the supremum over t yields $M \leq (1 + M)/2$, hence $M \leq 1$, hence $f \equiv 1$, contradicting $f(y) > 1$. Therefore no such y exists and $f \equiv 1$, which is the case $c = 0$.

From now on assume $f(x_0) < 1$ for some $x_0 > 0$.

Step 3: h is bounded and $f(x) \to 0$ as $x \to \infty$.

Let $I = \inf_{x>0} f(x) < 1$, choose $\varepsilon \in \left(0, \frac{1-I}{2}\right)$, and choose $y_\varepsilon > 0$ with $f(y_\varepsilon) < I + \varepsilon$. If $h(x) = y_\varepsilon$ for some x, then (D) gives

$$1 + I \leq 1 + f(2x) = 2f(y_\varepsilon) < 2I + 2\varepsilon,$$

so $1 - I < 2\varepsilon$, contradicting the choice of ε. Hence h omits the value y_ε. Since h is continuous on the interval $(0, \infty)$ and $h(x) \to 0 < y_\varepsilon$ as $x \to 0^+$, the intermediate value theorem forces $h(x) < y_\varepsilon$ for every x. Thus $xf(x) < y_\varepsilon$ for all x, so $f(x) < y_\varepsilon/x \to 0$.

Step 4: normalisation.

By Step 1 and Step 3, f is continuous with $f(0^+) = 1$ and $f(\infty) = 0$, so there is $\delta > 0$ with $f > 3/4$ on $(0, \delta)$ and $R > \delta$ with $f < 1/4$ on (R, ∞). The set $f^{-1}(1/2)$ is therefore a nonempty (intermediate value theorem) closed subset of the compact interval $[\delta, R]$, so it has a least element x_1. Since $f > 1/2$ near 0 and f omits the value $1/2$ on $(0, x_1)$, the intermediate value theorem gives $f > 1/2$ on $(0, x_1)$. Replacing f by $x \mapsto f(x_1 x)$, which is again a solution by the scaling symmetry, we may and do assume

$$f(1) = \tfrac{1}{2} \qquad \text{and} \qquad f(t) > \tfrac{1}{2} \text{ for } 0 < t < 1.$$

It now suffices to prove $f(x) = 1/(1 + x)$; undoing the normalisation then gives $f(x) = 1/(1 + cx)$ with $c = 1/x_1 > 0$.

Step 5: $\lim_{x\to\infty} h(x) = 1$.

By (D) and positivity of f, $f(h(x)) = \frac{1}{2}(1 + f(2x)) > \frac{1}{2} = f(1)$, so $h(x) \neq 1$ for every x. As h is continuous with $h(x) \to 0$ as $x \to 0^+$, the intermediate value theorem gives $h(x) < 1$ for all x.

Suppose $h(x) \not\to 1$. Then there are $\delta \in (0, 1)$ and $x_n \to \infty$ with $h(x_n) \leq 1 - \delta$. Choose $\delta'' \in (0, 1 - \delta)$ with $f > 3/4$ on $(0, \delta'']$; on the compact interval $[\delta'', 1 - \delta]$ the continuous function f attains a minimum, which exceeds $1/2$ because $f > 1/2$ on $(0, 1)$. Hence

$$\mu := \inf\{f(t) : 0 < t \leq 1 - \delta\} > \tfrac{1}{2},$$

and so $f(h(x_n)) \geq \mu > 1/2$ for all n. But $f(2x_n) \to 0$ by Step 3, so (D) gives $f(h(x_n)) \to \frac{1}{2}$, a contradiction. Therefore $h(x) \to 1$ as $x \to \infty$.

Step 6: the reflection identity $f(y) + f(1/y) = 1$ for all $y > 0$.

Fix $y > 0$ and replace y by xy in E:

$$f(xf(xy)) + f(xyf(x)) = 1 + f(x(1+y)).$$

Let $x \to \infty$. Then $xy \to \infty$, so $xf(xy) = \frac{1}{y}\, h(xy) \to \frac{1}{y}$ and $xyf(x) = y\, h(x) \to y$, while $x(1+y) \to \infty$ gives $f(x(1+y)) \to 0$. Since f is continuous at the positive numbers $1/y$ and y, the three limits give

$$f(1/y) + f(y) = 1.$$

Step 7: conclusion.

Apply Step 6 to the positive number $yf(x)$: $f(yf(x)) = 1 - f\left(\frac{1}{yf(x)}\right)$. Substituting into $E(x,y)$ and cancelling the 1,

$$f(xf(y)) = f(x+y) + f\left(\frac{1}{yf(x)}\right).$$

Fix y, multiply through by $xf(y)$, and let $x \to \infty$, examining each term.

The left side is $vf(v) = h(v)$ with $v = xf(y) \to \infty$, so it tends to 1 by Step 5.

The first term on the right is

$$xf(y)f(x+y) = f(y) \cdot \frac{x}{x+y} \cdot h(x+y) \longrightarrow f(y) \cdot 1 \cdot 1 = f(y).$$

For the second term, put $u = \dfrac{1}{yf(x)}$, which tends to ∞ because $f(x) \to 0$. Then $f(u) = h(u)/u = h(u)\, yf(x)$, so

$$xf(y)\, f(u) = f(y)\, y\, h(u)\, \big(xf(x)\big) = f(y)\, y\, h(u)\, h(x) \longrightarrow f(y)\, y.$$

Therefore $1 = f(y) + yf(y)$, that is $f(y) = \dfrac{1}{1+y}$ for every $y > 0$, as required. Undoing the normalisation of Step 4 completes the classification.

Layer 1: WHAT

Block 1: verify the family, and record the scaling symmetry $f \mapsto f(c\,\cdot)$, which tells you the answer is a one parameter family and that a normalisation will be needed.

Block 2: put $y = x$. The diagonal identity $2f(xf(x)) = 1 + f(2x)$ names the true unknown $h(x) = xf(x)$ and is used four separate times.

Block 3: bound f near 0 straight from the equation, deduce $h \to 0$ at 0^+, and then run the diagonal identity through lim sup and lim inf to get $f(0^+) = 1$. The only technical point is that h covers a punctured neighbourhood of 0, so limits along h see the same limsup as limits along the identity.

Block 4: dispose of the case $f \geq 1$. The equation gives $f(\lambda x) \leq f(x+y)$, an affine contraction inequality; iterating to its fixed point bounds f, and the diagonal identity then squeezes sup f to 1.

Block 5: in the remaining case, the diagonal identity shows h cannot hit a value where f is nearly minimal, so h is bounded and $f(x) \to 0$ at infinity.

Block 6: spend the scaling symmetry: put the least solution of $f(t) = 1/2$ at $t = 1$.

Block 7: sharpen the bound on h to a limit: h omits 1, hence $h < 1$, and $f(h(x)) \to 1/2$ forces $h \to 1$.

Block 8: substitute $y \mapsto xy$ and let $x \to \infty$; the three terms converge to $f(1/y)$, $f(y)$ and 0, giving the reflection identity.

Block 9: feed the reflection identity back into E, multiply by $xf(y)$ so that every term converges, and let $x \to \infty$ to get $f(y)(1+y) = 1$.

The nominal route treats this as a substitution puzzle and spends the exam hunting for values of x and y that make two arguments coincide. The bifurcation point is Block 2, at the decision of what the unknown is. Once you write $h(x) = xf(x)$ and read the diagonal identity as a statement about h, every later step is a question about a single continuous function on an interval: what is its limit at each end, and which values does it omit. Those are questions the intermediate value theorem answers.

Layer 2: HOW

Block 1. **Cue.** The equation is unchanged when x and y are swapped, and it is unchanged under $x \mapsto cx$ in the sense that $f(c\,\cdot)$ is again a solution. **Why natural.** Symmetries are free information. The swap symmetry tells you the diagonal is the degenerate direction and therefore the cheapest substitution. The scaling symmetry tells you in advance that the answer contains a free parameter, so any attempt to compute a specific value such as $f(1)$ is doomed until you normalise. **Link to the next block.** Take the diagonal.

Block 2. **Cue.** $y = x$ makes the two left hand terms identical. **Why natural.** It converts a two variable equation into a one variable identity, and the argument that appears, $xf(x)$, is exactly the product versus sum hybrid that the whole equation is about. Naming it h is not decoration: everything you can prove is a statement about the limits and the image of h. **Link to the next block.** You need boundary information about f before the identity can be exploited.

Block 3. **Cue.** The term $f(xf(y))$ with y fixed and $x \to 0^+$ sends its argument to 0^+, and the right side stays bounded by continuity at y. **Why natural.** That is a free upper bound for f near 0, and boundedness near 0 is what makes $h(x) \to 0$. With $h \to 0$ at 0^+, the diagonal identity relates the behaviour of f near 0 to itself, and a self referential relation between lim sup and lim sup is an equation: $2L_+ = 1 + L_+$. The one honest technicality is that lim sup along the curve h equals lim sup along the identity, and that needs h to cover a punctured neighbourhood of 0, which is the intermediate value theorem. **Link to the next block.** The value 1 at the left end is now known unconditionally.

Block 4. **Cue.** The right side of E contains $f(x+y)$ and the left contains $f(xf(y))$; if $f \geq 1$, the other left hand term is at least 1 and can be discarded. **Why natural.** Discarding a term turns an equation into an inequality, and the inequality $f(\lambda x) \leq f(x+y)$ compares a multiplicative shift with an additive one. Multiplication beats addition, so iterating drives the argument to the fixed point of $u \mapsto u/\lambda + y$, and continuity converts the infinite chain into a single bound $f \leq f(c_0)$. This is the standard way to prove boundedness from a one sided functional inequality. **Link to the next block.** With f bounded, the diagonal identity plus surjectivity of h pins the supremum.

Block 5. **Cue.** The diagonal identity says $f(h(x)) = \frac{1}{2}(1 + f(2x)) \geq \frac{1}{2}(1 + \inf f)$, so f can never be close to its infimum at a point of the image of h. **Why natural.** Values near the infimum are therefore omitted by h, and an omitted value plus continuity on an interval plus values below it forces h to stay below it. That is a boundedness statement about $xf(x)$, which is exactly a decay statement about f. Note how little was used: only that $\inf f < 1$. **Link to the next block.** You now have both boundary behaviours, 1 at 0^+ and 0 at ∞.

Block 6. **Cue.** A one parameter family of solutions and a proof that must end with a specific formula. **Why natural.** Normalise, using the least point where $f = 1/2$, because $1/2$ is the value

the diagonal identity refers to: (D) says $f(h(x)) > 1/2$ always. Choosing the least such point is what makes the interval $(0, 1)$ a region where $f > 1/2$, which is the fact used in Step 5. **Link to the next block.** The normalisation is chosen to serve the next step, not for aesthetics.

Block 7. Cue. $f(h(x)) > f(1)$ for every x. **Why natural.** So h omits the value 1, hence $h < 1$ everywhere by the same omission argument as before. Then $f(h(x)) \to 1/2$ as $x \to \infty$, and since f is bounded away from $1/2$ on any interval $(0, 1 - \delta]$, the point $h(x)$ must approach 1. This upgrades a bound into a limit, which is what the final computation needs. **Link to the next block.** You now know the exact asymptotic $f(x) \sim 1/x$.

Block 8. Cue. You have an asymptotic at infinity and an equation with two free variables. **Why natural.** Substituting $y \mapsto xy$ makes both left hand arguments converge to finite positive limits as $x \to \infty$, namely $1/y$ and y, while the right hand argument runs to infinity where f vanishes. That is the only substitution for which all three terms have limits, and it produces the reflection identity for free. **Link to the next block.** The reflection identity lets you replace an awkward term by a computable one.

Block 9. Cue. Every term of the resulting identity behaves like $1/(\text{argument})$ for large x, so multiplying by the right factor makes each term converge. **Why natural.** The only asymptotic information you have is $uf(u) \to 1$, so any term of the form $f(\text{large})$ must be multiplied by that large quantity before it can be used. Multiplying by $xf(y)$ is exactly what turns the left side into $h(xf(y)) \to 1$; the same factor happens to make both right hand terms converge, to $f(y)$ and $yf(y)$. The final equation $1 = f(y)(1 + y)$ then falls out.

Layer 3: WHY

The mathematical content is a clash between multiplication and addition. On the left the function is evaluated at products, on the right at a sum. A continuous function that reconciles the two must convert the multiplicative structure near infinity into the additive structure, and the only continuous way to do that with a bounded positive function is $f(x) \sim 1/(cx)$: then $xf(y) \approx x/(cy)$ and $f(xf(y)) \approx cy/x \cdot c$, and the additive constant 1 on the right is what forces the exact Mobius form rather than a mere asymptotic. Substituting $g = 1/f - 1$ turns the answer into $g(x) = cx$, so the equation is secretly asserting that a certain transform of f is linear, which is the Cauchy equation in disguise, and the hypothesis of continuity is doing precisely the job it does for Cauchy: excluding pathological solutions.

The organising principle for the proof is that a functional equation on an open half line with only continuity has no interior structure to exploit. All the information lives at the two ends. So the programme is forced: pin $f(0^+)$, pin $f(\infty)$, upgrade the second to the sharp asymptotic $xf(x) \to 1$, and then propagate inwards by taking limits inside the equation. Every single step of the ideal solution is an instance of that programme, and each one uses the same two instruments: the diagonal identity (D) and the fact that a continuous function on an interval which omits a value and takes a smaller one stays below it.

The discernment mechanism to train is the question *what is the true unknown?* Here it is not f but $h(x) = xf(x)$, because h is the object that appears as an argument in the diagonal identity, and because the target family is characterised by h being bounded with limit $1/c$. Whenever a functional equation feeds a composite expression into the unknown function, promote that composite to the status of the unknown and ask the three standard questions about it: what are its limits at the ends of the domain, what is its image, and which values does it omit.

Knowledge gaps, stated plainly. This problem is not approachable at all without the following installed, and none of it is on the A-level Further Mathematics syllabus except the intermediate

value theorem. First, lim sup and lim inf of a function at a point, their existence for bounded functions, and their behaviour under adding constants and under composition with a curve that covers a punctured neighbourhood. Second, the intermediate value theorem in the form "the continuous image of an interval is an interval", used three times to convert omitted values into global inequalities. Third, compactness: a continuous positive function on a closed bounded interval attains a positive minimum. Fourth, iteration of an affine contraction and convergence to its fixed point. Fifth, the discipline of taking limits inside an identity only after checking that every term converges. If any of the five is missing, the solution cannot be written, and no amount of algebraic ingenuity substitutes for them.

Revolving concepts

Functional equations with only continuity. The standard toolkit: exploit symmetry (the diagonal substitution), look for injectivity or monotonicity to cancel f, and if neither is available, work with asymptotics at the boundary of the domain. Continuity is used to move limits inside f and to upgrade pointwise information to interval information. See Book 1, field: Functional equations, topic: continuous solutions.

The Cauchy equation and its relatives. $g(x+y) = g(x) + g(y)$ with g continuous (or merely bounded on a set of positive measure, or monotone on an interval) forces $g(x) = cx$. The present answer is $1/f - 1 = cx$, so the problem is a Mobius twisted Cauchy equation. See Book 1, field: Functional equations, topic: Cauchy's equation.

Omitted values and the intermediate value theorem. If g is continuous on an interval, g never takes the value v, and g takes some value less than v, then $g < v$ everywhere. This single fact is used to prove $h < y_\varepsilon$ (Step 3) and $h < 1$ (Step 5), and its surjectivity twin is used in Step 2 and in the covering lemma of Step 1. See Book 1, field: Real analysis, topic: intermediate value theorem.

Affine contractions and their iteration, generalised. The recursion $u_{k+1} = au_k + b$ with $a \neq 1$ has the closed form

$$u_k = a^k \left(u_0 - \frac{b}{1-a} \right) + \frac{b}{1-a},$$

so when $|a| < 1$ every orbit converges to the fixed point $b/(1-a)$ geometrically. Here $a = 1/\lambda$ and $b = y$, giving the fixed point $c_0 = \lambda y/(\lambda - 1)$. The reusable pattern is: a one sided inequality $f(u) \leq f(T(u))$ with T a contraction plus continuity of f yields the global bound $f \leq f(\text{fixed point})$, by iterating and passing to the limit. The same pattern with T expanding, or with the inequality reversed, gives lower bounds instead. See Book 1, field: Sequences and recurrences, topic: affine recursions and fixed points.

The auxiliary function $h(x) = xf(x)$. For the answer $f = 1/(1+cx)$ one has $h(x) = x/(1+cx)$, which is strictly increasing from 0 to $1/c$; after the normalisation $c = 1$ this is exactly the content of Steps 5 and the bound $h < 1$. Promoting a composite argument to an unknown in its own right is the general move. See Book 1, field: Functional equations, topic: auxiliary functions.

Limits inside functional equations. If every term of an identity has a limit along a chosen path, the identity passes to the limit. The craft is choosing the path (here $y \mapsto xy$ with $x \to \infty$) and the normalising factor (here $xf(y)$) so that no term diverges. See Book 1, field: Real analysis, topic: passing to the limit.

Circuit Bank

- **WHEN** a functional equation is symmetric in its two variables **DO** substitute $y = x$ before anything else, and name whatever composite argument appears **BECAUSE** the diagonal is where the two sides collapse to a single unknown combination, and that combination is almost always the true unknown of the problem

- **WHEN** a functional equation on an open half line comes with continuity and nothing else **DO** go to the two ends of the domain and pin the limits there, then propagate inwards by taking limits inside the equation **BECAUSE** with no regularity beyond continuity there is no interior structure to exploit, so all the information lives at the boundary
- **WHEN** a continuous function on an interval is shown to omit a value v and to take some value below v **DO** conclude that it stays below v everywhere **BECAUSE** the continuous image of an interval is an interval, so a function cannot jump over a value it omits
- **WHEN** a functional equation yields a one sided inequality $f(u) \leq f(T(u))$ with T an affine contraction **DO** iterate to the fixed point and use continuity **BECAUSE** the infinite chain of inequalities collapses to the single global bound $f \leq f(\text{fixed point})$, which is how boundedness is obtained without any compactness
- **WHEN** an equation is invariant under rescaling the variable **DO** normalise before attempting uniqueness, and choose the normalisation to serve the next step **BECAUSE** the solution set is a one parameter family, so no argument can pin a value until the symmetry has been spent
- **WHEN** you want to take a limit in an identity whose terms individually vanish or diverge **DO** multiply through by the factor that makes each term converge, using the asymptotic you have already proved **BECAUSE** $uf(u) \to 1$ is usable only against a factor of size u, so the multiplier is dictated by the asymptotic, not chosen

Backtesting

1. First useful cue: the symmetry in x and y, which makes $y = x$ the cheapest substitution and produces $2f(xf(x)) = 1 + f(2x)$, the identity that carries the entire solution.
2. Method the expression suggested: study the auxiliary function $h(x) = xf(x)$, since the product versus sum mismatch means the answer must decay like $1/x$ and h is the normalised unknown.
3. Where you would hesitate, and the thought to reinforce: at the case $f \geq 1$, which looks like an irrelevant technicality and is in fact where the constant solution hides. Reinforce this: a classification must account for every solution, and the family $1/(1 + cx)$ includes $c = 0$ only because the constant 1 is genuinely a solution; the proof of the general case assumes $\inf f < 1$ at Step 3 and would be circular without the dichotomy.
4. The trap the question sets: assuming differentiability. Only continuity is given, and the entire proof is built to avoid derivatives; a solution that differentiates the equation scores nothing regardless of the answer it reaches.
5. What makes you recognise it faster next time: a functional equation that mixes $f(xf(y))$ with $f(x + y)$ is a multiplication against addition problem, and its solutions are always of Mobius or exponential type; compute the asymptotics first and the algebra afterwards.
6. The reusable decision rule: diagonal substitution, name the composite argument, pin both boundary limits, normalise the scaling symmetry, then take limits inside the equation with a multiplier chosen to make every term converge.

Mirror drills

1. Show that if f solves the equation then so does $x \mapsto f(cx)$ for every $c > 0$, and deduce that no substitution can determine $f(1)$. *Answer:* write the original equation at the pair (cx, cy); the value $f(1)$ can be moved to any value in the range of a solution by rescaling, so it is not determined by the equation alone.

2. Let g be continuous and positive on $(0,\infty)$ with $g(x)\to 0$ as $x\to 0^+$, and suppose g never takes the value 1. Prove $g(x)<1$ for all x. *Answer:* if $g(x_1)>1$ for some x_1, pick $x_0<x_1$ with $g(x_0)<1$ and apply the intermediate value theorem on $[x_0,x_1]$.

3. For $\lambda>1$ and $y>0$ let $T(u)=u/\lambda+y$. Find the fixed point c_0, show $T^k(u)\to c_0$ for every $u>0$, and deduce that a continuous f satisfying $f(u)\le f(T(u))$ for all u obeys $f\le f(c_0)$. *Answer:* $c_0=\lambda y/(\lambda-1)$ and $T^k(u)-c_0=(u-c_0)\lambda^{-k}$; iterate the inequality and use continuity at c_0.

Chapter 5

Putnam 2021

5.1 A1. Grasshopper hops of length five to the diagonal

A grasshopper starts at the origin in the coordinate plane and makes a sequence of hops. Each hop has length 5*, and after each hop the grasshopper is at a point whose coordinates are both integers; thus, there are* 12 *possible locations for the grasshopper after the first hop. What is the smallest number of hops needed for the grasshopper to reach the point* $(2021, 2021)$*?*

Read it out loud

This is a lattice walk with a finite step set and a single target, and the command term is *smallest number*. Read that phrase as a two-part contract before you touch the numbers: a lower bound that holds for every walk, and one explicit walk that attains it. Any answer that has only one half of the contract is worth at most half the marks, and the half that most candidates skip is the bound.

The fields wired in are elementary number theory (which integer vectors have length 5: the representations $25 = 0^2 + 5^2 = 3^2 + 4^2$, hence the twelve vectors $(0, \pm 5)$, $(\pm 5, 0)$, $(\pm 3, \pm 4)$, $(\pm 4, \pm 3)$), and, for the bound, the geometry of a finite step set: a potential function that no hop can change by more than a fixed amount. That second field is the one being tested. The examiners tell you the step count (twelve) to make sure you list the vectors; the real question is which measurement of progress you choose.

Within thirty seconds you should have written the twelve vectors, the target, and the observation that the target sits on the diagonal direction $(1, 1)$ while no hop of length 5 points along that diagonal (the vector $(5/\sqrt{2}, 5/\sqrt{2})$ is not a lattice point). That single observation already tells you the Euclidean bound $2021\sqrt{2}/5 \approx 571.6$ cannot be tight, and it tells you where to look instead: measure progress along the diagonal with the linear functional $x + y$, whose value each hop changes by at most $3 + 4 = 7$.

Approaches, ranked worst to best

1. **Greedy construction only.** *Cue:* the target is on the diagonal, so alternate $(3, 4)$ and $(4, 3)$. *Circuit:* aim along the target direction. (this finds a walk of 578 hops in a minute but proves nothing about minimality; without a bound you have half a solution and no way to know whether 577 is possible)

2. **Euclidean distance bound.** *Cue:* each hop has length 5. *Circuit:* distance over step length. (gives $\lceil 2021\sqrt{2}/5 \rceil = 572$; it stalls because the best hop in the diagonal direction only advances $7/\sqrt{2} < 5$ along that direction, so the bound is loose by six hops and you cannot close the gap by constructions)

3. **Parity of $x + y$.** *Cue:* every hop changes $x + y$ by $\pm 1, \pm 5$ or ± 7, all odd. *Circuit:* invariant modulo 2. (a genuine circuit, and it proves the hop count is even, but on its own it only pushes an odd bound up by one; here it is a supporting check, not the main argument)

4. **Linear functional $x + y$ plus construction.** *Cue:* the target is symmetric in x and y and no hop is diagonal. *Circuit:* potential that changes by a bounded amount per step. (each hop moves $x + y$ by at most 7, the target has $x + y = 4042 > 577 \cdot 7$, and $288(3, 4) + 288(4, 3) + (5, 0) + (0, 5)$ attains 578; both halves of the contract in four lines)

The ideal solution

The lattice points at distance 5 from a given lattice point are obtained by adding one of the twelve vectors $(0, \pm 5)$, $(\pm 5, 0)$, $(\pm 3, \pm 4)$, $(\pm 4, \pm 3)$, because $25 = a^2 + b^2$ with a, b integers forces $\{|a|, |b|\} = \{0, 5\}$ or $\{3, 4\}$.

Lower bound. Let $s = x + y$ be the sum of the grasshopper's coordinates. Initially $s = 0$, and at the target $s = 4042$. For each of the twelve hop vectors (a, b) we have $a + b \leq 7$, with equality exactly for $(3, 4)$ and $(4, 3)$. Thus after n hops, $s \leq 7n$. If $n \leq 577$ then $s \leq 7 \cdot 577 = 4039 < 4042$, so the target has not been reached. Hence at least 578 hops are needed.

Construction. Since $2021 = 288 \cdot 7 + 5$, take 288 hops of $(3, 4)$, 288 hops of $(4, 3)$, one hop of $(5, 0)$ and one hop of $(0, 5)$. The position after these 578 hops is $(288 \cdot 7 + 5,\ 288 \cdot 7 + 5) = (2021, 2021)$, and every intermediate position is a lattice point because every hop vector is a lattice vector.

Therefore the smallest number of hops is $\boxed{578}$.

Layer 1: WHAT

Block 1: list the hop set from $25 = 0^2 + 5^2 = 3^2 + 4^2$. Block 2: choose the potential $s = x + y$ and compute its maximum change per hop, 7. Block 3: divide, $4042 > 577 \cdot 7$, so at least 578. Block 4: write $2021 = 288 \cdot 7 + 5$ and assemble the walk.

Juxtapose with the nominal route. The zombie writes the same Block 1, then reaches for the only measure of progress that A-level trained into him, Euclidean distance, gets 572, and spends twenty minutes failing to build a 572-hop walk without understanding why. The bifurcation point is the choice of potential in Block 2: the zombie measures progress with the norm that the hop length is stated in; you measure progress with the linear functional aligned with the target, because a bound is only tight when the extremal hop for the potential is one you can actually use in the construction.

Layer 2: HOW

Block 1. **Cue.** "Each hop has length 5" and "coordinates are both integers" together with the planted number twelve. **Why natural.** A lattice vector of length 5 is a solution of $a^2 + b^2 = 25$; the only nonnegative solutions are $(0, 5), (5, 0), (3, 4), (4, 3)$ and the signs give $2 + 2 + 4 + 4 = 12$. Matching the examiners' count of twelve confirms you have the right set, which is the whole reason they printed it. **Link to the next block.** Now you know the step set is finite and fixed, so any linear function of position changes by a bounded amount per hop, which is precisely what a lower bound needs.

Block 2. **Cue.** "Smallest number of hops" to a target on the diagonal. **Why natural.** A step count lower bound comes from a quantity that the walk must move from 0 to some value V while each step moves it by at most M; then $n \geq V/M$. The tightest such quantity is the linear functional whose level lines are the supporting line of the hop set in the target direction. The hop set's convex hull has an edge from $(3, 4)$ to $(4, 3)$, and that edge lies on $x + y = 7$. So $x + y$ is not a guess: it is the supporting line of the hull that faces the target. **Link to the next block.** With $M = 7$ and $V = 4042$, division is the next line.

Block 3. **Cue.** $4042/7 = 577.4\ldots$. **Why natural.** If 577 hops sufficed, the sum would be at most 4039; it must be 4042. This is the complete argument and it needs no case analysis, because the bound is on a single scalar. **Link to the next block.** A bound of 578 is only worth its ink if 578 is attainable; the construction must use the extremal hops $(3, 4)$ and $(4, 3)$ almost exclusively, with slack of $578 \cdot 7 - 4042 = 4$ to absorb.

Block 4. **Cue.** The slack is 4 and the two hops $(5, 0)$, $(0, 5)$ each lose exactly 2 relative to the maximum 7. **Why natural.** $2021 = 288 \cdot 7 + 5$; the diagonal part $(2016, 2016)$ is 288 copies of $(3, 4) + (4, 3)$, and the remainder $(5, 5)$ is one axis hop in each direction. The lost slack is $2 + 2 = 4$, matching the arithmetic exactly, which is your check that nothing is wasted. **Link to the next block.** Bound equals construction; box the answer.

Layer 3: WHY

The truth underneath is weak linear programming duality. If T is the target and V the step set, then for any linear functional φ, the number of steps satisfies $n \geq \varphi(T) / \max_{v \in V} \varphi(v)$, and the best φ is a supporting functional of the convex hull of V in the direction of T. Ignoring integrality, the minimal number of steps is exactly the gauge of T with respect to that hull, and the gauge equals the maximum of $\varphi(T)$ over all φ with $\max_V \varphi \leq 1$. Here the hull is a dodecagon and the ray through $(1, 1)$ crosses the edge $x + y = 7$, so the Euclidean bound (which corresponds to a circle of radius 5 containing the hull) is strictly weaker: the circle bulges beyond the hull in the diagonal direction.

The discernment mechanism: whenever a bound is not attainable, ask which constraint of the bound is violated by every construction. The Euclidean bound assumes a hop of length 5 pointing at the target; none exists; therefore replace the norm by the functional that the actual extremal hops saturate. Train this by always writing the hop set's hull edges before choosing a potential, and by keeping parity as a second, independent check: the sum $x + y$ changes by an odd number every hop, so the hop count has the parity of 4042, namely even, and 578 is even. The examiners nudged with the diagonal target and the specific number 5 (small enough that the hull is a dodecagon with a non-axis edge facing the diagonal); a target like $(2021, 0)$ would have made the axis hop extremal and the problem trivial.

Revolving concepts

Taxicab metric. $d((x_1, y_1), (x_2, y_2)) = |x_1 - x_2| + |y_1 - y_2|$; each hop moves taxicab distance at most 7, and the target is at taxicab distance 4042. It is the norm whose unit ball is a diamond, and the diamond's edges are exactly the supporting lines $\pm x \pm y = \text{const}$. See Book 1, field: Inequalities, topic: norms and dual norms.

Supporting functionals and the gauge of a convex set. For a compact convex set K containing 0 in its interior, the gauge $\|T\|_K = \min\{t \geq 0 : T \in tK\}$ satisfies $\|T\|_K = \max\{\varphi(T) : \varphi \text{ linear}, \max_K \varphi \leq 1\}$. A step-count problem with step set V has fractional answer $\|T\|_{\text{conv}(V)}$. See Book 1, field: Inequalities, topic: convexity and duality.

Sums of two squares. $25 = 0^2 + 5^2 = 3^2 + 4^2$; in general the number of lattice points on $x^2 + y^2 = n$ is governed by the prime factorisation of n (four times the excess of divisors $\equiv 1 \pmod 4$ over divisors $\equiv 3 \pmod 4$). See Book 1, field: Number theory, topic: sums of two squares.

Parity invariants. Every hop vector has odd coordinate sum, so the parity of $x + y$ alternates; after n hops it equals the parity of n. Generalised recursion: if a quantity q_n satisfies $q_n \equiv q_{n-1} + 1 \pmod 2$, then $q_n \equiv q_0 + n$; the same one-line induction runs every invariant argument on walks. See Book 1, field: Games and invariants, topic: parity.

Circuit Bank

- **WHEN** you see "smallest number of steps" with a finite step set **DO** write the two-part contract: a scalar potential bound and an explicit construction, and refuse to box an answer until both exist **BECAUSE** because minimality is a universal statement and a construction is an existential one; a single argument never proves both
- **WHEN** you see a target direction along which no step is extremal for the Euclidean norm **DO** replace the norm by the linear functional of the hull edge facing the target **BECAUSE** because a bound is tight exactly when its extremal step is usable, and the supporting line of the hull is where that happens
- **WHEN** you see every step change some integer quantity by an odd amount **DO** record the parity of the step count as a free second constraint **BECAUSE** because parity can raise a divisional bound by one at zero cost and is a cheap check on your construction

Backtesting

1. First useful cue: the planted count of twelve hop vectors, forcing you to list them and notice that $(3,4)$ and $(4,3)$ are the diagonal-facing ones.
2. Method the expression suggested: bound a linear potential per hop, then divide; the diagonal target picks $x+y$.
3. Where you would hesitate, and the thought to reinforce: after getting 572 from Euclid and failing to construct it. Reinforce: an unattainable bound means the wrong potential, not a harder construction.
4. The trap the question sets: hop length 5 in the Euclidean norm invites a Euclidean bound; the answer lives in the taxicab norm.
5. What makes you recognise it faster next time: a target on a symmetry line of the step set plus a step set with no step on that line.
6. The reusable decision rule: choose the potential whose maximum per step is attained by a step you will actually use in the construction.

Mirror drills

1. (drill) Same hop set, target $(2021,0)$. *Answer:* the potential x changes by at most 5, so $n \geq 405$ since $404 \cdot 5 = 2020$; parity of $x+y$ forces n odd; construction $403(5,0)+(3,4)+(3,-4)$ gives 405.
2. (drill) A knight on an infinite board starts at the origin; least number of moves to $(2022,2022)$. *Answer:* $x+y$ changes by at most 3 per move and $4044 = 3 \cdot 1348$; 674 copies of $(1,2)+(2,1)$ attain 1348.
3. (drill) Steps of length $\sqrt{5}$ between lattice points, target $(0,2021)$. *Answer:* y changes by at most 2, so $n \geq 1011$; parity of $x+y$ makes n odd, consistent; $504(1,2)+506(-1,2)+(2,1)$ attains 1011.

5.2 A2. A one to the infinity limit hiding a geometric mean

For every positive real number x, let

$$g(x) = \lim_{r \to 0} \left((x+1)^{r+1} - x^{r+1} \right)^{\frac{1}{r}}.$$

Find $\lim_{x \to \infty} \dfrac{g(x)}{x}$.

Read it out loud

Two nested limits: an inner limit in r that defines a function g, then an outer limit in x of $g(x)/x$. Read the inner one structurally before you compute anything. At $r = 0$ the base is $(x+1)^1 - x^1 = 1$ and the exponent $1/r$ is unbounded; this is the 1^∞ indeterminate form, the same form as $(1 + 1/n)^n$. The field is single-variable real analysis: indeterminate forms, the definition of the derivative, and the number e. Underneath sits a second field, means of a function over an interval, which explains why the answer is what it is.

What is being tested is discipline about the order of limits. The zombie sees $x \to \infty$ and $r \to 0$, guesses that the order does not matter, expands $(x+1)^{r+1} - x^{r+1} \approx (r+1)x^r$, and gets $(r+1)^{1/r} \to e$. The number is right and the argument is worthless without a uniform estimate, because iterated limits do not commute in general. The ideal route removes the issue entirely by computing the inner limit in closed form for each fixed x.

The first useful cue is the exponent $r+1$, planted so that the base equals 1 at $r = 0$. Within thirty seconds you should have written $h(r) = (x+1)^{r+1} - x^{r+1}$, noted $h(0) = 1$, and written $\ln h(r)/r$, because a 1^∞ form is always resolved by taking logarithms, and $\ln h(r)/r$ with $\ln h(0) = 0$ is literally the difference quotient of $\ln h$ at 0.

Approaches, ranked worst to best

1. **Interchange the limits.** *Cue:* $x \to \infty$ looks like it should simplify the base to $(r+1)x^r$. *Circuit:* leading-term asymptotics. (gives $(r+1)^{1/r} \to e$ in two lines, but the interchange is unjustified; making it rigorous needs Taylor's theorem with a remainder bounded uniformly in $x \geq 1$, which is the official second solution and costs more ink than any other route)

2. **Expand in r.** *Cue:* the base is $1 + (\text{something}) \cdot r + O(r^2)$. *Circuit:* $(1 + ar + O(r^2))^{1/r} \to e^a$. (correct: $(x+1)^{r+1} = (x+1)e^{r\ln(x+1)}$ and similarly for x, so the base is $1 + r\left[(x+1)\ln(x+1) - x\ln x\right] + O(r^2)$; this is the derivative route written with more symbols, and the $O(r^2)$ bookkeeping is where a careless line loses a mark)

3. **Integral representation and squeeze.** *Cue:* u^{r+1} evaluated at $x+1$ and x is a definite integral of $(r+1)u^r$. *Circuit:* fundamental theorem of calculus, then a mean of u^r over $[x, x+1]$. (gives $x \leq \left(\int_x^{x+1} u^r\, du\right)^{1/r} \leq x+1$ and hence $e \leq g(x)/x \leq e(1+1/x)$; no closed form for g needed, but the direction of the inequality flips with the sign of r and both cases must be written)

4. **Derivative of $\ln h$ at zero.** *Cue:* $h(0) = 1$, exponent $1/r$. *Circuit:* 1^∞ resolves as $\exp((\ln h)'(0))$. (three lines to $g(x) = (x+1)^{x+1}/x^x$ and then $g(x)/x = (1+1/x)^{x+1} \to e$ is the definition of e; explicit, no cases, no interchange)

The ideal solution

Fix $x > 0$ and set $h(r) = (x+1)^{r+1} - x^{r+1}$. For $r > -1$ the function $u \mapsto u^{r+1}$ is strictly increasing on $(0, \infty)$, so $h(r) > 0$; hence for $0 < |r| < 1$ the quantity $h(r)^{1/r} = \exp(\ln h(r)/r)$ is well defined.

The function h is differentiable, with

$$h'(r) = (x+1)^{r+1}\ln(x+1) - x^{r+1}\ln x, \qquad h(0) = 1, \qquad h'(0) = (x+1)\ln(x+1) - x\ln x.$$

Since $h(0) = 1$ we have $\ln h(0) = 0$, and $\ln h$ is differentiable at 0 with $(\ln h)'(0) = h'(0)/h(0) = h'(0)$. By the definition of the derivative,

$$\lim_{r\to 0}\frac{\ln h(r)}{r} = \lim_{r\to 0}\frac{\ln h(r) - \ln h(0)}{r-0} = (x+1)\ln(x+1) - x\ln x.$$

By continuity of exp,

$$g(x) = \exp\big((x+1)\ln(x+1) - x\ln x\big) = \frac{(x+1)^{x+1}}{x^x}.$$

Therefore

$$\frac{g(x)}{x} = \frac{(x+1)^{x+1}}{x^{x+1}} = \left(1+\frac{1}{x}\right)^{x+1} = \left(1+\frac{1}{x}\right)^{x}\left(1+\frac{1}{x}\right).$$

As $x\to\infty$, $\ln\big((1+1/x)^x\big) = \ln(1+t)/t$ with $t = 1/x \to 0^+$, which tends to 1 (the derivative of ln at 1), so $(1+1/x)^x \to e$, while $1+1/x \to 1$. Hence

$$\lim_{x\to\infty}\frac{g(x)}{x} = \boxed{e}.$$

Layer 1: WHAT

Block 1: freeze x, name the base $h(r)$, observe $h(0) = 1$: a 1^∞ form. Block 2: take logarithms; $\ln h(r)/r$ is the difference quotient of $\ln h$ at 0. Block 3: differentiate h at 0 and exponentiate: $g(x) = (x+1)^{x+1}/x^x$. Block 4: $g(x)/x = (1+1/x)^{x+1} \to e$.

The nominal route never freezes x. It looks at the double limit as one object, simplifies the base for large x, and reads off e from $(r+1)^{1/r}$. The bifurcation point is Block 1: the decision to treat the inner limit as a function of r alone, with x a constant. Once x is frozen, $h(0) = 1$ is visible and the derivative circuit fires; the interchange question never arises because g is computed exactly.

Layer 2: HOW

Block 1. **Cue.** The exponent $r+1$ on both terms, and the outer power $1/r$. **Why natural.** The examiners could have written $(x+1)^r - x^r$; they wrote $r+1$ so that the base is exactly 1 at $r = 0$. A base tending to 1 with an exponent tending to infinity is the 1^∞ form, and the only question about such a form is the rate at which the base leaves 1, which is a derivative. **Link to the next block.** A rate is read from the logarithm, so Block 2 takes ln.

Block 2. **Cue.** $\ln h(r)/r$ with $\ln h(0) = 0$. **Why natural.** $\ln h(r)/r = (\ln h(r) - \ln h(0))/(r-0)$ is, by definition, the difference quotient of $\ln h$ at 0. No l'Hôpital is needed: l'Hôpital applied to a $0/0$ quotient whose denominator is r is the definition of the derivative with an extra name attached. Knowing this saves you from checking the hypotheses of a theorem you do not need. **Link to the next block.** Differentiability of $\ln h$ at 0 reduces to differentiability of h, and h is a sum of exponentials in r.

Block 3. **Cue.** $h(r) = e^{(r+1)\ln(x+1)} - e^{(r+1)\ln x}$. **Why natural.** Differentiating an exponential in r brings down the logarithm of the base, so $h'(0) = (x+1)\ln(x+1) - x\ln x$, and $g(x) = \exp(h'(0)) = (x+1)^{x+1}/x^x$. Note what this is: $g(x) = e\cdot\exp\big(\int_x^{x+1}\ln u\,du\big)$, that is e times the geometric mean of u over $[x, x+1]$, because $\int_x^{x+1}\ln u\,du = (x+1)\ln(x+1) - x\ln x - 1$. That is the structural meaning of the closed form. **Link to the next block.** The outer limit is now a limit of an explicit elementary function, and dividing by x is a one-line algebraic simplification.

Block 4. Cue. $(x+1)^{x+1}/x^{x+1}$. **Why natural.** The exponents match, so the ratio is $(1+1/x)^{x+1}$, and $(1+1/x)^x \to e$ is the definition of e (or the derivative of ln at 1 if you want it from first principles). The geometric-mean reading predicts this: $g(x)$ is e times a number between x and $x+1$, so $g(x)/x$ is squeezed between e and $e(1+1/x)$. **Link to the next block.** Box e.

Layer 3: WHY

The truth beneath the route is the general rule for 1^∞: if $f(0)=1$ and f is differentiable at 0, then $f(t)^{1/t} \to e^{f'(0)}$ as $t \to 0$, because $\ln f(t)/t \to (\ln f)'(0) = f'(0)$. Every 1^∞ limit on the Putnam is this rule, sometimes with t replaced by $1/n$. The discernment mechanism is to identify the base, evaluate it at the limit point, and if it is 1, compute the derivative of its logarithm there; if it is not 1, the limit is 0 or ∞ and there is nothing to compute.

The deeper truth is the power-mean identity. By the fundamental theorem of calculus, $(x+1)^{r+1} - x^{r+1} = (r+1)\int_x^{x+1} u^r\,du$, so

$$h(r)^{1/r} = (r+1)^{1/r}\left(\int_x^{x+1} u^r\,du\right)^{1/r}.$$

The second factor is the power mean of order r of the identity function on $[x, x+1]$, and power means tend to the geometric mean as the order tends to 0; the first factor tends to e. So $g(x) = e \cdot \mathrm{GM}_{[x,x+1]}(u)$, and since any mean of values in $[x, x+1]$ lies in $[x, x+1]$, the answer e is forced before any calculation. This is how a 120 scorer reads the problem: the answer is visible from the structure, and the derivative computation is the verification you write down.

The examiners nudged with two planted details: the exponent $r+1$ (base equals 1 at $r=0$) and the division by x in the outer limit (removing the factor GM $\approx x$ so that only e survives). Train the discernment by refusing to interchange limits without a uniform estimate: the official second solution shows the price of the interchange, a Taylor remainder bounded by $(r+1)r/x$ uniformly for $x \geq 1$. Iterated limits genuinely differ in general; $\lim_{r\to 0^+}\lim_{x\to\infty} x^{-r} = 0$ while $\lim_{x\to\infty}\lim_{r\to 0^+} x^{-r} = 1$. When the inner limit has a closed form, compute it and the question disappears.

Revolving concepts

The 1^∞ form and the derivative. If $f(0)=1$ and $f'(0)$ exists then $\lim_{t\to 0} f(t)^{1/t} = e^{f'(0)}$. Equivalently, if $a_n \to 0$ and $a_n b_n \to L$ then $(1+a_n)^{b_n} \to e^L$. Generalised recursion: none, but the rule iterates: for $f(t)^{1/t^2}$ with $f(0)=1$, $f'(0)=0$, the limit is $e^{f''(0)/2}$ by the same logarithm argument one order higher. See Book 1, field: Real analysis, topic: indeterminate forms.

Power means of a function. For a probability measure μ on an interval and positive f, $M_r(f) = (\int f^r\,d\mu)^{1/r}$ for $r \neq 0$, and $M_0(f) = \exp(\int \ln f\,d\mu)$ (the geometric mean). Lemma: $M_r(f) \to M_0(f)$ as $r \to 0$, since $\int f^r\,d\mu = 1 + r\int \ln f\,d\mu + O(r^2)$ for bounded f; and M_r is nondecreasing in r. On $[x, x+1]$ with $f(u)=u$, $M_0 = e^{-1}(x+1)^{x+1}/x^x$. See Book 1, field: Inequalities, topic: power means.

The number e. $(1+1/x)^x \to e$ as $x \to \infty$ through the reals, because $x\ln(1+1/x) = \ln(1+t)/t \to 1$; moreover $(1+1/x)^x$ increases to e and $(1+1/x)^{x+1}$ decreases to e. Also $(1+r)^{1/r} \to e$ as $r \to 0$ from either side. See Book 1, field: Real analysis, topic: the exponential function.

Interchange of limits. Iterated limits agree when one of the limits is uniform in the other variable (the Moore-Osgood theorem); without uniformity they may differ. The Putnam grader expects either a uniform bound or an exact computation of the inner limit. See Book 1, field: Real analysis, topic: uniform convergence.

Circuit Bank

- **WHEN** you see $\lim_{t\to 0} f(t)^{1/t}$ and $f(0) = 1$ **DO** write $\exp\big((\ln f)'(0)\big) = \exp(f'(0))$ and differentiate **BECAUSE** because the logarithm turns the form into a difference quotient at 0, and the difference quotient of a differentiable function is its derivative by definition
- **WHEN** you see $b^{r+1} - a^{r+1}$ with the exponent shifted by one **DO** write $(r+1)\int_a^b u^r\,du$ **BECAUSE** because the shifted exponent is the antiderivative of u^r up to the factor $r+1$, and an integral of u^r over an interval is a power mean that stays between the endpoints
- **WHEN** you see two nested limits in different variables **DO** freeze the outer variable and compute the inner limit exactly **BECAUSE** because interchanging limits needs a uniform estimate you would have to prove, while an exact inner limit makes the question vanish

Backtesting

1. First useful cue: the exponent $r+1$, which makes the base equal to 1 at $r = 0$.
2. Method the expression suggested: logarithm, then the derivative of $\ln h$ at 0; the outer limit then reduces to $(1 + 1/x)^{x+1}$.
3. Where you would hesitate, and the thought to reinforce: the temptation to send $x \to \infty$ first because the base looks simpler. Reinforce: freeze the outer variable; the inner limit is elementary.
4. The trap the question sets: the correct number e is obtainable by an illegal interchange, so a wrong argument produces a right answer and a low score.
5. What makes you recognise it faster next time: any base that evaluates to 1 at the limit point with a reciprocal exponent; and any difference $b^s - a^s$ is $s\int_a^b u^{s-1}\,du$.
6. The reusable decision rule: 1^∞ means differentiate the log of the base at the limit point; never interchange limits you can compute.

Mirror drills

1. (drill) For $a, b > 0$ find $\lim_{r\to 0}\big((a^r + b^r)/2\big)^{1/r}$. *Answer:* the base is 1 at $r = 0$ with log-derivative $(\ln a + \ln b)/2$, so the limit is $\sqrt{ab}$: the power mean tends to the geometric mean.
2. (drill) Find $\lim_{r\to 0}\big(\int_0^1 (1+t)^r\,dt\big)^{1/r}$. *Answer:* $\exp\big(\int_0^1 \ln(1+t)\,dt\big) = \exp(2\ln 2 - 1) = 4/e$.
3. (drill) Find $\lim_{n\to\infty}(n!)^{1/n}/n$. *Answer:* ln of it is $\frac{1}{n}\sum_{k=1}^n \ln(k/n) \to \int_0^1 \ln t\,dt = -1$, so the limit is $1/e$; a geometric mean again.

5.3 A3. Regular tetrahedron with lattice vertices on a sphere

Determine all positive integers N for which the sphere

$$x^2 + y^2 + z^2 = N$$

has an inscribed regular tetrahedron whose vertices have integer coordinates.

Read it out loud

A sphere centred at the origin of radius $\sqrt{N}$; four lattice points on it forming a regular tetrahedron; "determine all N". The last phrase is again a two-part contract: a family of constructions and a proof that nothing outside the family works. The fields wired in are the metric geometry of the regular tetrahedron (equal edges, circumcentre at the centroid), lattice integrality (dot products of integer vectors are integers, determinants of integer matrices are integers), and the algebra that connects lengths to determinants, which is the Gram matrix. The answer will be a Diophantine condition on N, and Diophantine conditions come from integer invariants, so the question is really: which integer invariant of the configuration is pinned down by regularity?

Read the objects by their definitions. A vertex is a position vector $P_i \in \mathbb{Z}^3$ with $|P_i|^2 = N$ (the sphere is centred at the origin, so position vectors are the natural objects, another planted cue). "Regular" means all six edges are equal, and since $|P_i - P_j|^2 = 2N - 2P_i \cdot P_j$, this says exactly that all six dot products $P_i \cdot P_j$ ($i \neq j$) are equal to one integer c. So within thirty seconds you should have written: $|P_i|^2 = N$, $P_i \cdot P_j = c \in \mathbb{Z}$ for $i \neq j$. That is the whole configuration encoded in two integers, and from here the problem is linear algebra.

The examiners were nudging toward $N = 3m^2$: the cube $[-m, m]^3$ has the sphere of radius $m\sqrt{3}$ as circumsphere and its alternate vertices form a regular tetrahedron. The construction is a memory retrieval. The necessity is the test.

Approaches, ranked worst to best

1. **Coordinates and distance equations.** *Cue:* four points, six equal distances. *Circuit:* write down everything. (ten polynomial equations in twelve integer unknowns; no structure, no progress; this is zombie maths and it stalls on the first page)
2. **Volume of the tetrahedron.** *Cue:* regular tetrahedron with known circumradius. *Circuit:* memorised formulas $R = a\sqrt{6}/4$ and $V = a^3/(6\sqrt{2})$, plus $V = |\det|/6$. (correct: $a^2 = 8N/3$ and $\det(P_2 - P_1, P_3 - P_1, P_4 - P_1) = a^3/\sqrt{2}$ is an integer, so $256N^3/27$ is a perfect square; but two formulas must be recalled exactly, and a slip in either is fatal)
3. **Complete to a cube.** *Cue:* a regular tetrahedron is the alternate vertices of a cube. *Circuit:* antipodal points $-P_i$ are lattice points on the same sphere. (the official route: the eight points $\pm P_i$ form a cube of edge $2\sqrt{N/3}$ once you know the centroid is the origin; its volume $8(N/3)^{3/2}$ is an integer determinant; one geometric picture and one formula)
4. **Gram matrix.** *Cue:* integer coordinates plus equal dot products. *Circuit:* $G = MM^{\mathsf{T}}$, $\det G = (\det M)^2$, and four vectors in $\mathbb{R}^3$ have singular Gram matrix. (no geometric formulas at all: singularity of the 4×4 Gram matrix forces $c = -N/3$, and the 3×3 Gram determinant $16k^3$ must be a perfect square; the laziest route because every step is forced by a definition)

The ideal solution

Answer. Exactly the integers $N = 3m^2$ with m a positive integer.

Construction. Let $N = 3m^2$ and take

$$P_1 = (m, m, m), \quad P_2 = (m, -m, -m), \quad P_3 = (-m, m, -m), \quad P_4 = (-m, -m, m).$$

Each has $|P_i|^2 = 3m^2 = N$. The six differences are $(0, 2m, 2m)$, $(2m, 0, 2m)$, $(2m, 2m, 0)$, $(2m, -2m, 0)$, $(2m, 0, -2m)$, $(0, 2m, -2m)$ up to sign, each of squared length $8m^2$. A tetrahedron with six equal edges is regular, so these four lattice points are the vertices of a regular tetrahedron inscribed in the sphere.

Necessity. Suppose $P_1, \dots, P_4 \in \mathbb{Z}^3$ are the vertices of a regular tetrahedron inscribed in $x^2 + y^2 + z^2 = N$, with common edge length $a > 0$. Then $|P_i|^2 = N$ and, for $i \neq j$, $P_i \cdot P_j = \frac{1}{2}\big(|P_i|^2 + |P_j|^2 - |P_i - P_j|^2\big) = N - a^2/2 =: c$, which is an integer because it is a dot product of integer vectors.

Let M be the 4×3 matrix whose rows are $P_1, \dots, P_4$ and let $G = MM^{\mathsf{T}}$ be the 4×4 Gram matrix, so $G_{ij} = P_i \cdot P_j$. Then $G = (N - c)I + cJ$, where J is the all-ones matrix. Since $\operatorname{rank} G \leq \operatorname{rank} M \leq 3$, $\det G = 0$. The matrix J has eigenvalue 4 on $(1, 1, 1, 1)$ and eigenvalue 0 on the three-dimensional orthogonal complement, so G has eigenvalues $N + 3c$ (once) and $N - c$ (three times), and

$$0 = \det G = (N - c)^3(N + 3c).$$

As $N - c = a^2/2 > 0$, we get $c = -N/3$. Since c is an integer, $3 \mid N$; write $N = 3k$ with $k \geq 1$, so $c = -k$.

Now let M_3 be the 3×3 integer matrix with rows P_1, P_2, P_3 and $G_3 = M_3 M_3^{\mathsf{T}}$ its Gram matrix. Then $G_3 = (N - c)I + cJ = 4kI - kJ$ (with I, J now 3×3), whose eigenvalues are $4k - 3k = k$ (once) and $4k$ (twice), so $\det G_3 = 16k^3$. On the other hand $\det G_3 = (\det M_3)^2 = D^2$ with $D \in \mathbb{Z}$. Hence $D^2 = 16k^3$, so $4 \mid D$ and $k^3 = (D/4)^2$ is a perfect square. For every prime p, $3\, v_p(k)$ is even, hence $v_p(k)$ is even; therefore $k = m^2$ for some positive integer m, and $N = 3m^2$.

Combining both parts, the sphere $x^2 + y^2 + z^2 = N$ has an inscribed regular tetrahedron with integer vertices if and only if

$$\boxed{N = 3m^2 \text{ for some positive integer } m.}$$

Layer 1: WHAT

Block 1: encode the configuration as $|P_i|^2 = N$, $P_i \cdot P_j = c \in \mathbb{Z}$. Block 2: four vectors in $\mathbb{R}^3$ have a singular Gram matrix; $(N - c)^3(N + 3c) = 0$ gives $c = -N/3$, hence $3 \mid N$. Block 3: the 3×3 Gram determinant $16k^3$ equals an integer squared, hence k is a square. Block 4: the cube's alternate vertices give the construction.

The nominal route writes coordinates and distances (square roots, twelve unknowns) or reaches for memorised tetrahedron formulas. The bifurcation point is Block 1: the decision to describe the tetrahedron by its Gram matrix (two integers) rather than by its coordinates or its edge length. Every later step is then a definition: rank of a product, eigenvalues of $aI + bJ$, determinant of a product. The zombie route can also arrive, via the circumradius formula, but it arrives carrying formulas it cannot verify in the arena.

Layer 2: HOW

Block 1. **Cue.** "Integer coordinates", "regular", and a sphere centred at the origin. **Why natural.** Integrality lives in dot products and determinants, not in lengths: $|P_i - P_j| = a$ is irrational in general, but $P_i \cdot P_j$ is an integer. Regularity is a statement about six equal lengths, and the polarisation identity $|P_i - P_j|^2 = |P_i|^2 + |P_j|^2 - 2P_i \cdot P_j$ converts it into six equal integers. The centre at the origin makes position vectors the objects, so $|P_i|^2 = N$ is given for free. **Link to the next block.** A set of vectors with all pairwise dot products known is exactly a Gram matrix, and a Gram matrix carries two invariants you can compute: its rank and its determinant.

Block 2. **Cue.** Four vectors in three-dimensional space. **Why natural.** $G = MM^{\mathsf{T}}$ has rank at most 3, so the 4×4 determinant vanishes; and $G = (N - c)I + cJ$ has explicitly known eigenvalues because J is rank one. The factor $N + 3c$ is the eigenvalue on $(1, 1, 1, 1)$, and its vanishing says

$M^{\mathsf{T}}(1,1,1,1) = P_1 + P_2 + P_3 + P_4 = 0$: the centroid is the origin. That is the algebraic form of the geometric fact that a regular tetrahedron's circumcentre is its centroid, derived rather than recalled. **Link to the next block.** $c = -N/3$ is an integer, so $3 \mid N$; the remaining question is what constrains $k = N/3$, and the only unused integer invariant is a determinant.

Block 3. Cue. Three of the vectors are linearly independent (their Gram determinant is $16k^3 \neq 0$) and integral. **Why natural.** $\det(M_3 M_3^{\mathsf{T}}) = (\det M_3)^2$ and $\det M_3 \in \mathbb{Z}$. So the number $16k^3$ must be a perfect square, and a cube that is a square is a sixth power up to units: exponent bookkeeping prime by prime. **Link to the next block.** Necessity is complete: $N = 3m^2$. Now the construction must exist for every m, and the Block 2 fact $\sum P_i = 0$ tells you where to look: the antipodes $-P_i$ are on the sphere too, and eight points of the form $\pm P_i$ with all dot products $\pm k$ form a cube.

Block 4. Cue. $N = 3m^2 = m^2 + m^2 + m^2$. **Why natural.** The cube $[-m, m]^3$ has circumradius $m\sqrt{3}$, and its alternate vertices (those with an even number of minus signs) are pairwise at squared distance $8m^2$. Verifying six differences is faster and safer than any other check. **Link to the next block.** Both halves done; box the family.

Layer 3: WHY

Three truths. First, a configuration of vectors is determined up to isometry by its Gram matrix, so any question "which lengths are possible" is a question about which Gram matrices are realisable, and realisability of a Gram matrix over $\mathbb{Z}^3$ imposes integrality on det and rank conditions on size. Second, for $aI + bJ$ the spectrum is a (with multiplicity $n-1$) and $a + nb$ (once); every equal-dot-product configuration, that is every regular simplex, has a Gram matrix of this form, so the eigenvalue computation is a reflex. Third, an integer matrix E with $EE^{\mathsf{T}} = \lambda I$ in odd dimension n forces $\lambda^{n/2} = |\det E| \in \mathbb{Z}$, hence λ a perfect square; the cube's three edge vectors are such an E with $\lambda = 4k$. This is why the dimension matters: in the plane, $\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}$ has $EE^{\mathsf{T}} = 2I$ with 2 not a square, and indeed a square with integer vertices can be inscribed in $x^2 + y^2 = 2$.

The discernment mechanism: when the answer must be a Diophantine condition, list the integer invariants of the configuration (dot products, determinants, areas times 2, volumes times 6) and ask which one regularity pins to a formula in N. Train it by solving every lattice-geometry problem twice, once with coordinates and once with the Gram matrix, until the second becomes the reflex. The examiners nudged with "integer coordinates" (integrality), "regular" (equal dot products), and the origin-centred sphere (position vectors), and they chose the tetrahedron rather than the octahedron precisely because the tetrahedron needs the extra step $c = -N/3$ before the determinant fires.

Revolving concepts

Regular tetrahedron. Edge a, circumradius $R = a\sqrt{6}/4$, volume $a^3/(6\sqrt{2})$; circumcentre, centroid and incentre coincide; the vertex vectors from the centre satisfy $P_i \cdot P_j = -R^2/3$; it is the set of alternate vertices of a cube of edge $a/\sqrt{2}$. See Book 1, field: Geometry, topic: regular polyhedra.

Gram matrix. For vectors $v_1, \ldots, v_m \in \mathbb{R}^n$ with matrix M of rows v_i, $G = MM^{\mathsf{T}}$; G is positive semidefinite, $\operatorname{rank} G = \operatorname{rank} M \leq n$, and for $m = n$, $\det G = (\det M)^2$ is the squared volume of the parallelepiped. Generalised recursion: a regular n-simplex centred at the origin with circumradius squared N has $c = -N/n$ and the Gram determinant of n of its vertex vectors is $N^n(n+1)^{n-1}/n^n$; for $n = 3$ this is $16N^3/27 = 16k^3$. Schoenberg (1937) used exactly this to classify the n for which a regular n-simplex has vertices in $\mathbb{Z}^n$: always for $n \equiv 3 \pmod 4$, for $n \equiv 1 \pmod 4$ if and only if $n+1$ is a sum of two squares, and for even n if and only if $n+1$ is a perfect square. See Book 1, field: Matrices and determinants, topic: Gram matrices.

Spectrum of $aI + bJ$. Eigenvalues a (multiplicity $n-1$, on the sum-zero hyperplane) and $a+nb$ (on the all-ones vector); determinant $a^{n-1}(a+nb)$. See Book 1, field: Linear algebra, topic: rank-one updates.

Perfect-square bookkeeping. If k^3 is a perfect square then k is, by parity of p-adic valuations; if $D^2 = 16k^3$ then $4 \mid D$. See Book 1, field: Number theory, topic: valuations.

Circuit Bank

- **WHEN** you see lattice points with prescribed equal lengths or angles **DO** write the Gram matrix of the position vectors and record that its entries and its determinant are integers **BECAUSE** because integrality lives in dot products and determinants, while lengths are irrational; the Gram matrix is the smallest complete integer description of the configuration
- **WHEN** you see $n+1$ vectors in $\mathbb{R}^n$ with all dot products known **DO** write $\det G = 0$ and solve for the unknown dot product **BECAUSE** because the Gram matrix of dependent vectors is singular, and for equal dot products the determinant factors as $(N-c)^n(N+nc)$
- **WHEN** you see "for which N" in a lattice problem **DO** compute an integer determinant in terms of N and demand it be an integer or a perfect square **BECAUSE** because a Diophantine condition on N can only come from an integer invariant, and $(\det M)^2 = \det G$ is the invariant that regularity makes computable

Backtesting

1. First useful cue: "integer coordinates" together with "regular", signalling integer dot products that are all equal.
2. Method the expression suggested: Gram matrix $(N-c)I + cJ$; singularity in size four, squared determinant in size three.
3. Where you would hesitate, and the thought to reinforce: whether $c = -N/3$ needs the circumradius formula. Reinforce: it falls out of $\det G = 0$; derive, do not recall.
4. The trap the question sets: reaching for coordinates and distances, which produce square roots and twelve unknowns.
5. What makes you recognise it faster next time: any regular figure inscribed in an origin-centred sphere is a Gram matrix of the form $aI + bJ$.
6. The reusable decision rule: encode lattice configurations by Gram matrices; extract integrality from determinants.

Mirror drills

1. (drill) For which positive integers N does the circle $x^2 + y^2 = N$ have an inscribed square with integer vertices? *Answer:* exactly the N that are sums of two squares; if (a, b) is on the circle, so are $(-b, a), (-a, -b), (b, -a)$, and conversely any lattice point on the circle gives $N = a^2 + b^2$.
2. (drill) Show no equilateral triangle has all three vertices in $\mathbb{Z}^2$. *Answer:* with edge vectors u, v, $|u|^2 = |v|^2 = a^2$ and $u \cdot v = a^2/2$, the Gram determinant is $3a^4/4$, which must be the square of an integer; that makes $\sqrt{3}$ rational.
3. (drill) For which N does $x^2 + y^2 + z^2 = N$ have an inscribed regular octahedron with integer vertices? *Answer:* the vertices are $\pm u_1, \pm u_2, \pm u_3$ with u_i mutually orthogonal integer vectors of squared length N; the matrix E with rows u_i has $EE^{\mathsf{T}} = NI$, so $N^{3/2} = |\det E| \in \mathbb{Z}$ and N is a perfect square; conversely $(m, 0, 0), (0, m, 0), (0, 0, m)$ and their negatives work.

5.4 A4. Two divergent integrals whose difference survives

Let

$$I(R) = \iint_{x^2+y^2\le R^2} \left(\frac{1+2x^2}{1+x^4+6x^2y^2+y^4} - \frac{1+y^2}{2+x^4+y^4} \right) dx\, dy.$$

Find

$$\lim_{R\to\infty} I(R),$$

or show that this limit does not exist.

Read it out loud

Say the structure back to yourself before you touch a pen. This is a double integral over an expanding disc of a difference of two rational functions, each of which is a quadratic in x and y divided by a quartic in x and y. Both denominators are strictly positive everywhere, so for each fixed R there is nothing improper about the integral: the integrand is continuous on a compact set and $I(R)$ is an honest number. The only analysis lives in the limit.

Now the size count, which is the single most informative thing you can do in the first thirty seconds. At radius r, each numerator is of size r^2 and each denominator is of size r^4, so each integrand is of size r^{-2}, and the area element in polar coordinates is $r\,dr\,d\theta$. Each of the two pieces on its own therefore contributes $\int^R r^{-1}\,dr$, which is $\log R$: each piece diverges. The problem is engineered so that two logarithmic divergences cancel and a finite residue is left behind. That tells you the whole shape of the answer immediately. You are not being asked to integrate anything; you are being asked to match the two pieces against each other so exactly that only the residue remains. Any route that computes the two integrals separately is dead on arrival, because there is nothing to compute separately.

The fields wired into this question: symmetry of a domain under a finite group and averaging over that group, linear changes of variable in the plane with their Jacobians, polar coordinates for quartic forms, elementary rational trigonometric integrals, and a controlled interchange of a limit with an integral. The examiners' guiding lights are the two denominators. Look at them side by side: $x^4+6x^2y^2+y^4$ and x^4+y^4, with constants 1 and 2. The binomial identity $(x+y)^4+(x-y)^4 = 2(x^4+6x^2y^2+y^4)$ is A-level algebra, and the constants 1 and 2 sit in exactly the ratio that identity produces. That is the plant. The map $(x,y)\mapsto(x+y,x-y)$ is what the question is about, and everything else is bookkeeping.

The only obstruction is that the numerators do not match: $1+2x^2$ against $1+y^2$. So within thirty seconds you should have written down the two denominators, the identity for $(x+y)^4+(x-y)^4$, and the note "numerators asymmetric, domain symmetric, so average them".

Approaches, ranked worst to best

1. **Evaluate each term separately.** *Cue:* two fractions joined by a minus sign. *Circuit:* split the integral, integrate term by term. Ranks last because it is not merely long, it is invalid. Each term integrates to something of order $\log R$, so neither limit exists and the difference of the two limits is meaningless. A candidate who writes $\lim(A(R)-B(R)) = \lim A(R) - \lim B(R)$ here has lost the problem in one line.

2. **Polar coordinates then grind the radial integrals.** *Cue:* quartic forms become r^4 times a function of θ. *Circuit:* substitute $u=r^2$, partial fractions, arctan plus log. This does work. For fixed θ the radial integrals are elementary and give an arctangent plus a logarithm, so you obtain $I(R) = \alpha(\theta)$ terms integrated over θ with a $\log R$ coefficient $\int_0^{2\pi}(\frac{\cos^2\theta}{2(1+\sin^2 2\theta)} - \frac{\sin^2\theta}{4(1-\frac{1}{2}\sin^2 2\theta)})\,d\theta$ that you must then prove vanishes, plus a finite part you must then assemble.

Two pages of arctangent bookkeeping, and you still have to discover the angular shift to evaluate the finite part. It costs far more ink than the structure demands.

3. **Discover the reroute inside polar coordinates.** *Cue:* $x^4 + 6x^2y^2 + y^4 = r^4(1 + \sin^2 2\theta)$ and $x^4 + y^4 = r^4(1 - \frac{1}{2}\sin^2 2\theta)$. *Circuit:* compare the two angular profiles and look for a shift. In polar the two denominators are $1 + r^4(1 + \sin^2 2\theta)$ and $2 + r^4(1 - \frac{1}{2}\sin^2 2\theta)$, and replacing (r, θ) by $(r\sqrt{2}, \theta + \pi/4)$ in the second gives exactly twice the first. This is the same reroute seen through polar spectacles and it is a good route, but you have to notice a simultaneous rotation and dilation inside trigonometric identities rather than in one line of binomial algebra.

4. **Kill the odd parts, then apply the linear map to the first term only.** *Cue:* $D_1 = 1 + x^4 + 6x^2y^2 + y^4$ is invariant under both $x \leftrightarrow y$ and $x \mapsto -x$. *Circuit:* odd numerators integrate to zero over a symmetric domain. Replace $1 + 2x^2$ by $1 + (x - y)^2$, legal because $(1 + 2x^2) - (1 + (x - y)^2) = x^2 + 2xy - y^2$ is a sum of two pieces each odd under a symmetry of the disc and of D_1. Then the first term is exactly twice the second term composed with $(x, y) \mapsto (x + y, x - y)$. Excellent and short, but it needs two separate parity arguments rather than one principle.

5. **Symmetrise the whole integrand, then apply the linear map.** *Cue:* the disc and both denominators are invariant under $x \leftrightarrow y$, the numerators are not. *Circuit:* average a function over the symmetry group of the domain before doing anything else. The best route. One legal move makes both numerators symmetric, after which the two terms become the same function up to the map $(x, y) \mapsto (x + y, x - y)$, and the whole difference collapses to a single integral over an annulus. Three lines to the collapse, and the surviving integral is a standard rational trigonometric integral.

The ideal solution

The limit exists and equals

$$\boxed{\frac{\pi \log 2}{\sqrt{2}}} = \frac{\sqrt{2}}{2}\,\pi \log 2,$$

where log is the natural logarithm.

Write $D_R = \{(x, y) : x^2 + y^2 \le R^2\}$. Both denominators are at least 1, so the integrand is continuous on D_R and $I(R)$ exists for each R.

Step 1: symmetrise. The reflection $\sigma(x, y) = (y, x)$ maps D_R onto itself and preserves area, so $\iint_{D_R} F = \iint_{D_R} F \circ \sigma$ for any continuous F, and hence

$$\iint_{D_R} F = \iint_{D_R} \tfrac{1}{2}\,(F + F \circ \sigma)\,.$$

Both denominators are σ invariant, so only the numerators feel the averaging:

$$I(R) = \iint_{D_R} \big(f(x, y) - g(x, y)\big)\,dx\,dy,$$

$$f(x, y) = \frac{1 + x^2 + y^2}{1 + x^4 + 6x^2y^2 + y^4}, \qquad g(x, y) = \frac{1 + \frac{1}{2}(x^2 + y^2)}{2 + x^4 + y^4}.$$

Step 2: the linear map. Let $T(x, y) = (x + y, x - y)$. Then

$$(x + y)^2 + (x - y)^2 = 2(x^2 + y^2), \qquad (x + y)^4 + (x - y)^4 = 2(x^4 + 6x^2y^2 + y^4),$$

the second being the even part of the binomial expansion doubled. Therefore

$$g(T(x, y)) = \frac{1 + x^2 + y^2}{2\,(1 + x^4 + 6x^2y^2 + y^4)} = \tfrac{1}{2} f(x, y).$$

The map T is linear with $\det T = -2$, and $|T(p)| = \sqrt{2}\,|p|$, so T carries D_R bijectively onto $D_{R\sqrt{2}}$. The change of variables formula gives

$$\iint_{D_R} f = 2\iint_{D_R} g\circ T = 2\cdot\frac{1}{|\det T|}\iint_{D_{R\sqrt{2}}} g = \iint_{D_{R\sqrt{2}}} g.$$

Subtracting the integral of g over D_R,

$$I(R) = \iint_{D_{R\sqrt{2}}} g - \iint_{D_R} g = \iint_{A_R} g, \qquad A_R = \{R \le |p| \le R\sqrt{2}\}.$$

Step 3: polar coordinates and an exact radial integral. Since $x^4+y^4 = r^4\left(1-\frac{1}{2}\sin^2 2\theta\right)$, write $c = c(\theta) = 1-\frac{1}{2}\sin^2 2\theta$, so that $\frac{1}{2} \le c \le 1$. Then

$$I(R) = \int_0^{2\pi}\int_R^{R\sqrt{2}} \frac{\left(1+\frac{1}{2}r^2\right) r}{2+r^4 c}\,dr\,d\theta.$$

For fixed θ substitute $u = r^2$:

$$\int_R^{R\sqrt{2}} \frac{\left(1+\frac{1}{2}r^2\right) r\,dr}{2+r^4c} = \frac{1}{2}\int_{R^2}^{2R^2}\frac{du}{2+cu^2} + \frac{1}{4}\int_{R^2}^{2R^2}\frac{u\,du}{2+cu^2},$$

$$= \frac{1}{2\sqrt{2c}}\left[\arctan\left(u\sqrt{\tfrac{c}{2}}\right)\right]_{R^2}^{2R^2} + \frac{1}{8c}\log\frac{2+4cR^4}{2+cR^4}.$$

Step 4: uniform error bounds. For $B > 0$ we have $\arctan(2B) - \arctan(B) \le \frac{\pi}{2} - \arctan B = \operatorname{arctan}\frac{1}{B} \le \frac{1}{B}$. Taking $B = R^2\sqrt{c/2}$ and using $c \ge \frac{1}{2}$, the arctangent term is at most

$$\frac{1}{2\sqrt{2c}}\cdot\frac{1}{R^2\sqrt{c/2}} = \frac{1}{2cR^2} \le \frac{1}{R^2}.$$

For the logarithm, $\frac{2+4cR^4}{2+cR^4} = 4\,(1-t)$ with $t = \frac{3}{2(2+cR^4)}$, and $2+cR^4 \ge \frac{1}{2}R^4$ gives $0 \le t \le 3R^{-4}$. For $R \ge 2$ this puts $t \le \frac{1}{2}$, where $|\log(1-t)| \le 2t$, so

$$\left|\frac{1}{8c}\log\frac{2+4cR^4}{2+cR^4} - \frac{\log 4}{8c}\right| \le \frac{1}{8c}\cdot\frac{6}{R^4} \le \frac{3}{2R^4}.$$

Since $\frac{\log 4}{8c} = \frac{\log 2}{4c}$, the bracketed radial integral differs from $\frac{\log 2}{4c(\theta)}$ by at most $R^{-2} + \frac{3}{2}R^{-4}$, uniformly in θ. Integrating over $\theta \in [0, 2\pi]$,

$$\left|I(R) - \frac{\log 2}{4}\int_0^{2\pi}\frac{d\theta}{c(\theta)}\right| \le 2\pi\left(\frac{1}{R^2} + \frac{3}{2R^4}\right) \longrightarrow 0.$$

Step 5: the angular integral. Put $\phi = 2\theta$. The function $\phi \mapsto (1-\frac{1}{2}\sin^2\phi)^{-1}$ has period π and is symmetric about $\phi = \frac{\pi}{2}$, so

$$\int_0^{2\pi}\frac{d\theta}{1-\frac{1}{2}\sin^2 2\theta} = \frac{1}{2}\int_0^{4\pi}\frac{d\phi}{1-\frac{1}{2}\sin^2\phi} = 2\int_0^{\pi}\frac{d\phi}{1-\frac{1}{2}\sin^2\phi} = 4\int_0^{\pi/2}\frac{d\phi}{1-\frac{1}{2}\sin^2\phi}.$$

Divide numerator and denominator by $\cos^2\phi$ and set $t = \tan\phi$:

$$\int_0^{\pi/2}\frac{d\phi}{1-\frac{1}{2}\sin^2\phi} = \int_0^{\infty}\frac{dt}{1+\frac{1}{2}t^2} = \sqrt{2}\int_0^{\infty}\frac{ds}{1+s^2} = \frac{\pi}{\sqrt{2}}.$$

Hence $\int_0^{2\pi} c(\theta)^{-1}\,d\theta = 2\sqrt{2}\,\pi$, and

$$\lim_{R\to\infty} I(R) = \frac{\log 2}{4}\cdot 2\sqrt{2}\,\pi = \frac{\pi\log 2}{\sqrt{2}} \approx 1.5398.$$

Layer 1: WHAT

Block 1: size count in polar shows each half diverges like $\log R$, so the problem is a cancellation problem, not an evaluation problem.

Block 2: the disc and both denominators are invariant under $x \leftrightarrow y$, so replace each numerator by its average under that swap at no cost.

Block 3: the symmetrised numerators are $1 + x^2 + y^2$ and $1 + \frac{1}{2}(x^2 + y^2)$, and the symmetrised denominators are linked by $(x+y)^4 + (x-y)^4 = 2(x^4 + 6x^2y^2 + y^4)$; hence $f = 2\,(g \circ T)$ with $T(x,y) = (x+y, x-y)$.

Block 4: T scales lengths by $\sqrt{2}$ and has $|\det T| = 2$, so integrating f over D_R equals integrating g over $D_{R\sqrt{2}}$; the difference telescopes to g over the annulus between the two discs.

Block 5: on the annulus, polar coordinates plus $u = r^2$ give an exact arctangent plus logarithm; the arctangent part is $O(R^{-2})$ and the logarithm tends to $\frac{\log 2}{4c(\theta)}$ uniformly.

Block 6: the surviving angular integral is a standard tan substitution.

The nominal route is Block 1 followed by "polar coordinates, integrate in r, then in θ". It reaches the answer but only after separating a divergent $\log R$ coefficient, proving that coefficient is zero by a second angular integral, and then reassembling a finite part from four arctangent limits. The bifurcation point is the line immediately after the size count: the zombie writes $x = r\cos\theta$, the ideal route writes "the domain has a symmetry the numerators lack". Everything downstream follows from that single choice, because symmetrising is what turns two unrelated fractions into one fraction seen at two scales.

Layer 2: HOW

Block 1. **Cue.** Two rational functions, numerator degree 2, denominator degree 4, integrated over a disc of growing radius, with the answer allowed not to exist. **Why natural.** The examiners explicitly offer "or show that this limit does not exist", which is a confession that convergence is the delicate point. A degree count in polar coordinates settles it in one line: integrand of size r^{-2} against area element $r\,dr\,d\theta$ gives $\int^R \frac{dr}{r}$. Both halves diverge, so the minus sign between them is load bearing and must never be split. **Link to the next block.** Knowing the two halves must cancel to leading order tells you they are the same object in disguise, so your job is to find the disguise.

Block 2. **Cue.** The region $x^2 + y^2 \le R^2$ is symmetric in x and y, and so are both denominators, but the numerators $1 + 2x^2$ and $1 + y^2$ are not. **Why natural.** If a measure preserving map σ of the domain onto itself exists, then $\iint F = \iint F \circ \sigma$, and averaging the two is free. You use it whenever the asymmetry sits in a part of the expression you wish were symmetric. Here it is a targeted move: the denominators already have the symmetry, so averaging changes nothing except to repair the numerators. This is not a guess about what might help; it is removing the only structural mismatch in sight. **Link to the next block.** With symmetric numerators, each fraction is now a function of $x^2 + y^2$ over a quartic form, so the two fractions can only differ by a change of variable.

Block 3. **Cue.** $x^4 + 6x^2y^2 + y^4$ next to $x^4 + y^4$, with constant terms 1 and 2. **Why natural.** $(x+y)^4 + (x-y)^4$ is twice the even part of $(x+y)^4$, that is $2(x^4 + 6x^2y^2 + y^4)$. The constants match the same factor of two, and $(x+y)^2 + (x-y)^2 = 2(x^2+y^2)$ makes the numerators match as well. Three separate matches under the same substitution is not a coincidence; it is the problem being read backwards from g. That is the reason the move is natural rather than inspired. **Link to the next block.** You now have $f = 2(g \circ T)$ as a pointwise identity, so all that is left is to see what T does to the domain.

Block 4. **Cue.** $T(x,y) = (x+y, x-y)$ is a rotation by $\frac{\pi}{4}$ composed with a dilation by $\sqrt{2}$. **Why natural.** $|T(p)| = \sqrt{2}|p|$ maps the disc of radius R onto the disc of radius $R\sqrt{2}$, and $|\det T| = 2$

exactly cancels the factor 2 in $f = 2(g \circ T)$. So $\iint_{D_R} f$ and $\iint_{D_{R\sqrt{2}}} g$ are literally the same number. A difference of integrals of the same function over nested discs is an integral over the annulus between them. **Link to the next block.** The annulus $R \le |p| \le R\sqrt{2}$ has inner and outer radii in fixed ratio, which is exactly the shape that makes a log finite: $\int_R^{R\sqrt{2}} \frac{dr}{r} = \log\sqrt{2}$ for every R.

Block 5. Cue. An annulus integral of a fraction whose denominator is $2 + r^4 c(\theta)$. **Why natural.** Substituting $u = r^2$ makes both pieces standard: $\int \frac{du}{a+cu^2}$ is an arctangent and $\int \frac{u\,du}{a+cu^2}$ is a logarithm. Do the radial integral exactly rather than passing to the limit inside, because then the limit interchange needs no theorem, only two inequalities. The arctangent difference over $[R^2, 2R^2]$ is $O(R^{-2})$ because both endpoints are near $\frac{\pi}{2}$; the logarithm of a ratio of two quartics tends to $\log 4$ at rate R^{-4}, and both bounds hold uniformly in θ because $c(\theta) \ge \frac{1}{2}$. **Link to the next block.** A uniform bound integrated over a set of finite measure is still a bound, so the limit passes through the θ integral with an explicit error, not an appeal to a convergence theorem.

Block 6. Cue. $\int_0^{2\pi} \frac{d\theta}{1-\frac{1}{2}\sin^2 2\theta}$. **Why natural.** A rational function of $\sin^2$ and $\cos^2$ over a full period is handled by $t = \tan\phi$ after reducing to a quarter period using periodicity and reflection. That substitution turns $1 - \frac{1}{2}\sin^2\phi$ into $1 + \frac{1}{2}t^2$ and the integral into an arctangent. **Link to the next block.** Multiply by the radial factor $\frac{\log 2}{4}$ and box the answer.

Layer 3: WHY

The mathematical truth underneath this problem is that a conditionally convergent improper integral in the plane has no value until you specify the exhaustion, and that the value you get is a statement about how the integrand's tail is arranged around infinity, not about the integrand near the origin. Both halves here have tails of size r^{-2}, the exact borderline in two dimensions. Because the disc exhaustion is prescribed, and because the two halves are the same function viewed at scales differing by the factor $\sqrt{2}$, the difference of the two tails is exactly the mass sitting in one annulus of ratio $\sqrt{2}$, and that mass has a finite limit because the integrand is asymptotically homogeneous of degree -2. In other words: a function that behaves like $\frac{h(\theta)}{r^2}$ at infinity assigns the mass $\log(\lambda)\int h$ to every annulus of ratio λ, independently of where the annulus sits. That is the entire content of the answer, and $\log\sqrt{2} = \frac{1}{2}\log 2$ is where the $\log 2$ comes from.

The discernment mechanism is a two part test that you should run on every symmetric looking integral. First, list the symmetries of the domain (here the dihedral group of the square acts on the disc, and the swap $x \leftrightarrow y$ is the element you need). Second, ask which factors of the integrand already have that symmetry. If the denominator has it and the numerator does not, symmetrising is free and strictly simplifying, so it is never wrong to do it. This is the same instinct as replacing an expression by its symmetric average in an inequality: you are choosing a representative of an equivalence class, and you choose the most symmetric one.

The second discernment is scale matching. When two expressions differ by a constant factor in the same places (the 1 against the 2 in the constant terms, the $x^2 + y^2$ against $\frac{1}{2}(x^2 + y^2)$ in the numerators, the two quartic forms), the correct hypothesis is that a single linear map converts one into the other, and the factor is a power of the dilation ratio. Test that hypothesis by comparing degrees: a linear map scaling lengths by λ multiplies a form of degree d by λ^d. Here the quartics differ by $2 = (\sqrt{2})^2$ on the constants and by $(\sqrt{2})^4/2$ on the quartic parts, consistently with $\lambda = \sqrt{2}$. You do not search blindly for a substitution; you read off λ from the constants and then look for the rotation that fixes the angular profile.

To train this, take any Putnam integral with a quartic or biquadratic denominator and write the denominator in polar form. The angular profile of $x^4 + y^4$ is $1 - \frac{1}{2}\sin^2 2\theta$ and that of $x^4 + 6x^2y^2 + y^4$ is $1 + \sin^2 2\theta$; these are the same profile up to a quarter turn and a factor of two, and once you have seen that once you will see it forever. What the examiners nudged toward: they wrote the two

denominators in the same variables, in the same order, with constants in the ratio $1:2$, and offered you the escape hatch of nonexistence to force you to do the size count first.

Revolving concepts

Symmetrisation over a group action. If a finite group G acts on a domain D by measure preserving bijections, then $\iint_D F = \iint_D F^G$ where $F^G = \frac{1}{|G|}\sum_{\sigma\in G} F\circ\sigma$. For the disc, G can be taken to be any subgroup of $O(2)$; here the two element group generated by $x \leftrightarrow y$ is enough. The same principle in one dimension is $\int_{-a}^{a} F = \int_{-a}^{a} \frac{F(x)+F(-x)}{2}$, which is how odd parts are killed. See Book 1, field: Calculus and integrals, topic: symmetry reductions.

Linear change of variables in the plane. For an invertible linear T, $\iint_{T(E)} h = |\det T| \iint_E h\circ T$. The two facts you need about $T(x,y) = (x+y, x-y)$ are $|\det T| = 2$ and $|T(p)| = \sqrt{2}|p|$; the latter says T is $\sqrt{2}$ times an orthogonal map, namely a reflection composed with a rotation by $\frac{\pi}{4}$. See Book 1, field: Linear algebra, topic: similarity transformations and Jacobians.

Quartic forms in polar coordinates. $x^2+y^2 = r^2$, $x^4+y^4 = r^4(1-\frac{1}{2}\sin^2 2\theta)$, $x^4+6x^2y^2+y^4 = r^4(1+\sin^2 2\theta)$, $x^4-6x^2y^2+y^4 = r^4\cos 4\theta$, and $x^2y^2 = \frac{1}{4}r^4\sin^2 2\theta$. Every symmetric biquadratic is a combination of r^4 and $r^4\cos 4\theta$, which is why quarter turns keep appearing. See Book 1, field: Geometry, topic: invariant forms of the dihedral group.

Homogeneous tails and annulus mass. If $F(r,\theta) = \frac{h(\theta)}{r^2} + O(r^{-3})$ then $\iint_{\rho\le|p|\le\lambda\rho} F \to \log\lambda \int_0^{2\pi} h(\theta)\,d\theta$ as $\rho\to\infty$. This is the exact statement that borderline r^{-2} tails are logarithmic, and it is worth carrying as a formula. The general recursion hiding here is geometric rather than arithmetic: the annuli $A_R, A_{R\sqrt{2}}, A_{2R}, \dots$ each carry the same limiting mass, and the sum of k of them is $\iint_{|p|\le R2^{k/2}} \sim \frac{k}{2}\log 2 \int h$, which is the $\log R$ divergence you diagnosed in Block 1, recovered as a telescoping sum. Generalise: a function asymptotically homogeneous of degree $-n$ in n dimensions has scale invariant annulus mass, and the divergence of its integral is exactly logarithmic; for degree $-n-\varepsilon$ the annulus masses form a geometric sequence with ratio $\lambda^{-\varepsilon}$ and the total converges. See Book 1, field: Real analysis, topic: convergence tests for multiple integrals.

Standard rational trigonometric integrals. $\int_0^{\pi/2} \frac{d\phi}{a\cos^2\phi + b\sin^2\phi} = \frac{\pi}{2\sqrt{ab}}$ by $t = \tan\phi$, and $\int_0^{2\pi} \frac{d\theta}{a+b\cos\theta} = \frac{2\pi}{\sqrt{a^2-b^2}}$ for $a > |b|$ by the half angle substitution $t = \tan\frac{\theta}{2}$ or by a residue computation. Either gives the $2\sqrt{2}\pi$ above. See Book 1, field: Calculus and integrals, topic: trigonometric rational integrals.

Interchanging a limit with an integral. The clean tools are dominated convergence and uniform convergence on a set of finite measure. On the Putnam, the version that costs least ink and needs no citation is an explicit uniform error bound: show $|F_R(\theta) - F(\theta)| \le \varepsilon(R)$ for all θ with $\varepsilon(R)\to 0$, then integrate the inequality. See Book 1, field: Real analysis, topic: uniform convergence.

Circuit Bank

- **WHEN** a difference of two expressions is integrated over an unbounded or growing region **DO** check the size of each piece separately before splitting anything **BECAUSE** if each piece diverges, the minus sign is structural and splitting the integral is not merely inefficient, it is false
- **WHEN** the domain of integration has a symmetry that the denominator has but the numerator lacks **DO** replace the numerator by its average over that symmetry **BECAUSE** averaging over a measure preserving symmetry of the domain never changes the integral, so it is a free simplification that can only reveal structure
- **WHEN** two expressions in the same variables differ by matching constant factors on matching degrees **DO** hypothesise a single linear map and read the dilation ratio off the constants

BECAUSE a linear map scaling lengths by λ multiplies each homogeneous part of degree d by λ^d, so matching ratios across degrees determines λ before you have found the map

- **WHEN** you have $\iint_{D_1} h - \iint_{D_2} h$ with $D_2 \subset D_1$ **DO** write it as one integral over $D_1 \setminus D_2$ **BECAUSE** a difference of integrals of the same function over nested regions is the mass of the region between them, and that region is usually where all the analysis lives
- **WHEN** a limit must pass through an integral in a competition write up **DO** do the inner integral exactly first, then bound the error uniformly in the outer variable **BECAUSE** an explicit bound is shorter than invoking a convergence theorem and leaves a grader nothing to object to

Backtesting

1. First useful cue: numerator of degree 2 over denominator of degree 4 in the plane, which is the borderline r^{-2} tail, so each half diverges logarithmically and only the difference can converge.
2. Method the expression suggested: the pair $x^4 + 6x^2y^2 + y^4$ and $x^4 + y^4$ with constants 1 and 2 is the signature of $(x+y)^4 + (x-y)^4 = 2(x^4 + 6x^2y^2 + y^4)$, so the intended move is the linear map $(x, y) \mapsto (x+y, x-y)$.
3. Where you would hesitate, and the thought to reinforce: the numerators do not match, which stalls the substitution. The thought to install is that the disc is symmetric in x and y and so are both denominators, so you are entitled to average the numerators first. Asymmetry in a numerator over a symmetric domain is repairable, always, at zero cost.
4. The trap the question sets: computing the two integrals separately, or passing to the limit inside the θ integral without a bound. The first is fatal, the second loses presentation marks and can be fixed in two lines.
5. What makes you recognise it faster next time: memorise the polar profiles of the three symmetric biquadratics, and treat a growing disc with a r^{-2} tail as an annulus problem from the first line.
6. The reusable decision rule: symmetrise before you substitute, then match constants across degrees to find the linear map, then telescope nested discs into an annulus.

Mirror drills

1. Evaluate $\lim_{R\to\infty} \iint_{x^2+y^2\le R^2} \left(\frac{x^2}{1+x^4+y^4} - \frac{y^2}{1+x^4+y^4} \right) dx\, dy$. *Answer:* zero, and you should write one line: the integrand is odd under $x \leftrightarrow y$ while the disc is invariant, so every $I(R)$ is zero, no analysis needed.
2. Show that $\int_0^{2\pi} \frac{d\theta}{1+\sin^2 2\theta} = \sqrt{2}\,\pi$. *Answer:* periodicity reduces to $4\int_0^{\pi/2}$, then $t = \tan\phi$ gives $4\int_0^{\infty} \frac{dt}{1+2t^2} = 4 \cdot \frac{\pi}{2\sqrt{2}}$.
3. Let F be continuous on $\mathbb{R}^2$ with $F(x,y) = \frac{h(\theta)}{r^2}$ exactly, for $r \ge 1$, where h is continuous and 2π periodic. Compute $\iint_{1\le|p|\le\lambda} F$ and deduce for which h the integral over the whole plane converges. *Answer:* it equals $\log\lambda \int_0^{2\pi} h$, so it converges as $\lambda \to \infty$ only when $\int_0^{2\pi} h = 0$, which is precisely the cancellation the main problem engineers.

5.5 A5. Power sums over the units modulo 2021

Let A be the set of all integers n such that $1 \leq n \leq 2021$ and $\gcd(n, 2021) = 1$. For every nonnegative integer j, let

$$S(j) = \sum_{n \in A} n^j.$$

Determine all values of j such that $S(j)$ is a multiple of 2021.

Read it out loud

Read the definition of A before you read the sum. The set of integers in $[1, 2021]$ coprime to 2021 is not a random set: it is a complete set of representatives for the unit group $(\mathbb{Z}/2021\mathbb{Z})^\times$. So $S(j)$ is not a sum over an interval, it is a sum over a group, and $n \mapsto n^j$ is a group homomorphism on that group. Once you have said that sentence out loud, the problem has already changed character: sums of a homomorphism over a whole group are governed by one principle, and the only question left is when the homomorphism is trivial.

Second, read the modulus. The question asks when $S(j)$ is a multiple of 2021, and $2021 = 43 \times 47$ with both factors prime. A divisibility statement modulo a product of distinct primes is exactly the conjunction of two divisibility statements, one modulo each prime. So the answer will have the shape "condition at 43 and condition at 47", and any answer of a single condition is suspect on sight. The competition year being 43×47 is not decoration; it is the instruction.

The fields wired in: elementary number theory (Fermat's little theorem, primitive roots), the Chinese remainder theorem as an isomorphism of unit groups rather than a solving device, and the orthogonality principle for characters of finite abelian groups, which at this level is one three line argument rather than a theory. What is really being tested is whether you know that $(\mathbb{Z}/p\mathbb{Z})^\times$ is cyclic while $(\mathbb{Z}/2021\mathbb{Z})^\times$ is not, and whether you therefore refuse to run the argument at the composite modulus.

Within thirty seconds you should have written three things: $2021 = 43 \cdot 47$; $|A| = \phi(2021) = 42 \cdot 46 = 1932$; and the note "$A$ is the unit group, so multiply the index by a unit". The two numbers 42 and 46 are the orders of the two cyclic factors, and they are going to be the answer.

Approaches, ranked worst to best

1. **Compute $S(j)$ for many j and pattern match.** *Cue:* a free parameter j and a request to determine all values. *Circuit:* tabulate and guess. Ranks last: $S(1)$ alone is a sum of 1932 terms, and even a correct guess of the answer earns nothing without the two congruence arguments, which are the whole problem. Experiment is fine as private reconnaissance, but here the structure is visible without it.

2. **Faulhaber and inclusion exclusion.** *Cue:* a sum of jth powers over a range. *Circuit:* write $\sum_{k=1}^{m} k^j$ by Bernoulli polynomials. By inclusion exclusion, $S(j) = T_j(2021) - 43^j T_j(47) - 47^j T_j(43) + 2021^j$ where $T_j(m) = \sum_{k=1}^{m} k^j$. Each T_j is a Bernoulli polynomial evaluation with rational coefficients whose denominators must then be tracked modulo 43 and 47. It is correct and it is punishing, and the answer never becomes visible. It also fights the structure: the coprimality condition is a group condition, and inclusion exclusion is the wrong tool for a group.

3. **Multiply the summation variable by a unit, working modulo 2021.** *Cue:* A is closed under multiplication by any u coprime to 2021. *Circuit:* reindex a sum over a group by a translation. Take u coprime to 2021; then $n \mapsto un$ permutes A modulo 2021, so $S(j) \equiv u^j S(j)$ and $2021 \mid (u^j - 1)S(j)$. If you can choose u with $u^j - 1$ invertible modulo 2021, you get $S(j) \equiv 0$ at once. This delivers the sufficient direction in three lines and is a genuinely good move. It

cannot deliver necessity, because when $u^j - 1$ is a zero divisor the conclusion is empty, and $(\mathbb{Z}/2021\mathbb{Z})^\times \cong C_{42} \times C_{46}$ is not cyclic, so zero divisors are unavoidable. Ranks here: half a solution, correctly reasoned.

4. **Pair n with $2021 - n$.** *Cue:* the set A is symmetric about $2021/2$. *Circuit:* involution pairing to kill odd powers. Since $n \in A$ if and only if $2021 - n \in A$, and $(2021 - n)^j \equiv (-n)^j \pmod{2021}$, every odd j gives $S(j) \equiv -S(j)$, and 2021 being odd gives $S(j) \equiv 0$. One line, and it settles every odd j for free. It is not a full solution because 42 and 46 are even, so all the interesting cases are even j, but it is the correct first probe and it is your check on the final answer.

5. **Split by the Chinese remainder theorem, then a primitive root at each prime.** *Cue:* "multiple of 2021" with $2021 = 43 \cdot 47$. *Circuit:* reduce a composite modulus to prime moduli, where the unit group is cyclic. The complete route. Reducing A modulo 43 hits each nonzero class exactly 46 times, so $S(j) \equiv 46 \sum_{a=1}^{42} a^j \pmod{43}$, and the classical power sum lemma at a prime finishes it. Shortest, fully rigorous, and the only route that explains why the answer has two conditions.

The ideal solution

The answer: $S(j)$ is a multiple of 2021 precisely for those j that are

$$\boxed{\text{neither a multiple of 42 nor a multiple of 46.}}$$

Lemma 5.5.1. *Let p be a prime and let $j \geq 0$ be an integer. Then*

$$\sum_{n=1}^{p-1} n^j \equiv \begin{cases} -1 \pmod{p} & \text{if } (p-1) \mid j, \\ 0 \pmod{p} & \text{if } (p-1) \nmid j. \end{cases}$$

Proof. Suppose $(p-1) \mid j$. If $j = 0$ then $n^j = 1$ for every n; if $j > 0$ then $n^j = (n^{p-1})^{j/(p-1)} \equiv 1 \pmod{p}$ for $1 \leq n \leq p-1$ by Fermat's little theorem. Either way the sum is congruent to $p - 1 \equiv -1$.

Suppose $(p-1) \nmid j$. Write $S = \sum_{n=1}^{p-1} n^j$ and let g be a primitive root modulo p, which exists because $(\mathbb{Z}/p\mathbb{Z})^\times$ is cyclic of order $p-1$. Multiplication by g permutes the nonzero residue classes modulo p, so reindexing the sum by $n \mapsto gn$ gives

$$S \equiv \sum_{n=1}^{p-1} (gn)^j = g^j S \pmod{p}.$$

Since g has order exactly $p-1$ and $(p-1) \nmid j$, we have $g^j \not\equiv 1$, so $g^j - 1$ is a nonzero element of the field $\mathbb{Z}/p\mathbb{Z}$ and hence invertible. From $(g^j - 1)S \equiv 0$ we conclude $S \equiv 0$. □

Now $2021 = 43 \cdot 47$ with 43 and 47 distinct primes, so $2021 \mid S(j)$ if and only if both $43 \mid S(j)$ and $47 \mid S(j)$.

The set A is a complete set of representatives for $(\mathbb{Z}/2021\mathbb{Z})^\times$, and by the Chinese remainder theorem the reduction map

$$(\mathbb{Z}/2021\mathbb{Z})^\times \longrightarrow (\mathbb{Z}/43\mathbb{Z})^\times \times (\mathbb{Z}/47\mathbb{Z})^\times, \qquad n \longmapsto (n \bmod 43,\ n \bmod 47)$$

is a bijection. Consequently, for each fixed nonzero class a modulo 43, the number of $n \in A$ with $n \equiv a \pmod{43}$ equals the number of choices of the second coordinate, namely $\phi(47) = 46$. Therefore

$$S(j) \equiv 46 \sum_{a=1}^{42} a^j \pmod{43},$$

and symmetrically, each nonzero class modulo 47 is hit $\phi(43) = 42$ times, so

$$S(j) \equiv 42 \sum_{b=1}^{46} b^j \pmod{47}.$$

Since $46 \equiv 3 \pmod{43}$ and $42 \not\equiv 0 \pmod{47}$, both multipliers are invertible, so they may be cancelled. Applying the lemma with $p = 43$ (where $p - 1 = 42$) and with $p = 47$ (where $p - 1 = 46$):

$$43 \mid S(j) \iff 42 \nmid j, \qquad 47 \mid S(j) \iff 46 \nmid j.$$

Hence $2021 \mid S(j)$ if and only if $42 \nmid j$ and $46 \nmid j$, as claimed.

Two checks. First, $j = 0$: here $S(0) = |A| = \phi(2021) = 42 \cdot 46 = 1932$, which is not a multiple of 2021, and indeed $42 \mid 0$, so $j = 0$ is correctly excluded. Second, every odd j should qualify, since 42 and 46 are even; and indeed the involution $n \mapsto 2021 - n$ on A gives $S(j) \equiv -S(j) \pmod{2021}$ for odd j, so $S(j) \equiv 0$. The two arguments agree.

Layer 1: WHAT

Block 1: read A as the unit group $(\mathbb{Z}/2021\mathbb{Z})^\times$, not as a filtered interval.

Block 2: factor $2021 = 43 \cdot 47$ and replace the single divisibility by the pair of divisibilities.

Block 3: use the Chinese remainder theorem as a counting statement: reducing A modulo 43 covers each nonzero class exactly 46 times, so $S(j) \equiv 46 \sum_{a=1}^{42} a^j \pmod{43}$.

Block 4: check that the multiplicity 46 is invertible modulo 43, so it can be cancelled and the problem reduces to a power sum over $\mathbb{F}_{43}^\times$.

Block 5: prove the prime power sum lemma by reindexing with a primitive root, which gives $(g^j - 1)S \equiv 0$ with $g^j - 1$ invertible exactly when $(p - 1) \nmid j$.

Block 6: assemble the two conditions and verify against $j = 0$ and against the odd j pairing.

The nominal route treats the phrase "multiple of 2021" literally and stays at the composite modulus: pick a unit u, get $S(j) \equiv u^j S(j)$, and then stall, because $u^j - 1$ can be divisible by exactly one of 43 and 47 and there is nothing to cancel. The bifurcation point is the very first line after reading the modulus: the zombie writes "let u be a unit modulo 2021", the ideal route writes "$2021 = 43 \cdot 47$, so it suffices to work modulo each prime". Everything else is identical machinery; the only decision is where you run it. The reason the split is mandatory rather than convenient is that the multiplier argument needs a field, and $\mathbb{Z}/2021\mathbb{Z}$ is not one.

Layer 2: HOW

Block 1. **Cue.** "all integers n such that $1 \le n \le 2021$ and $\gcd(n, 2021) = 1$". **Why natural.** That is the verbatim definition of a reduced residue system: a complete set of representatives for the unit group modulo 2021. The size is $\phi(2021) = 1932$. Naming the object correctly is the whole first move, because sums over groups obey a rule that sums over intervals do not. **Link to the next block.** Since the index set is a group and $n \mapsto n^j$ respects multiplication, the sum is a character sum, and the only structural question about a character sum is whether the character is trivial.

Block 2. **Cue.** "is a multiple of 2021" with $2021 = 43 \cdot 47$. **Why natural.** An integer is divisible by a product of distinct primes if and only if it is divisible by each. That converts one hard statement into two statements over fields, and fields are where cancellation is legal. You should test any four digit modulus for factorisation on reflex; $2021 = 2025 - 4 = 45^2 - 2^2 = 43 \cdot 47$ is a difference of two squares, so the factorisation costs one line. **Link to the next block.** Working modulo 43 requires knowing how often each residue class modulo 43 appears among the elements of A.

Block 3. **Cue.** A is defined by a coprimality condition at both primes at once. **Why natural.** The Chinese remainder theorem says the reduction map from $(\mathbb{Z}/2021\mathbb{Z})^\times$ to the product of the two prime unit groups is a bijection. Bijection means the fibre over each class modulo 43 has size exactly $|(\mathbb{Z}/47\mathbb{Z})^\times| = 46$. This is the theorem used as a counting device, which is its most useful form on the Putnam, rather than as a machine for solving simultaneous congruences. **Link to the next block.** The sum therefore factorises as a constant multiplicity times a power sum over $\mathbb{F}_{43}^\times$, and the constant must be checked for invertibility before it is cancelled.

Block 4. **Cue.** The multiplier 46 modulo 43, and 42 modulo 47. **Why natural.** $46 \equiv 3$ and $42 \equiv 42$, both nonzero, hence invertible in the respective fields. This is not a formality: if the multiplicity were divisible by the prime, the congruence would hold for every j and the answer at that prime would be vacuous. Take $m = 21 = 3 \cdot 7$ and the same computation gives $S(j) \equiv 6 \sum_{a=1}^{2} a^j \equiv 0 \pmod 3$ for every j, so the condition at 3 disappears entirely. The examiners chose 43 and 47 so that neither multiplicity collapses. **Link to the next block.** With the multiplier cancelled, everything rests on the classical power sum over the nonzero residues of a prime.

Block 5. **Cue.** $\sum_{n=1}^{p-1} n^j$ with j free. **Why natural.** The index set is the cyclic group $\mathbb{F}_p^\times$, so multiplication by a primitive root g permutes it. Reindexing gives $S \equiv g^j S$, hence $(g^j - 1)S \equiv 0$. In a field one of the two factors must vanish, and $g^j - 1$ vanishes exactly when the order $p - 1$ divides j. The case $(p-1) \mid j$ is then Fermat's little theorem term by term, giving $p - 1 \equiv -1$. The two cases are the two ways a character can behave: trivial characters sum to the size of the group, nontrivial characters sum to zero. **Link to the next block.** Both prime conditions are now explicit, so the final answer is their conjunction.

Block 6. **Cue.** The answer is a pair of noncongruences, and $j = 0$ sits at the boundary. **Why natural.** $j = 0$ is divisible by everything, so the answer must exclude it, and the direct computation $S(0) = 1932$ confirms this. The odd j pairing gives an independent confirmation on an infinite family. Always spend the two lines: a boundary case and an independent family, checked, is what separates a graded ten from a graded eight. **Link to the next block.** Write the answer as "j not divisible by 42 and not divisible by 46" and stop.

Layer 3: WHY

The mathematical truth beneath the whole problem is orthogonality. Let G be a finite abelian group and let $\chi : G \to k^\times$ be a group homomorphism into the multiplicative group of a field. Then

$$\sum_{g \in G} \chi(g) = \begin{cases} |G| \cdot 1_k & \text{if } \chi \text{ is trivial,} \\ 0 & \text{otherwise,} \end{cases}$$

and the proof is the reindexing you have just used: pick h with $\chi(h) \neq 1$, then $\Sigma = \chi(h)\Sigma$ forces $\Sigma = 0$. Here $G = \mathbb{F}_p^\times$, $k = \mathbb{F}_p$ and $\chi(n) = n^j$. So the lemma is not an isolated fact about power sums; it is the single most used identity in all of finite abelian character theory, and it will reappear on the Putnam as roots of unity filters, as Gauss sums, and as the evaluation of $\sum_{k=0}^{n-1} \omega^{jk}$.

The discernment mechanism is a three part test that you should run whenever a sum is taken over a set defined by a coprimality or a congruence condition. First: is the index set a group? If yes, name it. Second: is the summand a homomorphism on that group? If yes, you are looking at a character sum and orthogonality settles it. Third: is the coefficient ring a field? If not, split the modulus until it is, because cancellation is the engine of the whole argument and cancellation needs the absence of zero divisors. The composite modulus 2021 fails the third test, and that is the only reason the problem is not a two line exercise.

There is a real generalisation worth carrying, and it comes with a caveat that this problem was designed to dodge. Let $m = p_1 p_2 \cdots p_k$ be squarefree with distinct odd primes, and let $S(j)$ be the sum of jth powers over the reduced residues modulo m. The same counting gives, for each i,

$$S(j) \equiv \left(\prod_{l \neq i} (p_l - 1)\right) \sum_{a=1}^{p_i - 1} a^j \pmod{p_i},$$

so by the lemma $S(j) \equiv 0 \pmod{p_i}$ unless both $(p_i - 1) \mid j$ and $p_i \nmid \prod_{l \neq i}(p_l - 1)$. In other words the clean answer "$S(j)$ is divisible by m exactly when no $p_i - 1$ divides j" holds only when each multiplicity $\phi(m)/(p_i - 1)$ is invertible modulo p_i. For $m = 2021$ that check passes at both primes. For $m = 21$ it fails at $p = 3$ and the true answer becomes "$21 \mid S(j)$ if and only if $6 \nmid j$". Train the habit by always writing the multiplicity down explicitly and reducing it modulo the prime before cancelling anything.

One more structural fact worth lighting up: as a function of j, $S(j)$ mod 2021 is periodic. Its period divides the exponent of the group $(\mathbb{Z}/2021\mathbb{Z})^\times \cong C_{42} \times C_{46}$, which is $\operatorname{lcm}(42, 46) = 966$, the Carmichael function $\lambda(2021)$. So the answer set is a union of arithmetic progressions, and the exceptional j (those where $2021 \nmid S(j)$) form the union of the multiples of 42 and the multiples of 46, a set of density $\frac{1}{42} + \frac{1}{46} - \frac{1}{966}$. Notice that the naive guess "exceptional exactly when $966 \mid j$", which is what you would predict if you wrongly assumed $(\mathbb{Z}/2021\mathbb{Z})^\times$ were cyclic of order $\phi(2021)$ or of order $\lambda(2021)$, is far too small a set. The failure of cyclicity is the content of the problem.

Revolving concepts

Reduced residue system and Euler's function. The integers in $[1, m]$ coprime to m form a complete set of representatives for $(\mathbb{Z}/m\mathbb{Z})^\times$, of size $\phi(m)$, and ϕ is multiplicative on coprime arguments with $\phi(p) = p - 1$. Here $\phi(2021) = 42 \cdot 46 = 1932$. See Book 1, field: Number theory, topic: Euler's totient and unit groups.

Fermat and Euler. $a^{p-1} \equiv 1 \pmod{p}$ for $p \nmid a$, and more generally $a^{\phi(m)} \equiv 1 \pmod{m}$ for $\gcd(a, m) = 1$. The sharp exponent is the Carmichael function $\lambda(m)$, the exponent of the unit group; $\lambda(2021) = \operatorname{lcm}(42, 46) = 966 < 1932$. See Book 1, field: Number theory, topic: orders and Fermat Euler.

Primitive roots. $(\mathbb{Z}/p\mathbb{Z})^\times$ is cyclic of order $p - 1$ for every prime p; a generator is a primitive root. Cyclicity holds for p^e and $2p^e$ and fails for products of two distinct odd primes, which is exactly the failure exploited here. See Book 1, field: Number theory, topic: primitive roots and cyclicity.

Chinese remainder theorem as an isomorphism. For coprime m_1, m_2, the reduction map $\mathbb{Z}/m_1 m_2\mathbb{Z} \to \mathbb{Z}/m_1\mathbb{Z} \times \mathbb{Z}/m_2\mathbb{Z}$ is a ring isomorphism, and it restricts to a group isomorphism on unit groups. Used as a counting device it says: each class modulo m_1 is realised by exactly $\phi(m_2)$ units modulo $m_1 m_2$. See Book 1, field: Number theory, topic: Chinese remainder theorem.

Character sums and orthogonality. For a finite abelian group G and a homomorphism $\chi : G \to k^\times$, $\sum_{g \in G} \chi(g)$ is $|G|$ if χ is trivial and 0 otherwise. The special case $\chi(n) = n^j$ on $\mathbb{F}_p^\times$ is the lemma above; the special case $G = \mathbb{Z}/n\mathbb{Z}$ and $\chi(k) = \omega^k$ with ω a primitive nth root of unity is the roots of unity filter $\frac{1}{n}\sum_{t=0}^{n-1} \omega^{-tr}\omega^{tk} = [k \equiv r]$. See Book 1, field: Abstract algebra, topic: characters of finite abelian groups.

Periodicity of $S(j)$ in j, as a recursion. Because $n^{j+\lambda(m)} \equiv n^j$ for every unit n, the sequence $j \mapsto S(j) \bmod m$ satisfies the exact recursion $S(j + \lambda(m)) \equiv S(j)$, that is, a linear recurrence of order $\lambda(m)$ with characteristic polynomial $x^{\lambda(m)} - 1$. Generalise: any sequence of the form $j \mapsto \sum_{n \in G} \chi_n(g)^j$, where each χ_n takes values in a finite group of exponent e, is periodic in j with period dividing e, and its values are determined by the multiset of orders appearing. This is the same phenomenon as the eventual periodicity of $j \mapsto a^j \bmod m$, of Fibonacci residues modulo m (the Pisano period), and of any linear recurrence over a finite ring: the state space is finite and the transition is invertible, so the orbit is purely periodic. See Book 1, field: Sequences and recurrences, topic: periodicity modulo m.

The involution pairing. If a summation set A is closed under $n \mapsto m - n$ then for odd j the sum of jth powers is $\equiv 0$ modulo m when m is odd. This is the additive twin of the multiplicative reindexing and is worth trying first on any symmetric index set. See Book 1, field: Number theory, topic: symmetry of residue systems.

Circuit Bank

- **WHEN** a divisibility question modulo a composite m **DO** factor m and replace the statement by the conjunction of the statements at each prime power **BECAUSE** cancellation, primitive roots and the field axioms are available at a prime and not at a composite, and $x \equiv 0 \pmod{m}$ is equivalent to the conjunction of $x \equiv 0$ at each prime power dividing m
- **WHEN** a sum whose index set is all of a finite group **DO** reindex by multiplying the index by a fixed group element **BECAUSE** the reindexing is a bijection so the sum is unchanged, which forces $(\chi(h) - 1)\Sigma = 0$ and kills Σ whenever $\chi(h) \neq 1$ is invertible
- **WHEN** a set defined by coprimality to a composite, reduced modulo one prime factor **DO** count the multiplicity by the Chinese remainder theorem and reduce that multiplicity modulo the prime before cancelling it **BECAUSE** the multiplicity is $\phi(m)/(p-1)$ and it can itself be divisible by p, in which case the congruence is vacuous and the condition at that prime disappears
- **WHEN** a sum of jth powers with j a free parameter **DO** ask only whether $x \mapsto x^j$ is the trivial character on the relevant group **BECAUSE** orthogonality makes every nontrivial character sum vanish, so the answer is always a statement about which j make the exponent map trivial, that is about $\mathrm{ord}(x) \mid j$
- **WHEN** an index set symmetric under $n \mapsto m - n$ with m odd **DO** test odd j separately in one line **BECAUSE** the pairing gives $S \equiv -S$, hence $S \equiv 0$, which both solves an infinite family for free and gives an independent check on your final answer

Backtesting

1. First useful cue: A is the reduced residue system modulo 2021, so the sum is over a group and $n \mapsto n^j$ is a homomorphism on it.
2. Method the expression suggested: reindex the sum by multiplication by a unit, which is the standard proof of character orthogonality; but this only bites over a field, so factor $2021 = 43 \cdot 47$ first and run it at each prime.
3. Where you would hesitate, and the thought to reinforce: at the composite modulus the argument stalls because $u^j - 1$ may be a zero divisor. The thought to install is that $(\mathbb{Z}/m\mathbb{Z})^\times$ is cyclic only for m equal to $1, 2, 4, p^e, 2p^e$, so for $m = 43 \cdot 47$ there is no primitive root and you must split before you cancel.
4. The trap the question sets: answering "all j not divisible by 966", which is what you get if you assume a primitive root exists modulo 2021. The exceptional set is much larger than that.
5. What makes you recognise it faster next time: any four digit modulus in a Putnam number theory problem should be factored on sight ($2021 = 45^2 - 2^2$), and any sum over a coprimality condition should be renamed as a sum over a unit group in the first line.
6. The reusable decision rule: split the modulus into primes, count multiplicities by the Chinese remainder theorem, check the multiplicity is invertible, then apply orthogonality at each prime and take the conjunction.

Mirror drills

1. Evaluate $\sum_{n=1}^{p-1} n^{p-1} \bmod p$ and $\sum_{n=1}^{p-1} n \bmod p$ for an odd prime p. *Answer:* the first is $p-1 \equiv -1$ by Fermat; the second is $\frac{p(p-1)}{2} \equiv 0$, consistent with the lemma since $(p-1) \nmid 1$ for $p > 2$.
2. Determine all $j \geq 0$ for which 15 divides $\sum_n n^j$, the sum being over $n \in [1, 15]$ coprime to 15. *Answer:* modulo 3 the multiplicity is $\phi(5) = 4 \equiv 1$, so $3 \mid S(j)$ iff $2 \nmid j$; modulo 5 the multiplicity is $\phi(3) = 2$, invertible, so $5 \mid S(j)$ iff $4 \nmid j$. Since j odd already forces $4 \nmid j$, the answer is exactly the odd j. Check $j = 1$: the sum is 60.
3. Show in one line that $2021 \mid S(j)$ for every odd j, without using the lemma. *Answer:* $n \in A \iff 2021 - n \in A$, and $(2021 - n)^j \equiv -n^j \pmod{2021}$ for odd j, so $S(j) \equiv -S(j)$ and 2021 is odd.

5.6 A6. Binary digit polynomials and composite values

Let $P(x)$ be a polynomial whose coefficients are all either 0 or 1. Suppose that $P(x)$ can be written as a product of two nonconstant polynomials with integer coefficients. Does it follow that $P(2)$ is a composite integer?

Read it out loud

Say what the two hypotheses actually are. The coefficients of P lie in $\{0,1\}$, so $P(2)$ is precisely the integer whose binary digits are those coefficients; the problem is a statement about binary expansions dressed as a statement about polynomials. And P factors over $\mathbb{Z}$ as QR with both factors nonconstant, so $P(2) = Q(2)R(2)$ is a factorisation of that integer into two integers. The question is whether that integer factorisation is nontrivial.

That reduction is the whole problem, and you should reach it in one line. $P(2) = Q(2)R(2)$ is automatic. What is not automatic is that neither factor is ± 1. So the target is: for every nonconstant integer factor Q of P, show $|Q(2)| \geq 2$. Note that $Q(2) = 0$ is impossible, since P is monic with nonnegative coefficients and hence $P(2) \geq 2^{\deg P} > 0$.

Now the obstruction. You cannot argue that Q inherits nonnegative coefficients from P, because it does not: $x^4 + x^3 + x + 1 = (x+1)^2(x^2 - x + 1)$ has the factor $x^2 - x + 1$ with a negative middle coefficient. So the constraint on the coefficients of P has to be transmitted to Q through something invariant under factorisation, and there is exactly one such thing: the roots. Every root of Q is a root of P, and $Q(2) = \prod_\alpha (2 - \alpha)$ over the roots of Q counted with multiplicity, once Q has been normalised to be monic. So the question becomes: how far can the complex roots of a polynomial with 0 and 1 coefficients travel towards the point 2?

The fields wired in are therefore root location for polynomials with bounded coefficients (a geometric series estimate applied to the reversed polynomial), factorisation of real polynomials into real linear and quadratic factors, the easy half of stability theory (a monic real polynomial with all roots in the open left half plane has all coefficients positive), and the intermediate value theorem for an integrality argument at the end. Within thirty seconds you should have written: "answer is yes; suffices to prove $|Q(2)| \geq 2$; $Q(2) = \prod(2 - \alpha)$; so bound the roots of P". Everything after that is the execution of a plan you already have.

Approaches, ranked worst to best

1. **Search for a counterexample.** *Cue:* a yes or no question. *Circuit:* test small cases before committing. As reconnaissance this is fine and it is quick: run over all 0 and 1 polynomials of degree at most twelve, factor each, and check that no reducible one has prime value at 2. It buys you the correct answer and nothing else, and a Putnam grader awards nothing for it. Ranks last as a solution, first as a five minute sanity check.

2. **Force the factors to have nonnegative coefficients.** *Cue: P* has nonnegative coefficients. *Circuit:* positivity is inherited by factors. It ranks here because it is false and the counterexample is small: $x^4 + x^3 + x + 1 = (x+1)^2(x^2 - x + 1)$. Knowing the counterexample immediately is worth more than the attempt, because it tells you the constraint must travel through the roots and not through the coefficients.

3. **Work modulo small primes or with the binary digits directly.** *Cue:* $P(2)$ is the integer with those binary digits. *Circuit:* congruence conditions on digit sums. There is no congruence obstruction available: $P(2)$ can be any integer with the right digits, and reducibility of P is not detectable modulo any fixed integer. A digit based argument cannot see the factorisation of P over $\mathbb{Z}$, so this route has nowhere to go.

4. **Show no root of P lies in the closed disc $|z-2| \le 1$.** *Cue:* $|Q(2)| = \prod |2-\alpha|$, so $|Q(2)| > 1$ would follow. *Circuit:* make every factor of a product exceed one. This is the right instinct and the right shape of conclusion, and it deserves respect. The problem is the bookkeeping: the bounds you can actually prove are $|\alpha| < \frac{1+\sqrt{5}}{2}$ when $\mathrm{Re}(\alpha) \ge 0$, and $|\alpha| < 1.466$ when $|\arg\alpha| \le \frac{\pi}{4}$, and neither excludes points such as $1.4 + 0.3i$, which lies inside that disc. So the route stalls short of the finish for want of a sharper estimate, and you should abandon it rather than push.

5. **Locate the roots, then shift to the midpoint of 1 and 2.** *Cue:* you want to compare $Q(2)$ with $Q(1)$, and $Q(1)$ is an integer you can bound below. *Circuit:* translate so that a comparison becomes a parity statement about coefficients. The complete route. Prove every root of P has real part less than $\frac{3}{2}$, shift by $\frac{3}{2}$ so that all roots move into the open left half plane, deduce that the shifted factor has strictly positive coefficients, and read off $Q(2) > Q(1)$ because 2 and 1 are symmetric about $\frac{3}{2}$. Then $Q(1) \ge 1$ by the intermediate value theorem, so $Q(2) \ge 2$. Every constant in this argument is forced by the previous line, none is guessed.

The ideal solution

The answer is $\boxed{\text{yes}}$: $P(2)$ must be composite. Recall that 1 is neither prime nor composite, so ruling out both $P(2) = 1$ and $P(2)$ prime is the content.

Write $P(x) = a_0 + a_1 x + \cdots + a_n x^n$ with every $a_i \in \{0, 1\}$ and $a_n = 1$, so P is monic. Since P is a product of two nonconstant polynomials, $n \ge 2$.

Lemma 5.6.1. *$P(x) > 0$ for every real $x > 0$. In particular P has no positive real root.*

Proof. All terms are nonnegative and the leading term x^n is positive. □

Lemma 5.6.2. *If α is a root of P with $\mathrm{Re}(\alpha) \ge 0$, then $|\alpha| < \varphi := \frac{1+\sqrt{5}}{2}$.*

Proof. If $|\alpha| \le 1$ there is nothing to prove, so assume $|\alpha| > 1$ and put $t = |\alpha|^{-1} \in (0, 1)$. Dividing $P(\alpha) = 0$ by α^n and separating the first two terms,

$$1 + a_{n-1}\alpha^{-1} = -\sum_{k=2}^{n} a_{n-k}\alpha^{-k}.$$

Since $\alpha^{-1} = \overline{\alpha}/|\alpha|^2$, the hypothesis $\mathrm{Re}(\alpha) \ge 0$ gives $\mathrm{Re}(\alpha^{-1}) \ge 0$, so the left side has real part at least 1 and therefore modulus at least 1. The right side has modulus at most

$$\sum_{k=2}^{n} t^k < \sum_{k=2}^{\infty} t^k = \frac{t^2}{1-t}.$$

Hence $1 - t < t^2$, that is $t^2 + t - 1 > 0$, that is $t > \frac{\sqrt{5}-1}{2} = \varphi^{-1}$, and so $|\alpha| = t^{-1} < \varphi$. □

Lemma 5.6.3. *Every root α of P satisfies $\mathrm{Re}(\alpha) < \frac{3}{2}$.*

Proof. If $\mathrm{Re}(\alpha) \le 0$ there is nothing to prove, so assume $\mathrm{Re}(\alpha) > 0$; in particular $\alpha \ne 0$.

Suppose first that $|\arg\alpha| > \frac{\pi}{4}$. Then $\mathrm{Re}(\alpha) = |\alpha|\cos(\arg\alpha) < \frac{|\alpha|}{\sqrt{2}} < \frac{\varphi}{\sqrt{2}} = \frac{1+\sqrt{5}}{2\sqrt{2}} < 1.145 < \frac{3}{2}$, using the previous lemma.

Suppose instead that $|\arg\alpha| \le \frac{\pi}{4}$. Then $\arg(\alpha^{-1}) \in [-\frac{\pi}{4}, \frac{\pi}{4}]$ and $\arg(\alpha^{-2}) \in [-\frac{\pi}{2}, \frac{\pi}{2}]$, so both $\mathrm{Re}(\alpha^{-1}) \ge 0$ and $\mathrm{Re}(\alpha^{-2}) \ge 0$. Dividing $P(\alpha) = 0$ by α^n and separating the first three terms,

$$1 + a_{n-1}\alpha^{-1} + a_{n-2}\alpha^{-2} = -\sum_{k=3}^{n} a_{n-k}\alpha^{-k},$$

where an empty sum is 0. The left side has real part at least 1, hence modulus at least 1. If $|\alpha| \le 1$ then $\mathrm{Re}(\alpha) \le 1 < \frac{3}{2}$ and we are done; otherwise put $t = |\alpha|^{-1} \in (0,1)$ and bound the right side by

$$\sum_{k=3}^{n} t^k < \frac{t^3}{1-t},$$

so $t^3 + t - 1 > 0$. The function $x \mapsto x^3 + x - 1$ is strictly increasing on $[0,\infty)$ and has a unique positive root τ; since $(\frac{2}{3})^3 + \frac{2}{3} - 1 = -\frac{1}{27} < 0$ we get $\tau > \frac{2}{3}$. Therefore $t > \tau > \frac{2}{3}$ and $|\alpha| < \frac{3}{2}$, so $\mathrm{Re}(\alpha) \le |\alpha| < \frac{3}{2}$. □

Lemma 5.6.4. *Let $F \in \mathbb{R}[x]$ be monic of degree $d \ge 1$ with every root of negative real part. Then every coefficient of F is strictly positive.*

Proof. Over $\mathbb{R}$, F factors into monic linear factors $x - \beta$ with β a real root, and monic quadratic factors $(x-\alpha)(x-\overline{\alpha}) = x^2 - 2\mathrm{Re}(\alpha)x + |\alpha|^2$ coming from conjugate pairs of nonreal roots. In the first case $-\beta > 0$, so $x - \beta = x + (-\beta)$ has positive coefficients; in the second case $-2\mathrm{Re}(\alpha) > 0$ and $|\alpha|^2 > 0$, so again all coefficients are positive. Finally, if $G = \sum_{i=0}^{p} g_i x^i$ and $H = \sum_{j=0}^{q} h_j x^j$ have all coefficients positive, then for each $0 \le m \le p+q$ the coefficient of x^m in GH is $\sum g_i h_j$ over the pairs with $i + j = m$, $0 \le i \le p$, $0 \le j \le q$; this index set is nonempty (take i between $\max(0, m-q)$ and $\min(p, m)$) and every term is positive. Induction over the factors completes the proof. □

Proof of the statement. Let $P = QR$ with $Q, R \in \mathbb{Z}[x]$ nonconstant. Since P is monic, the leading coefficients of Q and R are integers whose product is 1, so both equal 1 or both equal -1; replacing the pair (Q, R) by $(-Q, -R)$ if necessary, we may assume both are monic. This changes neither P nor the product $Q(2)R(2)$.

Every root of Q is a root of P, so by the third lemma every root of Q has real part less than $\frac{3}{2}$. Hence $\widetilde{Q}(x) := Q\left(x + \frac{3}{2}\right)$ is a monic real polynomial of degree $d = \deg Q \ge 1$ all of whose roots have negative real part, so by the fourth lemma $\widetilde{Q}(x) = \sum_{i=0}^{d} c_i x^i$ with every $c_i > 0$. Since 2 and 1 are symmetric about $\frac{3}{2}$,

$$Q(2) - Q(1) = \widetilde{Q}\left(\tfrac{1}{2}\right) - \widetilde{Q}\left(-\tfrac{1}{2}\right) = \sum_{i=0}^{d} c_i \left(2^{-i} - (-2)^{-i}\right) = 2\sum_{i \text{ odd}} c_i 2^{-i} > 0,$$

the last inequality because $d \ge 1$ guarantees that the index $i = 1$ occurs and $c_1 > 0$.

Next, $Q(1) \ge 1$. Indeed $Q(1)$ is an integer, and if $Q(1) \le 0$ then, Q being monic, $Q(x) \to +\infty$ as $x \to +\infty$, so by the intermediate value theorem Q would have a real root in $[1,\infty)$; that root would be a positive real root of P, contradicting the first lemma.

Therefore $Q(2) > Q(1) \ge 1$, and $Q(2)$ is an integer, so $Q(2) \ge 2$. The identical argument gives $R(2) \ge 2$. Hence $P(2) = Q(2)R(2)$ is a product of two integers each at least 2, so $P(2)$ is composite. ■

Layer 1: WHAT

Block 1: $P(2) = Q(2)R(2)$, so the task is exactly to prove $|Q(2)| \ge 2$ for every nonconstant integer factor Q; the value 0 is impossible since $P(2) \ge 2^n$.

Block 2: the coefficient hypothesis is not inherited by Q (counterexample $x^2 - x + 1$), so it must be transmitted through the roots, since $Q(2) = \prod(2 - \alpha)$ once Q is monic.

Block 3: to get a lower bound on $Q(2)$, compare it with $Q(1)$, because $Q(1)$ is an integer that is at least 1 by the intermediate value theorem plus the absence of positive real roots of P.

Block 4: a comparison of F at two points symmetric about c is a statement about the parity of the coefficients of $F(x + c)$; here $c = \frac{3}{2}$ is forced as the midpoint of 1 and 2.

Block 5: those shifted coefficients are all positive exactly when all roots of Q lie to the left of $\frac{3}{2}$, which is the easy half of stability theory.

Block 6: prove that every root of P has real part below $\frac{3}{2}$ by dividing $P(\alpha) = 0$ by α^n, using that α^{-1} and α^{-2} have nonnegative real part in the sector $|\arg \alpha| \leq \frac{\pi}{4}$, and bounding the tail by a geometric series.

Block 7: assemble, $Q(2) > Q(1) \geq 1$ forces $Q(2) \geq 2$, same for R, done.

The nominal route runs Block 1 and then tries to prove something about the coefficients of Q, or hunts for a congruence obstruction on the binary digits. Both are dead: the coefficients of a factor are not constrained in sign, and reducibility over $\mathbb{Z}$ is invisible modulo any fixed integer. The bifurcation point is the sentence immediately after Block 1: the zombie asks "what do I know about the coefficients of Q", the ideal route asks "what do I know about the roots of Q". That single reframing is the entire problem, because roots are the only data that a factor is forced to share with P.

Layer 2: HOW

Block 1. **Cue.** "product of two nonconstant polynomials with integer coefficients" and the target "$P(2)$ composite". **Why natural.** Substituting $x = 2$ into the factorisation is free and it converts a polynomial statement into an integer statement. Composite means a product of two integers each of absolute value at least 2, so the only failure mode is a factor equal to ± 1 at $x = 2$. Naming the single failure mode is what turns a vague question into a target. **Link to the next block.** You now need a lower bound on $|Q(2)|$ that uses only the hypothesis on P, so ask what Q is forced to inherit from P.

Block 2. **Cue.** P has nonnegative coefficients, but its factors need not. **Why natural.** Test it: $x^4 + x^3 + x + 1 = (x+1)^2(x^2 - x + 1)$. Producing this in ten seconds saves you ten minutes of trying to prove a false lemma. What a factor does inherit is its root set: every root of Q is a root of P. And since P is monic, Q can be taken monic, whereupon $Q(2) = \prod_\alpha (2 - \alpha)$ exactly. **Link to the next block.** So the coefficient hypothesis must be converted into a region of the complex plane containing all roots of P, and the region must be positioned relative to the point 2.

Block 3. **Cue.** $Q(2)$ and $Q(1)$ are both integers, and $Q(1)$ has an easy lower bound. **Why natural.** P has no positive real root because all its coefficients are nonnegative and the leading one is 1. A monic Q with $Q(1) \leq 0$ would have a real root in $[1, \infty)$ by the intermediate value theorem, which P forbids. So $Q(1) \geq 1$ costs three lines. That is your anchor: if you can prove the single inequality $Q(2) > Q(1)$, integrality upgrades it to $Q(2) \geq 2$ and the problem is finished. **Link to the next block.** Comparing a polynomial at two points is easiest when they are symmetric about a centre, so shift to the midpoint.

Block 4. **Cue.** The two points are 1 and 2, with midpoint $\frac{3}{2}$. **Why natural.** Write $\widetilde{Q}(x) = Q(x + \frac{3}{2})$. Then $Q(2) = \widetilde{Q}(\frac{1}{2})$ and $Q(1) = \widetilde{Q}(-\frac{1}{2})$, so the difference is twice the sum of the odd degree terms of $\widetilde{Q}$ evaluated at $\frac{1}{2}$. If all coefficients of $\widetilde{Q}$ are positive, the difference is positive. This is why the constant $\frac{3}{2}$ appears: not because of any root bound, but because it is the midpoint of the two values you decided to compare. The root bound is then reverse engineered to fit it. **Link to the next block.** Positivity of the coefficients of a shifted polynomial is a statement about the location of its roots.

Block 5. **Cue.** Monic real polynomial, want all coefficients positive. **Why natural.** Factor over $\mathbb{R}$ into linear and quadratic pieces. A linear piece $x - \beta$ has positive coefficients exactly when $\beta < 0$; a quadratic piece $x^2 - 2\mathrm{Re}(\alpha)x + |\alpha|^2$ has positive coefficients exactly when $\mathrm{Re}(\alpha) < 0$. Products of polynomials with all coefficients positive have all coefficients positive. So "all roots in the open left half plane" implies "all coefficients positive", and this is the direction that is elementary. The converse is false and you must not use it. **Link to the next block.** Translating back, you need every root of P to satisfy $\mathrm{Re}(\alpha) < \frac{3}{2}$.

Block 6. Cue. Coefficients bounded by the leading coefficient, and a bound wanted on the roots. **Why natural.** The standard device is to divide $P(\alpha) = 0$ by α^n, which turns the polynomial into $1 + a_{n-1}\alpha^{-1} + a_{n-2}\alpha^{-2} + \cdots$ and makes the small quantity $|\alpha|^{-1}$ the variable. Keeping the first term alone gives the crude bound $|\alpha| < 2$; keeping two terms and using $\mathrm{Re}(\alpha^{-1}) \geq 0$ in the right half plane gives $|\alpha| < \varphi$; keeping three terms and using $\mathrm{Re}(\alpha^{-2}) \geq 0$ in the sector $|\arg \alpha| \leq \frac{\pi}{4}$ gives $|\alpha| < \frac{3}{2}$ there. Outside that sector, $\mathrm{Re}(\alpha) < \frac{|\alpha|}{\sqrt{2}} < \frac{\varphi}{\sqrt{2}} < \frac{3}{2}$ anyway. The number of terms you keep is dictated by the strength you need, and each extra term costs one sector condition. **Link to the next block.** All roots have real part below $\frac{3}{2}$, so the shifted factor has positive coefficients and the chain closes.

Block 7. Cue. $Q(2) > Q(1) \geq 1$ with $Q(2) \in \mathbb{Z}$. **Why natural.** A strict inequality between integers is an inequality with a gap of at least one, which is where the final ≥ 2 comes from. Applying the same to R gives two factors of size at least 2, which is the definition of composite. **Link to the next block.** Write the conclusion, and state explicitly that 1 is neither prime nor composite so the case $P(2) = 1$ is also excluded.

Layer 3: WHY

The mathematical truth under this problem is a theorem you should know by name. Cohn's irreducibility criterion says that if a prime p is written in base ten as $\sum a_i 10^i$, then $\sum a_i x^i$ is irreducible over $\mathbb{Z}$. Brillhart, Filaseta and Odlyzko proved in 1981 that the same holds in every base $b \geq 2$. Our problem is the contrapositive of that theorem in base 2, restricted to digits 0 and 1, with the extra observation that $P(2) = 1$ is also impossible. The mechanism in all these theorems is identical: bounded digits force all roots of the digit polynomial into a bounded region of the complex plane sitting strictly to the left of the base point b, and then no factor can be small at b.

Why does bounding the digits bound the roots? Because a root satisfies $|\alpha|^n = |\sum_{i<n} a_i \alpha^i| \leq \sum_{i<n} |\alpha|^i$, and once the coefficients are bounded the right side is a geometric series that cannot keep up with $|\alpha|^n$ unless $|\alpha|$ is small. That is the Cauchy bound and it gives $|\alpha| < 2$ in one line. The refinements used above are the same computation with more terms retained on the left, and each retained term buys a stronger bound at the price of a sector hypothesis that guarantees the retained terms do not cancel. This is a general and reusable pattern: to strengthen a modulus estimate, move terms from the estimated side to the exact side, and pay for it with an argument that the exact side cannot be small, usually a real part argument.

Now the discernment mechanism, which is what you must actually train. There are two decisions in this problem and both are forced by the question rather than invented.

First decision: coefficients or roots. The hypothesis is about coefficients of P; the conclusion is about a value of P at an integer, mediated by an unknown factor Q. Ask what survives factorisation. Coefficient positivity does not (one counterexample settles it). Root membership does, trivially and exactly. So the hypothesis must be pushed into root language, used there, and pulled back to a value. Whenever a problem gives you information about one polynomial and asks about its unknown factors, the transmission channel is the root set, every time.

Second decision: the choice of the comparison point and hence of the constant $\frac{3}{2}$. You want a lower bound on the integer $Q(2)$. The only cheap lower bound available anywhere in the problem is $Q(1) \geq 1$, which follows from P having no positive real root. So the natural target is $Q(2) > Q(1)$, and the natural technique for comparing a polynomial at two points is to centre at their midpoint, where the comparison becomes "the odd part is positive". The midpoint of 1 and 2 is $\frac{3}{2}$, and only now do you learn what you must prove about the roots. This is the opposite of the usual presentation, where $\frac{3}{2}$ appears first and looks like a rabbit from a hat. It is not a rabbit; it is $\frac{1+2}{2}$.

A concrete sanity example, worth running once by hand. Take $P(x) = x^4 + x^3 + x + 1 = (x+1)^2(x^2 - x + 1)$, so $P(2) = 27$. For $Q = x^2 - x + 1$, the shift gives $\widetilde{Q}(x) = (x + \frac{3}{2})^2 - (x + \frac{3}{2}) + 1 =$

$x^2 + 2x + \frac{7}{4}$, all coefficients positive as promised, and $\widetilde{Q}(\frac{1}{2}) = 3 = Q(2)$ while $\widetilde{Q}(-\frac{1}{2}) = 1 = Q(1)$. The inequality $3 > 1$ is exactly the mechanism, and $27 = 3 \times 9$ is the resulting factorisation of the binary number 11011_2.

Finally, the knowledge gaps, stated without disguise, because this is an A6 and pretending otherwise would be dishonest. To approach this problem at all you must have installed the following, none of which is in A-level Further Mathematics: (i) the fact that a real polynomial factors over $\mathbb{R}$ into linear and quadratic factors, with nonreal roots in conjugate pairs, which is the fundamental theorem of algebra plus a conjugation argument; (ii) the technique of bounding roots via the reversed polynomial and a geometric series, that is the Cauchy bound family; (iii) the easy direction of the Routh and Hurwitz stability circle of ideas, namely that all roots in the open left half plane forces all coefficients positive (and the knowledge that the converse fails, as $x^4 + x^3 + x^2 + x + 1$ shows, whose roots are primitive fifth roots of unity, two of which have positive real part); (iv) the habit of normalising a factorisation of a monic polynomial so that both factors are monic. If any of those four is missing you will not find this route, and no amount of cleverness substitutes for them. This is a knowledge gap, and knowledge gaps are named, not disguised.

Revolving concepts

Value at an integer as a base b expansion. If P has coefficients in $\{0, 1, \ldots, b-1\}$ then $P(b)$ is the integer with those base b digits, and conversely every positive integer determines a unique such P. Factorisations of P over $\mathbb{Z}$ therefore give factorisations of the integer, which is the bridge this problem is built on. See Book 1, field: Number theory, topic: base representations.

Cauchy bound and the reversed polynomial. For monic $P = x^n + \sum_{i<n} a_i x^i$, every root satisfies $|\alpha| \leq 1 + \max_i |a_i|$; for coefficients in $\{0, 1\}$ this gives $|\alpha| < 2$. The proof divides by α^n and sums a geometric series, which is the same as applying the trivial bound to the reversed polynomial $x^n P(1/x)$. Sharper versions retain more leading terms; the versions used above give $|\alpha| < \varphi$ in the closed right half plane and $|\alpha| < \frac{3}{2}$ in the sector $|\arg \alpha| \leq \frac{\pi}{4}$. See Book 1, field: Polynomials, topic: root location bounds.

Roots of unity and the golden ratio thresholds. The constants φ and τ (the real root of $x^3 + x - 1$) arise as the positive roots of $t^2 + t - 1$ and $t^3 + t - 1$, that is as the reciprocals of the thresholds at which a geometric tail $\sum_{k \geq m} t^k$ reaches 1. The general statement: $\sum_{k \geq m} t^k \geq 1$ if and only if $t^m + t \geq 1$, so retaining m leading terms produces the threshold polynomial $t^m + t - 1$, whose positive root increases to 1 as m grows. That is the general recursion behind the family of bounds, and it tells you in advance how much you can gain by retaining another term. See Book 1, field: Inequalities, topic: geometric series thresholds.

Real factorisation and conjugate pairs. A real polynomial factors over $\mathbb{R}$ as a product of monic linear factors and monic quadratics $x^2 - 2\mathrm{Re}(\alpha)x + |\alpha|^2$, one for each conjugate pair of nonreal roots. See Book 1, field: Polynomials, topic: factorisation over the reals.

Positivity from left half plane roots. If every root of a monic real F has negative real part then every coefficient of F is positive. The converse is false. This is the elementary half of the Routh and Hurwitz criterion, and on the Putnam it appears whenever a shift is used to convert a root location statement into a coefficient sign statement. See Book 1, field: Polynomials, topic: stability and coefficient signs.

Shifting to compare two values. For any polynomial F and any $a < b$, setting $c = \frac{a+b}{2}$ and $h = \frac{b-a}{2}$ gives $F(b) - F(a) = 2\sum_{i \text{ odd}} c_i h^i$ where $F(x + c) = \sum c_i x^i$. So sign information about the coefficients of the recentred polynomial is exactly sign information about $F(b) - F(a)$. See Book 1, field: Polynomials, topic: Taylor shifts.

Integrality upgrades strict inequalities. If $u > v$ with $u, v \in \mathbb{Z}$ then $u \geq v + 1$. Combined with a lower bound on v this is how a strict analytic inequality becomes an arithmetic conclusion, and it is the last step of this proof. See Book 1, field: Number theory, topic: integrality arguments.

Cohn, Brillhart, Filaseta and Odlyzko. If a prime is written in base $b \geq 2$ as $\sum a_i b^i$, then $\sum a_i x^i$ is

irreducible over $\mathbb{Z}$. The final step of the argument above, the shift by the midpoint, is due to Polya and Szego. See Book 1, field: Polynomials, topic: irreducibility criteria.

Circuit Bank

- **WHEN** a hypothesis about the coefficients of P and a question about an unknown factor Q **DO** translate the hypothesis into a region containing all roots of P, and use that every root of Q is a root of P **BECAUSE** sign and size conditions on coefficients are not inherited by factors, but the root set is inherited exactly, so the root set is the only reliable transmission channel
- **WHEN** you need a lower bound on $F(b)$ for a monic integer polynomial **DO** find a nearby point a where a lower bound is cheap, then prove $F(b) > F(a)$ by shifting to the midpoint of a and b **BECAUSE** after recentring, $F(b) - F(a)$ is twice the odd part of the shifted polynomial, so a coefficient sign statement settles the comparison, and integrality then buys a whole unit
- **WHEN** you want all coefficients of a real monic polynomial to be positive **DO** prove every root has negative real part and factor into linear and quadratic pieces **BECAUSE** each real factor $x + c$ with $c > 0$ and each quadratic $x^2 - 2\mathrm{Re}(\alpha)x + |\alpha|^2$ with $\mathrm{Re}(\alpha) < 0$ has positive coefficients, and products of such polynomials do too; the converse direction is false so never use it
- **WHEN** you need a bound on the roots of a polynomial with bounded coefficients **DO** divide by α^n and estimate the tail as a geometric series, retaining as many leading terms as the bound requires **BECAUSE** $|\alpha|^{-1}$ is the natural small variable, and retaining m terms yields the threshold $t^m + t - 1 > 0$, but only in a sector where the retained terms have nonnegative real part
- **WHEN** a polynomial has all coefficients nonnegative with a positive leading coefficient **DO** state at once that it has no positive real root **BECAUSE** every term is nonnegative and the leading term is strictly positive for $x > 0$, which is the cheapest root exclusion available and is often the anchor of an intermediate value argument

Backtesting

1. First useful cue: $P(2) = Q(2)R(2)$ is automatic, so the entire content is that neither factor equals ± 1 at $x = 2$.
2. Method the expression suggested: coefficients in $\{0, 1\}$ is a boundedness condition, and boundedness of coefficients is a root location condition via division by α^n; so bound the roots and then convert the bound into a statement about the value of a factor.
3. Where you would hesitate, and the thought to reinforce: the constant $\frac{3}{2}$ looks unmotivated. Install the reason: you chose to compare $Q(2)$ with $Q(1)$ because $Q(1) \geq 1$ is cheap, and the midpoint of 1 and 2 is $\frac{3}{2}$. The root bound is then designed to make the shift work, not the other way round.
4. The trap the question sets: assuming the factors inherit nonnegative coefficients. The counterexample $x^4 + x^3 + x + 1 = (x + 1)^2(x^2 - x + 1)$ should be in your hand within a minute.
5. What makes you recognise it faster next time: the phrase "coefficients all 0 or 1" together with an evaluation at an integer is a base representation problem, and base representation plus irreducibility is Cohn and Brillhart, Filaseta and Odlyzko. Recognising the family tells you the proof will be root location plus a shift.
6. The reusable decision rule: push coefficient hypotheses into root regions, compare values at two points by shifting to their midpoint, and finish with integrality.

Mirror drills

1. Show that every root of a monic P with coefficients in $\{0,1\}$ satisfies $|\alpha| < 2$, and explain in one sentence why this bound alone does not solve the problem. *Answer:* if $|\alpha| > 1$ then $|\alpha|^n \le \sum_{k<n} |\alpha|^k < \frac{|\alpha|^n}{|\alpha|-1}$, so $|\alpha| < 2$; it fails because it permits roots with real part close to 2, so the shift to $\frac{3}{2}$ would not send them into the left half plane.
2. Give a monic real polynomial with all coefficients positive that has a root with positive real part, and explain what it shows. *Answer:* $x^4 + x^3 + x^2 + x + 1$, whose roots are the primitive fifth roots of unity, two of which have real part $\cos 72^\circ > 0$; it shows the converse of the positivity lemma is false, so positivity of coefficients can never be used to deduce root location.
3. Deduce Cohn's criterion in base 2 from the result above: if P has coefficients in $\{0,1\}$ and $P(2)$ is prime, then P is irreducible over $\mathbb{Z}$. *Answer:* contrapositive, since a factorisation into nonconstant integer polynomials would make $P(2)$ composite.

5.7 B1. Random unit square avoiding lattice corners

Suppose that the plane is tiled with an infinite checkerboard of unit squares. If another unit square is dropped on the plane at random with position and orientation independent of the checkerboard tiling, what is the probability that it does not cover any of the corners of the squares of the checkerboard?

Read it out loud

Read the objects, not the story. The corners of the checkerboard are the integer lattice $\mathbb{Z}^2$. The dropped square is a rigid motion of $[0,1]^2$: a centre c and an angle θ. The phrase "position and orientation independent of the tiling" is the examiners handing you the probability measure: c uniform modulo the lattice, θ uniform, the two independent. So this is a geometric probability question of the Buffon family, wearing a tiling costume.

The fields wired in are integral geometry (a random rigid motion of a shape against a lattice), the symmetries of the two objects (the lattice under integer translations, the square under a quarter turn, which cut the sample space down to one cell times $[0, \pi/2)$), linearity of expectation, and one trigonometric integral at the end. What is really being tested is whether you compute a probability by drawing the forbidden region (zombie) or by counting.

The first useful cue: the dropped shape has area 1 and the lattice has one point per unit of area. Within thirty seconds you should have written $\mathbb{E}[N] = 1$, where N is the number of corners covered, for every fixed θ. That single line is the whole solution, because N can only be 0, 1 or 2, so $\mathbb{P}(N = 0) = \mathbb{P}(N = 2)$, and covering two adjacent corners is the intersection of two squares whose area is one line of trigonometry.

Approaches, ranked worst to best

1. **Forbidden-region bash.** *Cue:* "does not cover" read as "area of the safe region". *Circuit:* compute the complement of the union of four rotated unit squares centred at the corners of the cell. (Why worst: as soon as $\theta \neq 0$ the four squares overlap pairwise, since their width along each axis is $\sin\theta + \cos\theta > 1$; the union is a polygon whose vertices you would find by intersecting lines, a page of ink and three chances for a skipped-calculation error.)

2. **Circumscribed square.** *Cue:* "rotated square" suggests describing membership in the frame of the dropped square. *Circuit:* a square is $|s|, |t| \leq \frac{1}{2}$ in its own coordinates, so the safe centres form a smaller square inside the square circumscribed about the cell. (Why second: it is the official route and it is correct; for fixed θ the safe region is a square of side $\sin\theta + \cos\theta - 1$. But the converse direction, that every centre within $\frac{1}{2}$ of a side of the circumscribed square covers some corner, needs a case argument that must be self-generated in the arena.)

3. **Count, then compare.** *Cue:* area 1 dropped on a lattice of covolume 1. *Circuit:* the expected number of lattice points in a random translate of a set equals its area, and a variable in $\{0,1,2\}$ with known mean satisfies $\mathbb{P}(0) = \mathbb{P}(2) + 1 - \mathbb{E}$. (Why best: no region is drawn. The only geometry is the intersection of two unit squares offset by a unit vector, which in the rotated frame is a rectangle $(1-\cos\theta) \times (1-\sin\theta)$.)

The ideal solution

Let the dropped square have centre c and let θ be the angle its sides make with the axes. The square is invariant under rotation by $\pi/2$, so θ may be taken uniform on $[0, \pi/2)$; the lattice is invariant under integer translations, so c may be taken uniform on the cell $C = [0,1)^2$; the two are independent. Put $u = (\cos\theta, \sin\theta)$, $v = (-\sin\theta, \cos\theta)$ and $Q_0 = \{su + tv : \ |s| \leq \frac{1}{2},\ |t| \leq \frac{1}{2}\}$, the dropped square translated to the origin. A point P is covered iff $P - c \in Q_0$, iff (central symmetry,

$Q_0 = -Q_0$) $c \in P + Q_0$. Fix $\theta \in (0, \pi/2)$; the two endpoints carry probability zero. Let N be the number of lattice points covered.

Step 1: $\mathbb{E}[N] = 1$. Since $N = \sum_{P \in \mathbb{Z}^2} \mathbf{1}[c \in P + Q_0]$,

$$\mathbb{E}[N] = \sum_{P \in \mathbb{Z}^2} \operatorname{area}\big(C \cap (P + Q_0)\big) = \sum_{P \in \mathbb{Z}^2} \operatorname{area}\big((C - P) \cap Q_0\big) = \operatorname{area}(Q_0) = 1,$$

because the translates $C - P$, $P \in \mathbb{Z}^2$, partition the plane (only finitely many terms are nonzero).

Step 2: $N \le 2$. Let $P \neq P'$ be lattice points in the dropped square and $w = P' - P = (a, b)$. Then $|\langle w, u\rangle| \le 1$ and $|\langle w, v\rangle| \le 1$, that is, $|a\cos\theta + b\sin\theta| \le 1$ and $|-a\sin\theta + b\cos\theta| \le 1$. If $ab > 0$ the first quantity equals $|a|\cos\theta + |b|\sin\theta \ge \cos\theta + \sin\theta > 1$; if $ab < 0$ the second equals $|a|\sin\theta + |b|\cos\theta > 1$. So $ab = 0$. If $a = 0$ then $|b|\sin\theta \le 1$ and $|b|\cos\theta \le 1$, and $|b| \ge 2$ would force $\sin\theta, \cos\theta \le \frac{1}{2}$, impossible as $\sin^2\theta + \cos^2\theta = 1$; so $w = (0, \pm 1)$. Likewise $b = 0$ gives $w = (\pm 1, 0)$. Thus any two covered lattice points are adjacent. Three covered points would be pairwise adjacent, but $P_3 - P_2 = (P_3 - P_1) - (P_2 - P_1)$ is a difference of two distinct unit vectors, of squared length 2 or 4, hence not a unit vector. So $N \in \{0, 1, 2\}$.

Step 3. From $1 = \mathbb{E}[N] = \mathbb{P}(N = 1) + 2\,\mathbb{P}(N = 2)$ and $1 = \mathbb{P}(N = 0) + \mathbb{P}(N = 1) + \mathbb{P}(N = 2)$, subtraction gives $\mathbb{P}(N = 0) = \mathbb{P}(N = 2)$.

Step 4: the two-corner probability. By Step 2, $N = 2$ iff the square covers a pair $\{P, P + e_1\}$ or a pair $\{P, P + e_2\}$, and it covers at most one such pair. Hence

$$\mathbb{P}(N = 2) = \sum_P \mathbb{P}(P,\ P + e_1 \text{ covered}) + \sum_P \mathbb{P}(P,\ P + e_2 \text{ covered}).$$

Now P and $P + e_1$ are both covered iff $c - P \in K_1 := Q_0 \cap (e_1 + Q_0)$, and the partition argument of Step 1 gives $\sum_P \mathbb{P}(c - P \in K_1) = \operatorname{area}(K_1)$. In the coordinates (s, t) along u, v we have $Q_0 = [-\frac{1}{2}, \frac{1}{2}]^2$, and $e_1 + Q_0$ is Q_0 shifted by $(\langle e_1, u\rangle, \langle e_1, v\rangle) = (\cos\theta, -\sin\theta)$, so $K_1 = [\cos\theta - \frac{1}{2}, \frac{1}{2}] \times [-\frac{1}{2}, \frac{1}{2} - \sin\theta]$ has area $(1 - \cos\theta)(1 - \sin\theta)$. For e_2 the shift is $(\sin\theta, \cos\theta)$ and the area is the same. Therefore

$$\mathbb{P}(N = 0 \mid \theta) = \mathbb{P}(N = 2 \mid \theta) = 2(1 - \sin\theta)(1 - \cos\theta).$$

Step 5: average over θ.

$$\mathbb{P}(N = 0) = \frac{2}{\pi}\int_0^{\pi/2} 2(1 - \sin\theta)(1 - \cos\theta)\, d\theta = \frac{2}{\pi}\int_0^{\pi/2} \big(2 - 2\sin\theta - 2\cos\theta + \sin 2\theta\big)\, d\theta = \frac{2}{\pi}\big(\pi - 2 - 2 + 1\big).$$

The probability is $\boxed{2 - \dfrac{6}{\pi}} \approx 0.090$.

Layer 1: WHAT

Block 1: normalise the randomness (cell times quarter turn, independent) and fix the frame (u, v) of the dropped square, so that "covered" is two linear inequalities. Block 2: define the count N and compute $\mathbb{E}[N] = 1$ by the partition identity. Block 3: prove $N \le 2$ by the adjacency lemma. Block 4: two linear relations, one elimination: $\mathbb{P}(N = 0) = \mathbb{P}(N = 2)$. Block 5: $\mathbb{P}(N = 2)$ is the area of an intersection of two unit squares, read off in the (u, v) frame. Block 6: integrate the trigonometric polynomial over θ.

Juxtaposition. The nominal route: fix θ, draw the four forbidden squares around the corners of the cell, compute the area of the complement of their union. The ideal route: fix θ, ask for the mean of N. The bifurcation point is the sentence after "fix θ". The zombie writes "safe region equals cell minus a union of four squares"; you write "$\mathbb{E}[N] = 1$". Everything after that line is forced.

Layer 2: HOW

Block 1. Cue. "Position and orientation independent of the tiling", and a square (four-fold symmetric) against a lattice (translation invariant). **Why natural.** The probability space must be explicit before any computation, and the symmetries of both objects cut it to a unit cell times a quarter turn. The frame (u, v) is natural because the dropped square is the unit ball of the ℓ^∞ norm in its own coordinates, so membership is $|\langle P - c, u\rangle| \le \frac{1}{2}$, $|\langle P - c, v\rangle| \le \frac{1}{2}$, and nothing else in the problem is simpler than that. **Link to the next block.** With the frame fixed, every later membership test (Blocks 2, 3, 5) is a pair of linear inequalities.

Block 2. Cue. Unit area against unit covolume. **Why natural.** The expected number of lattice points in a uniformly random translate of a set K is $\mathrm{area}(K)$. It is double counting in continuous form: the pieces $(C - P) \cap Q_0$ over all P reassemble Q_0. There is no estimate here, only an identity, which is why it is free. **Link to the next block.** An expectation is not a probability. To convert, you need the range of N.

Block 3. Cue. A unit square has diameter $\sqrt{2} < 2$ and, in its own frame, width 1 in two perpendicular directions. **Why natural.** Two lattice points inside must differ by an integer vector whose projections on u and on v are both at most 1 in absolute value; integrality leaves only the four unit vectors. Pairwise adjacency of three lattice points is impossible. **Link to the next block.** With $N \in \{0, 1, 2\}$, the mean and total probability are two equations in three unknowns, and the unknown you want is expressible through the one you can compute.

Block 4. Cue. "N takes three values and I know its mean." **Why natural.** Total probability and expectation are independent linear relations; the combination eliminating $\mathbb{P}(N = 1)$ is their difference. The problem "probability of none" has become "probability of two", a much smaller and rigid event. **Link to the next block.** Compute the rigid event.

Block 5. Cue. "Two adjacent corners both covered." **Why natural.** That event is the intersection of two translates of Q_0. In the frame where Q_0 is axis-parallel the second translate is shifted by $(\cos\theta, -\sin\theta)$, and two axis-parallel unit squares offset by (p, q) with $|p|, |q| \le 1$ meet in a rectangle with sides $1 - |p|$ and $1 - |q|$. The lattice sum collapses by the same partition identity as Block 2. **Link to the next block.** The conditional probability is a trigonometric polynomial; integrate it.

Block 6. Cue. Uniform θ on $[0, \pi/2)$. **Why natural.** $\sin\theta\cos\theta = \frac{1}{2}\sin 2\theta$ and every term is elementary; the density is $2/\pi$. **Link to the next block.** Final answer.

Layer 3: WHY

The truth beneath the route: for a uniformly random translate (modulo the lattice) of a measurable set K, the number of lattice points inside has mean exactly $\mathrm{area}(K)$. This is Blichfeldt's observation, the engine of Minkowski's convex body theorem, and it is the same identity that makes Buffon's needle work (a random needle of length L on lines at spacing d crosses $2L/(\pi d)$ lines on average). It converts a hard probability into an easy expectation whenever the count is bounded by 2.

Why $N \le 2$ is the load-bearing fact: the examiners chose a unit square precisely because it has area 1 (so $\mathbb{E}[N] = 1$) and diameter $\sqrt{2}$ (too small to hold two non-adjacent lattice points, too large to be forced to hold one). The design of the problem is that the answer must be $\mathbb{P}(N = 2)$; everything else is decoration.

The discernment mechanism: the answer $2 - 6/\pi$ has a bare π in a denominator. That is the fingerprint of an average of a trigonometric polynomial over a uniformly random angle (Buffon's $2/\pi$ carries the same fingerprint). When the answer has this form the route is: conditional on

θ, obtain a polynomial in $\sin\theta$ and $\cos\theta$, then integrate. If your conditional expression is not a trigonometric polynomial, you have left the ideal route.

How to train it: every time a random shape meets a lattice or a family of lines, write the expected count first. Then ask for the maximal count. Then decide whether the probability of the maximal count is easier than the probability of zero; it almost always is, because the maximal configuration is rigid. The nudges were "unit" twice, "independent of the tiling", and "does not cover any": area equals expectation, uniform measure, and complement via the count.

Revolving concepts

Geometric probability and kinematic measure. A random rigid motion of the plane is (c, θ) with measure $dc\, d\theta$; Buffon's needle and Crofton's formula (the expected number of intersections of a random line with a curve is proportional to its length) live here. See Book 1, field: Probability, topic: geometric probability and expected counts.

Lattices and fundamental domains. The cell $[0,1)^2$ is a fundamental domain for $\mathbb{Z}^2$; its translates partition the plane; the covolume is 1. Blichfeldt's lemma: if $\text{area}(K) > 1$ then K contains two distinct points differing by a lattice vector. Minkowski: a convex, centrally symmetric K with $\text{area}(K) > 4$ contains a nonzero lattice point. See Book 1, field: Geometry, topic: lattices and fundamental domains.

The set of centres. The centres c for which $c + Q_0$ contains P form $P - Q_0$, which equals $P + Q_0$ by central symmetry. This is why the forbidden-region route exists and why it hurts: unions of translates overlap.

Norms. The square is the unit ball of the ℓ^∞ norm in its own frame; "covered" is a rotated ℓ^∞ inequality; the intersection of two ℓ^∞ balls is a box. See Book 1, field: Geometry, topic: norms and their unit balls.

Trigonometric integrals. $(\sin\theta + \cos\theta - 1)^2 = 2 + \sin 2\theta - 2\sin\theta - 2\cos\theta$, whose integral over $[0, \pi/2]$ is $\pi - 3$.

Recursion. None here; the only iteration is the sum over lattice translates, and the partition identity collapses it in one line.

Circuit Bank

- **WHEN** you see a random rigid motion of a shape relative to a lattice or a family of lines **DO** write $\mathbb{E}[\text{count}] = \text{area}$ (or length ratio) before anything else, then determine the possible values of the count **BECAUSE** the expectation is an identity, not an estimate, and a bounded count turns it into a probability
- **WHEN** you see "probability that none" with a count that is at most 2 **DO** write the total-probability and expectation equations and eliminate $\mathbb{P}(1)$ **BECAUSE** $\mathbb{P}(0) = \mathbb{P}(2) + 1 - \mathbb{E}$, and the event "two" is rigid and small
- **WHEN** you see a square in a rotated position **DO** describe membership as $|\langle P - c, u\rangle| \le \frac{1}{2}$, $|\langle P - c, v\rangle| \le \frac{1}{2}$ in the square's own frame; never in the ambient coordinates **BECAUSE** a square is an ℓ^∞ ball, and intersections of translates of ℓ^∞ balls are boxes with side $1 - |\text{shift}|$
- **WHEN** you see a final answer containing π in a denominator **DO** expect a trigonometric polynomial conditional on a uniform angle **BECAUSE** averaging sin and cos over a uniform angle is the only natural source of $1/\pi$

Backtesting

1. First useful cue: "unit square" dropped on "unit squares": area 1 against covolume 1, so $\mathbb{E}[N] = 1$.

2. Method the expression suggested: count, do not draw; $\mathbb{P}(N = 0) = \mathbb{P}(N = 2)$ plus one intersection area.
3. Where you would hesitate, and the thought to reinforce: proving $N \leq 2$ properly, and justifying $\sum_P \text{area}(C \cap (P + K)) = \text{area}(K)$. The thought: the translates of the cell partition the plane; write the one-line identity and move on. The adjacency lemma is two lines of integrality.
4. The trap the question sets: normalising θ on $[0, 2\pi)$ with density $1/(2\pi)$ and then integrating a formula valid on $[0, \pi/2)$ (fatal), and silently assuming only the four corners of one cell can be covered (true because the circumradius is $\sqrt{2}/2 < 1$, but it must be said, or better, avoided by summing over all P as above).
5. What makes you recognise it faster next time: the answer's fingerprint π in a denominator; the phrase "independent of the tiling"; any Buffon-type set-up with a shape of area equal to the covolume.
6. The reusable decision rule: random shape versus lattice: the expectation of the count is the area; the probability of the extreme count is the area of a rigid configuration; the probability of zero follows by elimination.

Mirror drills

1. A unit segment is dropped at random on a plane ruled with parallel lines one unit apart. Find the probability that it crosses a line. *Answer:* for fixed angle θ to the lines the expected number of crossings is $|\sin\theta|$; averaging over uniform θ gives $2/\pi$; since at most one line can be crossed, the probability is $2/\pi$.
2. A disc of radius $r < \frac{1}{2}$ is dropped at random on the unit lattice. Find the probability that it covers no lattice point. *Answer:* the count is 0 or 1 (diameter less than 1) and its mean is πr^2, so the probability is $1 - \pi r^2$.
3. A unit square is dropped at random on the unit lattice with its orientation fixed at $\pi/4$. Find the probability that it covers exactly two lattice points. *Answer:* $2(1 - \frac{\sqrt{2}}{2})^2 = 3 - 2\sqrt{2}$.

5.8 B2. Weighted geometric means and tuned AM-GM

Determine the maximum value of the sum

$$S = \sum_{n=1}^{\infty} \frac{n}{2^n} (a_1 a_2 \cdots a_n)^{1/n}$$

over all sequences $a_1, a_2, a_3, \ldots$ of nonnegative real numbers satisfying

$$\sum_{k=1}^{\infty} a_k = 1.$$

Read it out loud

Say what the object is, not what the story is. Write $G_n = (a_1 a_2 \cdots a_n)^{1/n}$, the running geometric mean of the first n terms, and $w_n = n/2^n$. The quantity is $S = \sum_{n \geq 1} w_n G_n$, and the only constraint is on the linear functional $\sum_k a_k$. So the entire problem reads: bound a weighted sum of prefix geometric means by a multiple of the sum of the terms, and make the multiple sharp. That is Carleman's inequality with a weight inserted, and the useful thing you own is not the statement of Carleman but its proof, which is a design procedure with a free parameter in it.

The fields wired in are inequalities (AM-GM, and specifically the fact that AM-GM has a free rescaling built into it), double series with nonnegative terms (so Tonelli lets you swap the order of summation with no hypothesis to check), the two geometric sums $\sum_{n \geq k} x^n$ and $\sum n x^n$, and equality analysis. What is really being tested is whether you treat AM-GM as a fixed statement (zombie) or as a one-parameter family of statements that you tune to fit the constraint (ideal).

The first useful cue is the phrase "geometric mean of a prefix", and the second is that $\sum a_k$, not $\sum$ of anything else, is what you are given. Within thirty seconds you should have written, with the positive numbers $c_1, c_2, \ldots$ left undetermined,

$$G_n = \Big(\prod_{k=1}^{n} c_k\Big)^{-1/n} \Big(\prod_{k=1}^{n} c_k a_k\Big)^{1/n} \leq \Big(\prod_{k=1}^{n} c_k\Big)^{-1/n} \cdot \frac{1}{n} \sum_{k=1}^{n} c_k a_k.$$

You do not choose the c_k yet. You choose them at the end, from a design equation.

There is a second free piece of information in the wording: the question says "maximum", not "supremum", so the value is attained. AM-GM is an equality exactly when all the $c_k a_k$ agree, so the extremal sequence must be $a_k \propto 1/c_k$. You therefore get the extremiser for free the moment you get the weights, and you can check the answer twice from one calculation.

Approaches, ranked worst to best

1. **Lagrange multipliers on infinitely many variables.** *Cue:* "maximise subject to a linear constraint". *Circuit:* stationary point of $S - \lambda \sum a_k$. (Why worst: $\partial S / \partial a_j = \frac{1}{a_j} \sum_{n \geq j} \frac{G_n}{2^n}$, so the stationarity conditions are an infinite nonlinear system coupling every variable to every later one. You would still have to justify that a maximiser exists in an infinite-dimensional set, and that no boundary point with some $a_j = 0$ beats it. Pages of ink, no certificate.)

2. **Carleman off the shelf.** *Cue:* prefix geometric means summed. *Circuit:* $\sum_{n \geq 1} G_n \leq e \sum_k a_k$. (Why second worst: it is a correct inequality and it is the right family, but $n/2^n \leq \frac{1}{2}$ only gives $S \leq e/2 \approx 1.359$. The weights $n/2^n$ are the whole content of the problem, and throwing them away throws away the sharp constant.)

3. **Guess the extremiser inside the geometric family.** *Cue:* equality cases of AM-GM are geometric because the weights are exponential. *Circuit:* set $a_k = (1-r)r^{k-1}$ and optimise one variable. (Why third: it is honest reconnaissance and it takes two minutes. With $s = \sqrt{r}$ you

get $S = \frac{2(1-s^2)}{(2-s)^2}$, whose derivative is $\frac{4-8s}{(2-s)^3}$, so $s = \frac{1}{2}$, $r = \frac{1}{4}$, $S = \frac{2}{3}$. This tells you the answer and the extremiser, but it proves nothing: you have optimised over one family, not over all sequences.)

4. **AM-GM with undetermined weights, then a design equation.** *Cue:* the constraint is $\sum a_k$, so the bound must collapse to a constant times $\sum a_k$. *Circuit:* insert c_k, apply AM-GM, swap the order of summation, then choose c_k so that every a_k receives the same coefficient. (Why best: the weights are not guessed, they are forced. One line of AM-GM, one Tonelli swap, one geometric series, and the equality case falls out of the same calculation. This is exactly the proof of Carleman's inequality with a different weight sequence, so it costs you nothing new.)

The ideal solution

The answer is $\boxed{\frac{2}{3}}$, attained at $a_k = 3 \cdot 4^{-k}$.

All terms are nonnegative, so S is a well defined element of $[0, \infty]$ and every rearrangement below is legitimate by Tonelli's theorem for nonnegative double series.

The bound. Since $1 + 2 + \cdots + n = \frac{n(n+1)}{2}$,

$$\Big(\prod_{k=1}^{n} 4^k\Big)^{1/n} = 4^{(n+1)/2} = 2^{n+1},$$

so for every $n \geq 1$, applying AM-GM to the n nonnegative numbers $4a_1, 4^2a_2, \ldots, 4^n a_n$,

$$2^{n+1}(a_1 a_2 \cdots a_n)^{1/n} = \Big(\prod_{k=1}^{n} 4^k a_k\Big)^{1/n} \leq \frac{1}{n}\sum_{k=1}^{n} 4^k a_k.$$

(The weight 4^k is not guessed; Block 2 of Layer 2 derives it from a single design equation.) Multiplying by $\frac{n}{2^n} \cdot \frac{1}{2^{n+1}}$,

$$\frac{n}{2^n}(a_1 \cdots a_n)^{1/n} \leq \frac{1}{2} \cdot \frac{1}{4^n}\sum_{k=1}^{n} 4^k a_k.$$

Summing over $n \geq 1$ and exchanging the two summations,

$$2S \leq \sum_{n=1}^{\infty}\sum_{k=1}^{n} 4^{k-n} a_k = \sum_{k=1}^{\infty} a_k \sum_{n=k}^{\infty} 4^{k-n} = \sum_{k=1}^{\infty} a_k \cdot \frac{1}{1-\frac{1}{4}} = \frac{4}{3}\sum_{k=1}^{\infty} a_k = \frac{4}{3}.$$

Hence $S \leq \frac{2}{3}$ for every admissible sequence.

Attainment. Take $a_k = 3 \cdot 4^{-k}$. Then $\sum_{k\geq 1} a_k = 3 \cdot \frac{1/4}{1-1/4} = 1$, so the sequence is admissible. Every AM-GM above is an equality because $4^k a_k = 3$ for all k. Explicitly,

$$(a_1 \cdots a_n)^{1/n} = \big(3^n 4^{-n(n+1)/2}\big)^{1/n} = 3 \cdot 2^{-(n+1)},$$

so

$$S = \sum_{n=1}^{\infty} \frac{n}{2^n} \cdot \frac{3}{2^{n+1}} = \frac{3}{2}\sum_{n=1}^{\infty} \frac{n}{4^n} = \frac{3}{2} \cdot \frac{1/4}{(1-1/4)^2} = \frac{3}{2} \cdot \frac{4}{9} = \frac{2}{3}.$$

Therefore the maximum value of S is $\frac{2}{3}$.

Layer 1: WHAT

Block 1: name the object as $\sum_n w_n G_n$ with G_n a prefix geometric mean, and note that the constraint is on $\sum a_k$, so the target of the bound is $C \sum a_k$. Block 2: rescale before applying AM-GM, with weights $c_k = t^k$ left undetermined, and read off the coefficient that each a_k finally receives. Block 3: impose

the design equation "that coefficient does not depend on k", which forces $t = 4$ and hands you $C = \frac{2}{3}$. Block 4: run the actual proof: AM-GM at fixed n, swap the order of the double sum, sum the geometric series in n. Block 5: equality in AM-GM for every n means $4^k a_k$ is constant; the constraint normalises it to $a_k = 3 \cdot 4^{-k}$; verify directly.

Juxtaposition. The nominal route fixes AM-GM in its bare form, $G_n \le \frac{1}{n}(a_1 + \cdots + a_n)$, and then discovers that $\sum_n \frac{n}{2^n} \cdot \frac{1}{n} \sum_{k \le n} a_k$ has the coefficient of a_k equal to $\sum_{n \ge k} 2^{-n} = 2^{1-k}$, which is not constant, is largest at $k = 1$, and yields only $S \le 1$. From there the zombie starts patching. The ideal route never applies AM-GM to a_k; it applies AM-GM to $c_k a_k$ and treats c_k as an unknown to be solved for.

The bifurcation point is the very first inequality you write. The zombie writes $(a_1 \cdots a_n)^{1/n} \le \frac{1}{n} \sum_{k \le n} a_k$. You write $(a_1 \cdots a_n)^{1/n} = \left(\prod_{k \le n} c_k\right)^{-1/n} \left(\prod_{k \le n} c_k a_k\right)^{1/n}$ and stop.

Layer 2: HOW

Block 1. **Cue.** A geometric mean of the first n terms, summed against weights, with a constraint on the plain sum $\sum a_k$. **Why natural.** AM-GM is the only elementary bridge from a product to a sum, and the constraint tells you what the right hand side has to look like when you are finished: a constant times $\sum_k a_k$. So the plan is fixed before any computation: turn each G_n into a linear combination of the a_k, swap the order of summation, and collect the coefficient of each a_k. **Link to the next block.** The bare AM-GM produces the coefficient 2^{1-k}, which decays in k. Something must be inserted to flatten it.

Block 2. **Cue.** The geometric mean is multiplicative and homogeneous of degree one: replacing a_k by $c_k a_k$ multiplies G_n by the purely n-dependent factor $(\prod_{k \le n} c_k)^{1/n}$. **Why natural.** That is exactly why weights cost nothing. The rescaling is invisible to the structure of the problem and gives you an infinite supply of AM-GM statements, one for each positive sequence (c_k). Since the weights $n/2^n$ are exponential in n, try $c_k = t^k$, for which $\left(\prod_{k \le n} t^k\right)^{1/n} = t^{(n+1)/2}$, and everything stays geometric. Then

$$S \le \sum_{n \ge 1} \frac{n}{2^n} t^{-(n+1)/2} \frac{1}{n} \sum_{k=1}^{n} t^k a_k = \frac{1}{\sqrt{t}} \sum_{k \ge 1} t^k a_k \sum_{n \ge k} (2\sqrt{t})^{-n},$$

and since $\sum_{n \ge k} (2\sqrt{t})^{-n} = \frac{2\sqrt{t}}{2\sqrt{t}-1} (2\sqrt{t})^{-k}$, the coefficient of a_k is

$$\frac{2}{2\sqrt{t}-1} \left(\frac{\sqrt{t}}{2}\right)^k.$$

Link to the next block. That coefficient must not depend on k, because the only information you have is $\sum a_k = 1$. This is the design equation.

Block 3. **Cue.** "The coefficient of a_k must be constant in k." **Why natural.** $(\sqrt{t}/2)^k$ is constant precisely when $\sqrt{t} = 2$, that is $t = 4$, and then the constant is $\frac{2}{2 \cdot 2 - 1} = \frac{2}{3}$. Nothing was guessed: the exponent base 4 is the unique solution of a linear equation in $\log t$. As a bonus, if $\sqrt{t} < 2$ the coefficient decays and the best you can extract is its value at $k = 1$, namely $\frac{\sqrt{t}}{2\sqrt{t}-1}$, which is a decreasing function of $\sqrt{t}$ on $(\frac{1}{2}, 2]$ and is therefore also minimised at $t = 4$; if $\sqrt{t} > 2$ the coefficients are unbounded and the method dies. So $t = 4$ is optimal within the family as well as forced by the design equation. **Link to the next block.** With t known, rewrite the argument cleanly with 4^k in place of t^k and hide the search.

Block 4. **Cue.** A double sum of nonnegative terms. **Why natural.** Tonelli's theorem says a double series of nonnegative terms may be summed in either order, with the same value in $[0, \infty]$. No

absolute convergence needs to be established first, which matters because you do not yet know S is finite. The inner sum $\sum_{n\ge k} 4^{k-n} = \frac{4}{3}$ is one geometric series. **Link to the next block.** The bound is now exactly $\frac{2}{3}\sum_k a_k = \frac{2}{3}$.

Block 5. Cue. The word "maximum" in the question. **Why natural.** AM-GM is an equality if and only if all its arguments agree, so equality in every step needs $4^k a_k$ independent of k, that is $a_k = C4^{-k}$; the constraint gives $C = 3$. The verification then needs $\sum_{n\ge1} nx^n = \frac{x}{(1-x)^2}$ at $x = \frac{1}{4}$, which is the derivative of the geometric series. **Link to the next block.** Both halves agree at $\frac{2}{3}$, so the bound is sharp and the maximum is attained.

Layer 3: WHY

The mathematical truth underneath is that AM-GM is not one inequality but a one-parameter family, indexed by the rescalings $a_k \mapsto c_k a_k$, and that the family is closed in exactly the right way: rescaling changes the geometric mean by a factor depending on n only. Whenever you must bound $\sum_n w_n (a_1 \cdots a_n)^{1/n}$ by a multiple of $\sum_k a_k$, the general machine is

$$\sum_{n\ge1} w_n G_n \le \sum_{k\ge1} \Big(c_k \sum_{n\ge k} \frac{w_n}{n} \Big(\prod_{j\le n} c_j \Big)^{-1/n} \Big) a_k,$$

and the whole art is choosing (c_k) to make the bracket constant, or at least bounded. That is a design equation, not a guess, and it is the reason the constant that comes out is sharp: the same choice that flattens the bracket is the choice that makes AM-GM tight along the extremiser $a_k \propto 1/c_k$.

Carleman's inequality is the same machine with $w_n = 1$. There the bracket is flattened by $c_k = (k+1)^k/k^{k-1}$, for which $\prod_{j\le n} c_j = (n+1)^n$ telescopes, so $\big(\prod_{j\le n} c_j\big)^{-1/n} = \frac{1}{n+1}$ and the bracket becomes $c_k \sum_{n\ge k} \frac{1}{n(n+1)} = \frac{c_k}{k} = \big(1+\frac{1}{k}\big)^k < e$. Notice the two telescopes: one in the product, one in the sum. Our problem is the easy member of the family because the weights are geometric, so both telescopes degenerate to geometric series.

The discernment mechanism is the shape of the constraint. You are told $\sum a_k = 1$ and nothing else, so the only inequality that can possibly be sharp is one whose right hand side is $C\sum a_k$. Any intermediate expression that still mentions n in a way you cannot sum is a sign you have applied AM-GM too early or with the wrong arguments. Train it like this: never write AM-GM with the given variables. Write it with $c_k a_k$, carry the c_k through to the end, and let the last line tell you what they are.

What the examiners nudged with: the weight $n/2^n$ is exponential with a polynomial factor, which is the signature of "the answer is a rational number and the extremiser is geometric". The factor n is there for cosmetics, since it cancels against the $\frac{1}{n}$ of AM-GM; strip it and the problem changes character, because then $\sum_{n\ge k} \frac{1}{n} x^n$ has no closed form and the bracket cannot be flattened. That cancellation of n against $\frac{1}{n}$ is the fingerprint that the intended route is AM-GM, and is the single most informative feature of the statement.

Revolving concepts

AM-GM and its equality case. For nonnegative $x_1, \ldots, x_n$, $(x_1 \cdots x_n)^{1/n} \le \frac{1}{n}(x_1 + \cdots + x_n)$, with equality if and only if all x_i are equal. The weighted form: for weights $\lambda_i > 0$ summing to 1, $\prod x_i^{\lambda_i} \le \sum \lambda_i x_i$. The rescaling $x_i \mapsto c_i x_i$ is what converts the unweighted form into a tunable tool. See Book 1, field: Inequalities, topic: AM-GM and rescaling.

Carleman's inequality. $\sum_{n\ge1} (a_1 \cdots a_n)^{1/n} \le e \sum_{k\ge1} a_k$ for nonnegative a_k, with e best possible. Its proof is the machine above with $c_k = (k+1)^k/k^{k-1}$. It is the limiting case $p \to \infty$ of Hardy's inequality $\sum_n \big(\frac{a_1+\cdots+a_n}{n}\big)^p \le \big(\frac{p}{p-1}\big)^p \sum_k a_k^p$. See Book 1, field: Inequalities, topic: Carleman and Hardy.

Tonelli for series. For $b_{n,k} \geq 0$, $\sum_n \sum_k b_{n,k} = \sum_k \sum_n b_{n,k}$ in $[0, \infty]$, with no convergence hypothesis. This is why you may swap before knowing $S < \infty$. See Book 1, field: Real analysis, topic: double series and interchange.

Geometric sums. $\sum_{n \geq k} x^n = \frac{x^k}{1-x}$ and $\sum_{n \geq 1} n x^n = \frac{x}{(1-x)^2}$ for $|x| < 1$, the second being $x \frac{d}{dx}$ of the first.

The generalised recursion. There is no recursion in the sequence (a_k), but there is a self-similar structure in the double sum, and it generalises cleanly. For $\lambda > 1$, the same argument with $c_k = \lambda^{2k}$ gives

$$\max\Big\{ \sum_{n \geq 1} \frac{n}{\lambda^n} (a_1 \cdots a_n)^{1/n} : \; a_k \geq 0, \; \sum_k a_k = 1 \Big\} = \frac{\lambda}{\lambda^2 - 1},$$

attained at $a_k = (\lambda^2 - 1)\lambda^{-2k}$; the case $\lambda = 2$ is this problem. The reason a geometric extremiser appears is that the weight sequence is geometric, so the design equation is linear in $\log t$ and has exactly one root. Whenever the weights are $w_n = n\lambda^{-n}$, expect the extremiser $a_k \propto \lambda^{-2k}$: the square is the footprint of the exponent $(n+1)/2$ in $\big(\prod_{k \leq n} t^k\big)^{1/n} = t^{(n+1)/2}$.

Circuit Bank

- **WHEN** you see a geometric mean of the first n terms summed over n against weights **DO** apply AM-GM to $c_k a_k$ with the c_k undetermined, swap the order of summation, and collect the coefficient of a_k **BECAUSE** the constraint is on $\sum a_k$, so the bound must end as a constant times $\sum a_k$, and the c_k are exactly the freedom that makes that possible
- **WHEN** you have an inequality with free parameters and a linear constraint **DO** write down the design equation "the coefficient of every constrained variable is equal" and solve it **BECAUSE** a bound of the form $\sum_k \gamma_k a_k$ can only be turned into a number using $\sum a_k = 1$ when the γ_k are constant, and the flattening choice is automatically the one that is sharp
- **WHEN** you see the word maximum rather than supremum **DO** find the equality case of your inequality first and use it as a check on the constant **BECAUSE** an attained maximum means the equality conditions are consistent with the constraint, so they determine the extremiser and give you a second, independent computation of the answer
- **WHEN** you see a double sum of nonnegative terms **DO** swap the order without justification beyond naming Tonelli **BECAUSE** nonnegativity makes both iterated sums equal in $[0, \infty]$, so you may swap before you know anything converges

Backtesting

1. First useful cue: prefix geometric mean summed against geometric weights, with only $\sum a_k$ given. That is the Carleman machine.
2. Method the expression suggested: AM-GM applied to rescaled terms $c_k a_k$, with c_k solved for at the end by demanding a k-independent coefficient.
3. Where you would hesitate, and the thought to reinforce: reaching for the bare AM-GM because it is the one you have memorised. Reinforce this: AM-GM is a family, and the member you want is the one whose equality case matches the answer you expect. Carry the parameter.
4. The trap the question sets: the factor n in $n/2^n$ looks like an obstacle and is actually the invitation, since it cancels the $\frac{1}{n}$ from AM-GM. A candidate who bounds $n/2^n \leq \frac{1}{2}$ and quotes Carleman gets $e/2$ and stops.
5. What makes you recognise it faster next time: the pair (geometric mean of a prefix, constraint on the plain sum). Also the cancellation of n against $\frac{1}{n}$, which tells you AM-GM is intended rather than convexity or rearrangement.

6. The reusable decision rule: rescale before you apply a symmetric inequality, keep the scaling free, and let the constraint pick it.

Mirror drills

1. For $\lambda > 1$, find the maximum of $\sum_{n\geq 1} \frac{n}{\lambda^n}(a_1 \cdots a_n)^{1/n}$ over nonnegative sequences with $\sum_k a_k = 1$. *Answer:* $\frac{\lambda}{\lambda^2-1}$, at $a_k = (\lambda^2-1)\lambda^{-2k}$; take $c_k = \lambda^{2k}$, so that $\left(\prod_{k\leq n} c_k\right)^{1/n} = \lambda^{n+1}$.
2. Prove Carleman's inequality $\sum_{n\geq 1}(a_1 \cdots a_n)^{1/n} \leq e \sum_k a_k$ by the same machine, discovering the weights rather than quoting them. *Answer:* demand $\prod_{j\leq n} c_j$ to be a clean function of n; $c_k = (k+1)^k/k^{k-1}$ gives $\prod_{j\leq n} c_j = (n+1)^n$, and then the coefficient of a_k is $c_k \sum_{n\geq k} \frac{1}{n(n+1)} = \frac{c_k}{k} = (1+\frac{1}{k})^k < e$.
3. Show that if the sequence is required to vanish beyond a fixed index N, the supremum of S is strictly less than $\frac{2}{3}$, yet tends to $\frac{2}{3}$ as $N \to \infty$. *Answer:* strictness because the chain $2S \leq \sum_{n\geq 1} 4^{-n} \sum_{k\leq n} 4^k a_k$ discards the terms with $n > N$, which are strictly positive while the corresponding G_n vanish; and $a_k = 4^{-k}/\sum_{j\leq N} 4^{-j}$ gives $S = \frac{c}{2}\sum_{n\leq N} n4^{-n} \to \frac{2}{3}$ with $c = 3/(1-4^{-N})$.

5.9 B3. Averaging a rotational derivative over circles

Let $h(x, y)$ be a real-valued function that is twice continuously differentiable throughout $\mathbb{R}^2$, and define

$$\rho(x, y) = yh_x - xh_y.$$

Prove or disprove: For any positive constants d and r with $d > r$, there is a circle $\mathcal{S}$ of radius r whose centre is a distance d away from the origin such that the integral of ρ over the interior of $\mathcal{S}$ is zero.

Read it out loud

There are two objects and you must name both before you compute anything. The first is the combination $yh_x - xh_y$. That is not a random pairing of derivatives: in polar coordinates $x = s\cos\theta$, $y = s\sin\theta$ the chain rule gives $\partial_\theta = -y\partial_x + x\partial_y$, so $\rho = -h_\theta$. The integrand is the angular derivative of h, with a sign. The second object is the family of candidate circles: their centres run over the circle $|c| = d$, so the family is a one-parameter loop, parametrised by the angle ϕ of the centre.

The fields wired in are vector calculus in polar coordinates, interchange of integrals (Fubini), the rotation symmetry of the family of discs, and the intermediate value theorem. What is really being tested is whether you recognise the question type. You are handed a continuum of candidate objects and asked to produce one on which a functional vanishes. You are not asked which one. That is the signature of an averaging argument: show the average of the functional over the family is zero, note the functional is continuous in the parameter, and let the intermediate value theorem manufacture the object. Constructing the circle explicitly is impossible, since h is arbitrary, and the impossibility of construction is itself the cue.

The first useful cue is the shape $yh_x - xh_y$, and within thirty seconds you should have written two lines: $\rho = -\partial_\theta h$, and therefore for every $s > 0$

$$\int_0^{2\pi} \rho(s\cos\theta, s\sin\theta)\, d\theta = -\big[h(s\cos\theta, s\sin\theta)\big]_{\theta=0}^{\theta=2\pi} = 0,$$

because h is a single-valued function of the point, hence 2π-periodic in θ. Every circle centred at the origin is a null set for ρ. That is the entire mechanism, and the rest of the solution is bookkeeping that converts "zero on origin-centred circles" into "zero on one of the given discs".

A caveat worth clarifying at once, because it decides between prove and disprove. If h were allowed to be the multivalued angle θ itself, then $\rho = -1$ and every disc integral would be $-\pi r^2$, never zero. So any counterexample would have to be built from a potential that is not single-valued, and single-valuedness is exactly what the hypothesis gives you. The statement is true, and you now know which line of the proof carries the weight.

Approaches, ranked worst to best

1. **Hunt for a counterexample.** *Cue:* the phrase "prove or disprove". *Circuit:* choose h making ρ of one sign. (Why worst: it cannot succeed, and finding out why costs the whole time budget unless you first notice $\rho = -h_\theta$. Every attempt fails for the same structural reason: ρ has zero mean on each circle about the origin, so it can never be of one sign on such a circle. Note that the failed attempt does hand you the mechanism, which is why it is worth ten seconds, not ten minutes.)

2. **Fourier expansion in the angle.** *Cue:* $\rho = -h_\theta$ and h is periodic in θ. *Circuit:* write $h(s, \theta) = \sum_m c_m(s)e^{im\theta}$, so $\rho = -\sum_m imc_m(s)e^{im\theta}$ and the mode $m = 0$ is annihilated. (Why third: it is correct and it explains everything, since averaging over the centre angle ϕ multiplies the m-th mode by $\frac{1}{2\pi}\int_0^{2\pi} e^{im\phi}d\phi = 0$ for $m \neq 0$, while $m = 0$ was already killed. But for a general twice

differentiable h you owe the grader a convergence argument for the series, and that is ink you do not need to spend.)

3. **Divergence theorem.** *Cue:* $\rho = \nabla h \cdot V$ with $V = (y, -x)$ and $\operatorname{div} V = 0$. *Circuit:* $\rho = \operatorname{div}(hV)$, so the area integral becomes a boundary integral. (Why second: it is a genuine simplification, reducing the double integral to $\oint_{\partial\mathcal{S}} h\, V \cdot n\, ds = rd \int_0^{2\pi} h\big(c_\phi + r(\cos t, \sin t)\big) \sin(\phi - t)\, dt$, and the averaging step then works by writing the point as a rotation of $(d + r\cos u, r\sin u)$ and using that $\sin u$ is odd while $|(d + r\cos u, r\sin u)|$ is even in u. Correct, and shorter in data, but the algebra of $V \cdot n$ costs more ink than the polar route and the averaging step is the same either way.)

4. **Average over the family, then intermediate value theorem.** *Cue:* a one-parameter family of candidates and a demand for existence, not construction. *Circuit:* $\int_0^{2\pi} I(\phi)\, d\phi = 0$ by Fubini plus rotation invariance, then IVT. (Why best: it uses the mechanism directly, needs only continuity of ρ, and every step is one line. The rotation invariance of the family is what makes the interchanged integral factor into a radial weight times an angular integral that is already known to vanish.)

The ideal solution

The statement is true. Fix $d > r > 0$ (in fact $d > r$ is never used; see the remark at the end).

For $\phi \in \mathbb{R}$ put $c_\phi = (d\cos\phi,\, d\sin\phi)$ and let $D_\phi = \{P \in \mathbb{R}^2 :\ |P - c_\phi| \le r\}$, so that the circles described in the problem are exactly the boundaries ∂D_ϕ, $\phi \in [0, 2\pi)$. Define

$$I(\phi) = \iint_{D_\phi} \rho\, dA.$$

We show $I(\phi_0) = 0$ for some ϕ_0. Since h is continuously differentiable, $\rho = yh_x - xh_y$ is continuous, which is all that is used.

Step 1: I is continuous. Translating, $I(\phi) = \iint_{|w|\le r} \rho(c_\phi + w)\, dA(w)$. The function ρ is uniformly continuous on the compact disc $\{|P| \le d + r\}$, which contains every D_ϕ; so given $\varepsilon > 0$ there is $\delta > 0$ with $|\rho(P) - \rho(P')| < \varepsilon/(\pi r^2)$ whenever $|P - P'| < \delta$ inside that disc, and $|\phi - \phi'|$ small forces $|c_\phi - c_{\phi'}| < \delta$. Hence $|I(\phi) - I(\phi')| < \varepsilon$.

Step 2: ρ has zero mean on every circle centred at the origin. Fix $s > 0$ and set $g(\theta) = h(s\cos\theta, s\sin\theta)$. Then g is continuously differentiable and 2π-periodic, and by the chain rule

$$g'(\theta) = -s\sin\theta\, h_x + s\cos\theta\, h_y = -\big(yh_x - xh_y\big) = -\rho$$

evaluated at $(s\cos\theta, s\sin\theta)$. Therefore

$$\int_0^{2\pi} \rho(s\cos\theta, s\sin\theta)\, d\theta = -\int_0^{2\pi} g'(\theta)\, d\theta = -\big(g(2\pi) - g(0)\big) = 0.$$

Step 3: the average of I over the family vanishes. Let $\chi(P, \phi)$ be the indicator of $|P - c_\phi| \le r$, so $I(\phi) = \iint_{\mathbb{R}^2} \chi(P,\phi)\rho(P)\, dA(P)$. The function $\chi\rho$ is bounded and measurable and vanishes outside the compact set $\{|P| \le d + r\} \times [0, 2\pi]$, so Fubini's theorem gives

$$\int_0^{2\pi} I(\phi)\, d\phi = \iint_{\mathbb{R}^2} \rho(P)\, M(P)\, dA(P), \qquad M(P) = \int_0^{2\pi} \chi(P,\phi)\, d\phi.$$

Write $P = (s\cos\theta, s\sin\theta)$ with $s > 0$. Then $|P - c_\phi|^2 = s^2 + d^2 - 2sd\cos(\theta - \phi)$, so

$$\chi(P,\phi) = 1 \iff \cos(\theta - \phi) \ \ge\ \lambda(s) := \frac{s^2 + d^2 - r^2}{2sd},$$

which depends on P only through s and on ϕ only through $\theta - \phi$. Integrating over a full period in ϕ therefore removes θ altogether:

$$M(P) = M(s) = \begin{cases} 0, & \lambda(s) > 1, \\ 2\arccos\lambda(s), & |\lambda(s)| \le 1, \\ 2\pi, & \lambda(s) < -1. \end{cases}$$

In particular M is a bounded measurable function of s alone. Passing to polar coordinates and using Step 2,

$$\int_0^{2\pi} I(\phi)\,d\phi = \int_0^{\infty} M(s)\left(\int_0^{2\pi} \rho(s\cos\theta, s\sin\theta)\,d\theta\right) s\,ds = \int_0^{\infty} M(s)\cdot 0\cdot s\,ds = 0.$$

(The single point $s = 0$ is a null set and may be ignored.)

Step 4: conclusion. I is continuous on $[0, 2\pi]$ and $\int_0^{2\pi} I(\phi)\,d\phi = 0$. If I had no zero then, being continuous on an interval, it would have constant sign by the intermediate value theorem, and its integral would be strictly positive or strictly negative. Hence $I(\phi_0) = 0$ for some $\phi_0 \in [0, 2\pi]$, and $\mathcal{S} = \partial D_{\phi_0}$ is a circle of radius r whose centre is at distance d from the origin with $\iint_{\operatorname{int}\mathcal{S}} \rho = 0$. This proves the statement.

Remark. Two hypotheses are stronger than needed. Only $h \in C^1$ was used, not C^2. And $d > r$ was never used: if $d \le r$ the discs contain the origin, but $\{0\}$ is a null set and Step 2 holds for every $s > 0$, so the conclusion holds for all $d, r > 0$. Also, the argument gives slightly more: unless $I \equiv 0$, the function I takes both signs, so at least two of the circles work.

Layer 1: WHAT

Block 1: read $yh_x - xh_y$ as $-\partial_\theta h$. Block 2: deduce that ρ has zero mean on every circle centred at the origin, using only that h is single valued. Block 3: recognise the question type (existence within a continuous family) and set up $I(\phi)$, the functional on the family. Block 4: average I over the family; interchange the two integrals. Block 5: observe that the interchanged weight $M(P)$ is rotation invariant, because the family of discs is a rotation orbit, hence M is radial and the inner integral is the one from Block 2. Block 6: continuity plus the intermediate value theorem.

Juxtaposition. The nominal route reads "prove or disprove" and starts hunting for an h that breaks it, or else tries to evaluate $I(\phi)$ for a general h, which is hopeless. A more sophisticated zombie does spot $\rho = -h_\theta$, and then tries to exploit it disc by disc: in polar coordinates the disc D_ϕ meets the circle of radius s in an arc, and the angular integral over that arc is $h(P_1) - h(P_2)$, the difference of h at the two endpoints. That is a genuine simplification and it is where the official write-up passes through, but on its own it goes nowhere, because the endpoint values of an arbitrary h are arbitrary.

The bifurcation point is the sentence immediately after "the arc integral is $h(P_1) - h(P_2)$". The zombie tries to control that difference. You integrate over ϕ first, so that the arcs sweep out the full circle and the endpoint differences telescope to zero. In one phrase: do not evaluate the member, average the family.

Layer 2: HOW

Block 1. **Cue.** The specific combination $y\,\partial_x - x\,\partial_y$ applied to h. **Why natural.** Any first order expression $ah_x + bh_y$ is a directional derivative, and the question is which coordinate system turns it into a single partial derivative. The vector field $(y, -x)$ is tangent to the circles about the origin, so the natural coordinates are polar and the operator is $-\partial_\theta$. There are only two such combinations you will ever meet: $xh_x + yh_y = s\,h_s$ (Euler) and $yh_x - xh_y = -h_\theta$. Recognising which one you have is a lookup, not an insight. **Link to the next block.** A pure θ derivative integrates to a difference of endpoint values, so closed loops in θ give zero.

Block 2. **Cue.** h is a function on the plane, so it is 2π-periodic in θ. **Why natural.** This is the only place in the whole proof where a hypothesis about h is used, and it is the place where a counterexample would have to attack. Periodicity is the statement that h is single valued, and the fundamental theorem of calculus over a full period gives zero. If you tried $h = \theta$, which is not single valued, you would get $\rho = -1$ and the statement would fail, so this line is exactly the pivot. **Link to the next block.** You now own a family of sets on which ρ integrates to zero, namely the origin-centred circles. The given discs are not among them, so you need a device to move information from one family to the other.

Block 3. **Cue.** "There is a circle $\mathcal{S}$ such that the integral is zero", with no hint as to which. **Why natural.** When a problem asks for existence inside a continuum of candidates and the data are arbitrary, no construction can exist, so the proof must be non constructive. The two standard non constructive engines are averaging (the mean lies between the extremes) and topology (degree, fixed points). Here the family is a loop parametrised by one angle and the functional is real valued, so averaging plus the intermediate value theorem is the cheapest engine available. **Link to the next block.** To average, you need $I(\phi)$ as an explicit integral so that you may interchange the order of integration.

Block 4. **Cue.** A double integral over (P, ϕ) with a bounded integrand on a compact set. **Why natural.** Fubini applies with nothing to check. The point of interchanging is that after doing the ϕ integral first, the geometry of the individual disc disappears and only the total time each point P spends inside the moving disc remains. **Link to the next block.** That occupation function $M(P)$ is the object to inspect.

Block 5. **Cue.** The family $\{D_\phi\}$ is a single orbit of the rotation group acting on one disc. **Why natural.** If a family is closed under rotation, then averaging over the family produces a rotation invariant weight, and a rotation invariant function of the plane is a function of s alone. The explicit computation confirms it: membership is $\cos(\theta - \phi) \geq \lambda(s)$, and integrating over a full period of ϕ erases θ. Once M is radial, the double integral factors in polar coordinates into a radial integral times the angular integral of Block 2, which is zero. **Link to the next block.** The average of I is zero, and only a compactness fact is missing.

Block 6. **Cue.** A continuous real function on an interval with zero mean. **Why natural.** A continuous function with no zero has constant sign by the intermediate value theorem, and then its integral is nonzero. The continuity of I is not a formality worth skipping: it follows from uniform continuity of ρ on a compact disc containing all the D_ϕ, and a grader will want that sentence. **Link to the next block.** Existence of the circle.

Layer 3: WHY

The mathematical truth underneath is a statement about the operator ∂_θ. Its kernel consists of the radial functions, and its image, among continuous functions, consists exactly of the functions with zero mean on every circle about the origin. So the map $h \mapsto \rho$ produces only functions that are invisible to circular averaging, and the whole problem is a test of whether you can convert an arbitrary disc into a superposition of circular averages. The conversion device is the rotation group: the family of allowed discs is a single orbit, so integrating over the orbit turns any set into a radially symmetric weight, and a radially symmetric weight pairs a function against circular averages only. In the language you will meet later, $\rho\, dA$ is an exact form and the circle integrals detect exactly the failure of $d\theta$ to be exact; the mode $m = 0$ in the Fourier expansion of h in θ is the entire obstruction, and ∂_θ kills it.

The discernment mechanism has two components. First, when an expression is a first order combination of partial derivatives with variable coefficients, identify the flow of the coefficient vector field and change coordinates so that the expression becomes one partial derivative. The two coordinate systems that matter on the plane are Cartesian and polar, and $(y, -x)$ is the rotational field, so it is polar. Second, when the conclusion is existential over a continuum and the data are arbitrary, compute the mean of the functional over the continuum. If the mean is the target value, IVT ends it. This is the continuous form of the pigeonhole principle: at least one member is at most the average, and at least one is at least the average.

How to train it. Take every problem that says "there exists a position, an angle, a chord, a line, a partition, such that some quantity is zero or equal" and try averaging first, before any construction. The list of classical results with this shape is long: the ham sandwich theorem, the existence of a chord of a given length on a graph closing at both ends, the pancake theorem, the mean value theorem for integrals. In every one of them the object is produced, not built.

What the examiners nudged with, and how. They wrote ρ as $yh_x - xh_y$ rather than as $\operatorname{div}(h \cdot (y, -x))$ or as $-h_\theta$, because writing it in polar form would have given the game away. They said "prove or disprove" to make you spend time on counterexamples, and they set $d > r$ so that the discs avoid the origin, removing a distraction about the singularity of polar coordinates. They asked only for existence, which is the tell for averaging. And they assumed C^2, one degree more than needed, which is a deliberate false clue: a candidate who assumes the extra derivative is there to be used will look for a second order identity that does not exist.

Revolving concepts

Polar differential operators. With $x = s\cos\theta$, $y = s\sin\theta$: $\partial_s = \cos\theta\,\partial_x + \sin\theta\,\partial_y$ and $\partial_\theta = -s\sin\theta\,\partial_x + s\cos\theta\,\partial_y = -y\partial_x + x\partial_y$. Hence $xh_x + yh_y = s\,h_s$ (the Euler operator, which detects homogeneity: $sh_s = kh$ if and only if h is homogeneous of degree k) and $yh_x - xh_y = -h_\theta$. See Book 1, field: Calculus and integrals, topic: polar coordinates and differential operators.

Divergence-free fields and the divergence theorem. $V = (y, -x)$ has $\operatorname{div} V = 0$, so $\rho = \nabla h \cdot V = \operatorname{div}(hV) - h\operatorname{div} V = \operatorname{div}(hV)$, and $\iint_D \rho\, dA = \oint_{\partial D} hV \cdot n\, ds$. This is the second reading of the integrand and gives the alternative route. See Book 1, field: Calculus and integrals, topic: Green and divergence theorems.

The averaging existence principle. If J is a continuous real function on a connected parameter space and the average of J over that space is v, then J takes the value v somewhere. Proof: otherwise $J - v$ is continuous and never zero, so it has constant sign on the connected space, contradicting the vanishing of its mean. See Book 1, field: Real analysis, topic: intermediate value theorem and averaging.

Fubini and the occupation measure. For a family $\{K_g\}$ indexed by a compact group G with $K_g = gK$, the weight $M(P) = \int_G \mathbf{1}[P \in gK]\, dg$ is G invariant. For $G = SO(2)$ that means M is radial, and pairing a radial weight against ρ reduces to circular means of ρ. This is the general form of Step 3 and it is worth remembering with the group in it, because the same line proves the analogous statement for any bounded measurable region in place of the disc.

Fourier modes in the angle. Writing $h(s, \theta) = \sum_m c_m(s) e^{im\theta}$, the operator ∂_θ multiplies the m-th mode by im and annihilates $m = 0$; averaging over the centre angle ϕ multiplies the m-th mode by $\frac{1}{2\pi}\int_0^{2\pi} e^{im\phi} d\phi$, which vanishes for $m \neq 0$. The two annihilations cover all modes, which is the conceptual reason the average is zero.

Recursion. There is none in this problem: no sequence, no induction, no self-similar object. What replaces it is the group average, and the correct generalisation to store is the one stated above with G compact and M invariant, not a recursion.

Circuit Bank

- **WHEN** you see a first order expression $ah_x + bh_y$ with variable coefficients **DO** identify the flow of the field (a, b) and change to coordinates in which it is one partial derivative; for $(y, -x)$ that is polar and the expression is $-h_\theta$ **BECAUSE** a single partial derivative integrates by the fundamental theorem of calculus, and along a closed coordinate loop the result is zero when the potential is single valued
- **WHEN** you see an existence claim over a continuous family with arbitrary data **DO** compute the average of the functional over the family and apply the intermediate value theorem **BECAUSE** no construction can exist when the data are arbitrary, so the proof must be non constructive, and averaging is the cheapest non constructive engine for a real-valued functional
- **WHEN** you average a set-indexed functional over a group orbit of sets **DO** interchange the integrals and note that the resulting occupation weight is invariant under the group, hence constant on orbits **BECAUSE** invariance forces the weight to be a function of the orbit parameter alone, and pairing against an invariant weight sees only orbit averages of the integrand
- **WHEN** you are asked to prove or disprove and the only candidate counterexample needs a multivalued potential **DO** stop hunting and prove it, using single-valuedness as the pivot line **BECAUSE** the failure mode of the counterexample tells you exactly which hypothesis is load bearing, and a proof that uses that hypothesis once is usually short

Backtesting

1. First useful cue: the combination $yh_x - xh_y$, which is the angular derivative $-h_\theta$ in polar coordinates.
2. Method the expression suggested: circular means of ρ vanish by periodicity, so average the disc integral over the rotation family and finish with the intermediate value theorem.
3. Where you would hesitate, and the thought to reinforce: after seeing $\rho = -h_\theta$ you will want to attack a single disc, whose arcs give endpoint differences you cannot control. Reinforce this: when a single member is intractable and the family is a symmetry orbit, integrate over the family first.
4. The trap the question sets: the phrase "prove or disprove" invites a counterexample hunt, and the C^2 hypothesis invites you to look for a second derivative identity. Neither exists. Only C^1 and single-valuedness are used.
5. What makes you recognise it faster next time: the pairing of a rotational vector field with a rotation-symmetric family of candidate regions. Whenever the integrand is adapted to the same symmetry group as the family, the average over the family is the intended route.
6. The reusable decision rule: existence over a symmetric family equals zero average plus continuity. Never try to name the member.

Mirror drills

1. Show that for every continuously differentiable h on $\mathbb{R}^2$ and every $s > 0$, $\int_0^{2\pi} (yh_x - xh_y)(s\cos\theta, s\sin\theta)\, d\theta = 0$. *Answer:* the integrand is $-\frac{d}{d\theta} h(s\cos\theta, s\sin\theta)$ and h is 2π-periodic in θ, so the integral is $h(s,0) - h(s,2\pi) = 0$.
2. Show that the analogous statement for $\sigma = xh_x + yh_y$ is false. *Answer:* $\sigma = s\, h_s$ and s is not a periodic coordinate; take $h = \frac{1}{2}(x^2 + y^2)$, so $\sigma = x^2 + y^2 > 0$ off the origin and every

disc integral is strictly positive. This isolates periodicity of θ, not smoothness of h, as the load-bearing feature.

3. Let $K \subset \mathbb{R}^2$ be any bounded measurable set and let R_ϕ be rotation by ϕ about the origin. Show there is ϕ_0 with $\iint_{R_{\phi_0}K} \rho \, dA = 0$, provided $\phi \mapsto \iint_{R_\phi K} \rho$ is continuous. *Answer:* identical proof; the occupation weight $M(P) = \int_0^{2\pi} \mathbf{1}[P \in R_\phi K] \, d\phi$ is rotation invariant hence radial, so the averaged integral reduces to radial weights against circular means of ρ, all zero.

5.10 B4. A twisted product modulo a Fibonacci prime

Let $F_0, F_1, \ldots$ be the sequence of Fibonacci numbers, with $F_0 = 0$, $F_1 = 1$, and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$. For $m > 2$, let R_m be the remainder when the product $\prod_{k=1}^{F_m-1} k^k$ is divided by F_m. Prove that R_m is also a Fibonacci number.

Read it out loud

The object is $A = \prod_{k=1}^{N-1} k^k$ reduced modulo N, where $N = F_m$. Say that again slowly: it is a product over every nonzero residue class modulo N, but each factor is raised to a power that depends on the class. Call that a twisted product. Untwisted products over the full residue system are computable, because Wilson's theorem hands you $\prod_{k=1}^{p-1} k \equiv -1 \pmod{p}$; twisted ones are not, because the exponent k has no multiplicative meaning. Every route to this problem must therefore be a route that removes the twist.

The fields wired in are elementary number theory (Wilson's theorem, the fact that $\mathbb{Z}/p$ is a field so a quadratic has at most two roots, the criterion for -1 to be a quadratic residue) and the standard identity theory of the Fibonacci sequence (the addition formula, Cassini's identity, the sequence modulo 4). What is really being tested is whether you can read the two halves of the problem as one: the arithmetic half produces some square root of -1 modulo F_m, and the Fibonacci half must produce the same square root by a completely different route.

The first cue is planted in the statement itself and most candidates walk past it. The problem takes the trouble to say $F_0 = 0$. There is only one reason to normalise like that: 0 has to count as a Fibonacci number, so $R_m = 0$ must occur. And $R_m = 0$ means F_m divides the product, which happens for the cheapest possible reason when F_m is composite. So within thirty seconds you should have written: split on whether F_m is prime; if composite, $N = ab$ with $a, b < N$ makes $a^a b^b$ a factor of the product, so $R_m = 0 = F_0$ and that half is finished.

That leaves F_m prime, and now the second cue fires. Modulo a prime p the only computable products are the untwisted ones, and the only symmetry available on the exponent is the involution $k \mapsto p - k$, which turns k^k into $(-k)^{p-k}$. Multiplying the two versions of A makes the exponents add to the constant p, and a constant exponent is Wilson territory. So the plan is fixed before any computation: compute A^2, not A. That is why the answer can only be pinned down up to sign, and it is why the problem says R_m is *a* Fibonacci number rather than naming which one.

Approaches, ranked worst to best

1. **Reduce the exponents and multiply out.** *Cue:* k^k with k up to $F_m - 1$. *Circuit:* Fermat: $k^k \equiv k^{k \bmod (p-1)}$, then grind. (Why worst: it produces $p - 1$ unrelated powers with no structure, it does not generalise past $p = 13$, and it never touches the Fibonacci half of the problem at all. It is pure zombie arithmetic.)

2. **Discrete logarithm.** *Cue:* a product over the multiplicative group. *Circuit:* fix a primitive root g and write $A = g^E$ with $E = \sum_{k=1}^{p-1} k \cdot \mathrm{ind}_g(k)$. (Why second worst: it is a genuine idea, since it converts the product into a sum, but the sum couples the additive size k to the multiplicative index of k, and there is no closed form for such a mixed sum. It stalls immediately and for a reason worth naming: k as an exponent is additive data, k as a factor is multiplicative data, and no single coordinate system linearises both.)

3. **Pair k with its inverse.** *Cue:* Wilson's proof pairs k with k^{-1}. *Circuit:* $\prod k^k = \prod (k^{-1})^{k^{-1}}$ suitably reindexed. (Why third: the reflex is right, since Wilson's theorem is proved by pairing, but this particular pairing sends the exponent k to k^{-1}, which is again multiplicative data, so nothing simplifies. Knowing why the natural pairing fails is what points you at the correct one.)

4. **Pair k with $p-k$ and square.** *Cue:* the exponents k and $p-k$ add to the constant p. *Circuit:* $A \cdot \prod k^{p-k} = \prod k^p = ((p-1)!)^p$, and $\prod k^{p-k} \equiv \pm A$. (Why best: the involution $k \mapsto p-k$ is additive on exponents and multiplicative by -1 on bases, which is exactly the mismatch the problem is built from. One reindexing, one Wilson, and A^2 is a sign. The remaining work is to recognise that sign as -1 and to find a Fibonacci square root of -1 modulo F_m, which is Cassini's identity.)

The ideal solution

Write $N = F_m$ and $A = \prod_{k=1}^{N-1} k^k$, so that $R_m \equiv A \pmod N$ with $0 \le R_m < N$.

Small cases. For $m = 3$, $N = 2$ and $A = 1^1 = 1$, so $R_3 = 1 = F_1$. For $m = 4$, $N = 3$ and $A = 1^1 \cdot 2^2 = 4 \equiv 1$, so $R_4 = 1 = F_1$. From now on $m \ge 5$.

Case 1: N composite. Write $N = ab$ with $2 \le a \le b$; then $b = N/a \le N/2 < N$, so both a and b lie in $\{1, \dots, N-1\}$. If $a < b$ then a^a and b^b are two distinct factors of the product A, so ab divides $a^a b^b$, which divides A. If $a = b$ then $N = a^2$ with $a \ge 2$, and the single factor a^a of A is already divisible by $a^2 = N$. In both cases $N \mid A$, so $R_m = 0 = F_0$, a Fibonacci number.

Case 2: $N = p$ prime. First, m is odd. Indeed, the addition formula $F_{a+b} = F_a F_{b+1} + F_{a-1} F_b$ with $a = b = n$ gives

$$F_{2n} = F_n\big(F_{n+1} + F_{n-1}\big),$$

so if $m = 2n$ with $n \ge 3$ then $F_n \ge F_3 = 2$ and $F_{n+1} + F_{n-1} \ge F_4 + F_2 = 4$, making F_m composite. As $m \ge 5$, m must be odd.

Next, $p \equiv 1 \pmod 4$. The Fibonacci sequence modulo 4 begins $0, 1, 1, 2, 3, 1$ and then repeats, because $(F_6, F_7) \equiv (0, 1) \pmod 4$; so its period is 6. For odd m we have $m \equiv 1, 3, 5 \pmod 6$ and hence $F_m \equiv 1, 2, 1 \pmod 4$. Since $p = F_m \ge F_5 = 5$ is an odd prime, the value 2 is impossible, so $p \equiv 1 \pmod 4$.

Now reindex A by $k \mapsto p-k$:

$$A = \prod_{k=1}^{p-1} (p-k)^{\,p-k} \equiv \prod_{k=1}^{p-1} (-k)^{\,p-k} = (-1)^{\sum_{k=1}^{p-1}(p-k)} \prod_{k=1}^{p-1} k^{\,p-k} \pmod p.$$

Here $\sum_{k=1}^{p-1}(p-k) = \sum_{j=1}^{p-1} j = \frac{p(p-1)}{2}$, which is even because $p \equiv 1 \pmod 4$ makes $\frac{p-1}{2}$ even. Hence

$$A \equiv \prod_{k=1}^{p-1} k^{\,p-k} \pmod p.$$

Multiplying this by the definition of A collapses the exponents:

$$A^2 \equiv \prod_{k=1}^{p-1} k^k \cdot k^{\,p-k} = \prod_{k=1}^{p-1} k^p = \big((p-1)!\big)^p \equiv (-1)^p = -1 \pmod p,$$

using Wilson's theorem $(p-1)! \equiv -1 \pmod p$ and the fact that p is odd.

On the other hand, Cassini's identity $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$ with $n = m-1$ gives

$$F_m F_{m-2} - F_{m-1}^2 = (-1)^{m-1} = 1,$$

since m is odd. Reducing modulo $p = F_m$,

$$F_{m-1}^2 \equiv -1 \pmod p.$$

Therefore $A^2 \equiv F_{m-1}^2 \pmod p$, that is, $p \mid (A - F_{m-1})(A + F_{m-1})$. As p is prime, $A \equiv F_{m-1}$ or $A \equiv -F_{m-1} \pmod p$. In the second case

$$-F_{m-1} \equiv F_m - F_{m-1} = F_{m-2} \pmod p.$$

Finally $0 \le F_{m-2} < F_{m-1} < F_m = p$, so each of F_{m-1} and F_{m-2} is already reduced. Hence

$$R_m \in \{F_{m-1},\, F_{m-2}\},$$

a Fibonacci number in either case. Combining with Case 1 and the small cases, R_m is always a Fibonacci number. ■

Remark. Both alternatives really occur, so the ambiguity is not an artefact of the proof: $R_5 = 3 = F_4$, $R_7 = 8 = F_6$ and $R_{13} = 144 = F_{12}$ give F_{m-1}, while $R_{11} = 34 = F_9$ gives F_{m-2}. Deciding which one occurs is a genuinely harder question and the problem was written so that you do not have to.

Layer 1: WHAT

Block 1: read $F_0 = 0$ in the statement as permission for the answer 0, and split on the primality of F_m. Block 2: composite case, $N = ab$ makes N a factor of the product; $R_m = 0 = F_0$. Block 3: prime case, use $F_{2n} = F_n(F_{n+1} + F_{n-1})$ to force m odd. Block 4: Fibonacci modulo 4 has period 6, so $F_m \equiv 1$ or 2 for odd m, hence $p \equiv 1 \pmod 4$. Block 5: apply the involution $k \mapsto p - k$ to A; the sign is $(-1)^{p(p-1)/2} = +1$. Block 6: multiply the two expressions for A; exponents become the constant p; Wilson gives $A^2 \equiv -1$. Block 7: Cassini at $n = m - 1$ gives $F_{m-1}^2 \equiv -1$. Block 8: a quadratic over a field has at most two roots, so $A \equiv \pm F_{m-1}$, and $-F_{m-1} \equiv F_{m-2}$.

Juxtaposition. The nominal route sees k^k and reaches for Fermat to shrink the exponents, then multiplies out. It never produces a statement about F_m as a Fibonacci number, because it never produces an algebraic identity at all, only a number. The ideal route never computes A; it computes A^2, and only ever as a sign.

The bifurcation point is the line immediately after writing $A = \prod_{k=1}^{p-1} k^k$. The zombie writes "$k^k \equiv k^{k \bmod (p-1)}$". You write "replace k by $p - k$", because the exponent is additive data and $k \mapsto p - k$ is the only additive symmetry of the index set that is also multiplicative on the base, up to the sign -1.

Layer 2: HOW

Block 1. Cue. The statement fixes $F_0 = 0$ although the product is never obviously zero. **Why natural.** A normalisation stated but apparently unused is never decoration on the Putnam. If 0 is declared Fibonacci, then 0 is an intended value of R_m, so there are values of m for which F_m divides the product. Divisibility of $\prod_{k<N} k^k$ by N is immediate when N is composite, so the split on primality is handed to you by the wording. **Link to the next block.** The composite case is finished in two lines and you may forget it.

Block 2. Cue. $N = ab$ with $2 \le a \le b < N$. **Why natural.** Both a and b appear as indices in the product, so a^a and b^b appear as factors. The only care needed is the degenerate case $a = b$, which happens for genuine Fibonacci numbers ($F_{12} = 144 = 12^2$), and there a^a with $a \ge 2$ already contains a^2. **Link to the next block.** What remains is F_m prime, where the product is a unit and a genuine computation is required.

Block 3. Cue. A Fibonacci number that is prime. **Why natural.** Fibonacci divisibility is index divisibility: $F_a \mid F_b$ whenever $a \mid b$. The instance you need is the cheapest one, $F_{2n} = F_n(F_{n+1} + F_{n-1})$, which follows from the addition formula, itself the $(1,2)$ entry of $Q^{a+b} = Q^a Q^b$ for $Q = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$. So F_m prime and $m > 4$ forces m odd. **Link to the next block.** Oddness of m will be used twice: once for the sign in Cassini and once, through the value modulo 4, for the sign in the Wilson computation.

Block 4. **Cue.** You will need to know a parity, namely whether $\frac{p(p-1)}{2}$ is even. **Why natural.** That parity is decided by p modulo 4, so compute the Fibonacci sequence modulo 4. It is eventually periodic because the pair (F_n, F_{n+1}) modulo 4 determines all later pairs and there are finitely many pairs; the period is 6 by direct inspection. Since a Fibonacci prime above 2 cannot be even, the residue 2 is excluded and $p \equiv 1 \pmod 4$. **Link to the next block.** Both signs in the main computation are now determined.

Block 5. **Cue.** The exponent k in k^k is additive data attached to a multiplicative index. **Why natural.** The set $\{1, \ldots, p-1\}$ has exactly one useful involution that acts additively on the index and multiplicatively on the value: $k \mapsto p-k$, which is $k \mapsto -k$ in the group and $k \mapsto p-k$ in the exponent. Applying it to the whole product gives a second expression for A with the exponent reflected. Nothing here is a leap: it is the unique symmetry available. **Link to the next block.** You now have two expressions for the same A, whose exponents are k and $p-k$.

Block 6. **Cue.** Two exponents that sum to a constant. **Why natural.** Multiply the expressions. The product has every base raised to the same power p, so it is $((p-1)!)^p$, and Wilson's theorem evaluates it. The price of the symmetrisation is that you learn A^2, not A. That price is exactly the ambiguity in the final answer, which is why the problem asks only that R_m be *a* Fibonacci number. **Link to the next block.** $A^2 \equiv -1$, so A is a square root of -1 modulo p.

Block 7. **Cue.** You need a square root of -1 modulo F_m built from Fibonacci numbers, because the conclusion must be a Fibonacci number. **Why natural.** That is a search with exactly one hit. Among the standard identities, the only one expressing a Fibonacci square modulo F_m is Cassini, $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$, which at $n = m-1$ reads $F_{m-1}^2 \equiv (-1)^{m-1} \pmod{F_m}$. With m odd that is -1. Cassini itself is not an identity to memorise: it is $\det(Q^n) = (\det Q)^n$. **Link to the next block.** Two square roots of -1 in a field differ by a sign.

Block 8. **Cue.** $x^2 \equiv y^2$ in a field. **Why natural.** $\mathbb{Z}/p$ is a field, so $(A - F_{m-1})(A + F_{m-1}) \equiv 0$ forces one factor to vanish. Then $-F_{m-1} \equiv F_m - F_{m-1} = F_{m-2}$ by the defining recurrence, so the negative root is again a Fibonacci number, and both candidates already lie in $[0, p)$. **Link to the next block.** Done, in both cases.

Layer 3: WHY

The truth beneath the route is that in a finite abelian group you can evaluate the product of all elements, and nothing else. Wilson's theorem is the case $(\mathbb{Z}/p)^\times$ of the general fact that the product of all elements of a finite abelian group equals the product of its elements of order dividing 2, which here is -1 because ± 1 are the only such elements. A product in which the exponent varies with the element carries information that the group structure alone cannot see. So the entire strategy space consists of ways to convert a twisted product into an untwisted one, and there are only two kinds: pair the index with its additive negative, or pair it with its multiplicative inverse. Pairing with the inverse leaves the exponent scrambled; pairing with the additive negative makes the exponents sum to p. The choice is forced, not inspired.

The second truth is that Cassini's identity is a determinant. With $Q = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$ one has $Q^n = \begin{pmatrix} F_{n+1} & F_n \\ F_n & F_{n-1} \end{pmatrix}$, and taking determinants of both sides gives $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$ because the determinant is multiplicative. Once you see it that way, you also see the generalisation for free, and you stop treating Fibonacci identities as a list to be memorised.

The discernment mechanism is a two-sided reading of the conclusion. The problem promises that the answer belongs to a specific sequence. That promise is an instruction: whatever you compute about A, you must be able to match it against an identity of that sequence. Since the computation can only reach A^2, the identity you need must be about a Fibonacci *square* modulo F_m. There is exactly one such identity. So the search is finite and the target is known before you look. Train this by always asking, before starting a computation, what shape of statement the conclusion forces you to produce.

What the examiners nudged with, and how. They set $F_0 = 0$, which tells you 0 is admissible and therefore that the composite case exists. They wrote "R_m is also a Fibonacci number" rather than naming F_{m-1}, which tells you the method will lose a sign. They chose the exponent k to equal the base, which is the minimal way to make a product twisted while keeping the involution $k \mapsto p - k$ effective. And they restricted to $m > 2$ so that $F_m \geq 2$ and the product is nonempty.

Revolving concepts

Wilson's theorem. For a prime p, $(p-1)! \equiv -1 \pmod{p}$. Proof: pair each k with k^{-1}; the only self-paired elements are the roots of $x^2 = 1$, namely ± 1. Converse: if N is composite and $N > 4$ then $(N-1)! \equiv 0 \pmod{N}$. See Book 1, field: Number theory, topic: Wilson and factorial congruences.

Quadratic residue -1. For an odd prime p, the congruence $x^2 \equiv -1 \pmod{p}$ is solvable exactly when $p \equiv 1 \pmod{4}$, and then it has exactly two solutions, which are negatives of each other. One explicit root is $\left(\frac{p-1}{2}\right)!$, by Wilson. See Book 1, field: Number theory, topic: quadratic residues and Euler's criterion.

Fibonacci identities as matrix identities. With $Q = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, $Q^n = \begin{pmatrix} F_{n+1} & F_n \\ F_n & F_{n-1} \end{pmatrix}$. Then $Q^{a+b} = Q^a Q^b$ gives the addition formula $F_{a+b} = F_a F_{b+1} + F_{a-1} F_b$ and hence $F_{2n} = F_n(F_{n+1} + F_{n-1})$, while $\det Q^n = (\det Q)^n$ gives Cassini $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$. See Book 1, field: Sequences and recurrences, topic: linear recurrences in matrix form.

Divisibility structure. $\gcd(F_a, F_b) = F_{\gcd(a,b)}$, so for $a \geq 3$, $F_a \mid F_b$ if and only if $a \mid b$. Consequently F_m prime with $m > 4$ forces m prime, which is stronger than the oddness used above. See Book 1, field: Number theory, topic: strong divisibility sequences.

Pisano periods. The sequence $(F_n \bmod N)$ is purely periodic because the pair (F_n, F_{n+1}) determines the sequence in both directions and there are finitely many pairs modulo N. The period modulo 4 is 6, giving $0, 1, 1, 2, 3, 1$ repeating; this is what pins $p \equiv 1 \pmod{4}$.

The generalised recursion. For any second order recurrence $u_{n+1} = Pu_n - Qu_{n-1}$, put $M = \begin{pmatrix} P & -Q \\ 1 & 0 \end{pmatrix}$, so $\det M = Q$ and

$$u_{n+1}u_{n-1} - u_n^2 = Q^{n-1}\left(u_2 u_0 - u_1^2\right).$$

Fibonacci is $P = 1$, $Q = -1$, $u_0 = 0$, $u_1 = 1$, giving $(-1)^{n-1}(0 - 1) = (-1)^n$, which is Cassini. The same determinant argument gives Catalan's identity $F_n^2 - F_{n-s}F_{n+s} = (-1)^{n-s}F_s^2$ and d'Ocagne's identity. Store the determinant, not the list.

Circuit Bank

- **WHEN** you see a product over all nonzero residues with an exponent depending on the residue **DO** multiply the product by its image under $k \mapsto p - k$, so that the exponents add to the constant p, and evaluate with Wilson **BECAUSE** only untwisted products over a finite abelian group are computable, and $k \mapsto -k$ is the one involution that is additive on exponents and multiplicative on bases
- **WHEN** you arrive at $x^2 \equiv c$ modulo a prime **DO** find a second, structurally meaningful square root of c and conclude that x equals it up to sign **BECAUSE** a quadratic over a field has at

most two roots, so two square roots of the same element differ by a sign; the ambiguity is genuine information, not a defect

- **WHEN** the conclusion promises membership in a named sequence **DO** decide first what algebraic shape your computation can reach, then search that sequence's identities for exactly that shape **BECAUSE** here the computation can only reach a square, so only an identity about a Fibonacci square modulo F_m can close the gap, and Cassini is the unique candidate
- **WHEN** a problem states a normalisation that appears unused, such as $F_0 = 0$ **DO** treat it as a declaration that the corresponding value is an intended answer and look for the case that produces it **BECAUSE** examiners do not spend a clause on decoration; here it announces the composite case where the modulus divides the product
- **WHEN** you must show a modulus divides a product of many consecutive integers or their powers **DO** factor the modulus and locate the factors as separate terms of the product, checking the repeated-factor case **BECAUSE** $N = ab$ with $a, b < N$ puts a^a and b^b inside the product, and $N = a^2$ is covered because a^a contains a^2 once $a \geq 2$

Backtesting

1. First useful cue: the explicit $F_0 = 0$, which announces the composite case and the answer 0; then the twisted exponent in $\prod k^k$, which announces the involution $k \mapsto p - k$.
2. Method the expression suggested: symmetrise the exponent by reflection, square, apply Wilson, then match against Cassini.
3. Where you would hesitate, and the thought to reinforce: the two sign computations, $(-1)^{p(p-1)/2}$ and $(-1)^{m-1}$. Reinforce this: a sign that depends on p modulo 4 is decided by computing the sequence modulo 4, which is a six term calculation, not a research problem.
4. The trap the question sets: assuming F_m is prime and forgetting the composite case entirely, which is most of the values of m; and, in the composite case, forgetting that F_m can be a perfect square, as $F_{12} = 144$ is.
5. What makes you recognise it faster next time: any product over a full residue system with a nonconstant exponent. The reflex is reflect and multiply. Any appearance of a square root of -1 modulo p alongside a named integer sequence should send you straight to that sequence's Cassini-type identity.
6. The reusable decision rule: you cannot compute a twisted product, only its symmetrisation, so expect to determine the answer up to the ambiguity that the symmetrisation introduces, and read the problem statement to confirm that the ambiguity is tolerated.

Mirror drills

1. Derive Cassini's identity $F_{n+1}F_{n-1} - F_n^2 = (-1)^n$ without induction. *Answer:* take determinants in $Q^n = \begin{pmatrix} F_{n+1} & F_n \\ F_n & F_{n-1} \end{pmatrix}$, using $\det(Q^n) = (\det Q)^n = (-1)^n$.
2. Let p be a prime with $p \equiv 3 \pmod 4$. Show that $\prod_{k=1}^{p-1} k^k \equiv \pm 1 \pmod p$. *Answer:* the same reflection gives $A \equiv (-1)^{p(p-1)/2} \prod k^{p-k}$, and now $\frac{p-1}{2}$ is odd so the sign is -1; multiplying by A gives $A^2 \equiv -((p-1)!)^p \equiv -(-1) = 1$. (Check: $p = 7$ gives $A \equiv 1$.)
3. Show that if F_m is prime and $m > 4$ then m is prime. *Answer:* if $m = ab$ with $1 < a \leq b < m$ then $F_a \mid F_m$ by the addition formula applied repeatedly (or by $\gcd(F_a, F_m) = F_{\gcd(a,m)} = F_a$), and $1 < F_a < F_m$ since $a \geq 3$ when $m > 4$ is composite and not 4; handle $m = 2b$ directly through $F_{2n} = F_n(F_{n+1} + F_{n-1})$.

5.11 B5. Very odd matrices are acyclic digraphs

Say that an n-by-n matrix $A = (a_{ij})_{1 \le i,j \le n}$ with integer entries is very odd *if, for every nonempty subset S of $\{1, 2, \ldots, n\}$, the $|S|$-by-$|S|$ submatrix $(a_{ij})_{i,j \in S}$ has odd determinant. Prove that if A is very odd, then A^k is very odd for every $k \ge 1$.*

Read it out loud

Two words in the statement do all the work. The first is "odd". A determinant is a polynomial in the entries with integer coefficients, so it commutes with reduction modulo 2; "odd determinant" means "invertible over the field $\mathbb{F}_2$". Reduction also commutes with matrix multiplication, so $\overline{A^k} = \bar{A}^k$. Rewrite the whole problem over $\mathbb{F}_2$ and nothing is lost: A is very odd exactly when every principal submatrix of $\bar{A}$ is nonsingular, and the claim is that this property survives taking powers.

The second word is "every". The condition ranges over all nonempty subsets S, not just the initial segments $\{1, \ldots, r\}$. That distinction matters more than it looks. A condition on leading principal minors is destroyed by relabelling the indices; a condition on all principal minors is invariant under simultaneous permutation of the rows and the columns, that is, under conjugation by a permutation matrix. So the hypothesis is a permutation-conjugation invariant, and so is the conclusion, since $(P^{-1}AP)^k = P^{-1}A^kP$.

The fields wired in are linear algebra over a finite field, the combinatorics of directed graphs, and one clean induction. What is really being tested is a choice of strategy. You may either propagate the hypothesis through the product (compute principal minors of A^k from those of A) or classify the hypothesis (describe every very odd matrix) and then check that the description is stable under multiplication. Propagation needs the Cauchy and Binet expansion, which involves the non-principal minors $\det(A_{S,T})$ that the hypothesis says nothing about. Classification is the ideal route.

The first useful cue is the pair (invariant under permutation conjugation, must be closed under multiplication). Within thirty seconds you should have written the only family of matrices you know that is closed under multiplication, closed under passing to principal submatrices, and visibly nonsingular: the upper triangular matrices with 1 down the diagonal. That is the target normal form, and the whole problem becomes: show that over $\mathbb{F}_2$, very odd means permutation-conjugate to unitriangular.

Approaches, ranked worst to best

1. **Cauchy and Binet on A^2.** *Cue:* a minor of a product. *Circuit:* $\det\big((A^2)_{S,S}\big) = \sum_T \det(A_{S,T}) \det(A_{T,S})$ over all T with $|T| = |S|$. (Why worst: the identity is correct and the term $T = S$ contributes $\det(A_{S,S})^2 = 1$ modulo 2, which looks encouraging for about ten seconds. Then you notice that the hypothesis controls only the principal minors $\det(A_{T,T})$, while the sum involves the mixed minors $\det(A_{S,T})$ with $S \ne T$, about which you know nothing and for which no pairing or cancellation exists. This route is a dead end, and it is worth ten seconds precisely so that you learn what the hypothesis does not give you.)
2. **Induct on k through $A^k = A^{k-1}A$.** *Cue:* the conclusion is a statement for all k. *Circuit:* propagate the minor condition one factor at a time. (Why second worst: identical obstruction, since one factor still requires mixed minors. The failure is structural, not a matter of effort.)
3. **Unipotence.** *Cue:* a conjugation-invariant condition preserved by powers. *Circuit:* show $\bar{A} + I$ is nilpotent, which is preserved under multiplication. (Why third: this is the nearest wrong idea and the most instructive. Unipotence really is preserved by powers, and every very odd matrix is unipotent. But unipotence is strictly weaker than very odd: over $\mathbb{F}_2$ the matrix $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ equals $I + N$ with $N = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ and $N^2 = 0$, so it is unipotent, yet its

$(1,1)$ entry is 0 and it is not very odd. A conjugation invariant is too coarse; you need a permutation-conjugation invariant.)

4. **Classify very odd matrices over $\mathbb{F}_2$.** *Cue:* the condition is invariant under permutation conjugation and must be closed under multiplication. *Circuit:* very odd if and only if permutation-conjugate to unitriangular; encode the support as a digraph and prove it is acyclic. (Why best: once the classification is in hand the theorem is one line, because unitriangular matrices are closed under multiplication and every principal submatrix of one is again unitriangular. The classification itself costs a minimal-cycle argument and a topological ordering, both short.)

The ideal solution

Throughout, bars denote reduction modulo 2, and we work over the field $\mathbb{F}_2$. Since the determinant is a polynomial with integer coefficients in the entries, $\det(M)$ is odd if and only if $\det(\bar{M}) \neq 0$ in $\mathbb{F}_2$; and since reduction is a ring homomorphism, $\overline{A^k} = \bar{A}^k$. So it suffices to prove: if every principal submatrix of a matrix over $\mathbb{F}_2$ is nonsingular, the same holds for all its powers. Call such a matrix *very odd* over $\mathbb{F}_2$, and call a matrix *unitriangular* if it is upper triangular with every diagonal entry equal to 1.

For a permutation σ of $\{1,\dots,n\}$ let P_σ be the permutation matrix with $(P_\sigma)_{ij} = 1$ exactly when $i = \sigma(j)$; a direct computation gives $(P_\sigma^{-1} A P_\sigma)_{ij} = a_{\sigma(i)\sigma(j)}$.

Lemma 5.11.1. *For every σ and every nonempty S, $\det\big((P_\sigma^{-1} A P_\sigma)_S\big) = \det\big(A_{\sigma(S)}\big)$. In particular A is very odd if and only if $P_\sigma^{-1} A P_\sigma$ is.*

Proof. The submatrix of $(a_{\sigma(i)\sigma(j)})$ indexed by S has entries $a_{\sigma(i)\sigma(j)}$ for $i, j \in S$, which is the submatrix of A indexed by $\sigma(S)$ after applying the same bijection $S \to \sigma(S)$ to its rows and to its columns. Permuting rows and columns by the same permutation multiplies the determinant by the square of a sign, hence leaves it unchanged. □

Lemma 5.11.2. *If U is unitriangular then U^k is unitriangular for every $k \geq 1$, and every unitriangular matrix is very odd.*

Proof. A product of upper triangular matrices is upper triangular and its diagonal is the entrywise product of the diagonals, so U^k is unitriangular. If $S = \{i_1 < \cdots < i_r\}$ then $(U^k)_S$ is upper triangular with diagonal entries 1, because the ordering inherited from $\{1,\dots,n\}$ is the same one; hence $\det\big((U^k)_S\big) = 1 \neq 0$. □

Lemma 5.11.3. *If A is very odd over $\mathbb{F}_2$ then there is a permutation σ for which $P_\sigma^{-1} A P_\sigma$ is unitriangular.*

Proof. Taking $S = \{i\}$ gives $a_{ii} = 1$ for every i. Let G be the loopless directed graph on the vertex set $\{1,\dots,n\}$ with an edge $i \to j$ exactly when $i \neq j$ and $a_{ij} = 1$.

G has no directed cycle. Suppose it has one, and choose a directed cycle $i_1 \to i_2 \to \cdots \to i_M \to i_1$ of minimal length M. Since G is loopless, $M \geq 2$, and the vertices $i_1,\dots,i_M$ are distinct: a repetition $i_a = i_b$ with $a < b$ would produce a closed directed walk of length $b-a$ with $2 \leq b-a < M$, contradicting minimality. Minimality also forbids chords. Precisely, suppose $a_{i_j i_k} = 1$ with $k \not\equiv j$ and $k \not\equiv j+1 \pmod{M}$. Then the edge $i_j \to i_k$ followed by the arcs of the cycle from i_k round to i_j is a directed cycle of length $1 + \big((j-k) \bmod M\big)$, and the hypotheses on k give $1 \leq (j-k) \bmod M \leq M-2$, so this cycle has length at least 2 and at most $M-1$, contradicting minimality. Together with $a_{ii} = 1$ we conclude, for the principal submatrix B of A on $S = \{i_1,\dots,i_M\}$ with rows and columns listed in the cyclic order,

$$b_{jk} = \begin{cases} 1, & k \equiv j \text{ or } k \equiv j+1 \pmod{M}, \\ 0, & \text{otherwise.} \end{cases}$$

Every row of B then has an even number of entries equal to 1: exactly two when $M \geq 3$, and for $M = 2$ every entry is 1, so again two. Hence $B\mathbf{1} = \mathbf{0}$ over $\mathbb{F}_2$, where $\mathbf{1}$ is the all-ones column, so $\det B = 0$. This contradicts very oddness, so G is acyclic.

An acyclic digraph has a topological ordering. This is the standard fact about directed acyclic graphs, and it has a two-line proof. A finite acyclic digraph has a vertex with no incoming edge: take a directed path of maximal length and let v be its initial vertex; an edge $u \to v$ would close a cycle if u lay on the path, and would otherwise lengthen the path, and both are impossible. Delete v and induct on the number of vertices, obtaining a listing $v_1, \ldots, v_n$ of the vertices such that every edge $v_i \to v_j$ has $i < j$.

Let $\sigma(i) = v_i$ and $U = P_\sigma^{-1} A P_\sigma$, so $u_{ij} = a_{v_i v_j}$. If $i > j$ then $a_{v_i v_j} = 1$ would be an edge $v_i \to v_j$ against the ordering, so $u_{ij} = 0$; and $u_{ii} = a_{v_i v_i} = 1$. Thus U is unitriangular. □

Theorem 5.11.4. *If A is very odd then A^k is very odd for every $k \geq 1$.*

Proof. Reduce modulo 2; then $\bar{A}$ is very odd over $\mathbb{F}_2$. By Lemma 3 there is a permutation σ with $U = P_\sigma^{-1} \bar{A} P_\sigma$ unitriangular. Then $P_\sigma^{-1} \bar{A}^k P_\sigma = U^k$, which is unitriangular by Lemma 2 and therefore very odd; by Lemma 1, $\bar{A}^k = \overline{A^k}$ is very odd. Hence every principal minor of A^k is odd. □

Remark. Two features are sharp. The condition really needs all subsets, not the small ones: the $n \times n$ circulant with first row $(1, 1, 0, \ldots, 0)$ has every proper principal minor odd but even determinant, because its digraph is a single n-cycle. And parity, not nonvanishing, is what is closed under powers: over $\mathbb{Q}$ the matrix $\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$ has all principal minors nonzero, yet its square $\begin{pmatrix} 0 & 2 \\ -2 & 0 \end{pmatrix}$ has vanishing diagonal entries. That matrix has det $= 2$, so it is not very odd, and the example shows the reduction modulo 2 is doing real work rather than being a convenience.

Layer 1: WHAT

Block 1: read "odd determinant" as "invertible over $\mathbb{F}_2$" and move the entire problem there. Block 2: read "for every nonempty subset" as "invariant under permutation conjugation", which tells you which normal forms are legitimate. Block 3: choose the target normal form by asking what family is closed under multiplication, closed under principal submatrices, and visibly nonsingular; unitriangular is the unique cheap answer. Block 4: prove the easy direction, that unitriangular is very odd and stays unitriangular under powers. Block 5: encode A as a digraph on its off-diagonal support, having first noted $a_{ii} = 1$. Block 6: show a minimal directed cycle would give a principal submatrix whose rows all have even weight, hence singular. Block 7: acyclic gives a topological order, which is exactly a relabelling making A unitriangular. Block 8: assemble.

Juxtaposition. The nominal route reduces modulo 2 and then asks for $\det\big((A^2)_S\big)$ in terms of the minors of A, reaching for Cauchy and Binet. It fails because the hypothesis is about principal minors only and the expansion introduces mixed minors.

The bifurcation point is the sentence right after the reduction to $\mathbb{F}_2$. The zombie asks "what is a principal minor of A^k?". You ask "what does a very odd matrix look like?". In one phrase: when a property is claimed to be closed under an operation, classify the property instead of propagating it.

Layer 2: HOW

Block 1. Cue. The word "odd" applied to a determinant of an integer matrix. **Why natural.** Determinant is a polynomial with integer coefficients in the entries, so it commutes with any ring homomorphism, in particular with reduction modulo 2; and reduction is multiplicative, so it commutes with taking powers. Nothing about the problem is lost and the coefficient ring becomes a field with exactly one nonzero element, which turns every linear algebra statement into a statement

about the $0/1$ pattern. **Link to the next block.** Over $\mathbb{F}_2$ the hypothesis reads: every principal submatrix is nonsingular.

Block 2. Cue. "For every nonempty subset S", rather than for initial segments. **Why natural.** Simultaneously permuting the rows and the columns permutes the family of principal submatrices among themselves and preserves each determinant, because the sign of the permutation appears twice. So the hypothesis and the conclusion are invariants of A up to conjugation by a permutation matrix, and any normal form you aim for is allowed to use exactly that freedom and no more. This is also why the third-ranked approach fails: unipotence is invariant under all conjugations, so it cannot see the extra information carried by the index labels. **Link to the next block.** The target must be a normal form under permutation conjugation.

Block 3. Cue. The conclusion is about all powers. **Why natural.** You need a family $\mathcal{F}$ of matrices, stable under multiplication, each of whose members is very odd, and containing every very odd matrix. Stability under multiplication plus closure under principal submatrices plus visible nonsingularity is a demanding list, and there is exactly one standard family that meets it: upper triangular with unit diagonal. This is a lookup from your stock of matrix families, not a discovery. **Link to the next block.** Half of the classification, that unitriangular implies very odd and stays so under powers, is now immediate; the work is in the converse.

Block 4. Cue. $S = \{i\}$ is a legal subset. **Why natural.** The smallest instances of a hypothesis quantified over all subsets are always the first thing to extract, and here they give the diagonal outright: $a_{ii} = 1$ for every i. That is precisely the diagonal condition of unitriangularity, so what remains is only to arrange the off-diagonal zeros by relabelling. **Link to the next block.** The off-diagonal support is now the only unknown, and a $0/1$ pattern on ordered pairs is a directed graph.

Block 5. Cue. A $0/1$ matrix over $\mathbb{F}_2$ with a distinguished diagonal. **Why natural.** Encode the off-diagonal ones as edges. Under this dictionary a principal submatrix is an induced subgraph, so the hypothesis becomes a statement about every induced subgraph, and the target (unitriangular after relabelling) becomes exactly the statement that the digraph admits a topological ordering. Translating a linear algebra hypothesis into graph language is worthwhile precisely when the hypothesis is closed under passing to induced substructures. **Link to the next block.** You now need: no directed cycles.

Block 6. Cue. A cycle is the smallest obstruction to an ordering, so test cycles first. **Why natural.** Take a shortest cycle, because minimality is what forces the induced submatrix to be exactly the cyclic pattern with no chords. Then every row has exactly two ones (and for $M = 2$ the two coincide with the whole row), so the all-ones vector lies in the kernel over $\mathbb{F}_2$. Over $\mathbb{F}_2$ "all rows have even weight" is the cheapest singularity certificate in existence, and you should reach for it whenever a $0/1$ matrix has a regular pattern. **Link to the next block.** Acyclicity established.

Block 7. Cue. Acyclic finite digraph. **Why natural.** Every finite acyclic digraph has a vertex of indegree zero, found at the start of a maximal path, and deleting it preserves acyclicity; induction produces a listing along which all edges point forwards. That listing is a permutation σ, and $P_\sigma^{-1} A P_\sigma$ has zeros strictly below the diagonal by construction. **Link to the next block.** Combine with the easy direction to finish.

Layer 3: WHY

The truth beneath the route is a structure theorem: over $\mathbb{F}_2$, a matrix has all principal minors nonzero if and only if it is conjugate by a permutation matrix to a unitriangular matrix. The reason such a clean classification exists is that over $\mathbb{F}_2$ a matrix is nothing but its support, so a determinant condition becomes a combinatorial condition, and the combinatorial condition is exactly acyclicity of the support digraph. Over any larger field the classification collapses. The rational matrix $\begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$ has all principal minors nonzero yet is not permutation-triangularisable, and its square has vanishing diagonal entries, so the analogous statement over $\mathbb{Q}$ is false. That single example is the best possible argument that the reduction modulo 2 is the substance of the problem and not a convenience.

The second truth is about strategy. There are two ways to prove that a property is preserved by an operation. Propagation asks how the property transforms, and needs a formula relating the invariants of A^k to those of A. Classification describes the set of objects with the property and then checks the set is closed. Propagation is the default reflex and it is correct when the invariant has an addition formula (traces, ranks, characteristic polynomials). It is hopeless when the hypothesis constrains only part of a family of invariants, as here: the principal minors of a product involve non-principal minors of the factors, which are unconstrained. Learn to detect that mismatch early; it is what sends you to classification.

The discernment mechanism, stated as a test you can run in the arena: write down the group of transformations under which the hypothesis is invariant. If that group is the full conjugation group, look for eigenvalue or characteristic polynomial invariants. If it is only the permutation subgroup, then the labels carry information and you should look for a combinatorial encoding, almost always a graph. Here the quantifier over all subsets S is what pins the group down to permutations, so the graph is announced by the statement itself.

How to train it. Take every hypothesis of the form "for every subset" and immediately extract the singletons, then the pairs, then ask what structure on the index set the hypothesis describes. Singletons gave the diagonal; the pairs give that no 2-cycle occurs, since $\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ is singular; the general case is the cycle argument. The problem is built so that the small cases point straight at the general one.

What the examiners nudged with. They chose the word "odd" rather than "nonzero", which is the whole reduction. They quantified over all nonempty subsets rather than initial segments, which is what makes permutation conjugation the right symmetry and rules out the leading-minor conditions of LU factorisation. They asked for all $k \geq 1$ rather than $k = 2$, so that a propagation argument tuned to squares would not be enough and you would be pushed towards a normal form. And they named the property with a nonsense phrase, "very odd", to stop you searching your memory for a theorem and to force you to classify from scratch.

Revolving concepts

Reduction modulo a prime. For any ring homomorphism φ, $\varphi(\det M) = \det(\varphi M)$ and $\varphi(MN) = \varphi(M)\varphi(N)$, since both determinant and product are polynomial in the entries. Hence integer statements about parity of determinants are statements of linear algebra over $\mathbb{F}_2$. See Book 1, field: Matrices and determinants, topic: determinants over rings and reduction.

Principal versus leading principal minors. All leading principal minors nonzero is equivalent to the existence of an LU factorisation without pivoting; all principal minors nonzero is a strictly different and permutation-invariant condition. Confusing the two is the standard error here. See Book 1, field: Linear algebra, topic: LU factorisation and minors.

Cauchy and Binet. For $|S| = |T| = r$, $\det\big((MN)_{S,T}\big) = \sum_{|R|=r} \det(M_{S,R}) \det(N_{R,T})$. It is the correct tool when you know all minors of the factors, and the wrong tool when you know only the

principal ones. See Book 1, field: Matrices and determinants, topic: Cauchy and Binet formula.

Triangular matrices. Upper triangular matrices are closed under products and the diagonal multiplies entrywise; principal submatrices of an upper triangular matrix are upper triangular in the inherited order; U unitriangular means $U - I$ is strictly upper triangular, hence nilpotent with $(U - I)^n = 0$, so $U^k = \sum_{j=0}^{n-1} \binom{k}{j} (U - I)^j$ is a finite sum. See Book 1, field: Linear algebra, topic: triangular and unipotent matrices.

Even weight rows over $\mathbb{F}_2$. If every row of a matrix over $\mathbb{F}_2$ has an even number of ones, the all-ones vector is in the kernel and the determinant is zero. The circulant with first row $(1, 1, 0, \ldots, 0)$ is the canonical example: even determinant, all proper principal minors odd. See Book 1, field: Discrete structures and graphs, topic: parity arguments over $\mathbb{F}_2$.

Digraphs, cycles and topological order. A finite directed graph is acyclic if and only if its vertices admit a linear ordering along which every edge points forwards. The proof is the peeling induction: an acyclic finite digraph has a vertex of indegree zero (the start of a maximal path), delete and recurse. See Book 1, field: Discrete structures and graphs, topic: directed acyclic graphs.

The generalised recursion. The peeling induction above is the general template for finite ordered structures and deserves to be stored in its abstract form: every finite strict partial order extends to a linear order, proved by repeatedly extracting a minimal element. The same induction proves that a finite nilpotent-plus-identity pattern can be triangularised, that a finite acyclic dependency graph can be evaluated in order, and that any well-founded recursion on a finite set terminates. Whenever you must convert "no cycles" into "there is an order", this is the induction to write.

Circuit Bank

- **WHEN** you see a parity condition on determinants of integer matrices **DO** reduce the entire problem modulo 2 and work over $\mathbb{F}_2$ **BECAUSE** determinant and product are polynomial in the entries, so they commute with reduction, and over $\mathbb{F}_2$ a matrix is determined by its support, turning linear algebra into combinatorics
- **WHEN** you are asked to show a property is preserved by an operation and the hypothesis constrains only part of a family of invariants **DO** classify the objects with the property instead of propagating the property through the operation **BECAUSE** propagation needs a formula expressing the invariants of the image in terms of those of the source, and here the minors of a product involve non-principal minors that the hypothesis does not control
- **WHEN** a hypothesis is quantified over every subset of the index set **DO** note that the hypothesis is invariant under simultaneous permutation of rows and columns, and look for a normal form up to permutation conjugation, hence for a combinatorial encoding **BECAUSE** a full conjugation invariant such as unipotence is too coarse; the labels carry information and the natural carrier of labelled information is a graph
- **WHEN** you have a $0/1$ matrix over $\mathbb{F}_2$ with a regular pattern **DO** check the parity of the row weights; if every row has even weight, the all-ones vector is a kernel certificate **BECAUSE** it is the shortest possible singularity proof over $\mathbb{F}_2$ and it is exactly what a directed cycle produces
- **WHEN** you must convert "no cycles" into "there is an ordering" **DO** peel a vertex of indegree zero, found at the start of a maximal directed path, and induct **BECAUSE** acyclicity is preserved by deletion, and a maximal path cannot be extended or closed, so its initial vertex has no incoming edge

Backtesting

1. First useful cue: "odd determinant" for integer matrices, which is $\mathbb{F}_2$; then "every nonempty subset", which is permutation-conjugation invariance.

2. Method the expression suggested: classify very odd matrices as the permutation conjugates of unitriangular matrices, then note that unitriangularity is closed under powers.
3. Where you would hesitate, and the thought to reinforce: after the reduction you will want to compute a minor of A^2. Reinforce this: check first whether the hypothesis controls the objects the expansion produces. It does not, so classify instead.
4. The trap the question sets: settling for unipotence, which is preserved by powers and is implied by the hypothesis, but is strictly weaker, as $\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}$ over $\mathbb{F}_2$ shows. A second trap is verifying only small subsets S; the circulant with first row $(1, 1, 0, \ldots, 0)$ has all proper principal minors odd and even determinant.
5. What makes you recognise it faster next time: a determinant condition over $\mathbb{F}_2$ indexed by subsets is a statement about induced subgraphs of the support digraph. Whenever both are present, draw the graph.
6. The reusable decision rule: closure statements are proved by classification when the hypothesis is a normal form condition, and by propagation only when the invariant has a genuine addition or multiplication formula.

Mirror drills

1. Show that a square matrix over $\mathbb{F}_2$ in which every row has an even number of entries equal to 1 is singular. *Answer:* the all-ones column vector lies in the kernel, since each coordinate of the product is the row sum modulo 2.
2. Show that the $n \times n$ circulant C with $c_{ij} = 1$ exactly when $j - i \equiv 0$ or $1 \pmod{n}$ has even determinant, while every proper principal minor of C is odd. *Answer:* every row has weight 2, so $\det C \equiv 0$; deleting any index breaks the single n-cycle of the support digraph, leaving a disjoint union of directed paths, which is acyclic, so the induced submatrix is permutation-conjugate to unitriangular and has odd determinant.
3. Exhibit an integer matrix all of whose principal minors are nonzero but for which some principal minor of A^2 vanishes, and say why the problem is not contradicted. *Answer:* $A = \begin{pmatrix} 1 & 1 \\ -1 & 1 \end{pmatrix}$ has principal minors $1, 1, 2$ and $A^2 = \begin{pmatrix} 0 & 2 \\ -2 & 0 \end{pmatrix}$ has zero diagonal. No contradiction: $\det A = 2$ is even, so A is not very odd. Parity, not nonvanishing, is the property that is closed under powers.

5.12 B6. Iterated median trimming and quantile coupling

Given an ordered list of $3N$ real numbers, we can trim *it to form a list of N numbers as follows: We divide the list into N groups of 3 consecutive numbers, and within each group, discard the highest and lowest numbers, keeping only the median.*

Consider generating a random number X by the following procedure: Start with a list of 3^{2021} numbers, drawn independently and uniformly at random between 0 and 1. Then trim this list as defined above, leaving a list of 3^{2020} numbers. Then trim again repeatedly until just one number remains; let X be this number. Let μ be the expected value of $\left|X - \frac{1}{2}\right|$. Show that

$$\mu \geq \frac{1}{4}\left(\frac{2}{3}\right)^{2021}.$$

Read it out loud

Strip the story. The list is a complete ternary tree of depth 2021: the leaves are independent uniforms, every internal node is the median of its three children, and X is the root. Because the triples are disjoint and the leaves independent, the numbers at any fixed level are independent and identically distributed. So the state of the process is not a random variable, it is a distribution, and each round is a deterministic map on distributions: the distribution of the median of three independent copies.

The fields wired in are order statistics (the law of the median of three), iterated maps of an interval and their fixed points, the quantile transform and monotone coupling, and one elementary calculus inequality. What is really being tested is whether you can turn a recursion on distributions into a recursion on a single number. Everything else in the problem is bookkeeping.

The first cue is that the operation is the same at every level and the inputs at every level are independent and identically distributed, so within thirty seconds you should have written the cdf recursion. If Y_1, Y_2, Y_3 are independent with common cdf F then their median is at most x exactly when at least two of the three are, so the median has cdf

$$3F^2(1-F) + F^3 = 3F^2 - 2F^3 = \phi(F), \qquad \phi(u) = 3u^2 - 2u^3.$$

Therefore the cdf after k rounds is $F_k = \phi^{\circ k}$, the k-fold iterate of one cubic, applied to $F_0(x) = x$.

The second cue is in the constants the question hands you, and this is the part almost nobody reads. $\frac{1}{4}$ is $E\left|U - \frac{1}{2}\right|$ for a uniform U, that is, μ_0. And ϕ fixes $\frac{1}{2}$ with $\phi'(\frac{1}{2}) = 6 \cdot \frac{1}{2} - 6 \cdot \frac{1}{4} = \frac{3}{2}$, so $\frac{2}{3} = 1/\phi'(\frac{1}{2})$. The inequality you are asked to prove is precisely

$$\mu_{2021} \;\geq\; \mu_0 \cdot \Big(\frac{1}{\phi'(1/2)}\Big)^{2021},$$

which announces both the method (one round at a time, losing at most the factor $1/\phi'(\frac{1}{2})$ each time) and the reason it is true (the multiplier of the iterated map at its unstable fixed point is the fastest possible rate of concentration). Read the constants; they are the statement of the lemma you must prove.

A caveat worth clarifying. Trimming concentrates: $\phi(u) < u$ on $(0, \frac{1}{2})$, so F_k decreases there and X_k moves towards $\frac{1}{2}$. The problem asks you to prove concentration is not too fast, which is a lower bound on a shrinking quantity, and lower bounds on shrinking quantities are where sloppy arguments die.

Approaches, ranked worst to best

1. **Compute F_k explicitly.** *Cue:* $F_k = \phi^{\circ k}$ is a polynomial. *Circuit:* iterate the cubic and estimate. (Why worst: $\phi^{\circ k}$ has degree 3^k, so for $k = 2021$ there is no expression to write down, and the coefficients carry no usable structure. This is zombie computation dressed as a plan.)

2. **Moment recursion.** *Cue:* bound $E\big[(X_k - \frac{1}{2})^2\big]$ and convert. *Circuit:* variance shrinks by a computable factor. (Why second worst: the moments of a median are not determined by the moments of the input, so there is no closed recursion; and even with a bound on the second moment you would need a reverse inequality such as Paley and Zygmund to descend to $E|X_k - \frac{1}{2}|$, because Cauchy and Schwarz points the wrong way.)
3. **Density comparison against a box.** *Cue:* $f_k(\frac{1}{2}) = (\frac{3}{2})^k$ and f_k is nondecreasing on $[0, \frac{1}{2}]$. *Circuit:* compare f_k with the density that is $(\frac{3}{2})^k$ on an interval of length $\frac{1}{2}(\frac{2}{3})^k$ and zero elsewhere, then integrate by parts. (Why third: correct, and it is the second official solution. But it needs a separate induction proving f_k nondecreasing on $[0, \frac{1}{2}]$, plus a rearrangement style comparison of two densities with the same total mass. Three moving parts where one suffices.)
4. **Cumulative distribution with the axes transposed.** *Cue:* $\mu_k = 2\int_0^{1/2} F_k$ by parts, and the area under a monotone curve equals the area to the left of its inverse. *Circuit:* turn μ_k into $2\int_0^{1/2}(\frac{1}{2} - F_k^{-1}(y))\,dy$ and iterate. (Why second: correct, and it is the first official solution. It contains the right idea, namely that the iteration must be moved into the argument, but it hides that idea inside an area transposition instead of naming it.)
5. **Quantile coupling.** *Cue:* a recursion on distributions with a monotone map. *Circuit:* realise every X_k on one probability space as $\psi^{\circ k}(U)$ with $\psi = \phi^{-1}$, so the distributional recursion becomes a pointwise one. (Why best: the distributional recursion becomes an identity between random variables, the required inequality becomes one calculus check on ψ, and taking expectations is free. Nothing needs to be interchanged, integrated by parts, or induced separately.)

The ideal solution

Write $\phi(u) = 3u^2 - 2u^3$ for $u \in [0,1]$.

Step 1: the model. For $k \geq 0$ let X_k be the number that survives when the trimming procedure is applied repeatedly to a list of 3^k independent uniform variables on $[0,1]$, so that $X = X_{2021}$, and put $\mu_k = E\big|X_k - \frac{1}{2}\big|$, so that $\mu = \mu_{2021}$. Let F_k be the cumulative distribution function of X_k; thus $F_0(x) = x$ on $[0,1]$.

If Y_1, Y_2, Y_3 are independent with common cdf F, their median is at most x if and only if at least two of them are, so its cdf is $3F(x)^2\big(1 - F(x)\big) + F(x)^3 = \phi\big(F(x)\big)$. After one round of trimming, the 3^{k-1} surviving numbers are medians of pairwise disjoint triples of independent variables, hence are themselves independent and identically distributed with cdf $\phi \circ F_0$. Applying this at every level,

$$F_k = \phi \circ F_{k-1}, \qquad \text{so} \qquad F_k = \phi^{\circ k} \text{ on } [0,1].$$

Step 2: properties of ϕ. On $[0,1]$ we have $\phi'(u) = 6u(1-u) \geq 0$, with strict positivity on $(0,1)$, and $\phi(0) = 0$, $\phi(1) = 1$. Hence $\phi : [0,1] \to [0,1]$ is a continuous, strictly increasing bijection. Also

$$\phi(1-u) = 3(1-u)^2 - 2(1-u)^3 = 1 - 3u^2 + 2u^3 = 1 - \phi(u),$$

so $\phi\big(\frac{1}{2}\big) = \frac{1}{2}$ and ϕ maps $[0, \frac{1}{2}]$ onto itself. Consequently each $F_k = \phi^{\circ k}$ is a continuous strictly increasing bijection of $[0,1]$ with $F_k(1-x) = 1 - F_k(x)$ and $F_k(\frac{1}{2}) = \frac{1}{2}$.

Let $\psi = \phi^{-1} : [0,1] \to [0,1]$, again a continuous strictly increasing bijection, with $\psi(\frac{1}{2}) = \frac{1}{2}$ and $\psi(1-u) = 1 - \psi(u)$ (apply ψ to the displayed symmetry of ϕ).

Step 3: the coupling. Let U be uniform on $[0,1]$ and set $Y_k = \psi^{\circ k}(U)$, so $Y_0 = U$ and $Y_k = \psi(Y_{k-1})$. Since F_k is a continuous strictly increasing bijection of $[0,1]$ and $\psi^{\circ k} = F_k^{-1}$,

$$P(Y_k \leq x) = P\big(U \leq F_k(x)\big) = F_k(x),$$

so Y_k has the same distribution as X_k, and in particular $\mu_k = E\big|Y_k - \frac{1}{2}\big|$. The point of this construction is that the recursion, which was a relation between distributions, is now an identity between random variables on one probability space.

Step 4: the key inequality. We claim that

$$\left|\psi(u)-\tfrac12\right| \;\ge\; \tfrac23\left|u-\tfrac12\right| \qquad \text{for all } u\in[0,1].$$

By the symmetry $\psi(1-u)=1-\psi(u)$ both sides are unchanged under $u\mapsto 1-u$, so it suffices to prove it for $u\in[0,\tfrac12]$, where $\psi(u)\le\tfrac12$ and the claim reads $\tfrac12-\psi(u)\ge\tfrac23\big(\tfrac12-u\big)$. Since ϕ maps $[0,\tfrac12]$ bijectively onto itself, write $u=\phi(x)$ with $x=\psi(u)\in[0,\tfrac12]$; the claim becomes

$$\tfrac12-x \;\ge\; \tfrac23\Big(\tfrac12-\phi(x)\Big), \qquad \text{that is} \qquad \tfrac12-\phi(x) \;\le\; \tfrac32\Big(\tfrac12-x\Big) \quad (x\in[0,\tfrac12]).$$

Put $g(x)=\tfrac32\big(\tfrac12-x\big)-\big(\tfrac12-\phi(x)\big)$. Then $g(\tfrac12)=0$ and

$$g'(x)=-\tfrac32+\phi'(x)=-\tfrac32+6x-6x^2=-\tfrac32\big(1-2x\big)^2 \;\le\; 0,$$

so g is nonincreasing on $[0,\tfrac12]$ and therefore $g(x)\ge g(\tfrac12)=0$ there. This proves the claim.

Step 5: iterate and take expectations. By Step 3 and Step 4, pointwise on the probability space,

$$\left|Y_k-\tfrac12\right| = \left|\psi(Y_{k-1})-\tfrac12\right| \;\ge\; \tfrac23\left|Y_{k-1}-\tfrac12\right|,$$

so by induction $\big|Y_k-\tfrac12\big|\ge\big(\tfrac23\big)^k\big|U-\tfrac12\big|$ for every $k\ge0$. Taking expectations and using $E\big|U-\tfrac12\big| = \int_0^1\big|u-\tfrac12\big|\,du=\tfrac14$,

$$\mu_k = E\left|Y_k-\tfrac12\right| \;\ge\; \left(\tfrac23\right)^k\cdot\tfrac14.$$

Taking $k=2021$ gives $\mu\ge\tfrac14\big(\tfrac23\big)^{2021}$, as required. ■

Remark on sharpness. The only loss is the inequality $\phi'\le\phi'(\tfrac12)$ on $[0,\tfrac12]$, which is an equality only at the fixed point, so the exponential rate is exactly right and only a bounded constant is given away: numerically $\mu_k/\big(\tfrac14(\tfrac23)^k\big)$ increases from 1 and settles near 1.28. For comparison, the exact median of the whole list of $n=3^{2021}$ numbers has $E\big|\mathrm{med}-\tfrac12\big|$ of order $n^{-1/2}=(1/\sqrt3)^{2021}\approx(0.577)^{2021}$, which is much smaller than $(2/3)^{2021}$. Local trimming is genuinely worse than global median taking, and the problem is asking you to quantify how much information the greedy procedure throws away.

Layer 1: WHAT

Block 1: read the process as a complete ternary tree, so that the numbers at each level are independent and identically distributed and the state is a distribution. Block 2: compute the one-round map on cumulative distribution functions, $F\mapsto\phi\circ F$ with $\phi(u)=3u^2-2u^3$, from the order statistics of three. Block 3: read the constants of the problem, $\tfrac14=\mu_0$ and $\tfrac23=1/\phi'(\tfrac12)$, and conclude that the target is $\mu_k\ge\tfrac23\mu_{k-1}$. Block 4: note that ϕ is an increasing bijection fixing $\tfrac12$ and symmetric about it, so $\psi=\phi^{-1}$ exists with the same properties. Block 5: couple, that is, realise every X_k as $\psi^{\circ k}(U)$ for a single uniform U, turning the distributional recursion into $Y_k=\psi(Y_{k-1})$. Block 6: prove $\big|\psi(u)-\tfrac12\big|\ge\tfrac23\big|u-\tfrac12\big|$ by one derivative computation. Block 7: iterate pointwise and take expectations.

Juxtaposition. The nominal route stops at Block 3 with $F_k=\phi\circ F_{k-1}$, writes $\mu_k=2\int_0^{1/2}F_k(x)\,dx$, and then tries to bound $\int_0^{1/2}\phi\big(F_{k-1}\big)$ from below by $\tfrac23\int_0^{1/2}F_{k-1}$. The pointwise sufficient condition for that would be $\phi(u)\ge\tfrac23u$ on $[0,\tfrac12]$, and that is simply false: near $u=0$ one has $\phi(u)\approx3u^2$, which is far below $\tfrac23u$. In fact $\phi(u)\ge\tfrac23u$ only for $u\ge\frac{3-\sqrt{11/3}}{4}\approx0.271$. So the nominal route does not merely cost more ink, it cannot work.

The bifurcation point is the sentence after $F_k=\phi\circ F_{k-1}$. The composition sits on the outside, acting on the values of the cdf, and an integral of a composed function is not controlled from below by the integral of the inner function. You must move the iteration into the argument. The device that does exactly that is the inverse function, since $\big(\phi\circ F\big)^{-1}=F^{-1}\circ\phi^{-1}$, and the probabilistic name of the inverse cdf is the quantile function.

Layer 2: HOW

Block 1. **Cue.** The same operation is applied at every stage to disjoint triples. **Why natural.** Functions of disjoint families of independent variables are independent, so at every level the surviving numbers are independent, and by symmetry identically distributed. That is the only reason a one-variable recursion exists at all; if the triples overlapped, the level would carry correlations and no scalar recursion would close. **Link to the next block.** The state of the process is one distribution, so find the map that advances it.

Block 2. **Cue.** "Median of three". **Why natural.** The distribution of an order statistic of independent identically distributed variables is a function of the common cdf alone. For the median of three, "median at most x" means "at least two of the three are at most x", giving $3F^2(1-F)+F^3 = 3F^2-2F^3$. Working with the cdf rather than the density is deliberate: the cdf turns a counting event into a polynomial and gives you a map ϕ that does not involve x at all. **Link to the next block.** The recursion is $F_k = \phi \circ F_{k-1}$, so F_k is the k-th iterate of one cubic.

Block 3. **Cue.** The two constants in the inequality you are asked to prove. **Why natural.** $\frac{1}{4}$ is $E\left|U-\frac{1}{2}\right|$, the value of the quantity at $k=0$. So the inequality is an iteration of a single step, and the single step must lose exactly the factor $\frac{2}{3}$. Now compute $\phi'(\frac{1}{2}) = \frac{3}{2}$: the factor is the reciprocal of the multiplier of ϕ at its interior fixed point. Nothing about this is a guess; a problem that hands you μ_0 and $1/\phi'(\text{fixed point})$ is telling you the lemma. **Link to the next block.** You must prove $\mu_k \geq \frac{2}{3}\mu_{k-1}$, and you must do it without integrating a composed cdf.

Block 4. **Cue.** ϕ is strictly increasing on $[0,1]$ with $\phi(1-u) = 1-\phi(u)$. **Why natural.** Monotonicity is what allows inversion, and the symmetry is what makes $\frac{1}{2}$ a fixed point and lets you reduce every statement to $[0,\frac{1}{2}]$. Both facts are one line each and both are used twice, so extract them before starting. **Link to the next block.** $\psi = \phi^{-1}$ exists and inherits both properties.

Block 5. **Cue.** A recursion between distributions that you need to iterate as an inequality. **Why natural.** Inequalities iterate cleanly between numbers, not between measures. The quantile transform realises every member of the family on one probability space: if F is a continuous strictly increasing cdf and U is uniform, then $F^{-1}(U)$ has cdf F. Applying it to every F_k with the same U gives a monotone coupling, and because $F_k^{-1} = \left(\phi^{\circ k}\right)^{-1} = \psi^{\circ k}$, the outer composition on cdfs has become an inner composition on quantiles. This is the whole solution in one move. **Link to the next block.** The recursion is now the pointwise identity $Y_k = \psi(Y_{k-1})$, and any pointwise bound on ψ will iterate.

Block 6. **Cue.** You need $\left|\psi(u)-\frac{1}{2}\right| \geq \frac{2}{3}\left|u-\frac{1}{2}\right|$, that is, ψ does not contract distance to $\frac{1}{2}$ by more than $\frac{2}{3}$. **Why natural.** Rewriting in terms of ϕ, this says ϕ does not expand distance to $\frac{1}{2}$ by more than $\frac{3}{2}$, which is exactly the mean value theorem statement $\phi' \leq \phi'(\frac{1}{2})$ on $[0,\frac{1}{2}]$. And that is true because $\phi'(x) = 6x(1-x)$ is increasing on $[0,\frac{1}{2}]$. The clean way to write it is to show $g(x) = \frac{3}{2}\left(\frac{1}{2}-x\right) - \left(\frac{1}{2}-\phi(x)\right)$ vanishes at $x=\frac{1}{2}$ and has $g' = -\frac{3}{2}(1-2x)^2 \leq 0$, which also shows where the loss is: only at $x=\frac{1}{2}$ is the inequality tight. **Link to the next block.** A pointwise inequality on a coupled family iterates by induction and survives taking expectations.

Block 7. **Cue.** A pointwise almost sure inequality. **Why natural.** Expectation is monotone, so no interchange, no convergence theorem and no integration by parts is required. The last ingredient is $E\left|U-\frac{1}{2}\right| = \frac{1}{4}$, a one line integral. **Link to the next block.** Set $k = 2021$.

Layer 3: WHY

The mathematical truth beneath the route is a statement about iterated maps. The cdf map $\phi(u) = 3u^2 - 2u^3$ has three fixed points on $[0,1]$: 0 and 1 with multiplier $\phi' = 0$, and $\frac{1}{2}$ with multiplier $\phi'(\frac{1}{2}) = \frac{3}{2} > 1$. So $\frac{1}{2}$ is repelling for ϕ and attracting for $\psi = \phi^{-1}$, and the iterates $F_k = \phi^{\circ k}$ converge to the step function at $\frac{1}{2}$. In probabilistic language, $X_k \to \frac{1}{2}$ in probability, and the natural length scale of the distribution around $\frac{1}{2}$ shrinks by the factor $1/\phi'(\frac{1}{2}) = \frac{2}{3}$ per round. The problem asks you to prove that the scale shrinks by no more than that, and the reason it is true is that the multiplier at the fixed point is the largest value of ϕ' on $[0, \frac{1}{2}]$. This is why the constant $\frac{2}{3}$ is exactly right and not merely sufficient.

The second truth is about the mechanics of iterating inequalities. A recursion $F_k = \phi \circ F_{k-1}$ places the iteration in the value slot; a recursion $q_k = \psi \circ q_{k-1}$ places it in the argument slot, where $q_k = F_k^{-1}$. Functionals like $\int F$ respond well to changes in the argument and badly to changes in the value, so the correct move is to invert. The probabilistic name of that inversion is the quantile transform, and the correct name of what it buys you is a monotone coupling: one probability space carrying every X_k at once, ordered by a deterministic increasing map. Once the family is coupled, distributional statements become pointwise statements and every inequality you prove about the map is inherited by every expectation you care about.

The discernment mechanism, in the form you can run in the arena. First, when a random object is built by repeating one operation on independent copies, write the recursion at the level of cumulative distribution functions and never at the level of densities, because order statistics are polynomial in the cdf. Second, read the constants in the target inequality: if one of them is the value of the quantity at step zero, the target is an iterated one step lemma, and if another is a derivative of the recursion map at a fixed point, the one step lemma is a mean value theorem statement. Third, if the recursion composes on the outside and you must iterate an inequality, invert.

How to train it. Take any problem in which a distribution evolves by a fixed rule and force yourself to write three objects before doing anything: the map on cdfs, its fixed points, and the derivative at each fixed point. Then compare those numbers with the constants in the question. On this problem that comparison finishes the analysis in under a minute.

What the examiners nudged with, and how. The exponent 2021 and the base 3^{2021} are there to make explicit computation absurd, which is a way of telling you the answer is one lemma iterated. The constant $\frac{1}{4}$ is μ_0 and is the starting value. The constant $\frac{2}{3}$ is $1/\phi'(\frac{1}{2})$ and is the loss per round. The word "median" with three elements is the smallest case in which the cdf map is nonlinear, so the problem is the minimal instance of its own idea. And the absolute value $\left|X - \frac{1}{2}\right|$ rather than a variance is chosen because $E\left|\cdot\right|$ respects pointwise inequalities directly, which is the only property the coupling proof needs.

Revolving concepts

Order statistics. For n independent variables with common cdf F, the j-th smallest has cdf $\sum_{i=j}^{n} \binom{n}{i} F^i (1-F)^{n-i}$. For the median of 3 this is $3F^2 - 2F^3$; for the median of $2m+1$ it is the regularised incomplete beta function $I_F(m+1, m+1)$, whose derivative in F is $\frac{(2m+1)!}{m!\,m!}\big(F(1-F)\big)^m$. See Book 1, field: Probability, topic: order statistics and the beta distribution.

Quantile function and the inverse transform. If F is a continuous strictly increasing cdf on an interval and U is uniform on $[0,1]$, then $F^{-1}(U)$ has cdf F. Consequently a whole family of distributions can be realised on one probability space, monotonically, by using the same U; this is the monotone coupling. Composition inverts: $\big(\phi \circ F\big)^{-1} = F^{-1} \circ \phi^{-1}$. See Book 1, field: Probability, topic: quantile transform and coupling.

Iterated maps and multipliers. For a C^1 map ϕ of an interval with fixed point p, the number $\phi'(p)$ is the multiplier; $|\phi'(p)| < 1$ makes p attracting, $|\phi'(p)| > 1$ repelling, and near a repelling fixed point the orbit of ϕ^{-1} approaches p at the geometric rate $1/\phi'(p)$. Here $\phi(u) = 3u^2 - 2u^3$ has fixed

points $0, \frac{1}{2}, 1$ with multipliers $0, \frac{3}{2}, 0$. See Book 1, field: Sequences and recurrences, topic: iterated maps and fixed point multipliers.

Symmetric distributions. If $F(1-x) = 1 - F(x)$ then the law is symmetric about $\frac{1}{2}$, $F(\frac{1}{2}) = \frac{1}{2}$, and $E|Z - \frac{1}{2}| = 2\int_0^{1/2} (\frac{1}{2} - x)\, dF(x) = 2\int_0^{1/2} F(x)\, dx$ after an integration by parts. The symmetry is preserved by ϕ because $\phi(1-u) = 1 - \phi(u)$.

Independence of blockwise functions. If $Z_1, \ldots, Z_n$ are independent and the index set is partitioned into blocks, then functions of the separate blocks are independent. This is why the surviving numbers at each level are independent and the whole reduction to a scalar recursion is legitimate; it is the single hypothesis that would fail if the triples overlapped.

The generalised recursion. Replace "median of three" by "median of $2m+1$" and repeat k times. The cdf map is $\phi_m = I_u(m+1, m+1)$, again an increasing bijection of $[0,1]$ with $\phi_m(1-u) = 1 - \phi_m(u)$ and $\phi_m(\frac{1}{2}) = \frac{1}{2}$, and its derivative $\frac{(2m+1)!}{m!\,m!}\big(u(1-u)\big)^m$ is maximised on $[0, \frac{1}{2}]$ at $u = \frac{1}{2}$, with value

$$c_m = \phi_m'\left(\tfrac{1}{2}\right) = \frac{(2m+1)!}{m!\,m!\,4^m}.$$

The identical argument then gives $\mu_k \geq \frac{1}{4}\, c_m^{-k}$, and $m = 1$ gives $c_1 = \frac{3}{2}$, which is this problem. More generally still: if ϕ is any C^1 increasing bijection of $[0,1]$ with $\phi(1-u) = 1 - \phi(u)$ and $\phi' \leq L$ on $[0, \frac{1}{2}]$, then iterating the associated median type operation k times gives $E\big|X_k - \frac{1}{2}\big| \geq L^{-k} E\big|X_0 - \frac{1}{2}\big|$. Store that statement, not the cubic.

Knowledge gaps, stated plainly. Four things here are past A-level Further Mathematics and must be installed, not improvised. First, the distribution of an order statistic, and specifically that it is a polynomial in the parent cdf. Second, the quantile transform and monotone coupling, which is the entire method. Third, the language of iterated maps: fixed point, multiplier, attracting and repelling, and the fact that the multiplier governs the rate. Fourth, the fact that functions of disjoint blocks of independent variables are independent, which is what makes the level-by-level recursion valid. None of these is difficult; all four are invisible if you have not met them, and the problem is unapproachable without the second.

Circuit Bank

- **WHEN** you see a random object built by repeating one operation on independent identically distributed copies **DO** write the recursion for the cumulative distribution function, not for the density and not for the variable **BECAUSE** order statistics are polynomials in the parent cdf, so the one round map is a fixed function ϕ with no dependence on the point x
- **WHEN** you have a recursion $F_k = \phi \circ F_{k-1}$ on distribution functions and you must iterate an inequality **DO** pass to quantile functions and couple: set $Y_k = \phi^{-1}$ applied k times to a single uniform U **BECAUSE** composition on the outside of a cdf becomes composition on the inside of its inverse, and inequalities iterate between numbers but not between integrals of composed functions
- **WHEN** the target inequality contains the value of the quantity at step zero and a derivative of the recursion map at a fixed point **DO** prove the one step lemma with that derivative as the loss factor and induct **BECAUSE** a constant equal to $1/\phi'(p)$ at a fixed point p is never a coincidence; it is the exact geometric rate at which the inverse map approaches p
- **WHEN** you need a lower bound on an expectation of an absolute value **DO** obtain a pointwise almost sure inequality first, then take expectations **BECAUSE** monotonicity of expectation is free, whereas lower bounds obtained from moment inequalities point the wrong way and need Paley and Zygmund style corrections
- **WHEN** a map ϕ of $[0,1]$ satisfies $\phi(1-u) = 1 - \phi(u)$ **DO** reduce every claim to $[0, \frac{1}{2}]$ and record that $\frac{1}{2}$ is a fixed point **BECAUSE** the symmetry is inherited by all iterates and by the

inverse, so it halves the work in every step of the argument

Backtesting

1. First useful cue: the same operation applied to disjoint triples at every level, which forces a recursion on cumulative distribution functions, $F \mapsto 3F^2 - 2F^3$.
2. Method the expression suggested: read $\frac{1}{4}$ as μ_0 and $\frac{2}{3}$ as $1/\phi'(\frac{1}{2})$, then prove the one step lemma $\mu_k \geq \frac{2}{3}\mu_{k-1}$ by coupling and iterate it.
3. Where you would hesitate, and the thought to reinforce: after obtaining $F_k = \phi \circ F_{k-1}$ you will want to integrate. Reinforce this: an outer composition is useless for lower bounds because $\phi(u) \geq \frac{2}{3}u$ is false near zero. Invert first, then integrate.
4. The trap the question sets: treating trimming as if it were the median of the whole list. The true median of 3^{2021} numbers sits at distance of order $(0.577)^{2021}$ from $\frac{1}{2}$, which is far smaller than the bound you are asked to prove, so any argument that silently replaces trimming by global median taking proves nothing.
5. What makes you recognise it faster next time: a deterministic map on distributions with a symmetric fixed point, and a target constant that equals the reciprocal of the derivative there. That pairing means coupling.
6. The reusable decision rule: distributional recursions become pointwise recursions under the quantile transform, and only pointwise inequalities iterate for free.

Mirror drills

1. Show that if Y_1, Y_2, Y_3 are independent with common cdf F, the median has cdf $3F^2 - 2F^3$, and that the density of the median of $2m+1$ such variables is $\frac{(2m+1)!}{m!\,m!}\big(F(1-F)\big)^m f$. *Answer:* the median is at most x exactly when at least $m+1$ of the variables are, so the cdf is $\sum_{i \geq m+1} \binom{2m+1}{i} F^i (1-F)^{2m+1-i}$; differentiating in F telescopes to the displayed beta density.
2. Let F be a continuous strictly increasing cdf on $[0,1]$ and U uniform. Show $F^{-1}(U)$ has cdf F, and deduce that if $G = \phi \circ F$ with ϕ an increasing bijection then $G^{-1} = F^{-1} \circ \phi^{-1}$. *Answer:* $P\big(F^{-1}(U) \leq x\big) = P\big(U \leq F(x)\big) = F(x)$; the second claim is the inverse of a composition, and it is the reason the iteration moves from the value slot to the argument slot.
3. Prove the generalisation: if each round keeps the median of $2m+1$ consecutive numbers and there are k rounds, then $\mu_k \geq \frac{1}{4}c_m^{-k}$ where $c_m = \frac{(2m+1)!}{m!\,m!\,4^m}$. *Answer:* the cdf map ϕ_m has derivative $\frac{(2m+1)!}{m!\,m!}(u(1-u))^m$, which on $[0, \frac{1}{2}]$ is maximised at $u = \frac{1}{2}$ with value c_m; run the same coupling and the same mean value argument. For $m = 1$, $c_1 = \frac{3}{2}$.

Chapter 6

Putnam 2020

6.1 A1. Repunits, the order of ten modulo 101

How many positive integers N satisfy all of the following three conditions?

(i) *N is divisible by* 2020.

(ii) *N has at most* 2020 *decimal digits.*

(iii) *The decimal digits of N are a string of consecutive ones followed by a string of consecutive zeros.*

Read it out loud

This is a counting question wearing a divisibility mask. Condition (iii) is not a condition at all, it is a parametrisation instruction: it hands you the shape of every candidate and asks you to name the two parameters. Say it out loud as "k ones then b zeros", and the object on the page is $N = \underbrace{11\ldots1}_{k}\underbrace{00\ldots0}_{b}$, which is $R_k \cdot 10^b$ where R_k is the repunit with k ones. Condition (ii) is $k + b \leq 2020$ and condition (i) is a divisibility test. The fields wired in are elementary number theory (factorisation, coprimality, multiplicative order) and a two variable count over a triangular region. Nothing else.

Within thirty seconds two lines should be on your page: $N = R_k \cdot 10^b$ with $k \geq 1$, $b \geq 0$, $k + b \leq 2020$, and $2020 = 2^2 \cdot 5 \cdot 101$. The second line is compulsory because you are being asked about divisibility by a composite, and a composite divisibility condition always splits into its prime power components. The moment those two lines sit next to each other the problem is already solved in outline, because R_k is odd and ends in the digit 1, so R_k is coprime to 10. Thus the factor $2^2 \cdot 5$ can only be supplied by 10^b, and the factor 101 can only be supplied by R_k. The two parameters separate completely and never speak to each other again.

What is really being tested is whether you take the number apart multiplicatively before you touch it, and whether you know that $10^2 \equiv -1 \pmod{101}$. That single congruence is the whole arithmetic content: it says the order of 10 modulo 101 is 4, which converts "101 divides N" into "4 divides k". The examiners chose 2020 precisely because 101 has that clean small order and because $2^2 \cdot 5$ forces the innocent looking condition $b \geq 2$ that most candidates get wrong by one.

Approaches, ranked worst to best

1. **Enumerate and hope.** *Cue:* "how many positive integers". *Circuit:* list candidates, look for a pattern. (There are more than two million candidate pairs (k, b) and no pattern is visible without the arithmetic, so this produces nothing at all.)

2. **Write $N = (10^a - 10^b)/9$ and clear the nine.** *Cue:* the repunit formula $R_k = (10^k - 1)/9$. *Circuit:* turn the divisibility into $9 \cdot 2020 \mid 10^b(10^{a-b} - 1)$. (Correct, and it is the official route, but you now carry a spurious factor 3^2 through every line and have to remark that it is automatic. Any fraction you introduce and then cancel is ink you paid for twice.)

3. **Factorise $N = R_k \cdot 10^b$ and split by prime.** *Cue:* R_k is coprime to 10. *Circuit:* each prime power of the modulus is supplied by exactly one factor. (Best on the arithmetic side. No fractions, no redundant primes, and the two conditions $b \geq 2$ and $4 \mid k$ appear one line apart.)

4. **Count by fixing the number of zeros b.** *Cue:* b is the parameter in the divisibility condition on 10^b. *Circuit:* for each b count the admissible a. (This is how the official solution counts, and it produces $\sum_b \lfloor (2020 - b)/4 \rfloor$, a sum with a floor function that then has to be regrouped into runs of four equal terms. Correct but it costs a paragraph.)

5. **Count by fixing the number of ones k.** *Cue:* the constraint on k is a congruence, the constraint on b is an interval. *Circuit:* substitute $k = 4j$ so the surviving count is an arithmetic progression. (Best. Once $k = 4j$ is substituted, b runs over $2 \leq b \leq 2020 - 4j$, so the inner count is $2019 - 4j$

with no floor function anywhere, and the outer sum is a single arithmetic series. Choose to fix the variable whose condition is a congruence, because a congruence is what creates floor functions when it is left in the inner loop.)

The ideal solution

Every N satisfying (ii) and (iii) is

$$N = R_k \cdot 10^b, \qquad R_k = \underbrace{11\ldots1}_{k} = \frac{10^k - 1}{9},$$

with integers $k \geq 1$ and $b \geq 0$ and $k + b \leq 2020$, and distinct pairs (k, b) give distinct N. Factor the modulus:

$$2020 = 2^2 \cdot 5 \cdot 101.$$

The repunit R_k has final digit 1, so it is odd and not divisible by 5; that is, $\gcd(R_k, 10) = 1$. Consequently

$$v_2(N) = v_2(10^b) = b, \qquad v_5(N) = v_5(10^b) = b,$$

so $2^2 \cdot 5 \mid N$ if and only if $b \geq 2$. Also $\gcd(10^b, 101) = 1$, so $101 \mid N$ if and only if $101 \mid R_k$, and since $\gcd(9, 101) = 1$ this holds if and only if $101 \mid 10^k - 1$.

Lemma 6.1.1. *The multiplicative order of* 10 *modulo* 101 *is* 4. *Hence* $101 \mid 10^k - 1$ *if and only if* $4 \mid k$.

Proof. $10^2 = 100 \equiv -1 \pmod{101}$, so $10^4 \equiv 1 \pmod{101}$ and the order divides 4. It is not 1 because $10 \not\equiv 1$, and not 2 because $-1 \not\equiv 1$ modulo 101. So the order is exactly 4, and $10^k \equiv 1$ if and only if the order divides k. □

So $2020 \mid N$ if and only if $b \geq 2$ and $4 \mid k$. Writing $k = 4j$ with $j \geq 1$, condition (ii) becomes $b \leq 2020 - 4j$, and b must satisfy $2 \leq b \leq 2020 - 4j$. The number of such b is $2019 - 4j$ when $2020 - 4j \geq 2$, that is when $j \leq 504$, and there are none for $j \geq 505$. Therefore the answer is

$$\sum_{j=1}^{504} (2019 - 4j) = 504 \cdot 2019 - 4 \cdot \frac{504 \cdot 505}{2} = 504(2019 - 1010) = 504 \cdot 1009.$$

Hence the count is

$$\boxed{508536}.$$

Layer 1: WHAT

Block 1: read (iii) as a parametrisation and write $N = R_k 10^b$ with $k + b \leq 2020$. Block 2: factor $2020 = 2^2 \cdot 5 \cdot 101$ and note $\gcd(R_k, 10) = 1$, so the modulus splits between the two factors with no overlap. Block 3: the $2^2 \cdot 5$ part is carried by 10^b alone, giving $b \geq 2$. Block 4: the 101 part is carried by R_k alone, and the order of 10 modulo 101 is 4, giving $4 \mid k$. Block 5: substitute $k = 4j$ and count lattice points, one arithmetic series.

Juxtaposition. The nominal route writes $N = (10^a - 10^b)/9$, multiplies up to $9N = 10^b(10^{a-b} - 1)$, and asks when 18180 divides that; it then has to dispose of the factor 3^2 as automatic, keep two variables a and b tangled by the inequality $b < a \leq 2020$, and finish with a floor function sum. The ideal route never introduces the nine and never uses a: it uses the two natural parameters, the count of ones and the count of zeros, because those are the parameters in which the arithmetic separates. The bifurcation point is the very first line you write down. If you write $(10^a - 10^b)/9$ you are describing the number; if you write $R_k \cdot 10^b$ you are factorising it, and factorising is what a divisibility question always wants.

Layer 2: HOW

Block 1. **Cue.** "The decimal digits are a string of ones followed by a string of zeros." **Why natural.** A description of digits is a description of a decimal expansion, and every decimal expansion of that shape is a repunit times a power of ten. You are not being clever here, you are reading the definition of the object out loud in symbols. The two counts, ones and zeros, are the only degrees of freedom, so they are the coordinates. **Link to the next block.** Once the object is a product, the divisibility condition can be tested factor by factor, which is exactly the next block.

Block 2. **Cue.** A divisibility condition by the composite 2020. **Why natural.** Divisibility by a composite is never one condition. It is one condition per prime power, and those conditions are independent by the Chinese remainder theorem. So factor first, always, before anything else. Here $2020 = 2^2 \cdot 5 \cdot 101$, and the presence of 101 is the reason the year was chosen. **Link to the next block.** The factorisation is only useful if you know which factor of N can supply each prime, and that is the coprimality observation.

Block 3. **Cue.** R_k ends in the digit 1. **Why natural.** A repunit is coprime to 10 by inspection of its last digit, so the entire 2 and 5 content of N sits in 10^b and the entire 101 content sits in R_k. In valuation language $v_2(N) = v_5(N) = b$. This is the structural reroute: the two parameters are decoupled, so the answer will be a product region, not a tangled one. **Link to the next block.** Decoupled means you may now treat the two conditions separately, and the harder one is the condition on R_k.

Block 4. **Cue.** $101 \mid R_k$, which is a statement about $10^k \bmod 101$. **Why natural.** The set of exponents k with $10^k \equiv 1 \pmod{101}$ is exactly the set of multiples of the multiplicative order, because the exponents form a subgroup of $\mathbb{Z}$ and every subgroup of $\mathbb{Z}$ is the multiples of one integer. So the only work is to compute that order, and $10^2 = 100 = 101 - 1$ makes it immediate. **Link to the next block.** A congruence condition on k and an interval condition on b: the count is now a lattice point count.

Block 5. **Cue.** One congruence condition, one interval condition, one inequality linking them. **Why natural.** Put the congruence in the outer loop. Substituting $k = 4j$ makes the constraint disappear into the index and leaves an interval of b values whose length is linear in j, so the sum is arithmetic. If instead you loop over b on the outside, the inner count is the number of multiples of 4 in an interval, which is a floor function, and floor functions must then be regrouped by residue class. **Link to the next block.** Nothing remains but one arithmetic series.

Layer 3: WHY

The mathematical truth beneath this problem is that divisibility is multiplicative and N has been handed to you as a product of two coprime pieces. When $N = AB$ with $\gcd(A, B) = 1$, the condition $m \mid N$ decomposes prime by prime, and each prime power of m must be found in whichever of A, B is not coprime to it. Here that decomposition is total: $2^2 \cdot 5$ lives only in 10^b and 101 lives only in R_k. That is why the answer is a clean lattice count and not a case analysis.

The second truth is the order. For $\gcd(a, m) = 1$, the set $\{k \in \mathbb{Z} : a^k \equiv 1 \pmod{m}\}$ is a subgroup of $\mathbb{Z}$, hence equal to $d\mathbb{Z}$ where $d = \mathrm{ord}_m(a)$, and $d \mid \varphi(m)$ by Lagrange. So any question of the form "for which k does m divide the repunit R_k" or "for which k does m divide $a^k - 1$" has the same shape of answer: the multiples of one number. The discernment mechanism is to notice that the unknown appears as an exponent, and an unknown in an exponent means order, never expansion.

Train the discernment with a habit: when a Putnam problem names a specific large year number as a modulus, factor it in the margin before reading the rest of the sentence. $2020 = 2^2 \cdot 5 \cdot 101$, and 101 is the prime whose reciprocal has period 4 because $10^2 \equiv -1$. The examiners nudged you twice: they chose 101 so that the order computation is a one liner, and they chose $2^2 \cdot 5$ rather than 10 so that the sloppy answer $b \geq 1$ is wrong. Every off by one in this problem lives at $b \geq 2$.

Revolving concepts

Repunits. $R_k = (10^k - 1)/9 = \sum_{i=0}^{k-1} 10^i$. They satisfy the recursion $R_{k+1} = 10R_k + 1$, and the general solution of $u_{k+1} = 10u_k + 1$ is $u_k = C \cdot 10^k - 1/9$, which is where the closed form comes from. Key divisibility fact: $R_d \mid R_k$ when $d \mid k$, because $\gcd(10^k - 1, 10^d - 1) = 10^{\gcd(k,d)} - 1$. See Book 1, field: Number theory, topic: repunits and $a^k - 1$.

Multiplicative order. For $\gcd(a, m) = 1$, $\operatorname{ord}_m(a)$ is the least $k \geq 1$ with $a^k \equiv 1$, the set of all such k is $\operatorname{ord}_m(a)\mathbb{Z}$, and $\operatorname{ord}_m(a) \mid \varphi(m)$. Here $\operatorname{ord}_{101}(10) = 4$, which is also the statement that $1/101 = 0.\overline{0099}$ has period 4. See Book 1, field: Number theory, topic: orders and primitive roots.

Chinese remainder theorem as a splitting device. If $m = \prod p_i^{e_i}$ then $m \mid N$ if and only if $p_i^{e_i} \mid N$ for every i, and these conditions are logically independent. Used here to break 2020 into three unrelated questions. See Book 1, field: Number theory, topic: CRT.

p-adic valuation of a product. $v_p(AB) = v_p(A) + v_p(B)$, so a coprime factor contributes nothing. Used to get $v_2(N) = v_5(N) = b$ in one line rather than by cases. See Book 1, field: Number theory, topic: valuations.

Lattice counting under a congruence. To count pairs with $k \equiv 0 \pmod{d}$ and $k + b \leq M$ and $b \geq b_0$: substitute $k = dj$ so the congruence vanishes, then the inner count is an arithmetic progression in j. The general sum is $\sum_{j=1}^{J} (M - dj - b_0 + 1)$ with $J = \lfloor (M - b_0)/d \rfloor$. See Book 1, field: Combinatorics, topic: counting lattice points in triangles.

Circuit Bank

- **WHEN** you see a condition on the decimal digits **DO** write the number as a sum or product of powers of ten and repunits **BECAUSE** a digit description is a decimal expansion, and only the algebraic form can be tested for divisibility
- **WHEN** you see divisibility by a composite **DO** factor the modulus into prime powers before anything else **BECAUSE** by the Chinese remainder theorem the condition splits into independent conditions, and each prime usually attaches to a different factor of the number
- **WHEN** the unknown sits in an exponent inside a congruence **DO** compute the multiplicative order of the base and convert to a divisibility condition on the exponent **BECAUSE** the exponents that give 1 form a subgroup of $\mathbb{Z}$, so they are exactly the multiples of the order
- **WHEN** you must count pairs with one congruence condition and one interval condition **DO** loop over the congruence variable by substituting $k = dj$ **BECAUSE** substituting kills the congruence, and the remaining count is an interval length, so no floor functions ever appear

Backtesting

1. First useful cue: condition (iii) is a parametrisation, giving $N = R_k \cdot 10^b$, and $2020 = 2^2 \cdot 5 \cdot 101$.
2. Method the expression suggested: split the divisibility prime by prime and use coprimality to attach each prime power to exactly one factor.
3. Where you would hesitate, and the thought to reinforce: at the condition on b. The thought is $v_2(N) = v_5(N) = b$, so $4 \cdot 5 \mid N$ needs $b \geq 2$, not $b \geq 1$. Write the valuation, not the sentence "N must end in enough zeros".

4. The trap the question sets: two off by one errors, $b \geq 2$ rather than $b \geq 1$, and the upper limit $j \leq 504$ rather than $j \leq 505$. Check the extreme case explicitly: $j = 504$ gives $k = 2016$ and $b \in \{2, 3, 4\}$, three values, and $2019 - 4 \cdot 504 = 3$ confirms the formula at the edge.
5. What makes you recognise it faster next time: any problem naming the year as a modulus wants its factorisation, and any problem about strings of repeated digits wants repunits and orders.
6. The reusable decision rule: factor the number, factor the modulus, match prime powers to coprime factors, then count with the congruence in the outer loop.

Mirror drills

1. For how many k with $1 \leq k \leq 100$ is R_k divisible by 41? *Answer:* $10^5 \equiv 1 \pmod{41}$ and no smaller exponent works, so the order is 5 and $41 \mid R_k$ if and only if $5 \mid k$; there are 20 such k.
2. How many positive integers with at most 30 decimal digits, whose digits are a block of ones followed by a block of zeros, are divisible by 1000? *Answer:* the condition is $b \geq 3$ with $k \geq 1$ and $k + b \leq 30$, giving $\sum_{k=1}^{27}(28 - k) = 378$.
3. Show that R_k is never divisible by 2, 5, or 3^3, and find all k with $27 \mid R_k \cdot 10^b$. *Answer:* the last digit rules out 2 and 5; $R_k \equiv k \pmod{9}$ shows $9 \mid R_k$ iff $9 \mid k$, and since $v_3(R_k) = v_3(k)$ by lifting the exponent, $27 \mid R_k$ iff $27 \mid k$; 10^b contributes no factor 3, so the answer is $27 \mid k$.

6.2 A2. A binomial sum that counts coin flip majorities

Let k be a nonnegative integer. Evaluate

$$\sum_{j=0}^{k} 2^{k-j} \binom{k+j}{j}.$$

Read it out loud

This is a single finite binomial sum with a weight that is a power of two, and the command term is "evaluate", which promises a closed form. Say the summand out loud: a binomial coefficient in which both indices move as j moves, multiplied by 2 to the power of the complementary index. Two structural features jump out. First, $\binom{k+j}{j}$ counts arrangements of k objects of one kind and j of another. Second, and this is the guiding light, the two "lengths" in the summand add to a constant:

$$(k+j) + (k-j) + 1 = 2k+1.$$

A summand built from a block of length $k+j$, a free block of length $k-j$, and one more slot, with the total independent of j, is the fingerprint of a decomposition of objects of fixed length $2k+1$ according to where something happens for the first time.

The fields wired in are combinatorics (counting binary strings, bijections, the complement symmetry) and, if you prefer the language, elementary probability (the negative binomial distribution and a race between two players). Nothing analytic is involved: the answer will be an integer.

Within thirty seconds you should have computed the first two or three values. $S_0 = 1$, $S_1 = 2 \cdot 1 + 1 \cdot \binom{2}{1} = 4$, $S_2 = 4 + 2\binom{3}{1} + \binom{4}{2} = 16$. So the answer is 4^k, and you now know the size of the set you must build. Guessing the target is legitimate reconnaissance; it is the proof that must never be guessed. A power of two as an answer is a loud instruction: find a set of size 2^{2k+1}, split it in half by a symmetry, and show the sum counts one half.

Approaches, ranked worst to best

1. **Compute values and assert the pattern.** *Cue:* "evaluate". *Circuit:* tabulate S_0, S_1, S_2 and write 4^k. (Zero marks. It identifies the target, which is useful, but a Putnam grader wants the reason a pattern continues.)

2. **Induction through Pascal's rule.** *Cue:* $\binom{k+j}{j}$ with both indices moving. *Circuit:* split $\binom{k+j}{j} = \binom{k-1+j}{j} + \binom{k-1+j}{j-1}$ and reassemble. (Correct, and it is the first official route: the two pieces reassemble into $2S_{k-1}$ and $\frac{1}{2}S_k$, giving $S_k = 2S_{k-1} + \frac{1}{2}S_k$, hence $S_k = 4S_{k-1}$. It works, but it needs two index shifts, a boundary convention $\binom{n}{-1} = 0$, and the nerve to recognise S_k itself on the right hand side. Under exam pressure that is three places to slip.)

3. **Generating functions.** *Cue:* a binomial with upper index $k+j$ is a coefficient of $(1+x)^{k+j}$. *Circuit:* write the sum as $[x^k] \sum_j 2^{k-j}(1+x)^{k+j}$, sum the geometric series. (Clean and mechanical, and it is the third official route, but it ends at $2^{k+1}\sum_{j\le k}\binom{k}{j} - \sum_{j\le k}\binom{2k+1}{j}$, and the second of those partial sums is evaluated by exactly the complement symmetry used in the best route. So the generating function is a longer path to the same idea.)

4. **Probabilistic race.** *Cue:* $\binom{n-1}{k}2^{-n}$ is the negative binomial law. *Circuit:* let X be the index of the $(k+1)$st head and compute $2^{2k+1}P[X \le 2k+1]$. (Very short and correct, and it makes the symmetry visible. Ranked below the next one only because it converts an integer identity into a probability and back again, which costs two sentences of translation and forces you to state the law of X carefully.)

5. **Direct bijection on binary strings.** *Cue:* the constant total length $2k+1$. *Circuit:* count strings of length $2k+1$ with a majority of heads by the position of the $(k+1)$st head, then halve by the complement symmetry. (Best. Everything is a finite count, the sum is produced rather than manipulated, and the answer 4^k appears as half of 2^{2k+1}. No conventions, no index shifts, no boundary cases.)

The ideal solution

Write $S_k = \sum_{j=0}^{k} 2^{k-j}\binom{k+j}{j}$. We claim

$$S_k = \boxed{4^k}.$$

Let Ω be the set of strings of length $2k+1$ over the alphabet $\{H, T\}$, so $|\Omega| = 2^{2k+1}$. Call a string *heavy* if it contains at least $k+1$ letters H, and *light* if it contains at least $k+1$ letters T.

Claim 6.2.1. *S_k is the number of heavy strings.*

Proof. Let w be heavy. Since w contains at least $k+1$ letters H, there is a well defined position p at which the $(k+1)$st H occurs. The first $p-1$ letters contain exactly k letters H, so $p-1 = k+j$ where $j \geq 0$ is the number of letters T before position p; and $p \leq 2k+1$ forces $j \leq k$. Thus w is determined by the triple: the value of j; the arrangement of the first $k+j$ letters, which consist of exactly k letters H and j letters T and may be arranged in $\binom{k+j}{j}$ ways; and the last $2k+1-p = k-j$ letters, which are arbitrary, giving 2^{k-j} ways. Position p itself is forced to be H.

Conversely, every such triple produces a heavy string whose $(k+1)$st H sits at position $k+j+1$, and different triples produce different strings, because j is recovered from the string as the number of letters T before the $(k+1)$st H. Hence the heavy strings are in bijection with the triples, and their number is $\sum_{j=0}^{k}\binom{k+j}{j}2^{k-j} = S_k$. □

Now exchange the two letters. The map σ that replaces every H by T and every T by H is an involution of Ω, and it maps heavy strings onto light strings and light strings onto heavy strings. Hence there are as many heavy strings as light ones. Moreover every string is exactly one of the two: a string of odd length $2k+1$ has $\#H + \#T = 2k+1$, so $\#H \geq k+1$ if and only if $\#T \leq k$, that is, if and only if the string is not light. So Ω is the disjoint union of the heavy and the light strings, and the two classes have equal size:

$$S_k = \#\{\text{heavy}\} = \tfrac{1}{2}|\Omega| = \tfrac{1}{2}\cdot 2^{2k+1} = 4^k.$$

For $k = 0$ this reads $S_0 = 1$, which matches the sum directly.

Layer 1: WHAT

Block 1: compute S_0, S_1, S_2 to learn the target 4^k, and read $4^k = \frac{1}{2}\cdot 2^{2k+1}$ as "half of all strings of length $2k+1$". Block 2: notice $(k+j)+(k-j)+1 = 2k+1$ is constant, so the summand describes objects of fixed length $2k+1$ cut at one distinguished position. Block 3: identify the cut as the position of the $(k+1)$st letter H, and verify the decomposition is a bijection. Block 4: apply the letter exchange involution and the parity of $2k+1$ to conclude that heavy strings are exactly half of all strings.

Juxtaposition. The nominal route treats the sum as an algebraic object: apply Pascal, shift indices, recognise a copy of S_k inside itself, solve the resulting equation for S_k, then induct. It is a valid route and it is short on paper, but every step is an index manipulation with a boundary convention attached. The ideal route never manipulates the sum at all; it builds the set the sum counts and then counts that set a second way. The bifurcation point is the moment you see 2^{k-j} next to $\binom{k+j}{j}$. The zombie reads 2^{k-j} as a coefficient to be absorbed; you read it as the number of ways to fill $k-j$ free slots, which forces the question "free slots in what object of total length $2k+1$".

Layer 2: HOW

Block 1. **Cue.** "Evaluate", a finite sum with a small parameter. **Why natural.** Three values cost fifteen seconds and remove all ambiguity about the target. Knowing the answer is 4^k changes the search: you are no longer looking for a manipulation, you are looking for a set of size 2^{2k+1} and a two to one reason. **Link to the next block.** A power of two answer points to a counting proof, so the next job is to identify what the summand counts.

Block 2. **Cue.** $(k+j)+(k-j)$ is independent of j. **Why natural.** Whenever a sum over j has one factor counting arrangements in a block of length $k+j$ and another factor counting free choices in a block of length $k-j$, the sum is counting objects of the constant total length by cutting them at a position that depends on the object. That is the only way a j dependent split can produce a j independent whole. So the sum equals the number of strings of length $2k+1$ in some family, and j records where the cut falls. **Link to the next block.** The cut position must be determined by the string itself, and the binomial $\binom{k+j}{j}$ tells you exactly which statistic is being fixed: k letters of one kind before the cut.

Block 3. **Cue.** $\binom{k+j}{j}$ with exactly k letters H and j letters T before a forced H. **Why natural.** The forced letter at the cut is what makes the decomposition a bijection rather than an overcount: the cut is the first time the count of H reaches $k+1$. Any first passage decomposition has this shape, one forced letter at the moment of passage, everything before it constrained, everything after it free. State the inverse map explicitly and the bijection is proved rather than asserted. **Link to the next block.** Having identified the family as "at least $k+1$ letters H in $2k+1$ letters", you need its size, and the family is defined by a majority condition.

Block 4. **Cue.** Odd total length, and a condition of the form "at least half". **Why natural.** In an odd number of trials there are no ties, so the majority condition splits the whole space into two classes, and the exchange of letters is an explicit bijection between them. Two classes, equal size, disjoint, covering everything: each is half. The oddness is not decoration, it is what makes the split clean, and it is why the problem was posed with $2k+1$ hidden inside the summand rather than $2k$. **Link to the next block.** The count is $2^{2k+1}/2 = 4^k$ and nothing remains.

Layer 3: WHY

The truth beneath this problem is that a binomial identity with a closed form as clean as 4^k is a statement that two ways of counting one set agree. The sum counts by conditioning on a first passage time; the closed form counts by a symmetry. Every proof of the identity is somewhere on the spectrum between those two, and the algebraic proofs (Pascal induction, generating functions) are the same argument with the combinatorial content compiled away. The generating function route makes this explicit: it terminates at $\sum_{j=0}^{k}\binom{2k+1}{j} = 2^{2k}$, which is precisely the complement symmetry $\binom{2k+1}{j} = \binom{2k+1}{2k+1-j}$ in disguise.

The discernment mechanism is the constant total. Before deciding how to attack a binomial sum, add up the lengths appearing in the summand. If they sum to a constant, the sum is a decomposition of a fixed size family and you should hunt for the family. If they do not, no such reading exists and you fall back on Pascal, generating functions, or creative telescoping. That is a one line diagnostic you can run on every binomial sum you ever meet.

Train it by rewriting powers as counts. The habit is to refuse to read 2^{k-j} as "a coefficient" and to read it as "the number of ways to fill $k-j$ binary slots". Once every factor in a summand is a count of something, the summand is a description of an object, and the sum is a family. The examiners nudged you toward this by choosing the weight 2^{k-j} rather than, say, 3^{k-j}: base two matches the two letters in the binomial, so the family is binary strings and the symmetry is available.

Revolving concepts

First passage decomposition. To count sequences in which a running statistic first reaches a threshold, condition on the position of that first passage: everything before it is constrained by the near threshold count, the position itself carries a forced letter, and everything after is free. In general, the number of binary strings of length $n+m+1$ with at least $n+1$ letters H is $\sum_{j=0}^{m}\binom{n+j}{j}2^{m-j}$, which is our identity when $m=n$. See Book 1, field: Combinatorics, topic: first passage and ballot type decompositions.

Negative binomial law. For fair coin flips, the index X of the $(k+1)$st head satisfies $P[X=n]=\binom{n-1}{k}2^{-n}$ for $n \geq k+1$. Writing $n=k+j+1$ turns our sum into $2^{2k+1}P[X \leq 2k+1]$, and $P[X \leq 2k+1]=\frac{1}{2}$ by the majority symmetry. See Book 1, field: Probability, topic: negative binomial and waiting times.

Complement symmetry. $\binom{n}{j}=\binom{n}{n-j}$, and for odd n the two halves of the row of Pascal's triangle each sum to 2^{n-1}: $\sum_{j=0}^{(n-1)/2}\binom{n}{j}=2^{n-1}$. This is the same involution as the letter exchange. See Book 1, field: Combinatorics, topic: binomial symmetry and partial row sums.

Pascal recursion and its solution. $\binom{n}{r}=\binom{n-1}{r}+\binom{n-1}{r-1}$ applied inside the sum yields the recursion $S_k=2S_{k-1}+\frac{1}{2}S_k$, that is, the first order linear recursion $S_k=4S_{k-1}$ with $S_0=1$, whose general solution is $S_k=S_0r^k$ with $r=4$. Any time a sum can be re expressed in terms of itself with shifted parameters you obtain such a recursion; the general template $u_k=\alpha u_{k-1}+\beta u_k$ solves to $u_k=\frac{\alpha}{1-\beta}u_{k-1}$, so always collect the copies of u_k before panicking. See Book 1, field: Sequences and recurrences, topic: first order linear recursions.

Mechanical proof of hypergeometric identities. Every identity of this shape, where the ratio of consecutive summands is a rational function of the index, is decidable by Zeilberger's creative telescoping (the WZ method), which outputs a recursion in k and a certificate. You do not need it here and you cannot run it by hand quickly, but it is the honest statement of why sums like this always have clean answers. This is a named knowledge gap: install the existence of the algorithm, not the algorithm. See Book 1, field: Generating functions, topic: hypergeometric summation.

Circuit Bank

- **WHEN** a binomial sum has lengths that add to a constant, such as $(k+j)+(k-j)$ **DO** read the summand as a description of one object of that constant length, cut at a position depending on j **BECAUSE** a j dependent split that produces a j independent total is exactly what conditioning on a first passage looks like
- **WHEN** a power such as 2^{k-j} sits next to a binomial coefficient **DO** read the power as the number of ways to fill $k-j$ free binary slots **BECAUSE** when every factor is a count, the summand describes an object and the sum describes a family
- **WHEN** a counting problem has odd total size and a condition of the form "at least half" **DO** apply the exchange involution and conclude the family is exactly half of everything **BECAUSE** oddness forbids ties, so the two complementary classes partition the space and the involution matches them
- **WHEN** a finite sum with a parameter asks you to evaluate **DO** compute the first three values before choosing a method **BECAUSE** knowing the closed form tells you which kind of proof to look for, and a power of two answer means look for a bijection rather than an algebraic identity
- **WHEN** a sum reappears inside its own transformation **DO** collect the copies and solve for the sum **BECAUSE** $S=\alpha T+\beta S$ is a linear equation, and $\beta \neq 1$ makes it a one line recursion

Backtesting

1. First useful cue: $(k+j)+(k-j)+1 = 2k+1$ is constant, and $S_0, S_1, S_2 = 1, 4, 16$.
2. Method the expression suggested: count binary strings of length $2k+1$ by the position of the $(k+1)$st head, then halve by the exchange symmetry.
3. Where you would hesitate, and the thought to reinforce: at "why is this a bijection". The thought to write is the inverse map: given a heavy string, j is the number of tails before the $(k+1)$st head. A decomposition without a stated inverse is an overcount waiting to happen.
4. The trap the question sets: the sum looks like something to manipulate. Pascal and generating functions both work and both cost more ink, and the generating function route ends up needing the same symmetry anyway.
5. What makes you recognise it faster next time: the pattern $\binom{a+b}{b}c^d$ with $a+b+d$ constant is the signature of a first passage count, and base two weight with a two letter binomial means the alphabet is $\{H, T\}$.
6. The reusable decision rule: if a binomial sum has constant total length, prove it by building the family; only if it does not, reach for Pascal, generating functions, or creative telescoping.

Mirror drills

1. Evaluate $\sum_{j=0}^{k} \binom{k+j}{j} 2^{-j}$. *Answer:* 2^k; divide the identity by 2^k, or repeat the argument with the free slots weighted by 2^{-1} each.
2. Show that $\sum_{j=0}^{k} \binom{2k+1}{j} = 4^k$. *Answer:* the involution $j \mapsto 2k+1-j$ pairs the terms with $j \le k$ against those with $j \ge k+1$, so the sum is half of $\sum_{j=0}^{2k+1} \binom{2k+1}{j} = 2^{2k+1}$.
3. Alice needs 3 more points and Bob needs 2 more to win a match decided by fair coin flips. Find the probability Alice wins. *Answer:* condition on the position of Alice's third point: $\sum_{j=0}^{1} \binom{2+j}{j} 2^{-(3+j)} = \frac{1}{8} + \frac{3}{16} = \frac{5}{16}$.

6.3 A3. Iterated sine and a neutral fixed point

Let $a_0 = \pi/2$, and let $a_n = \sin(a_{n-1})$ for $n \geq 1$. Determine whether

$$\sum_{n=1}^{\infty} a_n^2$$

converges.

Read it out loud

This is an iterated map feeding a series. The map is $f(x) = \sin x$, the orbit starts at $\pi/2$, and you are asked about the summability of the squares of the orbit. Say the structure out loud: $f(0) = 0$ and $f'(0) = 1$, so the origin is a fixed point at which the derivative is exactly 1. That is a neutral fixed point, and neutral fixed points are the whole subject of this problem. An attracting fixed point ($|f'| < 1$) would give geometric decay and every series in sight would converge; a neutral fixed point gives polynomial decay, and polynomial decay is exactly where convergence questions become interesting.

The fields wired in are real analysis (monotone sequences, Taylor bounds with a controlled sign, comparison tests) and discrete dynamics (asymptotics of orbits near a parabolic fixed point). The first cue is the mismatch between $\sin x$ and x: since $\sin x = x - x^3/6 + \cdots$, each step loses about $a_n^3/6$, so the decay of the orbit is governed by a cubic correction, and cubic correction to a neutral fixed point means the orbit decays like $n^{-1/2}$ and its square like n^{-1}. The harmonic series is the boundary case, which is why the problem is worth asking.

Within thirty seconds two things should be on your page: $a_1 = \sin(\pi/2) = 1$, so the orbit is really the orbit of 1 under repeated sine, and the heuristic $a_n \approx \sqrt{3/n}$, hence $a_n^2 \approx 3/n$, hence divergence. Now the only job is a proof, and the crucial observation is that divergence requires only a lower bound. You never need the sharp constant $\sqrt{3}$. That single realisation cuts the work by three quarters.

Approaches, ranked worst to best

1. **Note $a_n \to 0$ and guess.** *Cue:* the terms tend to zero. *Circuit:* hope that small terms sum. (Worthless: $a_n \to 0$ is necessary for convergence, never sufficient, and the harmonic series is the standing counterexample.)

2. **Ratio or root test.** *Cue:* a series with positive terms. *Circuit:* compute $\lim a_{n+1}^2/a_n^2$. (Silent by construction. The ratio tends to 1 because $a_{n+1}/a_n = \sin(a_n)/a_n \to 1$, and every test that compares against a geometric series returns no information for polynomially decaying terms. Recognising this in one line is worth more than attempting it.)

3. **Numerical evidence.** *Cue:* a concrete recursion. *Circuit:* compute partial sums. (Reconnaissance only. The partial sums grow like $3 \log n$, so after ten thousand terms you have about 28, which is suggestive and proves nothing.)

4. **Full asymptotic by Stolz and Cesaro.** *Cue:* the cubic correction $\sin x = x - x^3/6 + O(x^5)$. *Circuit:* set $b_n = a_n^{-2}$, show $b_{n+1} - b_n \to \frac{1}{3}$, conclude $a_n^2 \sim 3/n$. (Correct, sharp, and the right way to understand the problem, but it costs an expansion of $(1-t)^{-1}$ to second order, a Cesaro step, and careful error control. On the clock that is several places to lose a mark, and none of the sharpness is needed for the question actually asked.)

5. **One sided induction against $1/\sqrt{n}$.** *Cue:* you need divergence, so you need only a lower bound. *Circuit:* prove $\sin x \geq x - x^3/6$ once, then induct on $a_n \geq n^{-1/2}$. (Best. One inequality, one induction, one comparison with the harmonic series. The deliberate slack in the constant, 1 instead of the true $\sqrt{3}$, is what makes the induction step close with a factorisable quadratic.)

The ideal solution

The series diverges.

Lemma 6.3.1 (polynomial bounds for sine)**.** *For all $x \geq 0$: $\sin x \leq x$, $\cos x \geq 1 - \frac{x^2}{2}$, and $\sin x \geq x - \frac{x^3}{6}$.*

Proof. Let $u(x) = x - \sin x$. Then $u(0) = 0$ and $u'(x) = 1 - \cos x \geq 0$, so $u \geq 0$ on $[0, \infty)$, which is the first inequality. Let $v(x) = \cos x - 1 + \frac{x^2}{2}$. Then $v(0) = 0$ and $v'(x) = x - \sin x = u(x) \geq 0$, so $v \geq 0$, which is the second. Let $w(x) = \sin x - x + \frac{x^3}{6}$. Then $w(0) = 0$ and $w'(x) = \cos x - 1 + \frac{x^2}{2} = v(x) \geq 0$, so $w \geq 0$, which is the third. □

First, the orbit stays in $(0, 1]$. Indeed $a_1 = \sin(\pi/2) = 1$, and if $a_n \in (0, 1]$ then $a_{n+1} = \sin(a_n) \in (0, 1]$, because sin is positive on $(0, \pi)$ and $\sin x \leq x \leq 1$ there. In particular $a_n > 0$ for every $n \geq 1$, and sin is strictly increasing on $[0, 1]$ since $\cos > 0$ there.

Claim 6.3.2. *$a_n \geq n^{-1/2}$ for every $n \geq 1$.*

Proof. Induction on n. For $n = 1$, $a_1 = 1 = 1^{-1/2}$.

Assume $a_n \geq n^{-1/2}$ for some $n \geq 1$. Both a_n and $n^{-1/2}$ lie in $(0, 1]$ and sin is increasing there, so

$$a_{n+1} = \sin(a_n) \geq \sin\left(n^{-1/2}\right) \geq n^{-1/2} - \frac{1}{6}n^{-3/2} = \frac{6n - 1}{6n^{3/2}},$$

using the lemma. Since $n \geq 1$, the right hand side is positive, so it suffices to prove

$$\frac{6n - 1}{6n^{3/2}} \geq \frac{1}{\sqrt{n + 1}}.$$

Both sides are positive, so this is equivalent to the inequality obtained by squaring and clearing denominators:

$$(n + 1)(6n - 1)^2 \geq 36n^3.$$

Expanding the left side gives $36n^3 + 24n^2 - 11n + 1$, so the inequality is equivalent to

$$24n^2 - 11n + 1 \geq 0, \qquad \text{that is} \qquad (3n - 1)(8n - 1) \geq 0,$$

which holds for every $n \geq 1$. Hence $a_{n+1} \geq (n + 1)^{-1/2}$, completing the induction. □

Therefore $a_n^2 \geq 1/n$ for all $n \geq 1$, and since the harmonic series diverges, the comparison test for series with nonnegative terms gives

$$\sum_{n=1}^{\infty} a_n^2 \geq \sum_{n=1}^{\infty} \frac{1}{n} = \infty.$$

The series $\sum a_n^2$ **diverges**.

Layer 1: WHAT

Block 1: identify the fixed point structure, $\sin 0 = 0$ with derivative 1, so decay is polynomial rather than geometric. Block 2: estimate the decay rate from the cubic term, giving the heuristic $a_n \approx \sqrt{3/n}$ and $a_n^2 \approx 3/n$, so the answer is divergence. Block 3: notice that divergence needs only a lower bound, so discard the sharp constant and aim at $a_n \geq n^{-1/2}$. Block 4: prove the single inequality $\sin x \geq x - x^3/6$ by a three step derivative ladder. Block 5: run the induction, reducing the step to $(3n - 1)(8n - 1) \geq 0$. Block 6: compare with the harmonic series.

Juxtaposition. The nominal route sets $b_n = 1/a_n^2$, expands sin to fifth order, expands $(1 - t)^{-1}$, shows $b_{n+1} - b_n \to 1/3$, invokes Cesaro, deduces $a_n^2 \sim 3/n$, and finally compares. Every one of

those steps is correct and every one is optional. The ideal route uses the asymptotic only as a private prediction, never writing it down, and pays for a one sided bound instead. The bifurcation point is the question "what exactly do I need". The zombie answers "the asymptotic behaviour of a_n"; you answer "a lower bound whose square is not summable", and the entire second half of the work evaporates.

Layer 2: HOW

Block 1. **Cue.** An iteration $a_{n+1} = f(a_n)$ with $f(0) = 0$, $f'(0) = 1$, and $0 < f(x) < x$ near 0. **Why natural.** The orbit is positive and strictly decreasing, hence convergent, and its limit is a fixed point of the continuous map f, so the limit is 0. Because $f'(0) = 1$ the convergence is not geometric, so the decay is governed by the first nonlinear term of f at 0. Naming this before computing anything is what tells you the answer will sit near the harmonic boundary. **Link to the next block.** Knowing the decay is polynomial, you next need its exponent, and that comes from the cubic term.

Block 2. **Cue.** $\sin x = x - x^3/6 + O(x^5)$, so each step subtracts about $a_n^3/6$. **Why natural.** For a step law $x \mapsto x - cx^{p+1}$, the reciprocal power $b = x^{-p}$ increases by almost the constant cp each step, because $(x - cx^{p+1})^{-p} = x^{-p}(1 - cx^p)^{-p} \approx x^{-p} + cp$. Hence $b_n \approx cpn$ and $x_n \approx (cpn)^{-1/p}$. With $c = 1/6$ and $p = 2$ this predicts $a_n \approx \sqrt{3/n}$. This is the standard machine for neutral fixed points, and you should be able to run it in your head in fifteen seconds. **Link to the next block.** The prediction says $a_n^2 \approx 3/n$, which is not summable, so the target is a lower bound of order $n^{-1/2}$.

Block 3. **Cue.** The question is convergence or divergence, not the value. **Why natural.** Divergence of a series with nonnegative terms follows from a lower bound alone. So you may throw away the upper estimate, the error terms, and the sharp constant. Be lazy in the good sense: solve the question asked, not the question the structure suggests. **Link to the next block.** A lower bound of the form $a_n \geq cn^{-1/2}$ must be proved by induction, and induction needs a constant that leaves slack.

Block 4. **Cue.** You must compare sin against a polynomial, from below. **Why natural.** The truth is $a_n \sim \sqrt{3}\, n^{-1/2}$ and the inequality $a_n \geq \sqrt{3}\, n^{-1/2}$ is false for every n: the orbit approaches the sharp constant from below. So the induction can only close with a constant strictly smaller than $\sqrt{3}$, and 1 is the tidiest such constant. This is a general principle for asymptotic inductions: deliberately weaken the constant, because the slack is what the induction step spends. **Link to the next block.** With the target $a_n \geq n^{-1/2}$ fixed, you need exactly one analytic input, the lower Taylor bound for sin.

Block 5. **Cue.** A cubic lower bound for sin on $[0, 1]$. **Why natural.** Two standard justifications exist. Taylor with Lagrange remainder gives $\sin x = x - x^3/6 + \frac{\sin c}{24} x^4$ for some $c \in (0, x)$, and $\sin c \geq 0$ for $x \leq \pi$. Or, with no theorem quoted, a ladder of three derivative arguments: $x \geq \sin x$ implies $\cos x \geq 1 - x^2/2$ implies $\sin x \geq x - x^3/6$, each step being "value zero at zero, nonnegative derivative". Use the ladder in a write up: it is self contained and each line is checkable. **Link to the next block.** Now the induction step is an algebraic inequality between two positive quantities.

Block 6. **Cue.** An inequality between $\frac{6n-1}{6n^{3/2}}$ and $(n+1)^{-1/2}$, both positive. **Why natural.** Both sides positive means squaring is reversible, and squaring removes the square roots. What is left is a cubic inequality whose cubic terms cancel exactly, leaving the quadratic $24n^2 - 11n + 1 = (3n-1)(8n-1)$. The cancellation of the n^3 terms is not luck: it is the statement that the two sides agree to leading order, which is precisely the slack in the constant showing up as a lower order surplus. **Link to the next block.** Comparison with $\sum 1/n$ finishes.

Layer 3: WHY

The mathematical truth here is the classification of fixed points by the derivative. If $|f'(L)| < 1$ the orbit converges to L geometrically and every reasonable series built from it converges. If $f'(L) = 1$ the linear term gives no contraction and the first nonlinear term takes over: for $x_{n+1} = x_n - cx_n^{p+1} + O(x_n^{p+2})$ with $c > 0$ and $x_n \downarrow 0$, the substitution $b_n = x_n^{-p}$ converts the multiplicative iteration into an almost additive one, $b_{n+1} - b_n \to cp$, and the Cesaro theorem gives $b_n/n \to cp$, that is

$$x_n \sim (cpn)^{-1/p}.$$

Every sine, tangent, logarithm and arctangent iteration you will ever meet is an instance. For sin: $c = 1/6$, $p = 2$, so $a_n \sim \sqrt{3/n}$, and consequently $\sum a_n^q$ converges precisely when $q > 2$.

The discernment mechanism is a two question filter applied to any iteration problem. Question one: what is the limit and what is f' there? Question two: if $f' = 1$, what is the first nonlinear term? The answers give the decay exponent immediately, and the decay exponent decides every convergence question about the orbit. Train it by computing the exponent for $\sin x$, $\tan^{-1} x$, $\log(1 + x)$, $x - x^2$, and $x/(1 + x)$ until the substitution $b = x^{-p}$ is reflex.

The examiners nudged you twice. They started the orbit at $\pi/2$ rather than at 1 so that the first step, $a_1 = 1$, is free and puts the orbit inside $[0, 1]$ where all the estimates are comfortable. And they asked about $\sum a_n^2$ rather than $\sum a_n$ or $\sum a_n^3$: those two are decided by the crudest of bounds, whereas the square sits exactly at the harmonic boundary, so the answer depends on the constant in the decay being finite and positive rather than on any coarse estimate.

Revolving concepts

Neutral (parabolic) fixed point. For f smooth with $f(0) = 0$, $f'(0) = 1$ and $f(x) = x - cx^{p+1} + O(x^{p+2})$ with $c > 0$, any orbit starting small and positive decays like $(cpn)^{-1/p}$. Method: set $b_n = x_n^{-p}$, expand to get $b_{n+1} = b_n + cp + O(x_n^p)$, then apply the Cesaro theorem to the increments. See Book 1, field: Real analysis, topic: iteration near a neutral fixed point.

The generalised recursion. The orbit recursion $x_{n+1} = x_n - cx_n^{p+1}$ has no closed form, but its reciprocal power form $b_{n+1} = b_n + cp + o(1)$ is an additive recursion with the general solution $b_n = cpn + o(n)$. This is the standard method for every "slowly decaying" recursion: find the power of x that linearises the increment. Sharper versions add a $\log n$ correction: here $a_n^2 = \frac{3}{n} - \frac{27 \log n}{5n^2} + O(n^{-2})$, which is why $na_n^2 < 3$ for every n and approaches 3 from below. See Book 1, field: Sequences and recurrences, topic: asymptotics of slowly decaying recursions.

Stolz and Cesaro. If $b_{n+1} - b_n \to L$ then $b_n/n \to L$. This is the discrete analogue of l'Hopital, and it is the engine that converts an increment estimate into an asymptotic. See Book 1, field: Real analysis, topic: Cesaro means and Stolz.

Taylor bounds with a sign. For $x \geq 0$ the partial sums of the alternating series for sin and cos bracket the function: $x - \frac{x^3}{6} \leq \sin x \leq x$ and $1 - \frac{x^2}{2} \leq \cos x \leq 1$. Prove them by the derivative ladder rather than quoting a remainder theorem when the write up must be self contained. See Book 1, field: Calculus and integrals, topic: Taylor inequalities.

Comparison and p series. $\sum n^{-q}$ converges if and only if $q > 1$; for nonnegative terms, a lower bound by a divergent series forces divergence and an upper bound by a convergent series forces convergence. See Book 1, field: Real analysis, topic: comparison tests.

Circuit Bank

- **WHEN** a series is built from the orbit of an iteration **DO** find the limit, then the derivative of the map at the limit **BECAUSE** derivative less than one in modulus means geometric decay and instant convergence, derivative exactly one means polynomial decay and a real question

- **WHEN** the fixed point is neutral and $f(x) = x - cx^{p+1} + \cdots$ **DO** substitute $b_n = x_n^{-p}$ and read off $b_{n+1} - b_n \to cp$ **BECAUSE** the reciprocal power turns a multiplicative decay into an additive drift, and additive drift is linear growth
- **WHEN** you are asked only to decide convergence or divergence **DO** prove a one sided bound and throw away the other half of the analysis **BECAUSE** divergence of nonnegative series follows from a lower bound alone, so the sharp constant is wasted ink
- **WHEN** you must prove an asymptotic bound by induction **DO** use a constant strictly weaker than the true one **BECAUSE** the induction step spends the slack between your constant and the truth, and with the sharp constant the statement is usually false for every n
- **WHEN** an inequality compares two positive expressions with square roots **DO** square both sides, then check that the leading terms cancel **BECAUSE** cancellation of the leading terms is the visible form of the slack, and what remains factorises

Backtesting

1. First useful cue: $\sin 0 = 0$ with derivative 1, so the fixed point is neutral and decay is polynomial; and $a_1 = 1$.
2. Method the expression suggested: read the decay exponent off the cubic term ($a_n \approx \sqrt{3/n}$), then prove the weaker one sided bound $a_n \geq n^{-1/2}$ by induction.
3. Where you would hesitate, and the thought to reinforce: choosing the constant in the bound. The thought is that the true constant $\sqrt{3}$ is approached from below, so any induction must use a strictly smaller constant; take 1 and the algebra factorises.
4. The trap the question sets: computing the full asymptotic. It is the natural instinct, it is four times the work, and it answers a question that was not asked.
5. What makes you recognise it faster next time: any iteration by sin, $\tan^{-1}$, $\log(1+x)$ or $x - x^2$ is a neutral fixed point problem, and the exponent comes from the first nonlinear term.
6. The reusable decision rule: for convergence questions about orbits, get the decay exponent from the first nonlinear term, then prove only the inequality that the requested direction needs.

Mirror drills

1. With the same a_n, determine all real $q > 0$ for which $\sum a_n^q$ converges. *Answer:* $a_n \sim \sqrt{3/n}$, so $a_n^q \sim 3^{q/2} n^{-q/2}$ and the series converges exactly when $q > 2$.
2. Let $x_0 = 1/2$ and $x_{n+1} = x_n - x_n^2$. Show $\sum x_n$ diverges and $\sum x_n^2$ converges. *Answer:* here $p = 1$, $c = 1$, so $x_n \sim 1/n$; the first series is harmonic and the second is a convergent p series. For a written proof, induct on $x_n \geq \frac{1}{n+2}$.
3. Let $x_0 = 1$ and $x_{n+1} = \log(1 + x_n)$. Find the decay rate of x_n and decide whether $\sum x_n^2$ converges. *Answer:* $\log(1+x) = x - x^2/2 + \cdots$, so $p = 1$, $c = 1/2$ and $x_n \sim 2/n$; hence $\sum x_n^2$ converges while $\sum x_n$ diverges.

6.4 A4. Splitting a strip and a derangement limit

Consider a horizontal strip of $N+2$ squares in which the first and the last square are black and the remaining N squares are all white. Choose a white square uniformly at random, choose one of its two neighbours with equal probability, and colour this neighbouring square black if it is not already black. Repeat this process until all the remaining white squares have only black neighbours. Let $w(N)$ be the expected number of white squares remaining. Find

$$\lim_{N\to\infty} \frac{w(N)}{N}.$$

Read it out loud

This is a random process on a one dimensional lattice that stops when no two white squares are adjacent, and you are asked for the asymptotic density of survivors. Say the structure out loud: the state is a set of white squares inside a strip with black walls; each step picks a white square and one of its two neighbours; the step does something only if that neighbour is white. The stopping condition, "every remaining white square has only black neighbours", is exactly "the white squares form an independent set in the path". The fields wired in are probability (expectation, conditioning, linearity), self similar recursion (the strip splits), and finally the analysis of a first order recursion with a factorial solution.

Two features are planted for you. First, many steps do nothing: if the chosen neighbour is already black, the configuration is unchanged. Those null steps cannot possibly matter for the final answer, so the first structural move is to condition on the steps that do something. Second, the moment any interior square turns black, it becomes a wall, and the strip becomes two shorter strips of exactly the same type, each with black ends. That is a self similarity cue as loud as any you will meet: the object after one effective move is two smaller copies of the original object.

Within thirty seconds you should have written $w(0) = 0$, $w(1) = 1$, and the sentence "condition on the first square to change colour, then the strip splits". The requested limit $w(N)/N$ tells you that w is asymptotically linear, which is a hint about how to normalise later.

Approaches, ranked worst to best

1. **Simulate and guess.** *Cue:* a concrete random process. *Circuit:* estimate $w(N)/N$ numerically. (Gives about 0.37 and points at $1/e$, which is useful reconnaissance and no more. A Putnam answer of $1/e$ demands the exponential series.)
2. **Markov chain on configurations.** *Cue:* a random process with states. *Circuit:* write the transition matrix. (The state space is all subsets of the strip; nothing is computable and nothing simplifies.)
3. **Linearity of expectation over individual squares.** *Cue:* expectation of a count. *Circuit:* $w(N) = \sum_i \Pr[\text{square } i \text{ stays white}]$. (The right instinct in general, but here each of those probabilities depends on i and on N through the same self similar structure, so you must solve the recursion anyway. It postpones the work rather than doing it.)
4. **Condition on the first effective move, then guess the closed form and verify by induction.** *Cue:* the strip splits after one square changes colour. *Circuit:* recursion in N, then induction. (This is the official route and it is correct. Its weakness is the middle: after differencing you hold $w(N) = w(N-1) + w(N-2)/(N-1)$ and you must already know the answer $(N+1)\sum_{k\le N+1}(-1)^k/k!$ in order to verify it. A verification is not a derivation, and in the arena you may not have guessed the right closed form.)
5. **Condition on the first effective move, then solve the recursion by normalising.** *Cue:* the requested limit is $\lim w(N)/N$, so w is asymptotically linear. *Circuit:* substitute $w(N) =$

$(N+1)u_N$, difference, telescope. (Best. The substitution is chosen so that the only denominator cancels, the recursion collapses to $d_N = -d_{N-1}/(N+1)$ for the increments $d_N = u_N - u_{N-1}$, and telescoping produces the alternating factorial series with no guessing at all. The exponential appears as the limit of its own series, which is the only honest way to arrive at $1/e$.)

The ideal solution

The answer is $\boxed{1/e}$.

Number the squares 1 to $N+2$ from left to right, so squares 1 and $N+2$ are black and squares 2 to $N+1$ are white initially. Set $w(0) = 0$ and $w(1) = 1$, both immediate: with $N = 1$ the single white square already has two black neighbours.

Step 1: only effective moves matter. A move is the ordered pair (chosen white square, chosen neighbour). Call a move *effective* if the chosen neighbour is white. A move that is not effective leaves the configuration unchanged, so the final configuration is a function of the sequence of effective moves alone. If W white squares remain, every pair (white square, direction) has probability $\frac{1}{2W}$, so conditional on a move being effective it is uniformly distributed over the effective pairs available at that moment. As long as two white squares are adjacent there is at least one effective pair, so each move is effective with probability at least $\frac{1}{2N}$, and therefore effective moves occur almost surely; each one decreases the number of white squares by one, so the process terminates after at most N effective moves.

Step 2: the first square to change colour. Let $N \geq 2$. Initially the effective pairs are: square i with its left neighbour, for $3 \leq i \leq N+1$, colouring square $i-1$; and square i with its right neighbour, for $2 \leq i \leq N$, colouring square $i+1$. There are $2(N-1)$ of them, all equally likely. Square j is produced by exactly two of these pairs when $3 \leq j \leq N$, and by exactly one when $j = 2$ or $j = N+1$. Hence the first square to be coloured is j with probability

$$p_j = \frac{1}{N-1}\ (3 \leq j \leq N), \qquad p_2 = p_{N+1} = \frac{1}{2(N-1)}.$$

Step 3: the strip splits. Once square j is black, no future move involves a square on the left of j together with a square on the right of j, since every move acts on a pair of adjacent squares. So the configuration is a system on squares 1 to j (black at both ends, $j-2$ white inside) and a system on squares j to $N+2$ (black at both ends, $N+1-j$ white inside). By Step 1, the effective moves available at any later time are the disjoint union of those available in the two systems, and a uniform choice from a disjoint union restricts to a uniform choice within each part; hence, viewed only at its own effective moves, each system evolves exactly as a standalone system of its size. By linearity of expectation (independence of the two sides is not needed), the conditional expectation of the final number of white squares, given that square j was coloured first, is $w(j-2) + w(N+1-j)$.

Step 4: the recursion. For $N \geq 2$, using $w(0) = 0$ and the reversal $j \mapsto N+3-j$ inside the sum,

$$w(N) = \frac{w(N-1)+w(0)}{N-1} + \frac{1}{N-1}\sum_{j=3}^{N}\big(w(j-2)+w(N+1-j)\big),$$

so

$$(N-1)\,w(N) = w(N-1) + 2\sum_{m=1}^{N-2} w(m).$$

Step 5: remove the sum. Writing that identity with $N-1$ in place of N (valid for $N \geq 3$) and subtracting,

$$(N-1)w(N) - (N-2)w(N-1) = w(N-1) - w(N-2) + 2w(N-2),$$

that is $(N-1)w(N) = (N-1)w(N-1) + w(N-2)$, so

$$w(N) = w(N-1) + \frac{w(N-2)}{N-1} \qquad (N \geq 3),$$

and it also holds at $N = 2$, where both sides equal 1.

Step 6: solve it. Substitute $w(N) = (N+1)u_N$, the factor $N+1$ being chosen so that the coefficient $\frac{1}{N-1}$ cancels. Then the recursion becomes $(N+1)u_N = Nu_{N-1} + u_{N-2}$, and subtracting $(N+1)u_{N-1}$ from both sides,

$$(N+1)(u_N - u_{N-1}) = -(u_{N-1} - u_{N-2}).$$

So the increments $d_N = u_N - u_{N-1}$ satisfy $d_N = -\frac{d_{N-1}}{N+1}$ for $N \geq 2$, with $d_1 = u_1 - u_0 = \frac{1}{2} - 0 = \frac{1}{2}$. Telescoping the product,

$$d_N = (-1)^{N-1}\frac{1}{2}\prod_{i=2}^{N}\frac{1}{i+1} = (-1)^{N-1}\frac{1}{2}\cdot\frac{2}{(N+1)!} = \frac{(-1)^{N-1}}{(N+1)!},$$

and summing the increments from $u_0 = 0$,

$$u_N = \sum_{j=1}^{N}\frac{(-1)^{j-1}}{(j+1)!} = \sum_{k=2}^{N+1}\frac{(-1)^k}{k!} = \sum_{k=0}^{N+1}\frac{(-1)^k}{k!},$$

the last step because the terms $k = 0$ and $k = 1$ cancel. Hence

$$w(N) = (N+1)\sum_{k=0}^{N+1}\frac{(-1)^k}{k!} \qquad (N \geq 0).$$

Step 7: the limit. Since $\sum_{k\geq 0}(-1)^k/k! = e^{-1}$,

$$\lim_{N\to\infty}\frac{w(N)}{N} = \lim_{N\to\infty}\frac{N+1}{N}\sum_{k=0}^{N+1}\frac{(-1)^k}{k!} = 1\cdot e^{-1} = \frac{1}{e}.$$

Layer 1: WHAT

Block 1: discard null moves and work with the embedded sequence of effective moves. Block 2: compute the law of the first square to change colour, $2(N-1)$ equally likely pairs. Block 3: observe that a black interior square is a wall, so the strip becomes two strips of the same type, and apply linearity of expectation. Block 4: assemble the recursion $(N-1)w(N) = w(N-1) + 2\sum_{m\leq N-2} w(m)$. Block 5: difference two consecutive instances to kill the sum, giving $w(N) = w(N-1) + w(N-2)/(N-1)$. Block 6: substitute $w(N) = (N+1)u_N$, get $d_N = -d_{N-1}/(N+1)$ for the increments, telescope, recognise the alternating exponential series. Block 7: take the limit.

Juxtaposition. The nominal route conditions on the first move, including the moves that change nothing, and immediately has to describe the configuration after a null move, which is the same configuration; the recursion then refers to itself with no progress and the candidate abandons it for simulation. A second nominal route computes Pr[square i survives] for each i and drowns in boundary cases. The bifurcation point is the sentence "choose a white square and one of its neighbours". The zombie treats every choice as a step of the process; you separate the moves that alter the state from those that do not, and condition on the former. Everything downstream follows from that one sentence.

Layer 2: HOW

Block 1. **Cue.** "Colour this neighbouring square black *if it is not already black*." **Why natural.** The conditional clause is the examiners telling you that some moves are null. Null moves do not change the state, so they cannot change the final configuration; they only change how long the process takes, and you were not asked about the time. Conditioning on effectiveness is legitimate because all pairs are equally likely at any instant, so the conditional law of an effective move is uniform on the currently effective pairs. **Link to the next block.** With null moves removed, the process is a sequence of state changes and you may condition on the first one.

Block 2. **Cue.** A finite, explicitly listable set of possible first changes. **Why natural.** Count pairs, not squares. Each interior square j with $3 \le j \le N$ can be reached from the left or from the right, whereas squares 2 and $N+1$ can only be reached from one side, since square 1 and square $N+2$ are already black and are never chosen. That asymmetry at the two ends is exactly the factor $\frac{1}{2}$ in p_2 and p_{N+1}, and it is the standard place to lose the problem. **Link to the next block.** You now have a probability for each possible first wall, and the wall is what makes the recursion.

Block 3. **Cue.** Every move acts on two adjacent squares, and the newly black square sits between the two halves. **Why natural.** A black square is absorbing and is never chosen (only white squares are chosen), so it acts as a permanent wall identical to the walls at the two ends. The two halves are therefore instances of the same problem with smaller parameters. Note what you do *not* need: independence of the two halves. Expectation is linear, so the marginal law on each side is enough, and the marginal law is the standalone process by the uniform restriction argument. **Link to the next block.** Self similarity plus known probabilities gives a closed recursion.

Block 4. **Cue.** A conditional expectation summed against p_j. **Why natural.** The sum $\sum_{j=3}^{N}(w(j-2)+w(N+1-j))$ runs over the same list of values twice in opposite orders, so it is $2\sum_{m=1}^{N-2} w(m)$. Multiply through by $N-1$ so that no fractions survive. A recursion in which $w(N)$ depends on all earlier values is a signal, not an obstacle. **Link to the next block.** A recursion containing $\sum_{m<N} w(m)$ is always attacked the same way.

Block 5. **Cue.** The recursion contains a running sum with a coefficient linear in N. **Why natural.** Write the same identity for $N-1$ and subtract: the two running sums differ by exactly one term, so the sum disappears and a short recursion remains. This is the standard device for every recursion of quicksort type, and it is why such recursions have clean solutions. **Link to the next block.** You now hold a two term recursion with a variable coefficient, which needs a normalisation before it will telescope.

Block 6. **Cue.** The question asks for $\lim w(N)/N$, and the recursion has the single denominator $N-1$. **Why natural.** Since the answer is a limit of $w(N)/N$, the natural unknown is $w(N)$ divided by a linear factor. Choose the factor to kill the denominator: with $w(N) = (N+a)u_N$ the last term becomes $\frac{N-2+a}{N-1}u_{N-2}$, which is exactly u_{N-2} when $a = 1$. That is the whole reason for the substitution; it is a choice made to cancel, not a guess. The resulting relation is a first order recursion in the increments, and first order recursions with rational ratios telescope into factorials, which is where $\sum(-1)^k/k!$ and hence e^{-1} come from. **Link to the next block.** Summing the increments gives the closed form and the limit follows from the exponential series.

Layer 3: WHY

The truth beneath this problem is self similar fragmentation. A random process that cuts a structure into two independent smaller copies of itself produces a recursion of the form $E_N = \sum_j p_j(E_{L(j)} +$

$E_{R(j)}$), and when the cut position is close to uniform the recursion always has the same two features: a running sum, removable by differencing, and a variable coefficient, removable by a linear normalisation. Quicksort, random binary search trees, and the parking problems all live in this family. The only thing that changes from problem to problem is the law of the cut, and here it is uniform on $2(N-1)$ ordered pairs, which is uniform on the cut positions apart from a halving at the two extreme positions.

The second truth is the appearance of e^{-1}. The closed form $w(N) = (N+1)\sum_{k=0}^{N+1}(-1)^k/k!$ says exactly that

$$\frac{w(N)}{N+1} = \frac{D_{N+1}}{(N+1)!} = \Pr[\text{a uniform random permutation of } N+1 \text{ letters has no fixed point}],$$

where D_m is the number of derangements; indeed $w(N) = D_{N+1}/N!$ turns the recursion $w(N) = w(N-1) + \frac{w(N-2)}{N-1}$ into the classical $D_m = (m-1)(D_{m-1} + D_{m-2})$. That identity is proved here through the recursion, and a direct combinatorial explanation via random permutations is known and worth hunting for. Whenever a limit of $1/e$ appears in a discrete problem, look for an inclusion and exclusion or a derangement in the background; the alternating factorial series is its signature.

The discernment mechanism: when a random process is described procedurally, ask three questions in order. Which moves change the state? What does one state change do to the structure? Is the resulting structure a smaller copy of the original? If the answers are "some", "it splits it", and "yes", you write a recursion and stop thinking about probability. Train it on this problem by deliberately writing the recursion before computing a single probability, then filling the probabilities in.

Revolving concepts

Embedded (jump) process. If a Markov process has moves that leave the state unchanged, the sequence of states visited is unchanged by deleting them, and conditional on a move being a state change, its law is the original law restricted and renormalised to the state changing moves. Use this whenever a process description contains "if it is not already". See Book 1, field: Probability, topic: jump chains and uniformisation.

Linearity of expectation without independence. $E[X+Y] = E[X] + E[Y]$ always holds. In splitting arguments this means you need only each part's marginal law, never joint independence, which is usually the harder thing to justify. See Book 1, field: Probability, topic: linearity of expectation.

Fragmentation recursions and the differencing trick. A recursion of the form $c_N a_N = b_N + \sum_{m<N} \lambda_m a_m$ is solved by writing the same identity at $N-1$ and subtracting, which reduces the order to two or one. The general shape here is the quicksort recursion $Na_N = f(N) + 2\sum_{m<N} a_m$, giving $Na_N - (N-1)a_{N-1} = f(N) - f(N-1) + 2a_{N-1}$. See Book 1, field: Sequences and recurrences, topic: recursions with running sums.

First order recursions with rational ratios. $d_N = r(N)d_{N-1}$ has the solution $d_N = d_1 \prod_{i=2}^{N} r(i)$. When $r(N) = \pm\frac{1}{N+c}$ the product is a reciprocal factorial, so the increments are $\pm 1/(N+c)!$ and their sum is a truncated exponential series. This is the general mechanism by which e and e^{-1} enter recursion problems. See Book 1, field: Sequences and recurrences, topic: telescoping products.

Derangements. $D_m = m!\sum_{k=0}^{m}\frac{(-1)^k}{k!}$, $D_m = mD_{m-1} + (-1)^m$, $D_m = (m-1)(D_{m-1} + D_{m-2})$, and $D_m/m! \to e^{-1}$. Here $w(N) = D_{N+1}/N!$. See Book 1, field: Combinatorics, topic: inclusion and exclusion, derangements.

Independent sets in a path as a stopping condition. "Every white square has only black neighbours" means the white set is independent in the path graph. The maximum such set has size $\lceil N/2 \rceil$, so $w(N) \le \lceil N/2 \rceil$ and the limit is at most $1/2$; the answer $1/e \approx 0.368$ sits comfortably below. Use such crude bounds as a sanity check on any answer you produce. See Book 1, field: Discrete structures and graphs, topic: independent sets in paths.

Circuit Bank

- **WHEN** a process description contains "if it is not already" **DO** condition on the moves that actually change the state and ignore the rest **BECAUSE** null moves cannot affect the final configuration, and conditioning on effectiveness keeps the uniform law over the moves that matter
- **WHEN** one move turns a structure into two smaller copies of itself **DO** write a recursion by conditioning on where the split happens, using linearity of expectation **BECAUSE** expectation is additive over the two parts, so only the marginal law on each part is needed and independence never has to be proved
- **WHEN** a recursion contains a sum of all earlier terms **DO** write the identity at $N-1$ and subtract **BECAUSE** the two running sums differ by one term, so differencing lowers the order and leaves a short recursion
- **WHEN** you must solve $a_N = a_{N-1} + \frac{a_{N-2}}{N-1}$ or similar **DO** substitute $a_N = (N+c)u_N$ with c chosen so the fractional coefficient becomes 1 **BECAUSE** the substitution is engineered to cancel the denominator, after which the increments obey a first order recursion that telescopes
- **WHEN** a discrete limit turns out to be $1/e$ **DO** look for derangements, inclusion and exclusion, or a truncated alternating factorial series **BECAUSE** the alternating series $\sum(-1)^k/k!$ is the only common source of e^{-1} in a counting problem

Backtesting

1. First useful cue: "if it is not already black" announces null moves, and a newly blackened interior square is a wall that splits the strip.
2. Method the expression suggested: condition on the first effective move, use self similarity and linearity of expectation, then solve the recursion.
3. Where you would hesitate, and the thought to reinforce: the endpoints $j = 2$ and $j = N+1$ have half the probability of the interior positions, because they can be reached from only one side. Count ordered pairs (square, direction), never squares, and the halving appears automatically.
4. The trap the question sets: trying to compute the survival probability of each individual square, or trying to track the process in real time including null moves. Both are correct in principle and neither terminates.
5. What makes you recognise it faster next time: a random process on an interval where one event creates a permanent wall is always a fragmentation recursion, and fragmentation recursions are always attacked by differencing then normalising.
6. The reusable decision rule: delete null moves, condition on the first real move, exploit self similarity with linearity of expectation, difference away any running sum, normalise by a linear factor chosen to cancel the denominator.

Mirror drills

1. From $(N-1)w(N) = w(N-1) + 2\sum_{m=1}^{N-2} w(m)$, derive $w(N) = w(N-1) + \frac{w(N-2)}{N-1}$. *Answer:* write the identity at $N-1$, subtract, and note the two running sums differ by the single term $2w(N-2)$.
2. Solve $d_N = -\frac{d_{N-1}}{N+1}$ for $N \geq 2$ with $d_1 = \frac{1}{2}$. *Answer:* $d_N = \frac{1}{2}(-1)^{N-1}\prod_{i=2}^{N}(i+1)^{-1} = \frac{(-1)^{N-1}}{(N+1)!}$.

3. Show that $D_m = (m-1)(D_{m-1} + D_{m-2})$ with $D_1 = 0$, $D_2 = 1$ implies $D_m = mD_{m-1} + (-1)^m$, and hence $D_m/m! = \sum_{k=0}^{m} (-1)^k/k!$. *Answer:* set $e_m = D_m - mD_{m-1}$; the first recursion gives $e_m = -e_{m-1}$ with $e_2 = 1$, so $e_m = (-1)^m$; dividing $D_m = mD_{m-1} + (-1)^m$ by $m!$ telescopes.

6.5 A5. Counting Fibonacci representations by complementation

Let a_n be the number of sets S of positive integers for which

$$\sum_{k\in S} F_k = n,$$

where the Fibonacci sequence $(F_k)_{k\geq 1}$ satisfies $F_{k+2} = F_{k+1} + F_k$ and begins $F_1 = 1$, $F_2 = 1$, $F_3 = 2$, $F_4 = 3$. Find the largest integer n such that $a_n = 2020$.

Read it out loud

This is a representation counting function. The sets S are sets of *indices*, not of values, so F_1 and F_2 are two different summands that happen both to equal 1; that is why $a_1 = 2$ and not 1. The command term is "find the largest n with $a_n = 2020$", which carries two hidden assertions: that a is not monotone (otherwise you would say "the" n), and that $\{n : a_n \leq 2020\}$ is finite, so that some last n exists. Your solution must supply both an n with $a_n = 2020$ and a proof that everything beyond it is strictly larger.

The fields wired in are Fibonacci arithmetic, representation counting in a numeration system, and one involution. The structural cue is the identity

$$F_1 + F_2 + \cdots + F_{m-1} = F_{m+1} - 1.$$

Read that out loud: the total of all the summands available below index m is one less than a Fibonacci number. Whenever the total of your available weights is a known number T, the map "take the complementary subset" is an involution turning a representation of n into a representation of $T - n$. That is the entire engine of this problem: it converts the case "the top index is not used" into another instance of the same counting function.

Within thirty seconds you should have $a_0 = 1$ (empty set), $a_1 = a_2 = 2$, $a_3 = a_4 = 3$, and the observation that Fibonacci numbers cut the integers into blocks $[F_m, F_{m+1})$. Answers of the form $F_{4040} - 1$ live at block boundaries, so the whole game will be to evaluate a at the two landmarks F_m and $F_{m+1} - 1$ and to bound it from below in between.

Approaches, ranked worst to best

1. **Tabulate a_n and look for a pattern.** *Cue:* a concrete counting function. *Circuit:* compute by dynamic programming and stare. (Essential reconnaissance and worthless as a proof: the answer involves F_{4040}, a number with over eight hundred digits.)

2. **Generating function $\prod_{k\geq 1}(1 + x^{F_k})$.** *Cue:* counting subsets with a prescribed sum. *Circuit:* coefficients of a product. (Correct as a definition, useless as a tool: the product has no closed form and no functional equation you can exploit here.)

3. **Zeckendorf plus local rewriting.** *Cue:* every representation is reachable from the greedy one. *Circuit:* repeatedly replace F_k by $F_{k-1} + F_{k-2}$ and count reachable states. (This is the official second route. It genuinely works and even gives a product formula over the blocks of the Zeckendorf bitstring, but the bookkeeping of the strings $(10)^k 0^\ell$ and their rewritings is long, and the final inequality needs a separate argument.)

4. **Split on the largest usable index, then complement.** *Cue:* the top index is either used or not, and the total below it is $F_{m+1} - 1$. *Circuit:* $a_n = a_{n-F_m} + a_{F_{m+1}-1-n}$ for $F_m \leq n < F_{m+1}$. (Best. Two cases, one line each, and the second case is handled by an involution rather than by a new idea. The recursion evaluates the landmarks immediately and drives the lower bound by induction.)

The ideal solution

The answer is $\boxed{n = F_{4040} - 1}$.

Throughout, $a_0 = 1$, counting the empty set. Two elementary facts:

Lemma 6.5.1 (partial sums, and existence of a representation). $\sum_{k=1}^{m-1} F_k = F_{m+1} - 1$ *for every* $m \geq 1$, *and* $a_n \geq 1$ *for every* $n \geq 0$.

Proof. For the identity, $F_k = F_{k+2} - F_{k+1}$ makes the sum telescope to $F_{m+1} - F_2 = F_{m+1} - 1$; for $m = 1$ the empty sum is $0 = F_2 - 1$. For the second statement, induct strongly on n: for $n \geq 1$ choose the largest m with $F_m \leq n$; then $n - F_m < F_{m+1} - F_m = F_{m-1} \leq F_m$, so by induction $n - F_m$ has a representation, and every index k occurring in it satisfies $F_k \leq n - F_m < F_m$, hence $k < m$; adjoining m gives a representation of n. □

Lemma 6.5.2 (top index recursion). *Let* $m \geq 2$ *and* $F_m \leq n < F_{m+1}$. *Then*

$$a_n = a_{n-F_m} + a_{F_{m+1}-1-n}.$$

Proof. Every S counted by a_n satisfies $\max S \leq m$, since $F_k \geq F_{m+1} > n$ for $k \geq m+1$. Split on whether $m \in S$.

Suppose $m \in S$. Then $S \setminus \{m\} \subseteq \{1, \ldots, m-1\}$ has sum $n - F_m$. Conversely, let T be *any* set with sum $n - F_m$. Since $n - F_m < F_{m+1} - F_m = F_{m-1} \leq F_m$, every index $k \in T$ has $F_k \leq n - F_m < F_m$, so $k < m$ and $T \cup \{m\}$ is counted by a_n. These two maps are mutually inverse, so this case contributes a_{n-F_m}.

Suppose $m \notin S$. Then $S \subseteq \{1, \ldots, m-1\}$, and by the partial sums lemma the complementary set $\{1, \ldots, m-1\} \setminus S$ has sum $(F_{m+1} - 1) - n$. Conversely, let U be *any* set with sum $F_{m+1} - 1 - n$. Since $n \geq F_m$ we have $F_{m+1} - 1 - n \leq F_{m-1} - 1 < F_m$, so as before $U \subseteq \{1, \ldots, m-1\}$, and $\{1, \ldots, m-1\} \setminus U$ is counted by a_n and omits m. Complementation inside $\{1, \ldots, m-1\}$ is an involution, so this case contributes $a_{F_{m+1}-1-n}$. □

Write $c_m = a_{F_m - 1}$ for $m \geq 2$, so $c_2 = a_0 = 1$ and $c_3 = a_1 = 2$.

Lemma 6.5.3 (the two landmarks). *For* $m \geq 2$,

$$a_{F_m} = a_{F_{m+1}-1} = 1 + c_{m-1} \quad (m \geq 3), \qquad \text{and} \qquad c_m = \left\lfloor \frac{m+1}{2} \right\rfloor, \qquad a_{F_m} = \left\lfloor \frac{m+2}{2} \right\rfloor.$$

Proof. Apply the top index recursion at the two ends of the block $[F_m, F_{m+1})$. At $n = F_m$ it gives $a_{F_m} = a_0 + a_{F_{m-1}-1} = 1 + c_{m-1}$. At $n = F_{m+1} - 1$, which lies in the block because $F_{m+1} - 1 \geq F_m$, it gives $a_{F_{m+1}-1} = a_{F_{m-1}-1} + a_0 = 1 + c_{m-1}$. The second of these says $c_{m+1} = c_{m-1} + 1$ for $m \geq 3$, and one checks $c_4 = a_2 = 2 = c_2 + 1$ directly, so $c_{m+1} = c_{m-1} + 1$ for all $m \geq 3$. With $c_2 = 1$ and $c_3 = 2$, the even and odd subsequences are arithmetic with common difference 1, giving $c_m = \lfloor (m+1)/2 \rfloor$. Hence $a_{F_m} = 1 + \lfloor m/2 \rfloor = \lfloor (m+2)/2 \rfloor$ for $m \geq 3$, and also for $m = 2$ since $a_{F_2} = a_1 = 2$. □

Lemma 6.5.4 (block minimum). *For* $m \geq 2$ *and* $F_m \leq n < F_{m+1}$ *we have* $a_n \geq \lfloor (m+2)/2 \rfloor$.

Proof. Strong induction on m. For $m = 2, 3, 4$ the blocks are $\{1\}$, $\{2\}$, $\{3, 4\}$ and the values $a_1 = a_2 = 2$, $a_3 = a_4 = 3$ meet the bounds $2, 2, 3$.

Let $m \geq 5$ and put $j = n - F_m$ and $j' = F_{m+1} - 1 - n$, so that $j, j' \geq 0$, $j + j' = F_{m-1} - 1$, and $a_n = a_j + a_{j'}$ by the top index recursion. The expression is symmetric in j and j', so assume $j \leq j'$.

If $j = 0$, then $j' = F_{m-1} - 1$, and $F_{m-2} \leq F_{m-1} - 1 < F_{m-1}$ because $F_{m-3} \geq 1$. So j' lies in block $m - 2$ and the induction hypothesis gives $a_{j'} \geq \lfloor m/2 \rfloor$, whence $a_n = 1 + a_{j'} \geq 1 + \lfloor m/2 \rfloor = \lfloor (m+2)/2 \rfloor$.

If $j \geq 1$, then $a_j \geq 2$: indeed j lies in some block, so the top index recursion writes a_j as a sum of two values of a, each at least 1 by the partial sums lemma. For j', note $j' \geq \frac{1}{2}(F_{m-1} - 1) \geq F_{m-3}$, the

last inequality being equivalent to $F_{m-3}+1 \le F_{m-2}$, which holds since $m \ge 5$. Also $j' \le F_{m-1}-1 < F_{m-1}$. So j' lies in a block with index q satisfying $m-3 \le q \le m-2$, and the induction hypothesis gives $a_{j'} \ge \lfloor (m-1)/2 \rfloor$. Therefore

$$a_n \ge 2 + \left\lfloor \frac{m-1}{2} \right\rfloor \ge \left\lfloor \frac{m+2}{2} \right\rfloor,$$

with equality when m is even and strict inequality when m is odd. □

Now finish. By the landmark lemma, $a_{F_{4040}-1} = c_{4040} = \lfloor 4041/2 \rfloor = 2020$. If $n > F_{4040}-1$, then $n \ge F_{4040}$, so n lies in a block $[F_m, F_{m+1})$ with $m \ge 4040$, and the block minimum lemma gives

$$a_n \ge \left\lfloor \frac{m+2}{2} \right\rfloor \ge \left\lfloor \frac{4042}{2} \right\rfloor = 2021 > 2020.$$

Hence no n larger than $F_{4040}-1$ has $a_n = 2020$, and $n = F_{4040}-1$ is the largest such integer.

Layer 1: WHAT

Block 1: fix the convention $a_0 = 1$ and compute the first few values, noticing the equal pairs $a_1 = a_2$, $a_3 = a_4$. Block 2: prove $\sum_{k \le m-1} F_k = F_{m+1} - 1$, the identity that makes complementation available. Block 3: split a representation on whether it uses the largest usable index m, and convert the "not used" case by complementing inside $\{1, \dots, m-1\}$; this gives $a_n = a_{n-F_m} + a_{F_{m+1}-1-n}$. Block 4: evaluate at the two ends of the block to get $a_{F_m} = a_{F_{m+1}-1} = 1 + a_{F_{m-1}-1}$, a lag two recursion, hence $\lfloor (m+1)/2 \rfloor$ and $\lfloor (m+2)/2 \rfloor$. Block 5: prove that the block minimum is attained at the ends, by induction using the same recursion. Block 6: read off the last block where the value 2020 can occur.

Juxtaposition. The nominal route computes a_n for n up to a few hundred, spots that the small values of a occur at $n = F_m$ and $n = F_m - 1$, guesses the answer $F_{4040}-1$, and then has no mechanism to prove that a_n never returns to 2020 later. The gap is not the formula, it is the monotone lower envelope. The ideal route builds one recursion that serves both purposes: evaluating the landmarks and bounding the interior. The bifurcation point is the sentence "suppose $m \notin S$". The zombie treats that case as a new, harder counting problem over $\{1, \dots, m-1\}$; you complement it and land back inside the same function a.

Layer 2: HOW

Block 1. **Cue.** Sets of indices, with $F_1 = F_2 = 1$ both available. **Why natural.** The problem says "sets S of positive integers" and the sum runs over $k \in S$, so the objects are index sets. Computing $a_1 = 2$ immediately confirms it and stops you from silently deduplicating the two ones. Small values are also what tell you the answer will be phrased at Fibonacci landmarks. **Link to the next block.** To say anything about a_n you must control which indices can appear, and that is decided by the block containing n.

Block 2. **Cue.** $F_m \le n < F_{m+1}$ bounds the largest usable index. **Why natural.** If $k \ge m+1$ then $F_k > n$, so index k cannot appear at all. The Fibonacci blocks are therefore the natural coordinates for this function, exactly as powers of ten are the natural coordinates for a digit problem. Every statement you prove should be indexed by the block. **Link to the next block.** Inside a block, the only structural question left is whether the top index m is used.

Block 3. **Cue.** A binary choice: is m in S or not. **Why natural.** This is the standard first move for counting representations in any numeration system, because removing the largest weight reduces the problem to the same problem with a smaller target. The "used" branch is immediate. The "not used" branch would normally be a different problem (subsets of a restricted index range), and the

reason it is not is the identity $\sum_{k\le m-1} F_k = F_{m+1} - 1$: complementation inside $\{1,\dots,m-1\}$ is a bijection onto representations of $F_{m+1}-1-n$, and, crucially, every representation of that number automatically lives inside $\{1,\dots,m-1\}$, so the count is the unrestricted $a_{F_{m+1}-1-n}$. **Link to the next block.** You now have a self referential recursion whose two arguments add up to $F_{m-1}-1$.

Block 4. **Cue.** The two arguments j and j' satisfy $j + j' = F_{m-1} - 1$, so at the ends of the block one of them is 0. **Why natural.** Evaluating a recursion at the extreme cases is free information, and here both ends give the same value, which explains the observed pairs $a_1 = a_2$ and $a_3 = a_4$. The resulting relation $c_{m+1} = c_{m-1} + 1$ has lag two, so it splits by parity into two arithmetic progressions; that parity split is where the floor function in the answer is born. **Link to the next block.** The landmarks give candidates for $a_n = 2020$; to know they are the last ones you must bound the interior of every later block.

Block 5. **Cue.** A recursion that expresses a_n as a sum of two values at smaller arguments. **Why natural.** Any such recursion converts a lower bound on all earlier blocks into a lower bound on the current one, provided you can locate the blocks of j and j'. With $j \le j'$ you always have $j' \ge \frac{1}{2}(F_{m-1}-1)$, and since F_{m-3} is less than half of F_{m-1} (the ratio tends to $\varphi^{-2} \approx 0.382$), j' sits in block $m-3$ or later. The other term is at least 2 whenever $j \ge 1$, and at least 1 always. Those two crude facts are exactly enough. **Link to the next block.** The block minimum is $\lfloor (m+2)/2 \rfloor$, an increasing function of m, so the set of n with $a_n \le 2020$ is finite and its last element is at the end of block 4039.

Block 6. **Cue.** $\lfloor (m+2)/2 \rfloor = 2021$ first happens at $m = 4040$. **Why natural.** $c_m = 2020$ has the two solutions $m = 4039$ and $m = 4040$, so $a_n = 2020$ at $n = F_{4039} - 1$ and at $n = F_{4040} - 1$; the larger is the answer. Everything from F_{4040} onwards is at least 2021 by the block minimum. State both halves: the value is attained, and nothing later attains it. **Link to the next block.** Nothing remains.

Layer 3: WHY

The truth beneath this problem is that a representation counting function in a numeration system is governed by a recursion on the leading digit, and that Fibonacci weights have one extra gift: the partial sums of the weights are again (nearly) weights. In base ten, the analogue of the identity $\sum_{k\le m-1} F_k = F_{m+1} - 1$ is $\sum_{k\le m-1} 9 \cdot 10^k = 10^m - 1$, which is exactly what makes "nines complement" arithmetic work. Complementation is available whenever the total weight is a landmark, and it converts an "omitted" case into an "included" case at a reflected target. That is why the second branch of the recursion costs nothing.

The second truth is about extremal questions for non monotone functions. To find the last n with $a_n = c$, you need a monotone lower envelope: a function $\mu(m)$, increasing, with $a_n \ge \mu(m)$ throughout block m, and equality somewhere. Here $\mu(m) = \lfloor (m+2)/2 \rfloor$, and equality holds at both ends of each block. The general shape of such an argument is: recursion, then induction on the block index, then the observation that the minimum of a sum is attained when the split is as lopsided as possible. Lopsided splits give one term equal to $a_0 = 1$, so representations of F_m and of $F_{m+1} - 1$ are the scarcest, and that scarcity grows by one every two blocks because the recursion has lag two.

The discernment mechanism: whenever a problem counts representations, ask what the largest usable weight is, split on it, and then look for an identity that lets you reflect the omitted case. Train it by proving the same recursion for base b digits and for the Zeckendorf system, and by computing the first twenty values of a_n by hand so that the landmark pattern is familiar rather than surprising. The examiners nudged you by choosing $2020 = \lfloor 4041/2 \rfloor$, a value that lands exactly on a landmark, and by starting the Fibonacci indexing at $F_1 = F_2 = 1$, which doubles every count at the bottom and forces you to read the definition carefully.

Revolving concepts

Zeckendorf representation. Every positive integer is a sum of Fibonacci numbers with indices at least 2, no two consecutive, and this representation is unique; it is produced by the greedy algorithm. The count a_n is the number of *all* representations, of which the Zeckendorf one is a single member, and every other representation is obtained from it by repeatedly rewriting $F_k \to F_{k-1} + F_{k-2}$. See Book 1, field: Number theory, topic: numeration systems and Zeckendorf.

Fibonacci partial sums. $\sum_{k=1}^{m} F_k = F_{m+2} - 1$, proved by telescoping $F_k = F_{k+2} - F_{k+1}$. Companions worth banking: $\sum F_{2k-1} = F_{2m}$, $\sum F_{2k} = F_{2m+1} - 1$, $\sum F_k^2 = F_m F_{m+1}$. See Book 1, field: Sequences and recurrences, topic: Fibonacci identities.

Complementation involution. If the available weights have total T, then $S \mapsto S^c$ is an involution on subsets carrying sum n to sum $T - n$, so the two counts are equal. This is the only content of the second branch of the recursion, and it is a general tool: any subset sum count is symmetric about $T/2$. See Book 1, field: Combinatorics, topic: involutions and symmetric subset sums.

Leading weight recursion (the general form). Let $u_1 \le u_2 \le \cdots$ be weights with partial sums $U_m = u_1 + \cdots + u_m$, and let $r(n)$ be the number of subsets of indices with weight sum n. If $u_m \le n$ and every representation of n omits all indices above m, then

$$r(n) = r(n - u_m) + \#\{\text{subsets of } \{1, \ldots, m-1\} \text{ of sum } n\} = r(n - u_m) + r(U_{m-1} - n),$$

the last equality when every representation of $U_{m-1} - n$ automatically avoids indices above $m - 1$. This is the general digit dynamic programme; the Fibonacci case is special only because $U_{m-1} = F_{m+1} - 1$ is itself a landmark. See Book 1, field: Combinatorics, topic: counting representations in numeration systems.

Lag two recursions. $c_{m+1} = c_{m-1} + d$ splits into two independent arithmetic progressions along the even and odd indices, giving $c_m = c_{m_0} + d\lfloor (m - m_0)/2 \rfloor$ on each parity class, which is why floor functions of the form $\lfloor (m+1)/2 \rfloor$ appear in the answer. The same split handles $u_{m+2} = u_m + f(m)$ in general: sum f over the arithmetic progression of the appropriate parity. See Book 1, field: Sequences and recurrences, topic: recursions with lag greater than one.

Circuit Bank

- **WHEN** you must count representations of n as a sum of distinct weights **DO** split on whether the largest usable weight is used **BECAUSE** removing the largest weight reduces to the same problem at a smaller target, and everything above it is excluded by size
- **WHEN** the total of all weights below the top equals a landmark such as $F_{m+1} - 1$ or $10^m - 1$ **DO** complement the subset instead of treating the omitted case separately **BECAUSE** complementation is an involution sending sum n to sum (total minus n), so the omitted case is the same counting function at a reflected argument
- **WHEN** a recursion relates a value to two earlier values whose arguments sum to a constant **DO** evaluate at the extreme splits first **BECAUSE** extreme splits force one argument to zero, which gives the landmark values and usually the extremal behaviour
- **WHEN** a recursion has lag two **DO** split by parity and solve two arithmetic progressions **BECAUSE** even and odd indices never interact, and the floor function in the answer is exactly this parity split
- **WHEN** you are asked for the largest n with $f(n) = c$ and f is not monotone **DO** prove an increasing lower envelope for f on each block, then locate the last block where c is reachable **BECAUSE** without a lower envelope you can never exclude a later return to the value c, which is where most attempts fail

Backtesting

1. First useful cue: the sets are index sets, so $a_1 = 2$; and $\sum_{k \le m-1} F_k = F_{m+1} - 1$ makes complementation available.
2. Method the expression suggested: split on the top index, complement the other branch, get $a_n = a_{n-F_m} + a_{F_{m+1}-1-n}$.
3. Where you would hesitate, and the thought to reinforce: after complementing, do you count subsets of $\{1, \dots, m-1\}$ or all subsets. The thought is a size check: $F_{m+1} - 1 - n \le F_{m-1} - 1 < F_m$, so no index above $m-1$ can occur anyway, and the count is the unrestricted a.
4. The trap the question sets: producing the value $F_{4040} - 1$ and stopping. Half the marks are in the proof that $a_n \ge 2021$ for every larger n, and that needs the block minimum.
5. What makes you recognise it faster next time: any counting question about Fibonacci or base b representations is a leading weight recursion, and any answer of the form $F_M - 1$ says the extremal cases sit at block boundaries.
6. The reusable decision rule: for representation counts, recurse on the top weight, complement the omitted branch, evaluate the recursion at extreme splits, then induct for the envelope.

Mirror drills

1. Compute a_{12} with the recursion. *Answer:* $F_6 = 8 \le 12 < 13 = F_7$, so $a_{12} = a_{12-8} + a_{13-1-12} = a_4 + a_0 = 3 + 1 = 4$.
2. Prove $\sum_{k=1}^{m-1} F_k = F_{m+1} - 1$. *Answer:* write $F_k = F_{k+2} - F_{k+1}$ and telescope to $F_{m+1} - F_2$.
3. Let b_n count the sets of indices $S \subseteq \{2, 3, 4, \dots\}$ with $\sum_{k \in S} F_k = n$. Show $a_n = b_n + b_{n-1}$ and check it at $n = 4$. *Answer:* split on whether $1 \in S$; since $F_1 = 1$, the sets containing 1 correspond to representations of $n-1$ using indices at least 2. At $n = 4$: $b_4 = 1$, $b_3 = 2$, and $a_4 = 3$.

6.6 A6. A Fejer mean of the square wave in disguise

For a positive integer N, let f_N be the function defined by

$$f_N(x) = \sum_{n=0}^{N} \frac{N + 1/2 - n}{(N+1)(2n+1)} \sin\big((2n+1)x\big).$$

Determine the smallest constant M such that $f_N(x) \le M$ for all N and all real x.

Read it out loud

This is a trigonometric polynomial containing only odd harmonics and only sines, so it is odd and 2π periodic, and it vanishes at every multiple of π. The coefficient is the interesting object. Say it out loud: a numerator linear in n, over a product of $(N+1)$ and $(2n+1)$. A linear numerator over a product of two linear factors is a partial fractions cue, and here the split is forced:

$$\frac{N + \frac{1}{2} - n}{(N+1)(2n+1)} = \frac{\frac{1}{2}\big[2(N+1) - (2n+1)\big]}{(N+1)(2n+1)} = \frac{1}{2n+1} - \frac{1}{2(N+1)}.$$

So f_N is the partial sum $\sum_{n \le N} \frac{\sin((2n+1)x)}{2n+1}$ minus $\frac{1}{2(N+1)} \sum_{n \le N} \sin((2n+1)x)$. The first of those is the classical Fourier partial sum of the square wave, whose sum on $(0, \pi)$ is $\pi/4$. The second is a plain trigonometric sum with a closed form. The answer is going to be $\pi/4$, and the whole problem is the word "smallest": you must show the family never exceeds $\pi/4$ and gets arbitrarily close to it.

The fields wired in are Fourier analysis (the odd sine series of the square wave, Cesaro summation, the Gibbs phenomenon), trigonometric telescoping, and one alternating series evaluation. What is really being tested is whether you know that the plain partial sums $\sum_{n \le N} \frac{\sin((2n+1)x)}{2n+1}$ do *not* obey the bound $\pi/4$: their maxima decrease to $\frac{1}{2}\,\mathrm{Si}(\pi) \approx 0.926$ and never come near $\pi/4 \approx 0.785$. That is the Gibbs phenomenon. The subtracted term is exactly the correction that removes the overshoot, and the problem is asking you to prove that it does.

Within thirty seconds you should have the partial fractions split on the page, plus the two evaluations $f_1(\pi/2) = 1 - \frac{1}{3} = \frac{2}{3}$ and $f_2(\pi/2) = 1 - \frac{1}{3} + \frac{1}{5} - \frac{1}{6} = 0.7$, which sit below $\pi/4$ and climb towards it.

Approaches, ranked worst to best

1. **Numerical maximisation.** *Cue:* a concrete family of functions. *Circuit:* plot f_N and read off the peak. (Reveals that the maximum always sits at $x = \pi/2$ and climbs to 0.7854, which is priceless reconnaissance and no proof at all.)

2. **Bound each term.** *Cue:* a finite sum of sines. *Circuit:* triangle inequality. (Gives $\sum 1/(2n+1)$, a logarithmically divergent bound. Useless, but worth ten seconds to notice that the cancellation between terms is the entire content, so no term by term estimate can ever work.)

3. **Monotonicity in N plus Fourier convergence.** *Cue:* the correction term is $O(1/N)$. *Circuit:* show $f_{N+1} \ge f_N$ on $(0, \pi)$, then pass to the limit. (This is the official route and it is correct. It costs a two case argument, because the sign of the newly added harmonic $\sin((2N+3)x)$ is not controlled, and it finishes by quoting pointwise convergence of the square wave Fourier series. Two named theorems and a case split.)

4. **Recognise the Fejer mean.** *Cue:* coefficients tapered linearly in n. *Circuit:* f_N is the Cesaro mean of order $2N+1$ of the square wave series, and the Fejer kernel is nonnegative. (Conceptually the best: a Cesaro mean of a bounded function is an average of its values, so it is trapped between $\inf g$ and $\sup g$, and Fejer's theorem gives convergence to g at points of continuity. Three lines if you are allowed the machinery. It is ranked below the next route

only because it quotes a theorem most candidates cannot state precisely under pressure, and because the write up must still identify f_N with the Cesaro mean, which is a computation.)

5. **Differentiate and watch the collapse.** *Cue:* a finite trigonometric sum whose derivative is a kernel. *Circuit:* $f_N'(x) = \dfrac{\cos x \sin^2((N+1)x)}{2(N+1)\sin^2 x}$, so the maximum is at $x = \pi/2$ for every N. (Best. Self contained, no Fourier theorem quoted, and it localises the problem to a single point, after which the value is an alternating series whose distance from $\pi/4$ is controlled exactly. The collapse of the derivative is not luck: the correction term was built to cancel the oscillating kernel, which is the same statement as the Fejer identification.)

The ideal solution

The answer is $M = \boxed{\dfrac{\pi}{4}}$.

Step 1: split the coefficient. Since $N + \frac{1}{2} - n = \frac{1}{2}[2(N+1) - (2n+1)]$,

$$\frac{N+\frac{1}{2}-n}{(N+1)(2n+1)} = \frac{1}{2n+1} - \frac{1}{2(N+1)},$$

so, writing $S_N(x) = \sum_{n=0}^{N} \frac{\sin((2n+1)x)}{2n+1}$ and $T_N(x) = \sum_{n=0}^{N} \sin((2n+1)x)$,

$$f_N(x) = S_N(x) - \frac{T_N(x)}{2(N+1)}.$$

Step 2: close the second sum. Using $2\sin x \sin((2n+1)x) = \cos(2nx) - \cos((2n+2)x)$ and telescoping,

$$2\sin x \; T_N(x) = 1 - \cos((2N+2)x) = 2\sin^2((N+1)x),$$

so for $\sin x \neq 0$,

$$T_N(x) = \frac{\sin^2((N+1)x)}{\sin x} = \frac{1-\cos((2N+2)x)}{2\sin x}.$$

The same trick with $2\sin x \cos((2n+1)x) = \sin((2n+2)x) - \sin(2nx)$ gives

$$S_N'(x) = \sum_{n=0}^{N} \cos((2n+1)x) = \frac{\sin((2N+2)x)}{2\sin x} \qquad (\sin x \neq 0).$$

Step 3: reduce to $(0,\pi)$. Each f_N is 2π periodic and odd, and $f_N(0) = f_N(\pi) = 0$ because $\sin((2n+1)k\pi) = 0$. Oddness together with periodicity gives $f_N(2\pi - x) = f_N(-x) = -f_N(x)$, so the values on $(\pi, 2\pi)$ are the negatives of those on $(0,\pi)$. Hence

$$\sup_{x\in\mathbb{R}} f_N(x) = \max\Big(0, \max_{x\in(0,\pi)} f_N(x)\Big).$$

Step 4: differentiate. Fix $x \in (0,\pi)$, so $\sin x > 0$. By the coefficient split and the closed form for T_N,

$$f_N(x) = S_N(x) - \frac{1-\cos((2N+2)x)}{4(N+1)\sin x},$$

and differentiating the quotient,

$$\frac{d}{dx}\left[\frac{1-\cos((2N+2)x)}{\sin x}\right] = \frac{(2N+2)\sin((2N+2)x)\sin x - \big(1-\cos((2N+2)x)\big)\cos x}{\sin^2 x}.$$

Together with the closed form for S_N', the first pieces cancel exactly:

$$f_N'(x) = \frac{\sin((2N+2)x)}{2\sin x} - \frac{\sin((2N+2)x)}{2\sin x} + \frac{\big(1-\cos((2N+2)x)\big)\cos x}{4(N+1)\sin^2 x},$$

that is,

$$f_N'(x) = \frac{\cos x \, \sin^2((N+1)x)}{2(N+1)\sin^2 x} \qquad (0 < x < \pi).$$

Step 5: locate the maximum. By the derivative formula the sign of f_N' on $(0,\pi)$ is the sign of $\cos x$: nonnegative on $(0,\pi/2]$ and nonpositive on $[\pi/2,\pi)$. Since f_N is continuous on $[0,\pi]$, it increases on $[0,\pi/2]$ and decreases on $[\pi/2,\pi]$, so

$$\max_{x\in[0,\pi]} f_N(x) = f_N(\pi/2) \geq f_N(0) = 0.$$

Combining with Step 3, $\sup_{x\in\mathbb{R}} f_N(x) = f_N(\pi/2)$ for every N.

Step 6: evaluate at $\pi/2$. Since $\sin((2n+1)\pi/2) = (-1)^n$ and $\sin^2((N+1)\pi/2)$ equals 1 for N even and 0 for N odd,

$$f_N(\pi/2) = A_N - \frac{\varepsilon_N}{2(N+1)}, \qquad A_N = \sum_{n=0}^{N} \frac{(-1)^n}{2n+1}, \qquad \varepsilon_N = \begin{cases} 1 & N \text{ even}, \\ 0 & N \text{ odd}. \end{cases}$$

Step 7: bracket the alternating series. For $t \in [0,1]$, summing the finite geometric series with ratio $-t^2$,

$$\sum_{n=0}^{N} (-1)^n t^{2n} = \frac{1-(-1)^{N+1}t^{2N+2}}{1+t^2},$$

and integrating over $[0,1]$ gives, since $\int_0^1 \frac{dt}{1+t^2} = \frac{\pi}{4}$,

$$A_N = \frac{\pi}{4} + (-1)^N E_N, \qquad E_N = \int_0^1 \frac{t^{2N+2}}{1+t^2}\,dt \in \left(0,\ \frac{1}{2N+3}\right),$$

the bound on E_N coming from $0 < \frac{1}{1+t^2} \leq 1$ on $[0,1]$.

If N is odd then $f_N(\pi/2) = A_N = \frac{\pi}{4} - E_N < \frac{\pi}{4}$.

If N is even then

$$f_N(\pi/2) = \frac{\pi}{4} + E_N - \frac{1}{2N+2} < \frac{\pi}{4} + \frac{1}{2N+3} - \frac{1}{2N+2} < \frac{\pi}{4}.$$

Hence $f_N(x) \leq f_N(\pi/2) < \frac{\pi}{4}$ for every N and every real x, so $M = \pi/4$ is admissible.

Step 8: minimality. Along odd N, $f_N(\pi/2) = \frac{\pi}{4} - E_N$ and $0 < E_N < \frac{1}{2N+3} \to 0$, so $f_N(\pi/2) \to \pi/4$. Therefore no constant smaller than $\pi/4$ bounds the whole family, and $\pi/4$ is the smallest such constant. It is a supremum and not a maximum: the bound is strict for every individual N and x.

Layer 1: WHAT

Block 1: partial fraction the coefficient, exposing $f_N = S_N - \frac{1}{2(N+1)}T_N$, a Fourier partial sum minus a tapering correction. Block 2: close T_N in one line by multiplying by $2\sin x$ and telescoping, and get S_N' the same way. Block 3: use oddness and periodicity to reduce the supremum to the interval $(0,\pi)$. Block 4: differentiate; the oscillatory kernel cancels and $f_N'(x) = \frac{\cos x \sin^2((N+1)x)}{2(N+1)\sin^2 x}$. Block 5: conclude the maximum sits at $x = \pi/2$ for every N. Block 6: evaluate there, getting the Leibniz partial sum minus a parity dependent correction. Block 7: bracket the Leibniz partial sum with an integral remainder and check both parities. Block 8: let N run through odd values for minimality.

Juxtaposition. The nominal route recognises the Fourier series of the square wave, asserts that the sum is $\pi/4$ on $(0,\pi)$, and concludes that $\pi/4$ is the bound because "the partial sums converge to $\pi/4$". That is false reasoning and it fails on exactly this family's more famous cousin: the plain partial sums S_N converge to $\pi/4$ pointwise and yet their maxima tend to $\frac{1}{2}\,\mathrm{Si}(\pi) \approx 0.926$. Pointwise convergence never bounds a family uniformly. The bifurcation point is the moment you realise that the subtracted term is not an error term to be discarded but the entire point of the problem. The zombie throws away $\frac{T_N}{2(N+1)}$ because it is $O(1/N)$; you keep it, because it is the term that kills the Gibbs overshoot, and the fact that it is small pointwise is irrelevant near the peak, where T_N is of size N.

Layer 2: HOW

Block 1. **Cue.** A coefficient that is linear in n over a product of linear factors. **Why natural.** Partial fractions in the summation index is the standard reflex, and here the numerator is visibly a combination of the two denominators: $2(N+1)-(2n+1)=2N+1-2n$. Splitting is not a manipulation for its own sake; it separates a universal object (the square wave partial sum, independent of the taper) from an N dependent correction. Once separated, each piece can be recognised. **Link to the next block.** The correction is a bare trigonometric sum, and bare trigonometric sums with angles in arithmetic progression always close.

Block 2. **Cue.** $\sum \sin((2n+1)x)$, angles in arithmetic progression with common difference $2x$. **Why natural.** Multiply by $2\sin$ of half the common difference. The product to sum identity then turns each term into a difference of consecutive cosines and the sum telescopes. This is the real reason such sums have closed forms, and it is worth more than memorising the formula: the same move gives $\sum \cos((2n+1)x)$, which you will need in the next block for S_N'. **Link to the next block.** With T_N in closed form, f_N is an explicit elementary function plus S_N, and you can differentiate.

Block 3. **Cue.** Only odd harmonics and only sines, and the question says "all real x". **Why natural.** Odd plus 2π periodic means the graph on $(\pi, 2\pi)$ is the reflection of the graph on $(0,\pi)$ through the origin, so the supremum over the line equals the supremum over $(0,\pi)$, provided you also check the two boundary points where $\sin x = 0$ and every term vanishes. Doing this reduction first is what allows you to divide by $\sin x$ freely afterwards. **Link to the next block.** On $(0,\pi)$ the function is smooth and you may look for critical points.

Block 4. **Cue.** f_N is now S_N minus an explicit quotient, and S_N' is a known kernel. **Why natural.** Differentiating a sum of sines with coefficients $1/(2n+1)$ removes the coefficients entirely and produces the kernel $\frac{\sin((2N+2)x)}{2\sin x}$. The derivative of the correction contains the same kernel with coefficient $\frac{2N+2}{4(N+1)}=\frac{1}{2}$, so the two cancel identically. That cancellation is the fingerprint of the design: the coefficient $\frac{1}{2(N+1)}$ in the problem was chosen to make it happen. What survives is nonnegative apart from the factor $\cos x$. **Link to the next block.** A derivative whose sign is the sign of $\cos x$ pins the maximum at $\pi/2$ for every N simultaneously.

Block 5. **Cue.** f_N' has the sign of $\cos x$ on $(0,\pi)$. **Why natural.** Sign of the derivative is all you need; the zeros of $\sin^2((N+1)x)$ are isolated points where the derivative vanishes without changing sign, so monotonicity on each half is unaffected. The uniformity in N is the gift: the supremum over x has become the value at one fixed point, and the problem has collapsed from two variables to one. **Link to the next block.** At $x=\pi/2$ the sines take the values $(-1)^n$, so the value is an alternating series.

Block 6. **Cue.** $\sin((2n+1)\pi/2)=(-1)^n$. **Why natural.** This is why $\pi/2$ is where the answer lives: the odd harmonics all attain modulus one there and the series becomes the Madhava, Gregory and Leibniz series for $\pi/4$. The correction contributes only for even N, because $\sin^2((N+1)\pi/2)$ vanishes when N is odd, and that parity split is what makes the two cases of the final estimate. **Link to the next block.** An alternating series needs a sharp two sided bound, not just convergence.

Block 7. **Cue.** A partial sum of an alternating series compared with its limit. **Why natural.** The integral form is better than the alternating series test here because it gives the error with a sign and a bound in one line: expand $\frac{1}{1+t^2}$ as a finite geometric series with an exact remainder, integrate, and read off $A_N=\frac{\pi}{4}+(-1)^N E_N$ with $0<E_N<\frac{1}{2N+3}$. For even N the overshoot E_N must be beaten by

the correction $\frac{1}{2N+2}$, and $\frac{1}{2N+3} < \frac{1}{2N+2}$ settles it with room to spare. **Link to the next block.** Both parities give a strict inequality, and the odd case shows the bound is approached.

Layer 3: WHY

The structural truth is this. Let g be the square wave, equal to $\pi/4$ on $(0,\pi)$, to $-\pi/4$ on $(-\pi,0)$, and to 0 at 0 and π, extended with period 2π. Its Fourier series is $\sum_{n\geq 0} \frac{\sin((2n+1)x)}{2n+1}$, whose partial sums of order m are $T_m^{\mathrm{F}} = \sum_{n \text{ odd}, n\leq m} \frac{\sin(nx)}{n}$, so $T_{2k+1}^{\mathrm{F}} = T_{2k+2}^{\mathrm{F}} = S_k$. Now average:

$$\frac{1}{2N+2}\sum_{m=0}^{2N+1} T_m^{\mathrm{F}} = \frac{1}{2N+2}\left(2\sum_{k=0}^{N-1} S_k + S_N\right),$$

and the coefficient of $\frac{\sin((2n+1)x)}{2n+1}$ on the right is $\frac{2(N-n)+1}{2N+2} = \frac{2N+1-2n}{2N+2}$, which is exactly the coefficient in f_N after multiplying the coefficient split out. In other words

$$f_N = \sigma_{2N+1}(g), \qquad \sigma_m(g) = \frac{1}{m+1}\sum_{j=0}^{m} T_j^{\mathrm{F}},$$

the Fejer (Cesaro) mean of order $2N+1$ of the square wave. Fejer means are convolutions with the Fejer kernel $K_m(t) = \frac{1}{m+1}\left(\frac{\sin((m+1)t/2)}{\sin(t/2)}\right)^2 \geq 0$, of total mass 1, so $\sigma_m(g)(x)$ is a weighted average of the values of g and therefore lies between $-\pi/4$ and $\pi/4$ everywhere; and Fejer's theorem gives $\sigma_m(g)(x) \to g(x)$ at every point of continuity, in particular $\sigma_m(g)(\pi/2) \to \pi/4$. That is the whole problem in two sentences, and it is why the answer is the supremum of g rather than anything larger.

Contrast this with the plain partial sums, which are convolutions with the Dirichlet kernel; that kernel changes sign, its mass is not bounded, and the partial sums overshoot: $\max_x S_N(x) \to \frac{1}{2}\operatorname{Si}(\pi) = \frac{1}{2}\int_0^\pi \frac{\sin t}{t}\,dt \approx 0.926$, about nine per cent of the jump above the limiting value. The examiners took the classical Gibbs example and applied the classical cure, then asked you to prove the cure works without naming it. That is the design of the problem, and recognising the linear taper $1 - \frac{n}{N+1}$ style coefficient as a Cesaro weight is the single cue that unlocks it.

Knowledge gaps, named without disguise. To meet this problem at full speed you must have installed: the Fourier series of the square wave and its sum $\pi/4$ on $(0,\pi)$; the Dirichlet kernel and the fact that pointwise convergence gives no uniform bound; the Fejer kernel, its nonnegativity, and the Cesaro mean bound $\inf g \leq \sigma_m(g) \leq \sup g$; the Gibbs phenomenon and the constant $\frac{1}{2}\operatorname{Si}(\pi)$; and the integral remainder form of the Madhava, Gregory and Leibniz series. None of these is in A level Further Mathematics. The good news is that the ideal solution above uses none of them: it needs only the product to sum identities, one quotient derivative, and a finite geometric series under an integral. Install the Fourier theory because it explains the problem; solve the problem with calculus because that is what you can execute under time pressure.

The discernment mechanism: when a family of trigonometric polynomials carries coefficients that taper linearly to zero at the top, suspect a Cesaro mean and expect a bound by the supremum of the underlying function. When you are asked for a uniform bound over a family, never argue from pointwise convergence; either bound each member directly or find the location of the maximum. Train it by computing the maximum of the plain partial sums numerically for a few N so that the Gibbs overshoot is a fact you have seen, not a story you have heard.

Revolving concepts

Telescoping trigonometric sums. For angles in arithmetic progression, multiply by $2\sin(\frac{d}{2})$ where d is the common difference and use $2\sin A\sin B = \cos(A-B) - \cos(A+B)$ or $2\sin A\cos B = \sin(A+B) - \sin(A-B)$. Results worth banking: $\sum_{n=0}^{N}\sin((2n+1)x) = \frac{\sin^2((N+1)x)}{\sin x}$, $\sum_{n=0}^{N}\cos((2n+1)x) =$

$\frac{\sin((2N+2)x)}{2\sin x}$, and the Dirichlet kernel $\frac{1}{2}+\sum_{n=1}^{N}\cos(nx)=\frac{\sin((N+\frac{1}{2})x)}{2\sin(x/2)}$. See Book 1, field: Complex numbers, topic: sums of sines and cosines in arithmetic progression.

Fourier series of the square wave. $\sum_{n\geq 0}\frac{\sin((2n+1)x)}{2n+1}=\frac{\pi}{4}$ for $x\in(0,\pi)$, $-\frac{\pi}{4}$ for $x\in(-\pi,0)$, and 0 at $x=0,\pi$. At $x=\pi/2$ it is the Madhava, Gregory and Leibniz series. See Book 1, field: Real analysis, topic: Fourier series of elementary waveforms.

Fejer means and kernel positivity. $\sigma_m(g)=\frac{1}{m+1}\sum_{j=0}^{m}T_j^{\mathrm{F}}(g)$ equals the convolution of g with $K_m(t)=\frac{1}{m+1}\left(\frac{\sin((m+1)t/2)}{\sin(t/2)}\right)^2\geq 0$, whose integral is 1. Hence $\inf g\leq\sigma_m(g)\leq\sup g$ everywhere, and $\sigma_m(g)\to g$ at points of continuity, uniformly on compact sets of continuity. See Book 1, field: Real analysis, topic: Cesaro summability and Fejer's theorem.

Gibbs phenomenon. The Fourier partial sums of a function with a jump overshoot near the jump by about 8.95 per cent of the jump, and the overshoot does not decay with N; for the square wave, $\max_x S_N(x)\to\frac{1}{2}\,\mathrm{Si}(\pi)\approx 0.926>\frac{\pi}{4}$. Cesaro averaging removes it. See Book 1, field: Real analysis, topic: Gibbs phenomenon.

Alternating series with an integral remainder. $\sum_{n=0}^{N}(-1)^n t^{2n}=\frac{1-(-1)^{N+1}t^{2N+2}}{1+t^2}$ integrated over $[0,1]$ gives $\sum_{n=0}^{N}\frac{(-1)^n}{2n+1}=\frac{\pi}{4}+(-1)^N\int_0^1\frac{t^{2N+2}}{1+t^2}dt$, with the remainder in $\left(\frac{1}{2(2N+3)},\frac{1}{2N+3}\right)$. The same device with $\frac{1}{1+t}$ gives the partial sums of $\log 2$. This is the general recursion generalisation of "truncate a geometric series exactly, then integrate": always keep the exact remainder rather than an asymptotic one, because you usually need its sign. See Book 1, field: Calculus and integrals, topic: geometric series with exact remainder.

Uniform bounds versus pointwise limits. $f_N\to f$ pointwise says nothing about $\sup_N\sup_x f_N$. To bound a family you must either locate the maxima or bound each member by a kernel argument. See Book 1, field: Real analysis, topic: modes of convergence.

Circuit Bank

- **WHEN** a coefficient is linear in the summation index over a product of linear factors **DO** partial fraction it in that index **BECAUSE** the split separates a universal series from an N dependent correction, and each half is then recognisable
- **WHEN** you meet a trigonometric sum with angles in arithmetic progression **DO** multiply by $2\sin$ of half the common difference and telescope **BECAUSE** the product to sum identities turn each term into a difference of consecutive terms, which is why such sums always close
- **WHEN** a family of trigonometric polynomials has coefficients tapering linearly to zero **DO** suspect a Cesaro or Fejer mean and expect the bound $\sup g$ **BECAUSE** Fejer kernels are nonnegative with unit mass, so a Cesaro mean is an average of values of the function and cannot overshoot
- **WHEN** you must bound a whole family uniformly **DO** locate the maximum of each member, never argue from a pointwise limit **BECAUSE** Gibbs is the standing counterexample: partial sums converge pointwise to a function of supremum $\pi/4$ while their own maxima tend to 0.926
- **WHEN** a derivative of a designed expression looks like it might cancel **DO** compute it and expect the oscillatory kernel to vanish **BECAUSE** problem setters build such families so that the derivative has constant sign, which is the only way the maximum can sit at the same point for every N
- **WHEN** the value at the critical point is an alternating series **DO** expand the underlying geometric series with an exact remainder and integrate **BECAUSE** this gives the error with a sign and a sharp bound, which is what a strict inequality needs

Backtesting

1. First useful cue: the coefficient splits as $\frac{1}{2n+1} - \frac{1}{2(N+1)}$, exposing a square wave partial sum minus a taper.
2. Method the expression suggested: close the taper by telescoping, differentiate, watch the kernel cancel, evaluate at $x = \pi/2$.
3. Where you would hesitate, and the thought to reinforce: after seeing that the correction is $O(1/N)$, the instinct is to discard it. The thought to reinforce is that near the peak T_N has size of order N, so the correction is of order one exactly where it matters; that is the whole problem.
4. The trap the question sets: concluding $M = \pi/4$ from pointwise convergence alone. The same argument applied to S_N gives the wrong answer, since $\max_x S_N(x) \to \frac{1}{2}\,\mathrm{Si}(\pi) \approx 0.926$, not $\pi/4$.
5. What makes you recognise it faster next time: linear tapers in n mean Cesaro means, and Cesaro means never overshoot; the answer to such a problem is the supremum of the limiting waveform.
6. The reusable decision rule: for a supremum over a family of trigonometric polynomials, differentiate first and hope for a constant sign factor; if the derivative collapses, the two variable problem becomes a one variable series estimate.

Mirror drills

1. Prove $\sum_{n=0}^{N} \sin((2n+1)x) = \frac{\sin^2((N+1)x)}{\sin x}$ for $\sin x \neq 0$. *Answer:* multiply by $2\sin x$, use $2\sin x \sin((2n+1)x) = \cos(2nx) - \cos((2n+2)x)$, telescope to $1 - \cos((2N+2)x) = 2\sin^2((N+1)x)$.
2. Verify the derivative formula by hand at $N = 1$. *Answer:* $f_1(x) = \frac{3}{4}\sin x + \frac{1}{12}\sin 3x$, so $f_1'(x) = \frac{3}{4}\cos x + \frac{1}{4}\cos 3x = \frac{3}{4}\cos x + \frac{1}{4}(4\cos^3 x - 3\cos x) = \cos^3 x$, which is $\frac{\cos x \sin^2(2x)}{4\sin^2 x}$.
3. Show $\left|\sum_{n=0}^{N} \frac{(-1)^n}{2n+1} - \frac{\pi}{4}\right| < \frac{1}{2N+3}$ and identify the sign of the error. *Answer:* the error is $(-1)^N \int_0^1 \frac{t^{2N+2}}{1+t^2}dt$, positive for even N and negative for odd N, and bounded by $\int_0^1 t^{2N+2}dt$.

6.7 B1. Binary digit signs annihilate low degree polynomials

For a positive integer n, define $d(n)$ to be the sum of the digits of n when written in binary (for example, $d(13) = 1+1+0+1 = 3$). Let

$$S = \sum_{k=1}^{2020} (-1)^{d(k)} k^3.$$

Determine S modulo 2020.

Read it out loud

This is a signed power sum. The weight is a cubic, k^3, and the sign is $(-1)^{d(k)}$, the parity of the number of ones in the binary expansion of k. Say that sentence slowly, because it already contains the entire problem: the weight is a *polynomial of degree three*, and the sign is a *function of the binary digits*. Two objects from two different worlds have been multiplied together, and the question is asking whether the sign world can eat the polynomial world.

The fields wired in are: binary expansions and the Thue Morse sequence; finite difference calculus (equivalently, generating functions and multiplicities of roots at $x = 1$); and elementary modular arithmetic to finish. Nothing else. There is no number theory of 2020 beyond the last three lines.

What is really being tested is one structural observation, and it is the observation that $(-1)^{d(k)}$ is *multiplicative across disjoint binary blocks*. If c is a multiple of 2^m and $0 \le r < 2^m$, then the binary expansions of c and r never collide, so $d(c+r) = d(c) + d(r)$, and therefore $(-1)^{d(c+r)} = (-1)^{d(c)}(-1)^{d(r)}$. That single line is the guiding light. Once you have it, the sign over a block of 2^m consecutive integers starting at a multiple of 2^m is a fixed constant times the pattern $(-1)^{d(0)}, (-1)^{d(1)}, \ldots, (-1)^{d(2^m-1)}$, and that pattern is precisely the sign pattern $\prod_i (-1)^{\varepsilon_i}$ attached to the m independent binary choices $\varepsilon_0, \ldots, \varepsilon_{m-1}$.

Within thirty seconds you should have written three things: $(-1)^{d(2^m j + r)} = (-1)^{d(j)}(-1)^{d(r)}$ for $0 \le r < 2^m$; the word "difference" next to it, because summing a function against all 2^m sign patterns of m independent binary choices is by definition an m-fold finite difference; and $2016 = 16 \cdot 126$, because 2016 is the largest multiple of 16 below 2020 and $16 = 2^4$ is the first power of two whose exponent exceeds the degree 3. That last line is the whole solution compressed into an arithmetic fact.

Approaches, ranked worst to best

1. **Term by term evaluation.** *Cue:* no cue at all, just panic. *Circuit:* none. Two thousand and twenty cubes, each with a binary digit count. This is zombie maths in its purest form. It cannot be done in the arena and it teaches nothing, so it ranks last even though it would in principle succeed.

2. **Consecutive pairing.** *Cue:* $d(2k+1) = d(2k) + 1$, so adjacent terms carry opposite signs. *Circuit:* pair up terms whose signs are forced to be opposite. Grouping $k = 2j$ with $k = 2j+1$ turns the sum into $\sum_j (-1)^{d(j)} \left((2j)^3 - (2j+1)^3\right)$, which is a Thue Morse sum against a quadratic. Repeat and you get a linear one, then a constant one. It works, but every pass costs a full expansion and a re-indexing, and you must track the boundary terms three times. It costs more ink than it should, and each expansion is a fresh opportunity for a skipped-calculation error.

3. **Blocks of eight with an explicit constant.** *Cue:* blocks of length 2^m have a clean sign pattern. *Circuit:* tile the range by aligned power of two blocks. Compute $\sum_{r=0}^{7} (-1)^{d(r)} r^3 = -48$ by hand, deduce that the block starting at $8j$ contributes $(-1)^{d(j)} \cdot (-48)$, then note $2016 = 8 \cdot 252$

and finish by showing $\sum_{j=0}^{251}(-1)^{d(j)} = 0$ from consecutive pairing. This is a correct and honest route. It ranks here because it needs two separate arguments and one numerical evaluation that you must get right.

4. **Generating function with $x\,d/dx$.** *Cue:* a product $\prod(1-x^{2^i})$ generates exactly these signs. *Circuit:* convert a signed digit sum into a polynomial identity. Use $\prod_{i=0}^{m-1}(1-x^{2^i}) = \sum_{k=0}^{2^m-1}(-1)^{d(k)}x^k$, apply the operator $x\frac{d}{dx}$ three times, set $x=1$, and note the left side still vanishes because $x=1$ was a root of multiplicity m. Correct, and it is the same fact as the winner, only wearing a coat. It ranks just below because you must argue that three applications of $x\,d/dx$ cannot destroy a root of multiplicity $m>3$, which is a sentence of extra work.

5. **Finite difference annihilation.** *Cue:* the sign is a product of independent digit signs, and a cubic has vanishing fourth difference. *Circuit:* when a sign is a product over independent binary choices, the sum is an iterated difference operator. Blocks of length 16 aligned at multiples of 16, each contributing exactly zero because four difference operators annihilate a cubic. One lemma, one tiling, one five term tail. This is the ideal route: it is the shortest, it generalises immediately to any degree and any base, and it never asks you to compute a single cube until the final line.

The ideal solution

For $h \neq 0$ write $(\Delta_h P)(x) = P(x+h) - P(x)$ and $(S_h P)(x) = P(x+h)$, so $\Delta_h = S_h - I$.

Lemma 6.7.1. *If P is a polynomial of degree $D \geq 1$ then $\Delta_h P$ has degree exactly $D-1$; if P is constant then $\Delta_h P = 0$. Consequently, if $\deg P < m$ then $\Delta_{h_1}\Delta_{h_2}\cdots\Delta_{h_m}P = 0$ for any nonzero $h_1,\ldots,h_m$.*

Proof. Write $P(x) = a_D x^D + \cdots$ with $a_D \neq 0$. Then $P(x+h) - P(x) = a_D D h x^{D-1} + (\text{lower order})$, and $a_D D h \neq 0$, so the degree drops by exactly one. Each application therefore drops the degree by one until the polynomial is constant, and one further application returns zero. Applying m operators to a polynomial of degree at most $m-1$ produces zero. □

Lemma 6.7.2. *Let $m \geq 1$, let c be a multiple of 2^m, and let P be any polynomial of degree less than m. Then*

$$\sum_{r=0}^{2^m-1}(-1)^{d(c+r)}P(c+r) = 0.$$

Proof. Since $2^m \mid c$ and $0 \leq r < 2^m$, the binary expansions of c and of r use disjoint digit positions, so $c+r$ has the digits of c above position $m-1$ and the digits of r below it. Hence $d(c+r) = d(c)+d(r)$ and $(-1)^{d(c+r)} = (-1)^{d(c)}(-1)^{d(r)}$.

Set $Q(x) = P(c+x)$, a polynomial of degree less than m. Every r with $0 \leq r < 2^m$ has a unique expansion $r = \sum_{i=0}^{m-1}\varepsilon_i 2^i$ with $\varepsilon_i \in \{0,1\}$, and then $(-1)^{d(r)} = \prod_{i=0}^{m-1}(-1)^{\varepsilon_i}$. Expanding the operator product,

$$\prod_{i=0}^{m-1}\Delta_{2^i} \;=\; \prod_{i=0}^{m-1}(S_{2^i} - I) \;=\; \sum_{T\subseteq\{0,\ldots,m-1\}}(-1)^{m-|T|}\,S_{r_T}, \qquad r_T = \sum_{i\in T}2^i,$$

because the shift operators commute and $S_h S_{h'} = S_{h+h'}$. Evaluating at $x=0$ and noting that $T \mapsto r_T$ is a bijection from subsets onto $\{0,1,\ldots,2^m-1\}$ with $|T| = d(r_T)$,

$$\Big(\prod_{i=0}^{m-1}\Delta_{2^i}Q\Big)(0) \;=\; (-1)^m\sum_{r=0}^{2^m-1}(-1)^{d(r)}Q(r).$$

The left side is zero by Lemma 1, since $\deg Q < m$. Multiplying by $(-1)^m(-1)^{d(c)}$ gives the claim. □

Now take $m = 4$, so $2^m = 16$, and $P(x) = x^3$, of degree $3 < 4$. Since $2016 = 16 \cdot 126$, the interval $\{0, 1, \ldots, 2015\}$ is the disjoint union of the 126 blocks $\{16j, 16j+1, \ldots, 16j+15\}$ for $j = 0, 1, \ldots, 125$, and every one of these blocks starts at a multiple of 16. By Lemma 2 each block contributes zero, so

$$\sum_{k=0}^{2015} (-1)^{d(k)} k^3 = 0, \qquad \text{hence} \qquad \sum_{k=1}^{2015} (-1)^{d(k)} k^3 = 0$$

(the term $k = 0$ is zero anyway). Therefore

$$S = \sum_{k=2016}^{2020} (-1)^{d(k)} k^3.$$

It remains to evaluate five terms. Now $2016 = 11111100000_2$, so $d(2016) = 6$ and 2016 is a multiple of 32. For $0 \le r \le 4$ the same disjointness argument gives $d(2016 + r) = 6 + d(r)$, so the signs of $k = 2016, 2017, 2018, 2019, 2020$ are the signs of $d(0), d(1), d(2), d(3), d(4)$, namely

$$+, \ -, \ -, \ +, \ -.$$

Hence

$$S = 2016^3 - 2017^3 - 2018^3 + 2019^3 - 2020^3.$$

Reduce modulo 2020 by writing each base as a small negative number: $2016 \equiv -4$, $2017 \equiv -3$, $2018 \equiv -2$, $2019 \equiv -1$, $2020 \equiv 0 \pmod{2020}$. Then

$$S \equiv (-4)^3 - (-3)^3 - (-2)^3 + (-1)^3 - 0 = -64 + 27 + 8 - 1 = -30 \pmod{2020},$$

and $-30 \equiv 1990 \pmod{2020}$. Therefore

$$\boxed{S \equiv 1990 \pmod{2020}}.$$

Layer 1: WHAT

Block 1. Read the sign as a function of binary digits and record $d(2^m j + r) = d(j) + d(r)$ for $0 \le r < 2^m$.

Block 2. Recognise that summing over one aligned block of length 2^m means summing over all 2^m sign patterns of m independent binary choices, which is an m-fold finite difference.

Block 3. State the difference lemma: m difference operators annihilate any polynomial of degree less than m.

Block 4. Choose the smallest legal m: the degree is 3, so $m = 4$ and the block length is 16.

Block 5. Check the arithmetic that makes the tiling work: $16 \mid 2016$, with $2016 = 16 \cdot 126$. Conclude that everything below 2016 cancels.

Block 6. Evaluate the five leftover terms, reading their signs off $d(2016) = 6$ and $d(r)$ for $r \le 4$.

Block 7. Reduce modulo 2020 by replacing $2016, \ldots, 2020$ with $-4, \ldots, 0$, so that the cubes are tiny.

The nominal route runs: "the sum has 2020 terms, so I need the value of $(-1)^{d(k)}$ for each k; let me find a pattern in the Thue Morse sequence." That route spends its energy on the sequence itself and eventually reaches consecutive pairing, then a second pairing, then a third, and arrives at the answer exhausted.

The bifurcation point is the sentence you write immediately after reading the problem. The zombie writes "what is $(-1)^{d(k)}$?". The ideal solver writes "over which ranges does $(-1)^{d(k)}$ integrate a cubic to zero?". One question is about individual values; the other is about an annihilation property. The whole problem is decided there.

Layer 2: HOW

Block 1. **Cue.** The definition of d is given in terms of binary digits, and binary digits of a number below 2^m never interact with the digits of a multiple of 2^m. **Why natural.** You are not being asked to understand $(-1)^{d(k)}$ as a sequence; you are being asked to use it as a weight. A weight is only useful if it factorises, and the only factorisation available for a digit sum is across disjoint digit positions. So the first move is forced: split k at a power of two boundary. **Link to the next block.** Once the sign factorises across a block, the block sum is a fixed constant times a universal pattern, so all blocks look alike and you only need to understand one.

Block 2. **Cue.** The universal pattern is $(-1)^{d(r)} = \prod_i (-1)^{\varepsilon_i}$, a product over m independent choices. **Why natural.** Summing f over all choices with the sign $\prod(-1)^{\varepsilon_i}$ is, by definition, the alternating sum $\sum_T (-1)^{|T|} f(r_T)$, and an alternating sum over subsets of shifts is exactly an iterated difference operator. This is not a leap. It is the standard dictionary between "alternating over independent binary choices" and "iterated difference". If you have ever written $\Delta^n f(0) = \sum_j (-1)^{n-j} \binom{n}{j} f(j)$, you already know it. **Link to the next block.** Difference operators have one famous property, and it is exactly the property you need.

Block 3. **Cue.** A difference operator lowers the degree of a polynomial by one, with no exceptions when $h \neq 0$. **Why natural.** The leading terms cancel and the next term $a_D D h x^{D-1}$ survives with a nonzero coefficient. That is a one line proof, and it is the definition-level reason: differencing is a discrete derivative, and m derivatives kill degree less than m. **Link to the next block.** You now have a free parameter, m, and the smaller you take it the finer the tiling, so take it as small as the degree permits.

Block 4. **Cue.** Degree three, so you need $m \geq 4$. **Why natural.** Be lazy. A larger m means longer blocks, which means the leftover tail after tiling is longer and you compute more cubes by hand. Minimal m minimises ink. **Link to the next block.** The choice $m = 4$ is only useful if 16 divides something close to 2020.

Block 5. **Cue.** $2020 = 16 \cdot 126 + 4$. **Why natural.** This is the arithmetic check that decides whether the plan survives. It does: 2016 is a multiple of 16, so the tiling is exact and the tail has only five terms. **Link to the next block.** Five terms is a finite computation, and you now need their signs.

Block 6. **Cue.** 2016 is not just a multiple of 16; it is a multiple of 32, and $2020 - 2016 = 4 < 32$. **Why natural.** This means the same factorisation applies again on the tail: $d(2016 + r) = d(2016) + d(r)$ for $r \leq 4$. You never have to count binary digits of 2017, 2018, 2019 or 2020 separately. You count $d(2016) = 6$ once, then read $d(0), d(1), d(2), d(3), d(4) = 0, 1, 1, 2, 1$. **Link to the next block.** The five signed cubes are still large numbers, but the question only wants a residue.

Block 7. **Cue.** The modulus is 2020 and the terms sit at 2016 to 2020. **Why natural.** The bases are within 4 of the modulus, so replacing k by $k - 2020$ turns each cube into a single digit cube. Never cube a four digit number when the question hands you a modulus that makes it a one digit number. **Link to the next block.** The arithmetic collapses to $-64 + 27 + 8 - 1 = -30$, and normalising into $[0, 2020)$ gives 1990.

Layer 3: WHY

The mathematical truth beneath this route is that $(-1)^{d(\cdot)}$ is a character. On the block $\{c, c+1, \ldots, c+2^m-1\}$ with $2^m \mid c$, the map $r \mapsto (-1)^{d(r)}$ is the product of the m coordinate characters of the

group $(\mathbb{Z}/2)^m$ acting on the binary digits $\varepsilon_0, \ldots, \varepsilon_{m-1}$. Pairing a function against the full product of coordinate characters is the classical way of writing an m-fold difference, and an m-fold difference is a discrete m-th derivative. Polynomials of degree less than m have vanishing m-th derivative. That is the entire content of the problem, and everything else is arithmetic.

The same truth in generating function language is that

$$\prod_{i=0}^{m-1} \left(1 - x^{2^i}\right) = \sum_{k=0}^{2^m-1} (-1)^{d(k)} x^k$$

has a zero of order exactly m at $x = 1$, one order from each factor. Moments of order less than m read off derivatives of order less than m at $x = 1$, so they vanish. The two languages are the same lemma; the difference operator version is shorter because it never asks you to think about what $x \frac{d}{dx}$ does to a root of high multiplicity.

The discernment mechanism is a two part question that you should ask of every signed sum. First: does the sign factorise over some decomposition of the index? Second: over the block on which it factorises, what does the sign annihilate? Here the answers are "yes, over binary blocks" and "every polynomial of degree less than the block exponent". If the sign had been $(-1)^k$ the block length would be 2 and the annihilated class would be constants, which is why $\sum(-1)^k$ telescopes. If the sign had been a b-adic digit character in base b you would get $\prod_{i<m}(1 - x^{b^i}\zeta)$ style products and the same annihilation of low degrees. The general shape is always: character times polynomial, block length controls degree killed.

To train the discernment, do this. Every time you meet a weight that is defined by digits, immediately ask what happens over a block whose length is the relevant power of the base. Every time you meet an alternating sum over subsets, immediately write it as an iterated difference. These are two of the cheapest circuits on the Putnam and they fire together here.

The examiners nudged you with three planted signals. The modulus 2020 and the upper limit 2020 are the same number, which is the nudge to reduce k modulo 2020 at the end and turn the cubes into single digits. The exponent 3 is small and fixed, which is the nudge that a bounded number of difference operators will suffice. And the upper limit is four above $2016 = 16 \cdot 126$, which is the nudge that a power of two tiling is meant to be exact except for a very short tail. If the limit had been, say, 2019, the tail would have been four terms; if it had been a prime, you would have suspected a different structure entirely.

Revolving concepts

Binary digit sum and the Thue Morse sequence. Define $t_k = (-1)^{d(k)}$. The defining recursion is $t_0 = 1$, $t_{2k} = t_k$, $t_{2k+1} = -t_k$. This is a 2-automatic sequence: its value is a finite state function of the binary expansion. The generalised recursion, which you should install once and reuse, is: for a base b and a digit function g, the sequence $u_{bk+r} = \chi(r)u_k$ is determined by the multiplicative rule across digit blocks, and every sum of u against a function over an aligned block of length b^m becomes an m-fold operator applied to that function. Pointer: see Book 1, field: Number theory, topic: automatic sequences and digit characters.

Finite difference calculus. $\Delta_h P(x) = P(x + h) - P(x)$; $\deg \Delta_h P = \deg P - 1$ for nonconstant P; $\Delta_h^m P = 0$ if and only if $\deg P < m$; Newton's forward difference formula $P(x) = \sum_j \binom{x}{j} (\Delta_1^j P)(0)$ for the step one case. The composite $\Delta_{h_1} \cdots \Delta_{h_m}$ applied to x^m equals $m!\, h_1 h_2 \cdots h_m$, which is the exact analogue of the m-th derivative of x^m being $m!$. Pointer: see Book 1, field: Polynomials, topic: finite differences.

The exact first nonvanishing moment. From the composite formula with $h_i = 2^i$,

$$\sum_{k=0}^{2^m-1} (-1)^{d(k)} k^m = (-1)^m\, m! \prod_{i=0}^{m-1} 2^i = (-1)^m\, m!\, 2^{m(m-1)/2}.$$

For $m = 3$ this reads $-6 \cdot 8 = -48$, which is exactly the constant that route three in the ranking has to compute by hand. Knowing the general formula means you never compute such a constant again. Pointer: see Book 1, field: Number theory, topic: digit sums against polynomial weights.

Generating function for a digit character. $\prod_{i=0}^{m-1}(1 - x^{2^i}) = \sum_{k<2^m}(-1)^{d(k)}x^k$, with a zero of order m at $x = 1$. More generally $\prod_{i\geq 0}(1 - x^{2^i})$ is the Thue Morse generating series. The multiplicity of the root at $x = 1$ is exactly the number of factors, and multiplicity at $x = 1$ is the generating function encoding of moment vanishing. Pointer: see Book 1, field: Generating functions, topic: root multiplicity and moment vanishing.

Reduction of a residue tail. If a sum modulo M has its surviving terms in $[M - c, M]$ for small c, replace k by $k - M$ before doing any arithmetic. Pointer: see Book 1, field: Number theory, topic: modular normalisation.

Circuit Bank

- **WHEN** a sign of the form $(-1)^{\text{digit sum in base } b}$ multiplying a polynomial **DO** split the range into blocks of length b^m starting at multiples of b^m, with m greater than the degree **BECAUSE** the sign factorises across disjoint digit positions, so each block sum is an m-fold difference and m differences annihilate degree less than m
- **WHEN** an alternating sum over all subsets of a fixed finite set of shifts **DO** write it as the composite difference operator $\prod_i(S_{h_i} - I)$ evaluated at a point **BECAUSE** shift operators commute and expanding the product reproduces the subset sum with the same signs, and each factor lowers polynomial degree by exactly one
- **WHEN** a summation limit that is a small amount above a multiple of a power of two **DO** tile up to that multiple and treat the short tail by hand **BECAUSE** the structural cancellation is exact only on aligned blocks, so the tail is where all the actual arithmetic lives, and it should be as short as possible
- **WHEN** a sum to be evaluated modulo M whose surviving terms sit just below M **DO** replace each k by $k - M$ before raising to any power **BECAUSE** cubing -4 costs one line and cubing 2016 costs four, and the extra digits carry no information modulo M

Backtesting

1. First useful cue: the sign is a binary digit parity, which means it factorises across power of two blocks, so the very first thing written down is $d(2^m j + r) = d(j) + d(r)$ for $0 \leq r < 2^m$.
2. Method the expression suggested: a polynomial weight against a digit character means iterated finite differences; choose the block exponent to exceed the degree and tile the range.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at the step "an alternating sum over binary patterns is a difference operator". Reinforce it by writing $\Delta_{h_1}\Delta_{h_2} = S_{h_1+h_2} - S_{h_1} - S_{h_2} + I$ once by hand and seeing the four subsets with their signs. After that the general case is bookkeeping, not insight.
4. The trap the question sets: the number 2020 is dressed to look like a number theory problem about $2020 = 2^2 \cdot 5 \cdot 101$. It is not. The only role of 2020 is that 2016 is a multiple of 16 and that the tail sits within four of the modulus. Chasing the factorisation of 2020 wastes the whole session.
5. What makes you recognise it faster next time: the pairing "digit character times polynomial". Whenever those two appear multiplied, the answer is always "block of length base to the power (degree plus one), everything cancels, compute the tail".

6. The reusable decision rule: a weight that is a product over independent digit choices turns any sum over an aligned block into an iterated difference, and iterated differences annihilate polynomials of low degree. Choose the block length minimal, tile, and do arithmetic only on the remainder.

Mirror drills

1. Evaluate $\sum_{k=0}^{63}(-1)^{d(k)}k^5$. *Answer:* zero. The range is one aligned block of length 2^6, and six difference operators annihilate a polynomial of degree 5.
2. Evaluate $\sum_{k=0}^{2^m-1}(-1)^{d(k)}k^m$ in closed form. *Answer:* $(-1)^m m!\,2^{m(m-1)/2}$. The composite $\Delta_1\Delta_2\cdots\Delta_{2^{m-1}}$ applied to x^m equals $m!$ times the product of the steps, and the sign $(-1)^m$ comes from the operator expansion. Check $m=3$: $-6\cdot 8=-48$.
3. Determine $\sum_{k=1}^{2020}(-1)^{d(k)}k$ modulo 2020. *Answer:* zero. The degree is 1, so take blocks of length $2^2=4$; since $2020=4\cdot 505$, the range $\{0,1,\ldots,2019\}$ tiles exactly and contributes nothing, leaving only $k=2020$, whose contribution is $\pm 2020\equiv 0$.

6.8 B2. Adjacent pairing decides the peg sliding game

Let k and n be integers with $1 \leq k < n$. Alice and Bob play a game with k pegs in a line of n holes. At the beginning of the game, the pegs occupy the k leftmost holes. A legal move consists of moving a single peg to any vacant hole that is further to the right. The players alternate moves, with Alice playing first. The game ends when the pegs are in the k rightmost holes, so whoever is next to play cannot move and therefore loses. For what values of n and k does Alice have a winning strategy?

Read it out loud

This is an impartial game under the normal play convention: both players have the same moves available, and the player who cannot move loses. Positions are k-element subsets of $\{1, \ldots, n\}$, a move slides one peg rightwards to any vacant hole, and the unique position with no legal move is the block of k rightmost holes. The fields wired in are combinatorial game theory (normal play, winning and losing positions, pairing strategies) and nothing else; there is no arithmetic here beyond parity.

What is really being tested is whether you reach for a strategy instead of a computation. The answer to "for what n and k" is going to be a congruence condition, and the moment you accept that, you are being told that some parity object is conserved. The naive candidate is the number of moves in the game. Kill that candidate immediately: a peg may slide one hole or ten, so different lines of play from the same start can take different numbers of moves. Play $n = 4$, $k = 2$ once and you will see a two move game and a three move game from the identical start. So the parity is not in the move count; it must be in the *shape of the position*.

Once you know the parity lives in the position, the only cheap objects on a line of n holes that carry parity are pairings of the holes. And there is exactly one canonical pairing of a line: cut it into adjacent blocks $\{1,2\}, \{3,4\}, \ldots$. That pairing exists only when n is even, and the starting configuration respects it only when k is even. Two evenness conditions have appeared without any calculation, and the answer is already visible: Bob wins when n and k are both even, Alice wins otherwise.

Within thirty seconds you should have written: "normal play, so I want a pairing or a mirroring strategy"; "pair adjacent holes"; and "if n or k is odd, Alice's first move should destroy the obstruction and hand Bob an even by even board". The rest is verification.

Approaches, ranked worst to best

1. **Total move count parity.** *Cue:* normal play games are often decided by whether the total number of moves is fixed and odd. *Circuit:* count the moves, take the parity. This is the trap, and it is a good trap because the game does have a conserved quantity of exactly this flavour. It is false here: from $\{1,2\}$ on four holes, the line $1 \to 3$, $2 \to 4$ takes two moves while $2 \to 3$, $1 \to 2$, $2 \to 4$ takes three. Ranks last because it produces a confident wrong answer.

2. **Full Sprague Grundy tabulation.** *Cue:* impartial game under normal play. *Circuit:* compute the Grundy value of every position by the mex recursion. Correct in principle and it is what a machine would do, but the state space is all $\binom{n}{k}$ subsets, the game does not decompose into independent components, and you only need to know who wins the single starting position, not a Grundy value that could be added to other games. It costs enormous ink for information you do not need.

3. **Induction on n and k with a case analysis of Alice's first move.** *Cue:* the answer looks like a congruence in n and k. *Circuit:* strong induction on board size. This can be pushed through, but Alice has up to $k(n-k)$ first moves and you must handle all of them in the losing cases. You end up rediscovering the pairing anyway, in a disguised and longer form.

4. **Small case tabulation, then conjecture.** *Cue:* a two parameter family with a congruence answer. *Circuit:* compute the small table, read the pattern. Necessary reconnaissance and you should always do it. It is not a proof, so it cannot rank above a method that produces one, but it tells you which theorem to prove and it costs almost nothing.
5. **Adjacent pairing strategy with a one move trim.** *Cue:* normal play, a line of holes, an answer that is a parity condition. *Circuit:* when the loser is the player who runs out of moves, look for an involution of the board under which the second player can always mirror. Bob keeps every adjacent pair either full or empty; Alice's move always breaks one full pair and half fills a strictly later empty pair, and Bob completes it. When a parity obstruction exists, Alice spends one move retiring an end hole to remove it. Two short arguments, no case explosion, and both halves are checkable in the arena.

The ideal solution

Number the holes 1 to n from left to right. A position is a set $P \subseteq \{1, \dots, n\}$ with $|P| = k$; a move replaces P by $(P \setminus \{p\}) \cup \{v\}$ for some $p \in P$ and some vacant $v > p$. The initial position is $P_0 = \{1, \dots, k\}$ and the player who cannot move loses.

Lemma 6.8.1. *The game terminates, and the unique position with no legal move is* $\{n-k+1, \dots, n\}$.

Proof. The potential $\Phi(P) = \sum_{p \in P} p$ strictly increases with every move and is bounded above by $\sum_{i=n-k+1}^{n} i$, so only finitely many moves are possible. A position has no legal move exactly when no vacant hole lies to the right of any peg, that is, when every vacant hole precedes every peg, which for a set of size k means $P = \{n-k+1, \dots, n\}$. □

Theorem 6.8.2. *Alice has a winning strategy if and only if at least one of n and k is odd. If n and k are both even, Bob has a winning strategy.*

Step 1: both even, so the player to move loses.
Let n and k both be even, with $0 \le k \le n$. Split the holes into the $n/2$ adjacent pairs

$$Q_i = \{2i-1,\, 2i\}, \qquad i = 1, \dots, n/2.$$

Call a position P *balanced* if for every i either $Q_i \subseteq P$ or $Q_i \cap P = \emptyset$. Since k is even, $P_0 = \{1, \dots, k\}$ is balanced.

Claim 6.8.3. *If P is balanced and the player to move makes any legal move, the opponent has a legal reply after which the position is again balanced.*

Proof. Say the mover takes the peg at p to the vacant hole $v > p$. Let $Q_i \ni p$ and $Q_j \ni v$. Since P is balanced and $p \in P$, we have $Q_i \subseteq P$; since $v \notin P$, we have $Q_j \cap P = \emptyset$, so $Q_j \neq Q_i$ and hence $j \neq i$. If $j < i$ then $v \le 2j \le 2i-2 < 2i-1 \le p$, contradicting $v > p$. Therefore $j > i$, and consequently every element of Q_j is larger than every element of Q_i.

Let p' be the other hole of Q_i and v' the other hole of Q_j. After the mover's move, Q_i contains exactly the peg at p' and Q_j contains exactly the peg at v, while every other pair is untouched and hence still full or empty. The opponent moves the peg at p' to v': this is legal because $v' > p'$ (as Q_j lies entirely to the right of Q_i) and v' is vacant. Now Q_i is empty, Q_j is full, and every other pair is unchanged, so the position is balanced. □

Bob's strategy is: always restore balance. By the claim this is always available to him, so he never runs out of moves. By Lemma 1 the game terminates, and it can only terminate on a turn where the player to move has no move; that player is Alice. Hence when n and k are both even, the second player wins.

Step 2: two retirement facts.

(R1) If hole 1 is vacant, no peg can ever occupy it, because every move goes rightwards. The remaining play is therefore identical to the game on holes $2, \dots, n$ with the same k pegs, and the terminal position of the original game restricts to the terminal position of that game.

(R2) If hole n is occupied, that peg can never move, and no other peg can move into hole n. The remaining play is therefore identical to the game on holes $1, \dots, n-1$ with the other $k-1$ pegs, and again the terminal positions correspond.

Step 3: at least one of n, k odd, so Alice wins.

In each case Alice makes one move and hands Bob a game which, after retiring end holes, is an even by even game with Bob to move; Step 1 then says Bob loses.

Case A: n odd and k odd. Since $k < n$, hole n is vacant. Alice moves the peg at hole k to hole n, reaching $\{1, \dots, k-1\} \cup \{n\}$. By (R2) the remaining game is the game on the $n-1$ holes $1, \dots, n-1$ with the $k-1$ pegs in the leftmost $k-1$ holes. Both $n-1$ and $k-1$ are even, and Bob is to move, so Bob loses.

Case B: n odd and k even. Since $k < n$, hole $k+1$ exists and is vacant. Alice moves the peg at hole 1 to hole $k+1$, reaching $\{2, \dots, k+1\}$. By (R1) the remaining game is the game on the $n-1$ holes $2, \dots, n$ with the k pegs in the leftmost k of them. Both $n-1$ and k are even, and Bob is to move, so Bob loses.

Case C: n even and k odd. Alice moves the peg at hole 1 to hole n, reaching $\{2, \dots, k\} \cup \{n\}$. By (R1) and (R2) the remaining game is the game on the $n-2$ holes $2, \dots, n-1$ with the $k-1$ pegs in the leftmost $k-1$ of them. Both $n-2$ and $k-1$ are even, and Bob is to move, so Bob loses.

In Cases A and C with $k = 1$ the reduced game has no pegs at all, so Bob is immediately unable to move, which is the same conclusion. This proves the theorem, and the answer is:

$$\boxed{\text{Alice wins exactly when } n \text{ is odd or } k \text{ is odd}}.$$

Layer 1: WHAT

Block 1. Formalise: positions are k-subsets, a move slides one peg right into a vacancy, the terminal position is the right block, and the game is finite because $\sum_{p \in P} p$ strictly increases.

Block 2. Test the move count parity hypothesis on a small board and watch it fail. Conclude the parity lives in the position, not the length of the game.

Block 3. Look for a pairing of the board that the initial position respects. On a line there is only one natural one, adjacent blocks of two, and it exists exactly when n is even and is respected by $\{1, \dots, k\}$ exactly when k is even.

Block 4. Prove the mirroring claim: Alice must break a full pair and half fill a strictly later empty pair, so Bob can complete that pair.

Block 5. Record the two retirement facts: a vacant hole 1 is dead forever, and an occupied hole n is frozen forever.

Block 6. In each of the three odd cases, find a single Alice move that trims the board to even by even.

The nominal route runs: "impartial game, so compute Grundy values". That route builds a table of $\binom{n}{k}$ positions, finds no product structure to exploit, and never terminates by hand.

The bifurcation point is Block 2. The zombie asks "how many moves does this game take?" and, on seeing a plausible answer for the first small case, commits. The ideal solver asks "is the number of moves even determined by the start?", finds one three move line and one two move line on the four hole board, and immediately abandons counting for structure. Everything downstream follows from refusing to count.

Layer 2: HOW

Block 1. Cue. "The game ends when the pegs are in the k rightmost holes, so whoever is next to play cannot move and therefore loses." That is the normal play convention stated in words.

Why natural. Before anything else, confirm the game is finite and identify the unique terminal position, because every argument about who runs out of moves is vacuous if play could go on forever. The potential $\sum p$ is the obvious monotone quantity because every move increases exactly one coordinate. **Link to the next block.** A finite normal play game is decided by structure, and the two structural tools are Grundy values and pairing strategies.

Block 2. Cue. The question asks for which n and k, not for a Grundy value, so the answer is a congruence. **Why natural.** Congruence answers come from conserved parities. There are only two candidates: the parity of the number of moves, and the parity of some feature of the position. Test the cheap one first, on the smallest board where a choice exists. It fails, which is information, not a setback. **Link to the next block.** Having eliminated the move count, you are forced to look at the position itself.

Block 3. Cue. A line of n holes, and a losing condition of the form "you run out of moves". **Why natural.** In a normal play game the second player wins whenever the board admits an involution such that every move of the first player can be answered by the image move. On a line, the involution that respects the "rightwards only" constraint must map each hole to a neighbour, because the reply has to be legal in the same direction. Pairing $\{1,2\},\{3,4\},\ldots$ is the unique such involution. This is not a guess, it is the only candidate the geometry allows. **Link to the next block.** The pairing exists only for even n, and the start respects it only for even k, which is exactly the answer the small table predicted.

Block 4. Cue. The mirroring reply must be legal, which means it must go rightwards. **Why natural.** This is the one place where the proof could fail, so it is the one place that deserves a careful sentence. Alice takes a peg out of a full pair and puts it into an empty pair. A full pair cannot be the same as an empty pair, and the target must be to the right, which forces the target pair to lie entirely to the right of the source pair. That is what makes Bob's completing move legal. **Link to the next block.** Bob's strategy is now well defined and always available, so the game cannot end on his turn.

Block 5. Cue. Moves only go right, so extreme holes can become permanently irrelevant. **Why natural.** A vacant leftmost hole can never be filled, and an occupied rightmost hole can never be emptied. Each fact removes one hole from the board without changing the structure of the remaining game. This is the only way to change the parity of n inside the game, so it is what Alice must use. **Link to the next block.** Alice now needs one move that removes the right number of holes and pegs.

Block 6. Cue. Three parity cases, each needing a different trim: (n,k) odd odd needs one hole and one peg removed; odd even needs one hole and no peg; even odd needs two holes and one peg. **Why natural.** Match the trim to the deficit. Moving the rightmost peg to hole n freezes a peg and retires the right end. Moving the leftmost peg away from hole 1 retires the left end. Moving the leftmost peg all the way to hole n does both at once. There are exactly three trims and exactly three cases, and they line up because retiring the left end changes n only, while retiring the right end changes both n and k. **Link to the next block.** Each case lands on an even by even board with Bob to move, which Step 1 has already settled.

Layer 3: WHY

The mathematical truth beneath this route is that a pairing strategy is a homomorphism from the board to a quotient in which every position is already "done". Formally, when n and k are both even, the balanced positions form a set closed under "Alice moves then Bob restores", and the

terminal position is balanced. The second player therefore never faces a position from which he cannot answer, so the length of the game, whatever it is, ends on the first player's turn. No Grundy value is ever computed, because Grundy values carry information about sums of games and this game is never summed with another.

There is also a genuine conserved quantity in this game, and knowing it is what lets you kill the move count hypothesis without guessing. For a position P define

$$I(P) = \#\{(v, p) : v \text{ vacant},\ p \in P,\ v > p\}.$$

A move that carries a peg from p to v decreases I by exactly $v - p$. To see this, note that the pairs not involving p or v are untouched; the pair (v, p) is destroyed; each vacancy strictly between p and v loses its partner p and gains no new one; each peg strictly between p and v loses the vacancy v; and every hole outside $[p, v]$ exchanges one partner for another. Counting, the loss is $1 + (\text{holes strictly between}) = v - p$. Since $I(P_0) = k(n-k)$ and I is zero exactly at the terminal position, the total distance travelled by all pegs over the whole game is always $k(n-k)$, no matter how the game is played. What is *not* conserved is how that total is split into moves, which is precisely why the move count parity argument fails and why a genuine strategy is needed.

The discernment mechanism is a three question filter for any impartial game. First, is the game finite, and what is the unique terminal position? Second, is the total number of moves forced, and if so is its parity the answer? Third, if not, does the board carry an involution under which the second player can mirror, and does the starting position respect it? Only when all three fail do you reach for Sprague Grundy. On this problem the filter answers "yes, unique right block", "no, and here is a counterexample", "yes, adjacent pairing", and you are done.

To train it, take any normal play game you meet and immediately look for the involution before anything else. Mirroring beats computation whenever the board has a symmetry compatible with the move rule. The compatibility requirement is the part people forget: on this board a reflection through the centre is a perfectly good involution of the holes but it is useless, because the reply to a rightward move would have to be a leftward move.

The examiners nudged you deliberately. They wrote "a line of n holes", which is a board with an obvious block structure. They allowed a peg to move to *any* vacant hole further right, not just the adjacent one, which is exactly the modification that destroys the move count invariant and forces you to think. And they asked "for what values of n and k", which promises a congruence and therefore promises a parity object.

Revolving concepts

Impartial games under normal play. Both players have the same legal moves from any position and the player unable to move loses. Every position is either an $\mathcal{N}$-position (the player to move wins) or a $\mathcal{P}$-position (the previous player wins). The defining recursion is: P is a $\mathcal{P}$-position if and only if every move from P leads to an $\mathcal{N}$-position, and P is an $\mathcal{N}$-position if and only if some move leads to a $\mathcal{P}$-position. Generalised, this is the Sprague Grundy function $G(P) = \operatorname{mex}\{G(P') : P \to P'\}$ where mex of a set of nonnegative integers is the least nonnegative integer not in it; $\mathcal{P}$-positions are exactly those with $G = 0$, and for a disjoint sum of games $G(P \oplus Q) = G(P) \oplus G(Q)$ with $\oplus$ the bitwise exclusive or. Here the balanced positions are exactly a family of $\mathcal{P}$-positions and you never need the numerical values. Pointer: see Book 1, field: Games and invariants, topic: Sprague Grundy theory.

Pairing and mirroring strategies. If the board admits an involution σ with no fixed points, the starting position is σ-symmetric, and for every legal move m the "image" move is legal immediately afterwards, then the second player wins by always playing the image. The engineering content is always the legality of the reply, not the existence of σ. Pointer: see Book 1, field: Games and invariants, topic: pairing strategies and strategy stealing.

Monotone potentials and termination. A strictly monotone bounded integer functional on positions proves a game is finite. Here $\Phi(P) = \sum_{p \in P} p$ works. Pointer: see Book 1, field: Games and invariants, topic: termination arguments.

The inversion count and its recursion. $I(P) = \#\{(v, p) : v \text{ vacant}, p \text{ peg}, v > p\}$ satisfies $I(P') = I(P) - (v - p)$ for a move from p to v, with $I(P_0) = k(n-k)$ and $I = 0$ at the end. Written as a recursion, $I_{t+1} = I_t - d_t$ where d_t is the distance of the t-th move, so $\sum_t d_t = k(n-k)$ for every play. The general lesson is that when a game has a conserved total but a free decomposition of that total into moves, the move count carries no invariant and a strategy is required. Pointer: see Book 1, field: Combinatorics, topic: inversions of a binary word.

Related classical games. Sliding coins rightwards along a strip is the family containing the Silver Dollar Game and Mock Turtles; when jumping over occupied cells is forbidden, the game becomes staircase Nim on the gaps. The present rule permits jumping, which is exactly why the gap structure is not preserved and the pairing argument replaces the Nim analysis. Pointer: see Book 1, field: Games and invariants, topic: coin sliding games.

Circuit Bank

- **WHEN** an impartial normal play game whose answer is asked as a condition on parameters **DO** look for a pairing or mirroring strategy for the second player before computing any Grundy value **BECAUSE** a congruence answer signals a conserved parity, and the cheapest carrier of parity on a structured board is an involution the second player can imitate
- **WHEN** a game where you suspect the number of moves is fixed **DO** find two lines of play from the same start with different lengths, or prove the total is forced **BECAUSE** a move count argument is worthless unless the count is determined by the position, and one explicit pair of unequal lines settles it in ten seconds
- **WHEN** moves that go only in one direction along a line **DO** ask which cells become permanently dead, namely a vacancy with nothing to its left that can fill it and an occupied cell with nothing to its right **BECAUSE** dead cells can be deleted from the board without changing the game, and deleting cells is the only way a player can change the parity of the board size
- **WHEN** a mirroring strategy whose reply must obey a direction constraint **DO** check that the target block lies entirely on the correct side of the source block **BECAUSE** the existence of the involution is never the difficulty; the legality of the mirrored reply is, and it is the only step that can fail

Backtesting

1. First useful cue: normal play, plus an answer requested as a condition on n and k, which together say "find the parity object" and, on a line, "pair adjacent holes".
2. Method the expression suggested: a second player pairing strategy for the symmetric case, plus a single trimming move by the first player to remove any parity obstruction.
3. Where you would hesitate, and the thought to reinforce: the hesitation is whether Bob's reply is always legal. Reinforce with the sentence "a full pair and an empty pair are distinct, and a rightward move forces the empty one to lie entirely to the right". That single observation is the load bearing step of the whole proof.
4. The trap the question sets: allowing a peg to move to *any* vacant hole to the right, which looks like generosity but is there to destroy the move count invariant. If the rule were "move to the adjacent vacant hole", the total number of moves would be forced to $k(n-k)$ and the answer would be a different, simpler parity condition.

5. What makes you recognise it faster next time: the shape "line of cells, one directional moves, last player to move wins" should fire the adjacent pairing circuit instantly, and the shape "answer is a condition on parameters" should fire "conserved parity".
6. The reusable decision rule: in a finite normal play game, before computing anything, ask whether the board has a fixed point free involution compatible with the direction of play under which the starting position is symmetric; if it does, the second player wins, and the first player's only hope is a move that destroys the symmetry defect.

Mirror drills

1. Same game, but a legal move is to slide a peg into the vacant hole immediately to its right (no jumping). Who wins? *Answer:* pegs can no longer pass one another, so the i-th peg must travel from hole i to hole $n-k+i$, and the game always lasts exactly $k(n-k)$ moves. Alice wins if and only if $k(n-k)$ is odd, that is, if and only if k is odd and n is even.
2. Prove that in the original game the total distance travelled by all pegs is the same for every line of play. *Answer:* the inversion count $I(P) = \#\{(v,p) : v \text{ vacant}, p \text{ peg}, v > p\}$ drops by exactly $v-p$ on a move from p to v, starts at $k(n-k)$ and ends at 0.
3. Why must the pairs be adjacent? Show that pairing hole i with hole $n+1-i$ cannot support Bob's strategy. *Answer:* the starting position $\{1,\dots,k\}$ is not even symmetric under that pairing, since hole 1 is occupied while hole n is not; and even if it were, Bob's completing move could be forced to go leftwards, which is illegal.

6.9 B3. Take logarithms and count exponential arrivals

Let $x_0 = 1$, and let δ be some constant satisfying $0 < \delta < 1$. Iteratively, for $n = 0, 1, 2, \ldots$, a point x_{n+1} is chosen uniformly from the interval $[0, x_n]$. Let Z be the smallest value of n for which $x_n < \delta$. Find the expected value of Z, as a function of δ.

Read it out loud

This is a first passage time for a *multiplicative* random walk. Read the recursion literally: x_{n+1} is uniform on $[0, x_n]$, which is the same sentence as $x_{n+1} = x_n U_{n+1}$ with U_{n+1} uniform on $[0, 1]$ and independent of everything before. So $x_n = U_1 U_2 \cdots U_n$, a product of independent uniforms, and Z is the first index at which that product drops below a fixed level.

The fields wired in are: transformations of random variables, sums of independent exponentials and the Gamma distribution, the tail sum formula for the expectation of a nonnegative integer variable, and (as the one line reading of the whole thing) the Poisson process. There is no calculus of variations, no differential equation, and nothing that needs the word "martingale".

What is really being tested is one reflex: a product of independent factors becomes a sum of independent increments the moment you take logarithms. The problem has been engineered so that the resulting increments are the nicest possible ones. If U is uniform on $[0, 1]$ then $\mathbb{P}(-\log U > t) = \mathbb{P}(U < e^{-t}) = e^{-t}$, so $-\log U$ is exponential with rate one. The walk $-\log x_n$ is therefore a sum of independent rate one exponentials, which is by definition the sequence of arrival times of a rate one Poisson process, and the question "when does the walk first exceed the level $\log(1/\delta)$?" is the question "how many arrivals happen in an interval of length $\log(1/\delta)$?".

Within thirty seconds you should have written three lines: $x_n = U_1 \cdots U_n$; $-\log U_i \sim \mathrm{Exp}(1)$; $L := \log(1/\delta)$ and $Z = \min\{n : T_n > L\}$ where $T_n = -\log x_n$. At that point the answer $1 + L$ is already visible, because the expected number of rate one arrivals in a window of length L is L and Z is one more than that count.

Approaches, ranked worst to best

1. **Convolve the densities of products of uniforms.** *Cue:* a product of n independent uniforms has a computable density. *Circuit:* find the law of x_n, then sum $\mathbb{P}(x_n \geq \delta)$. Doable, since the density of the product is $(-\log x)^{n-1}/(n-1)!$ on $(0, 1)$, but you are deriving that by an integral recursion that is nothing other than the logarithm in disguise, and you are doing it without admitting it. It ranks last because it is the ideal method performed blindfolded.

2. **Condition on x_1 and solve an integral equation.** *Cue:* the problem is scale invariant, so conditioning on the first step reproduces the same problem with a rescaled δ. *Circuit:* self-similarity of the process. Write $f(\delta) = \mathbb{E}[Z]$ and $f(\delta) = 1 + \int_\delta^1 f(\delta/c)\, dc$, then substitute $h(x) = f(1/x)$ to get $xh(x) - x = \int_1^x h$, differentiate and solve $h'(x) = 1/x$. This is correct and it is the route the official solution takes, but it silently uses that h is continuous before differentiating an integral, and it asks you to solve a functional equation rather than to recognise a distribution you already know.

3. **Renewal reward with Wald's identity.** *Cue:* Z is a stopping time for a sum of independent identically distributed increments. *Circuit:* Wald plus memorylessness. After taking logarithms, Wald's identity gives $\mathbb{E}[T_Z] = \mathbb{E}[Z]\,\mathbb{E}[E_1] = \mathbb{E}[Z]$, and $T_Z = L + R$ where the overshoot R is again exponential with mean one by memorylessness. Hence $\mathbb{E}[Z] = L + 1$. Short and beautiful, and it explains where the "+1" comes from, but it quotes two theorems and needs $\mathbb{E}[Z] < \infty$ established first.

4. **Logarithms plus the Poisson process.** *Cue:* the increments are exactly rate one exponentials. *Circuit:* arrival times of a Poisson process. $Z = N(L) + 1$ where $N(L)$ is the number of arrivals

in $[0, L]$, and $N(L)$ is Poisson with mean L, so $\mathbb{E}[Z] = L + 1$. This is the correct mental picture and it is two lines. It ranks just below the winner only because it imports a named process.

5. **Logarithms plus the tail sum formula.** *Cue:* expectation of a nonnegative integer variable, and $\{Z > n\}$ is exactly $\{T_n \le L\}$. *Circuit:* sum the survival probabilities, swap by positivity. $\mathbb{E}[Z] = \sum_{n\ge 0} \mathbb{P}(T_n \le L)$, each term is the Gamma cumulative distribution, and swapping sum and integral collapses the exponential series against e^{-t} to give $1 + L$. Nothing is imported beyond the Gamma density, every interchange is justified by nonnegativity, and the finiteness of $\mathbb{E}[Z]$ comes out of the computation rather than being assumed. This is the ideal write up.

The ideal solution

Write $L = \log(1/\delta)$, so $L > 0$ because $0 < \delta < 1$.

Step 1: realise the process as a product. Saying that x_{n+1} is uniform on $[0, x_n]$ is exactly saying $x_{n+1} = x_n U_{n+1}$ where $U_1, U_2, \dots$ are independent and uniform on $[0, 1]$, independent of the past: if $x_n = t > 0$ then tU_{n+1} is uniform on $[0, t]$. Hence

$$x_n = U_1 U_2 \cdots U_n \qquad (n \ge 0),$$

with the empty product equal to $x_0 = 1$. Each $U_i > 0$ almost surely, so $x_n > 0$ almost surely and the logarithms below are defined.

Step 2: take logarithms. Put $E_i = -\log U_i$. For $t \ge 0$,

$$\mathbb{P}(E_i > t) = \mathbb{P}(U_i < e^{-t}) = e^{-t},$$

so the E_i are independent exponential random variables of rate one. Setting $T_0 = 0$ and $T_n = E_1 + \cdots + E_n = -\log x_n$, the stopping condition transforms as

$$x_n < \delta \iff -\log x_n > -\log \delta \iff T_n > L,$$

so

$$Z = \min\{n \ge 0 : T_n > L\}.$$

Since $T_0 = 0 \le L$ we have $Z \ge 1$, which matches $x_0 = 1 \ge \delta$.

Step 3: the survival events. The sequence T_n is nondecreasing, so $\{Z > n\} = \{T_n \le L\}$ for every $n \ge 0$. For a random variable Z taking values in $\{0, 1, 2, \dots\} \cup \{\infty\}$,

$$\mathbb{E}[Z] = \sum_{n\ge 0} \mathbb{P}(Z > n),$$

with both sides simultaneously finite or infinite; this is Tonelli applied to $\sum_{n\ge 0} \mathbb{E}[\mathbf{1}_{\{Z>n\}}]$.

Step 4: the law of T_n. For $n \ge 1$, T_n has density $g_n(t) = t^{n-1}e^{-t}/(n-1)!$ on $t > 0$. Indeed $g_1(t) = e^{-t}$, and if g_{n-1} has this form then

$$g_n(t) = \int_0^t g_{n-1}(s)\, e^{-(t-s)}\, ds = e^{-t} \int_0^t \frac{s^{n-2}}{(n-2)!}\, ds = \frac{t^{n-1}e^{-t}}{(n-1)!},$$

which is the induction step.

Step 5: sum. Using $\mathbb{P}(Z > 0) = 1$ and $\mathbb{P}(Z > n) = \mathbb{P}(T_n \le L) = \int_0^L \frac{t^{n-1}e^{-t}}{(n-1)!}\, dt$ for $n \ge 1$,

$$\mathbb{E}[Z] = 1 + \sum_{n=1}^{\infty} \int_0^L \frac{t^{n-1}e^{-t}}{(n-1)!}\, dt = 1 + \int_0^L e^{-t} \sum_{n=1}^{\infty} \frac{t^{n-1}}{(n-1)!}\, dt = 1 + \int_0^L e^{-t}e^{t}\, dt = 1 + L.$$

Every term of the double sum is nonnegative, so the interchange of $\sum$ and $\int$ is justified by Tonelli's theorem. In particular $\mathbb{E}[Z]$ is finite, so $Z < \infty$ almost surely.

Therefore

$$\boxed{\mathbb{E}[Z] = 1 + \log(1/\delta) = 1 - \log \delta}$$

with log the natural logarithm.

Remark (the same computation in one line). The points $T_1 < T_2 < \cdots$ are the arrival times of a rate one Poisson process, so $N(L) := \#\{n \geq 1 : T_n \leq L\}$ is Poisson with mean L, and $Z = N(L) + 1$. Hence $\mathbb{E}[Z] = L + 1$, and moreover $Z - 1$ is Poisson, so $\mathrm{Var}(Z) = L$ as well.

Check. As $\delta \to 1^-$, $L \to 0$ and $\mathbb{E}[Z] \to 1$: with δ near 1 the very first point falls below δ with probability near one. As $\delta \to 0^+$, $\mathbb{E}[Z] \to \infty$ logarithmically, which is right because the walk descends by a constant factor per step on average.

Layer 1: WHAT

Block 1. Rewrite "uniform on $[0, x_n]$" as "x_n times an independent uniform on $[0, 1]$", so x_n is a product of independent identically distributed factors.

Block 2. Take logarithms; the product becomes a sum, and the summands are rate one exponentials because $\mathbb{P}(-\log U > t) = e^{-t}$ directly from the definition.

Block 3. Translate the stopping rule: $x_n < \delta$ becomes $T_n > L$ with $L = \log(1/\delta)$, so Z is a first passage index above a fixed level.

Block 4. Use monotonicity of T_n to identify the survival events $\{Z > n\} = \{T_n \leq L\}$, and use the tail sum formula for $\mathbb{E}[Z]$.

Block 5. Write down the Gamma density of T_n and hence each survival probability as an integral over the fixed window $[0, L]$.

Block 6. Swap sum and integral, recognise $\sum t^{n-1}/(n-1)! = e^t$, and watch $e^{-t}e^t = 1$ collapse the integral to L.

The nominal route runs: "let $f(\delta)$ be the answer; condition on the first point; this gives $f(\delta) = 1 + \int_\delta^1 f(\delta/c)\, dc$; now solve the functional equation." It arrives at the same place after a substitution, a differentiation under an integral sign, and a first order differential equation.

The bifurcation point is the first line you write. The zombie writes "let $f(\delta) = \mathbb{E}[Z]$", that is, it starts from the *answer* and tries to characterise it. The ideal solver writes "$x_{n+1} = x_n U_{n+1}$", that is, it starts from the *process* and changes coordinates until the process is one it already knows. Naming the unknown is the reflex that costs you the elegant solution here, because the unknown is not the object with structure; the walk is.

Layer 2: HOW

Block 1. Cue. "Chosen uniformly from the interval $[0, x_n]$" has the previous value in the range of the uniform, not in the answer. **Why natural.** A uniform law on $[0, t]$ is the scaling by t of the uniform law on $[0, 1]$. That is the definition of the uniform distribution, not an insight. So the recursion is multiplication by an independent factor with a fixed law, and the process is a multiplicative random walk with independent increments. **Link to the next block.** A walk with independent multiplicative increments is a walk with independent additive increments after one change of coordinates.

Block 2. Cue. The state space is $(0, 1]$ under multiplication. **Why natural.** The logarithm is the isomorphism from the multiplicative group of positive reals to the additive reals. Whenever a recursion multiplies, the coordinate in which it is linear is the logarithm. The bonus is that the specific law here is the nicest one there is: $-\log U$ for U uniform is exponential, which you verify in one line from the definition of the distribution function. **Link to the next block.** Additive walks with independent increments are first passage problems, and first passage above a level is what the question asks.

Block 3. **Cue.** "The smallest value of n for which $x_n < \delta$." **Why natural.** A threshold in the original coordinate becomes a threshold in the logarithmic coordinate, because log is increasing. The parameter δ enters only through $L = \log(1/\delta)$, which is the first sign that the answer will be a function of L alone and probably an affine one. **Link to the next block.** You now need the distribution of the first passage index, and the cheapest handle on any stopping index is its tail.

Block 4. **Cue.** The quantity requested is an expectation of a nonnegative integer variable. **Why natural.** For such a variable the tail sum formula turns an expectation into a sum of probabilities of "not yet stopped" events. Here that event is $\{T_n \le L\}$, which is a statement about a sum of n known independent variables, and therefore computable. Computing the law of Z itself would require the joint statement $T_n \le L < T_{n+1}$, which is strictly harder for no benefit. **Link to the next block.** Each survival probability is a Gamma cumulative distribution over the same window.

Block 5. **Cue.** A sum of n independent rate one exponentials. **Why natural.** This is the Gamma distribution, and the density is obtained by one convolution induction that takes two lines. You should never look this up in the arena; the induction is faster than remembering. **Link to the next block.** You now have a sum of integrals over one fixed interval, all sharing the factor e^{-t}.

Block 6. **Cue.** A sum over n of $t^{n-1}/(n-1)!$. **Why natural.** That is the exponential series, and it is sitting next to e^{-t}. Be lazy: swap the order, let the series do the work, and the integrand becomes the constant 1. The swap needs no delicate justification because every term is nonnegative, so Tonelli applies with no hypotheses to check. **Link to the next block.** The integral of 1 over $[0, L]$ is L, so the answer is $1 + L$.

Layer 3: WHY

The mathematical truth beneath this route is the identification of two symmetries. The problem is *scale invariant*: conditioned on $x_1 = c$, the remaining process is a copy of the original with δ replaced by δ/c. That is what the official integral equation encodes. On the other side, the exponential distribution is the unique *memoryless* law on $[0, \infty)$: $\mathbb{P}(E > s + t \mid E > s) = \mathbb{P}(E > t)$. These are the same statement transported by the logarithm. Scale invariance in the multiplicative picture is memorylessness in the additive picture. So you may either derive and solve a functional equation expressing the symmetry, or change coordinates so that the symmetry is carried by a distribution whose properties you already know. The second is strictly less work, and it is available whenever the increments are multiplicative.

The truth also explains the shape of the answer. Split $\mathbb{E}[Z] = L + 1$ into its two pieces. The L is the expected number of complete crossings: the walk moves by an average of one unit per step, and it must cover a distance L. The 1 is the overshoot step, the one move that takes the walk past the level. Formally, Wald's identity says $\mathbb{E}[T_Z] = \mathbb{E}[Z]\,\mathbb{E}[E_1] = \mathbb{E}[Z]$, while $T_Z = L + R$ with R the overshoot; memorylessness makes R exponential with mean one regardless of L, so $\mathbb{E}[Z] = L + 1$. Every "+1" in a first passage answer is an overshoot, and knowing that lets you predict the shape of the answer before computing anything.

The discernment mechanism is a short filter. When you meet a random recursion, ask what group the increment acts by. If it is addition, work directly. If it is multiplication on $(0, \infty)$, take logarithms without hesitation. Then ask what the log increment law is: for uniform on $[0, 1]$ it is exponential, for uniform on $[a, b]$ with $0 < a < b$ it is uniform on $[\log(1/b), \log(1/a)]$, for a Beta law it is one of the classical families. Only when the log increment is not a named law do you fall back on functional equations. Training this is a matter of doing the substitution reflexively, even when the problem looks as though it wants a recursion in δ.

The examiners nudged you three times. They set $x_0 = 1$, which is the multiplicative identity and therefore the natural origin after taking logarithms. They chose the uniform distribution rather than

any other law on $[0, x_n]$, which is precisely the choice that makes the log increments exponential. And they asked for the answer "as a function of δ", which tells you the answer depends on δ alone and not on any further data, and hence that only the single number $\log(1/\delta)$ can appear.

Revolving concepts

Uniform scaling and the definition of a uniform law. If U is uniform on $[0,1]$ and $t > 0$ then tU is uniform on $[0,t]$; conversely any variable uniform on $[0,t]$ equals t times an independent uniform on $[0,1]$ in distribution. Pointer: see Book 1, field: Probability, topic: transformations of continuous random variables.

The probability integral transform. If F is a continuous strictly increasing distribution function and U is uniform on $[0,1]$ then $F^{-1}(U)$ has distribution function F; the case $F(t) = 1 - e^{-t}$ gives $-\log U \sim \mathrm{Exp}(1)$ and is the one you use here. Pointer: see Book 1, field: Probability, topic: inverse transform sampling.

Sums of exponentials and the Gamma law. If $E_1, \dots, E_n$ are independent $\mathrm{Exp}(\lambda)$ then $T_n = \sum E_i$ has density $\lambda^n t^{n-1} e^{-\lambda t}/(n-1)!$. The recursion is the convolution $g_n = g_{n-1} * g_1$, whose general solution for exponential increments is the Gamma density above; the same convolution recursion for a general increment density g defines the renewal density, and the general solution method is to pass to Laplace transforms, where $\widehat{g_n} = (\widehat{g})^n$ and $\sum_n \widehat{g_n} = \widehat{g}/(1-\widehat{g})$. Pointer: see Book 1, field: Probability, topic: convolutions and the Gamma family.

Poisson process. If T_n are the partial sums of independent $\mathrm{Exp}(\lambda)$ variables and $N(t) = \#\{n \geq 1 : T_n \leq t\}$, then $N(t)$ is Poisson with mean λt, and $\mathbb{P}(T_n > t) = \mathbb{P}(N(t) \leq n-1) = e^{-\lambda t} \sum_{j<n} (\lambda t)^j / j!$. This last identity is the same statement as the Gamma tail integral, and it converts the integral in Step 5 into a finite sum if you prefer. Pointer: see Book 1, field: Probability, topic: Poisson processes.

Tail sum formula. For Z with values in $\{0,1,2,\dots\} \cup \{\infty\}$, $\mathbb{E}[Z] = \sum_{n \geq 0} \mathbb{P}(Z > n)$, proved by writing $Z = \sum_{n \geq 0} \mathbf{1}_{\{Z > n\}}$ and applying Tonelli. Pointer: see Book 1, field: Probability, topic: expectations via tails.

Wald's identity and overshoot. If Z is a stopping time for independent identically distributed integrable increments E_i and $\mathbb{E}[Z] < \infty$, then $\mathbb{E}[T_Z] = \mathbb{E}[Z]\mathbb{E}[E_1]$. For exponential increments memorylessness makes the overshoot $T_Z - L$ exponential with the same rate, independent of L, which is why first passage answers of this type have the shape "level divided by mean drift, plus one". Pointer: see Book 1, field: Probability, topic: renewal theory and Wald's identity.

Memorylessness and scale invariance. The exponential law is the unique memoryless law on $[0,\infty)$; under $x \mapsto -\log x$ this is exactly the scale invariance of the uniform recursion. Pointer: see Book 1, field: Probability, topic: characterisations of the exponential distribution.

Circuit Bank

- **WHEN** a random recursion of the form $X_{n+1} = X_n \cdot (\text{independent factor})$ **DO** take logarithms immediately, before writing anything else **BECAUSE** log is the isomorphism from the multiplicative positives to the additive reals, so it converts a multiplicative walk into a sum of independent identically distributed increments, which is the only walk anyone knows how to analyse
- **WHEN** a uniform factor on $[0,1]$ inside a product **DO** write $-\log U \sim \mathrm{Exp}(1)$ **BECAUSE** $\mathbb{P}(-\log U > t) = \mathbb{P}(U < e^{-t}) = e^{-t}$ straight from the definition, and exponential increments turn the walk into a Poisson process whose counts you know exactly
- **WHEN** an expectation of a first passage index or any other nonnegative integer variable **DO** use $\mathbb{E}[Z] = \sum_{n \geq 0} \mathbb{P}(Z > n)$ **BECAUSE** the survival events are statements about the first n increments only, so they are computable, whereas the point probabilities of Z require a two sided sandwich

- **WHEN** a sum of integrals of nonnegative functions that you want to interchange **DO** cite Tonelli and swap without further comment **BECAUSE** nonnegativity makes the interchange unconditional, so there is no convergence hypothesis to check and no ink to spend
- **WHEN** a first passage answer that comes out as "level plus a constant" **DO** identify the constant as the expected overshoot **BECAUSE** Wald's identity splits $\mathbb{E}[Z]$ into the distance to be covered and the final partial step, and for memoryless increments the overshoot has the same law as a fresh increment

Backtesting

1. First useful cue: "uniform on $[0, x_n]$" with the previous state inside the interval, which says multiplicative increment, which says take logarithms.
2. Method the expression suggested: convert the product to a sum of rate one exponentials, read the stopping rule as a first passage above $L = \log(1/\delta)$, and evaluate the expectation by summing survival probabilities.
3. Where you would hesitate, and the thought to reinforce: the hesitation is whether you are allowed to realise the whole process with one independent uniform per step. Reinforce with "a uniform on $[0, t]$ *is* t times a uniform on $[0, 1]$", which is the definition and not a modelling assumption.
4. The trap the question sets: asking for the answer "as a function of δ" invites you to name $f(\delta)$ and hunt for an equation it satisfies. That route works but it costs a functional equation, a substitution and a differentiation under an integral, and it needs a continuity remark that most write ups omit.
5. What makes you recognise it faster next time: the pairing "product of independent factors" with "first time below a threshold". Those two phrases together mean Poisson process, always.
6. The reusable decision rule: change coordinates until the increments are independent and additive, then read off the answer from a distribution you already know, rather than characterising the answer by an equation.

Mirror drills

1. Same process, but Z' is the smallest n with $x_n < \delta^2$. Find $\mathbb{E}[Z']$. *Answer:* $1 + \log(1/\delta^2) = 1 - 2\log\delta$; only the level L changes, and it doubles.
2. Show that $Z - 1$ is Poisson distributed and deduce $\mathrm{Var}(Z)$. *Answer:* $\mathbb{P}(Z - 1 = j) = \mathbb{P}(T_j \leq L < T_{j+1}) = e^{-L}L^j/j!$, so $Z - 1 \sim \mathrm{Poisson}(L)$ and $\mathrm{Var}(Z) = L = -\log\delta$.
3. For independent uniforms $U_1, \ldots, U_n$ on $[0, 1]$, compute $\mathbb{P}(U_1U_2 \cdots U_n < \delta)$. *Answer:* $\delta \sum_{j=0}^{n-1} \frac{(\log(1/\delta))^j}{j!}$, since the event is $\{T_n > L\} = \{N(L) \leq n - 1\}$ and $N(L)$ is Poisson with mean L, while $e^{-L} = \delta$.

6.10 B4. Cyclic rotation makes the reciprocals sum to one

Let n be a positive integer, and let V_n be the set of integer $(2n+1)$-tuples $\mathbf{v} = (s_0, s_1, \ldots, s_{2n-1}, s_{2n})$ for which $s_0 = s_{2n} = 0$ and $|s_j - s_{j-1}| = 1$ for $j = 1, 2, \ldots, 2n$. Define

$$q(\mathbf{v}) = 1 + \sum_{j=1}^{2n-1} 3^{s_j},$$

and let $M(n)$ be the average of $\frac{1}{q(\mathbf{v})}$ over all $\mathbf{v} \in V_n$. Evaluate $M(2020)$.

Read it out loud

This is the average of the reciprocal of a weighted sum, taken over all closed lattice paths of length $2n$ with unit steps. The path starts and ends at height zero, and q collects the value 3^{height} at every intermediate time. The fields wired in are bijective combinatorics (paths and step words), finite group actions (the cyclic group of order $2n$ acting by rotation), and double counting. Nothing analytic and nothing about Catalan numbers.

Two features of the statement are shouting. The first is the reciprocal. Averages of reciprocals almost never have closed forms, because $\mathbb{E}[1/X]$ requires the whole distribution of X while $\mathbb{E}[X]$ requires only its mean. So if a clean answer exists, it cannot come from knowing the distribution of q; it must come from a structural identity that produces $1/q$ directly. The only such identity that exists is this one: if a group acts and q scales by a multiplier under the action, then the reciprocals over one orbit add up to something you can control.

The second is the shape of q itself. It is written as $1 + \sum_{j=1}^{2n-1}$, which is one short at the top and carries a stray additive constant. Since $s_{2n} = 0$, that constant is exactly $3^{s_{2n}}$. So

$$q(\mathbf{v}) = \sum_{j=1}^{2n} 3^{s_j},$$

a sum over a complete period. The problem was deliberately written in the asymmetric form to hide the cyclic symmetry from you. Undoing that disguise is the first move.

Within thirty seconds you should have written: $q = \sum_{j=1}^{2n} 3^{s_j}$; "rotate the step word"; and the guess $M(n) = 1/(2n)$. The case $n = 1$ confirms it in one line: the two paths give $q = 1 + 3 = 4$ and $q = 1 + \frac{1}{3} = \frac{4}{3}$, whose reciprocals $\frac{1}{4}$ and $\frac{3}{4}$ average to $\frac{1}{2}$. The tell that $1/(2n)$ is the right shape is that $4040 = 2 \cdot 2020$ is the length of the walk, that is, the order of the rotation group.

Approaches, ranked worst to best

1. **Find the distribution of q and average its reciprocal.** *Cue:* it is an expectation over a finite set. *Circuit:* transfer matrix or generating function in the height. You would need the joint law of all $2n$ heights, since q couples them all, and then you would need to invert. The reciprocal destroys every product structure a transfer matrix could give you. This is the definition of zombie maths on this problem.

2. **Expand $1/q$ as a geometric series.** *Cue:* q is a sum, and $1/(1+u) = 1 - u + u^2 - \cdots$. *Circuit:* series expansion of the reciprocal. The heights are both positive and negative, so q ranges over many orders of magnitude and no expansion point works uniformly. Even where it converges you are left with an infinite family of moment computations. It stalls immediately.

3. **Small cases then conjecture.** *Cue:* a clean numerical answer is demanded. *Circuit:* compute $n = 1, 2, 3$ and read the pattern. Cheap and essential. It produces $M(1) = \frac{1}{2}$, $M(2) = \frac{1}{4}$, $M(3) = \frac{1}{6}$, and, if you also try base 2 instead of base 3, it produces the far more valuable observation that the answer does not depend on the base at all. That observation is the real

content of the reconnaissance, because a parameter that cannot appear in the answer cannot appear in the mechanism.

4. **Rotate and analyse orbits.** *Cue:* closed path, so the step word can be rotated. *Circuit:* group action, then average orbit by orbit. Show q scales under rotation, deduce that the average of $1/q$ over each orbit is $1/(2n)$, then note that V_n is a disjoint union of orbits. Correct, and it is the official route, but it forces you to discuss orbits whose length m properly divides $2n$ and to argue that running k from 1 to $2n$ traverses such an orbit $2n/m$ times with each element equally often. That paragraph is pure bookkeeping.

5. **Rotate and double count.** *Cue:* the same cue, but finished by Fubini. *Circuit:* sum the orbit identity over the whole set, using only that each rotation is a bijection. Sum the identity $\sum_{k=1}^{2n} 1/q(\varphi^k \mathbf{w}) = 1$ over all $\mathbf{w}$, then interchange the two sums. Each φ^k is a bijection of the set, so the inner sum is the same total T every time, giving $2n\,T = |V_n|$. No orbit ever has to be examined. This is the ideal route: it is three lines shorter and it removes the only place where the official argument could be mis-stated.

The ideal solution

Fix any real $a > 0$ and define $q_a(\mathbf{v}) = 1 + \sum_{j=1}^{2n-1} a^{s_j}$. We prove that the average of $1/q_a$ over V_n equals $\frac{1}{2n}$, independently of a; the problem is the case $a = 3$.

Step 1: restore the symmetry. Since $s_{2n} = 0$ we have $a^{s_{2n}} = 1$, so

$$q_a(\mathbf{v}) = \sum_{j=1}^{2n} a^{s_j}.$$

Every term is positive, so $q_a(\mathbf{v}) > 0$ and $1/q_a$ is defined.

Step 2: pass to step words. Let W_n be the set of tuples $\mathbf{w} = (w_1, \ldots, w_{2n}) \in \{+1, -1\}^{2n}$ with $\sum_{i=1}^{2n} w_i = 0$, that is, with exactly n entries equal to $+1$. The map

$$\mathbf{w} \longmapsto \mathbf{v}, \qquad s_0 = 0, \quad s_j = w_1 + \cdots + w_j \ (1 \le j \le 2n),$$

is a bijection $W_n \to V_n$, with inverse $w_j = s_j - s_{j-1}$. Write $Q(\mathbf{w}) = q_a(\mathbf{v}) = \sum_{j=1}^{2n} a^{s_j}$.

Step 3: rotate. Extend the indices periodically by $w_{i+2n} = w_i$. Because $\sum_{i=1}^{2n} w_i = 0$, the partial sums satisfy $s_{j+2n} = s_j$ for all $j \ge 0$; this is exactly where the hypothesis $s_{2n} = 0$ is used, and it is the only place. Let $\varphi : W_n \to W_n$ be the rotation $\varphi(\mathbf{w})_i = w_{i+1}$, which is a bijection of W_n since it preserves the number of $+1$ entries. The i-th entry of $\varphi^k(\mathbf{w})$ is w_{i+k}, so the partial sums of $\varphi^k(\mathbf{w})$ are

$$s_j^{(k)} = \sum_{i=1}^{j} w_{i+k} = s_{j+k} - s_k.$$

Therefore

$$Q(\varphi^k \mathbf{w}) = \sum_{j=1}^{2n} a^{s_{j+k} - s_k} = a^{-s_k} \sum_{j=1}^{2n} a^{s_{j+k}} = a^{-s_k} \sum_{j=1}^{2n} a^{s_j} = a^{-s_k} Q(\mathbf{w}),$$

where the third equality holds because $j \mapsto j + k$ permutes a complete period and a^{s_j} has period $2n$ in j.

Step 4: the reciprocals over one rotation cycle add to one. For each fixed $\mathbf{w}$,

$$\sum_{k=1}^{2n} \frac{1}{Q(\varphi^k \mathbf{w})} = \sum_{k=1}^{2n} \frac{a^{s_k}}{Q(\mathbf{w})} = \frac{1}{Q(\mathbf{w})} \sum_{k=1}^{2n} a^{s_k} = \frac{Q(\mathbf{w})}{Q(\mathbf{w})} = 1.$$

The last step is not a coincidence: $Q(\mathbf{w})$ was defined to be the sum of the very multipliers a^{s_k} that the rotation produces.

Step 5: double count. Sum the identity of Step 4 over all $\mathbf{w} \in W_n$ and interchange the two finite sums:

$$|W_n| = \sum_{\mathbf{w}\in W_n} \sum_{k=1}^{2n} \frac{1}{Q(\varphi^k \mathbf{w})} = \sum_{k=1}^{2n} \sum_{\mathbf{w}\in W_n} \frac{1}{Q(\varphi^k \mathbf{w})} = \sum_{k=1}^{2n} \sum_{\mathbf{w}'\in W_n} \frac{1}{Q(\mathbf{w}')} = 2n \sum_{\mathbf{w}'\in W_n} \frac{1}{Q(\mathbf{w}')},$$

where the third equality uses only that each φ^k is a bijection of W_n. Hence

$$\frac{1}{|W_n|} \sum_{\mathbf{w}\in W_n} \frac{1}{Q(\mathbf{w})} = \frac{1}{2n},$$

and transporting along the bijection $W_n \to V_n$ gives $M(n) = \frac{1}{2n}$ for every $a > 0$.

Taking $n = 2020$ and $a = 3$,

$$\boxed{M(2020) = \frac{1}{4040}}.$$

Check, $n = 1$. $V_1 = \{(0,1,0),(0,-1,0)\}$, and $q(\mathbf{v}) = 1 + 3^{s_1}$ takes the values 4 and $1 + \frac{1}{3} = \frac{4}{3}$. The average of the reciprocals is $\frac{1}{2}(\frac{1}{4} + \frac{3}{4}) = \frac{1}{2}$, which is $1/(2 \cdot 1)$ as claimed. The same two paths with $a = 2$ give $\frac{1}{2}(\frac{1}{3} + \frac{2}{3}) = \frac{1}{2}$, confirming the independence of a.

Layer 1: WHAT

Block 1. Absorb the stray 1 into the sum using $s_{2n} = 0$, turning q into a sum over a full period $j = 1, \ldots, 2n$.

Block 2. Replace the path by its step word, because rotation is an operation on words and not on paths.

Block 3. Extend the word periodically. This is legitimate exactly because the path is closed, and it is the only use of $s_0 = s_{2n} = 0$.

Block 4. Compute the effect of a rotation on q: the heights shift by $-s_k$, so q is multiplied by a^{-s_k}.

Block 5. Sum the reciprocals over the $2n$ rotations. The multipliers are exactly the terms of q, so the sum telescopes to 1.

Block 6. Double count over all words and all rotations, using bijectivity of each rotation, to convert the per cycle identity into the global average $1/(2n)$.

The nominal route runs: "this is an expectation, so I need the distribution of q; let me set up a transfer matrix in the height and sum over paths". That route never recovers, because $1/q$ has no product structure in the path.

The bifurcation point is Block 1 and the sentence that follows it. The zombie sees $1 + \sum_{j=1}^{2n-1}$ and treats the 1 as a constant to be carried along. The ideal solver sees a sum that stops one short of a full period next to a stray constant, recognises that the constant is the missing term, and asks what symmetry the restored expression has. Everything after that is forced.

Layer 2: HOW

Block 1. Cue. A sum whose upper limit is $2n - 1$, one short of the natural $2n$, with a bare 1 added on. **Why natural.** Bare additive constants in a competition expression are almost never constants. They are values of the summand at a boundary index that has been hidden. Here $s_{2n} = 0$ is given in the hypothesis, so $a^{s_{2n}} = 1$, and the constant is the missing term. Restoring it makes the expression a sum over a complete period, which is the form in which a rotation can act. **Link to the next block.** A complete period invites a cyclic group action.

Block 2. **Cue.** The objects are paths, but the natural group action rotates *steps*, not heights. **Why natural.** Rotating heights makes no sense, since the height sequence is constrained by $s_0 = 0$. The steps, on the other hand, form an unconstrained word except for having n of each sign, and that condition is rotation invariant. So change coordinates from heights to steps first. The change is a bijection with an explicit inverse, so nothing is lost. **Link to the next block.** With the word in hand you can rotate, and you then need to know what rotation does to q.

Block 3. **Cue.** $s_0 = s_{2n} = 0$ is stated twice, once at each end. **Why natural.** A hypothesis stated at both ends of a sequence is a periodicity hypothesis. The total increment is zero, so the partial sums extend to a genuinely periodic sequence, and only then is "the height at time $j + k$" meaningful for $j + k > 2n$. Without this, rotation would change the heights by an amount depending on how far around you went. **Link to the next block.** Now compute the heights of the rotated word.

Block 4. **Cue.** The partial sums of a rotated word are differences of the original partial sums. **Why natural.** This is immediate: starting the walk at time k instead of time 0 shifts all heights down by s_k. A downward shift of all heights multiplies every term a^{s_j} by the same factor a^{-s_k}, so q scales rather than changing in any complicated way. A functional that merely scales under a group action is the best possible situation, because reciprocals then scale too. **Link to the next block.** You now have $2n$ values of q, all proportional to one another with known constants.

Block 5. **Cue.** The proportionality constants are a^{s_k}, and q is the sum of the a^{s_k}. **Why natural.** This is why the problem works, and it is not luck. The definition of q was engineered to be the sum over the group of the multipliers of the action. Whenever that is so, $\sum_k 1/q(\varphi^k \mathbf{w}) = \sum_k a^{s_k}/q(\mathbf{w}) = 1$. Be lazy and notice this before computing anything: the answer $1/(2n)$ is now inevitable, since $2n$ rotations share a total of 1. **Link to the next block.** One identity per rotation cycle now has to be turned into a statement about the whole set.

Block 6. **Cue.** You have an identity indexed by pairs (word, rotation) and want a statement about words alone. **Why natural.** That is a double count. Sum over both indices and interchange. The only property of φ^k you need is that it is a bijection of the finite set, which is obvious. Orbit lengths, stabilisers and Burnside never enter, because you are not counting orbits; you are counting pairs. **Link to the next block.** Divide by $|W_n|$ to get the average, then substitute $n = 2020$.

Layer 3: WHY

The mathematical truth beneath this route is that q is a *multiplier* function, or in the language of group cohomology a coboundary: $Q(\varphi^k \mathbf{w}) = \chi_k(\mathbf{w})\, Q(\mathbf{w})$ with $\chi_k(\mathbf{w}) = a^{-s_k}$. The general principle, which you should install once and reuse forever, is this. Let a finite group G of order N act on a finite set X, and suppose a positive function F on X satisfies $F(gx) = \lambda(g, x)^{-1} F(x)$ for some positive multipliers, with the compatibility

$$F(x) = \sum_{g \in G} \lambda(g, x).$$

Then $\sum_{g \in G} 1/F(gx) = 1$ for every x, and hence the average of $1/F$ over X is $1/N$. Both the rotation structure and the definition of q in this problem are exactly the hypotheses of that principle. The examiners built the function so that the multipliers are the summands.

The abstract skeleton, stripped of all lattice paths, is a two line identity you should keep in your head: for positive reals $x_1, \ldots, x_m$ with cyclic indexing and $S = \sum_j x_j$, set $Q_k = \sum_{j=1}^{m} x_{j+k}/x_k$. Then $Q_k = S/x_k$, so $\sum_{k=1}^{m} 1/Q_k = \sum_k x_k/S = 1$. The whole of B4 is this identity with $x_j = a^{s_j}$, plus the observation that rotating the word is what produces the ratios x_{j+k}/x_k.

The discernment mechanism is a two step filter for any average of a reciprocal. First: is there a group acting on the index set? Look for closed loops, cyclic conditions, and symmetric constraints; "starts and ends at zero" is the standard signal. Second: does the denominator scale under that action, and are the multipliers the summands of the denominator? If yes, the answer is one over the order of the group, and you never compute anything. If the answer is claimed to be independent of a parameter that appears in the definition, this is almost always the mechanism, because a computational proof would have to exhibit the cancellation of that parameter explicitly, while a symmetry proof never touches it.

To train it, treat every "average of $1/(\text{something})$" as a signal, not as an obstacle. Reciprocals are the sign that a partition of unity is hiding somewhere. The rotation trick used here is the same one behind the cycle lemma of Dvoretzky and Motzkin, which says that among the m rotations of a word of integers each at most 1 with total $r > 0$, exactly r have all partial sums positive. That lemma proves the ballot problem and the Catalan number formula in one line, and it is the same circuit: rotate a closed or nearly closed path and count what fraction of rotations have a property.

The examiners nudged you three times. They wrote q in the asymmetric form with a stray 1, which is the nudge (by concealment) toward restoring the period. They imposed $s_0 = s_{2n} = 0$, which is the closed loop condition and the only hypothesis rotation needs. And they asked for $M(2020)$ rather than a general formula, so that the answer $1/4040$ would visibly be one over the path length, which is one over the order of the rotation group.

Revolving concepts

Lattice paths and step words. A path $(s_0, \dots, s_{2n})$ with $s_0 = s_{2n} = 0$ and $|s_j - s_{j-1}| = 1$ corresponds bijectively to a word $\mathbf{w} \in \{\pm 1\}^{2n}$ with n entries of each sign, via $s_j = \sum_{i \le j} w_i$; hence $|V_n| = \binom{2n}{n}$. The recursion $s_j = s_{j-1} + w_j$ is the simplest first order linear recursion, and its general solution is the partial sum; the periodic extension $s_{j+2n} = s_j$ holds precisely when the total increment $\sum_i w_i$ vanishes, which is the general criterion for a linear recursion with an additive forcing term to be periodic. Pointer: see Book 1, field: Combinatorics, topic: lattice paths and Dyck words.

Group actions on finite sets and double counting. If a finite group G acts on a finite set X and $f : X \to \mathbb{R}$, then $\sum_x \sum_g f(gx) = |G| \sum_x f(x)$, because each g acts as a bijection. This is Fubini for finite sums, and it is strictly weaker and strictly easier than the orbit counting lemma; use it whenever you have an identity that holds along orbits. Pointer: see Book 1, field: Abstract algebra, topic: group actions and counting.

Multiplier (cocycle) functions. A positive F on X with $F(gx) = \lambda(g,x)^{-1} F(x)$ and $F(x) = \sum_g \lambda(g,x)$ satisfies $\sum_g 1/F(gx) = 1$. In this problem $\lambda(\varphi^k, \mathbf{w}) = a^{s_k}$. This is the abstract statement of the partition of unity that the problem conceals. Pointer: see Book 1, field: Abstract algebra, topic: cocycles and multiplier identities.

Cycle lemma of Dvoretzky and Motzkin. If $w_1, \dots, w_m$ are integers with each $w_i \le 1$ and $\sum w_i = r > 0$, then exactly r of the m cyclic rotations have all partial sums positive. Consequences include the ballot problem and $\frac{1}{n+1}\binom{2n}{n}$ for the Catalan numbers. Same circuit, different conclusion. Pointer: see Book 1, field: Combinatorics, topic: cycle lemma and ballot problems.

Reciprocal averages and partitions of unity. $\mathbb{E}[1/X]$ generally requires the full law of X, so a closed form for such an average signals an exact identity, typically $\sum_i 1/X_i = 1$ over some structured family. Pointer: see Book 1, field: Combinatorics, topic: partitions of unity in enumeration.

Circuit Bank

- **WHEN** an average of $1/q$ over a structured finite family **DO** look for a group acting on the family under which q scales, not for the distribution of q **BECAUSE** averages of reciprocals have no closed form unless an exact identity produces them, and the only common source of such identities is a family of reciprocals summing to one along an orbit

- **WHEN** a bare additive constant next to a sum whose index range stops one short of a natural period **DO** absorb the constant as the missing boundary term **BECAUSE** competition writers hide symmetry by peeling off an endpoint, and restoring it is what makes the group action visible
- **WHEN** a sequence constrained to have the same value at both ends **DO** extend it periodically and rotate **BECAUSE** equal endpoints means the total increment is zero, which is exactly the condition for the periodic extension of the partial sums to be consistent
- **WHEN** an answer that is claimed to be independent of a parameter appearing in the definition **DO** stop computing and look for a symmetry proof **BECAUSE** any computational proof would have to exhibit the cancellation of that parameter explicitly, whereas a symmetry proof never touches it
- **WHEN** an identity that holds for each orbit of a group action, and you want a global average **DO** sum over the whole set crossed with the group and interchange the two sums **BECAUSE** each group element acts as a bijection, so the inner sum is the same total every time, and orbit lengths never enter

Backtesting

1. First useful cue: the stray 1 in front of a sum that stops at $2n-1$, sitting next to the hypothesis $s_{2n}=0$. Absorbing the 1 is the entire opening move.
2. Method the expression suggested: cyclic rotation of the step word, under which q is multiplied by a^{-s_k}, followed by summation of reciprocals over the rotation cycle.
3. Where you would hesitate, and the thought to reinforce: the hesitation is whether the rotated path is still in V_n and whether the heights behave as claimed. Reinforce with "starting the same closed walk at a later time shifts every height by the same constant", which is a picture, not a computation.
4. The trap the question sets: the reciprocal invites you to hunt for the distribution of q, and the base 3 invites you to think the arithmetic of 3 matters. Neither is true: the answer holds for every positive base, and the distribution of q is never needed.
5. What makes you recognise it faster next time: the combination "closed path" plus "average of a reciprocal" plus "answer is one over the length". Those three together mean rotation, every time.
6. The reusable decision rule: if a positive functional scales under a finite group action and the multipliers are exactly the terms it is built from, then its reciprocals sum to one over each orbit, so its reciprocal averages to one over the order of the group.

Mirror drills

1. Let $x_1,\dots,x_m$ be positive reals indexed cyclically, and set $Q_k=\sum_{j=1}^{m} x_{j+k}/x_k$. Prove $\sum_{k=1}^{m} 1/Q_k = 1$. *Answer:* $Q_k = S/x_k$ with $S=\sum_j x_j$, so $\sum_k 1/Q_k = \sum_k x_k/S = 1$. This is B4 with the lattice path removed.
2. Replace 3 by any $a>0$ in the definition of q. What is the average of $1/q$? *Answer:* still $1/(2n)$; the proof never uses the value of a, only that every a^{s_j} is positive.
3. Where exactly does the argument need $s_{2n}=0$, and what breaks without it? *Answer:* it is needed for $s_{j+2n}=s_j$, so that rotating the word shifts all heights by the single constant s_k. If $s_{2n}=c\neq 0$ then $s_{j+2n}=s_j+c$, the multiplier depends on how far around you have gone, and q no longer merely scales.

6.11 B5. Reciprocal substitution pins four points to one line

For $j \in \{1,2,3,4\}$, let z_j be a complex number with $|z_j| = 1$ and $z_j \neq 1$. Prove that

$$3 - z_1 - z_2 - z_3 - z_4 + z_1 z_2 z_3 z_4 \neq 0.$$

Read it out loud

This is a nonvanishing statement about one specific symmetric polynomial in four unimodular variables. Say the polynomial out loud in symmetric function language: it is $3 - e_1 + e_4$, where e_k denotes the k-th elementary symmetric polynomial in $z_1, \ldots, z_4$. The middle terms e_2 and e_3 are absent, and the constant is $3 = 4 - 1$. Both facts are the same signal.

Substitute $z_1 = z_2 = z_3 = z_4 = 1$: the expression becomes $3 - 4 + 1 = 0$. So the polynomial vanishes at exactly the point that the hypothesis excludes. That is not a coincidence, it is the design of the problem, and it tells you the coordinate to work in. When an expression vanishes at a forbidden point, move that point to the origin. Set

$$u_j = 1 - z_j \neq 0.$$

The fields wired in are: complex numbers and the geometry of Mobius transformations; elementary symmetric polynomials and the reversal identity that inverts roots; and one elementary real inequality at the very end. Nothing requires multivariable calculus, compactness, or the geometry of hypocycloids, even though the published solutions use all three.

What is really being tested is whether you convert a quadratic constraint into a linear one. The hypothesis $|z| = 1$ is a circle, and circles are quadratic conditions that are painful to carry through algebra. But the excluded point $z = 1$ sits on that circle, so the Mobius map that sends $z = 1$ to infinity sends the circle to a straight line. That map is $z \mapsto 1/(1 - z)$, and the line it produces is

$$\operatorname{Re} \frac{1}{1 - z} = \frac{1}{2},$$

which you should verify inside ten seconds. Every hypothesis in the problem is now a single real linear equation per variable, and a nonvanishing claim under linear constraints is something you can finish by separating real and imaginary parts.

Within thirty seconds you should have written: $u_j = 1 - z_j$; $|u|^2 = 2 \operatorname{Re} u$ hence $\operatorname{Re}(1/u) = 1/2$; and $v_j = 1/u_j = \frac{1}{2} + it_j$ with t_j real.

Approaches, ranked worst to best

1. **Write $z_j = e^{i\theta_j}$ and separate into two real trigonometric equations.** *Cue:* four unimodular numbers, so four angles. *Circuit:* parametrise and grind. You obtain two real equations in four unknown angles and no structure whatsoever. There are far too many degrees of freedom to argue by inspection and no natural reduction. This is zombie maths.

2. **Constrained optimisation with Lagrange multipliers.** *Cue:* a nonvanishing claim on a compact set is a claim that a minimum is positive. *Circuit:* minimise $|3 - z_1 - z_2 - z_3|^2 - |1 - z_1 z_2 z_3|^2$ and locate critical points. This is the published first solution. It works: the gradient identity factorises as a product of three sines, forcing $z_i z_j = 1$ for some pair, after which the problem collapses. But it requires compactness, differentiability, a three variable gradient computation and a product to sum factorisation that has to be got exactly right. It costs a great deal of ink for a statement that is ultimately algebraic.

3. **Hypocycloid geometry.** *Cue:* sums of unimodular numbers with prescribed product fill the region bounded by a hypocycloid. *Circuit:* identify S_4, the set of $w_1 + \cdots + w_4$ with $|w_j| = 1$

and $\prod w_j = 1$, as the filled hypocycloid H_4. This is the published second solution and it is the most illuminating one conceptually. It requires an unsmoothing lemma proved by a singular matrix argument, a description of H_n as a rolling circle curve, and a boundary point analysis. Wonderful to know; far too heavy to be the route you take in the arena.

4. **Shift to $u_j = 1 - z_j$ and expand.** *Cue:* the expression vanishes when all $z_j = 1$. *Circuit:* move the forbidden point to the origin. This gets you the clean identity $3 - e_1(z) + e_4(z) = e_2(u) - e_3(u) + e_4(u)$, which already removes the constant term and the linear term. It is the right first move but it does not finish, because you still have to use $|z_j| = 1$, which in the u coordinate is the circle $|u - 1| = 1$.

5. **Shift, then invert: $v_j = 1/(1 - z_j)$.** *Cue:* the circle $|z| = 1$ passes through the excluded point $z = 1$. *Circuit:* a Mobius map sending a point of a circle to infinity turns that circle into a line. The circle condition becomes the line condition $\operatorname{Re} v_j = \frac{1}{2}$; dividing the shifted identity by $e_4(u)$ converts it into $e_2(v) - e_1(v) + 1$ by the reversal identity for elementary symmetric functions; writing $v_j = \frac{1}{2} + it_j$ turns that into $(\frac{1}{2} - \sigma) + i\frac{T}{2}$ with T and σ the first two elementary symmetric functions of the real numbers t_j; and vanishing would force $T = 0$ and $\sigma = \frac{1}{2}$, which contradicts $T = 0 \Rightarrow \sigma \leq 0$. Half a page, no calculus, no geometry, and it generalises. This is the ideal route.

The ideal solution

Lemma 6.11.1. *If $|z| = 1$ and $z \neq 1$, then $u = 1 - z$ satisfies $|u|^2 = 2\operatorname{Re}(u)$ and $\operatorname{Re}(u) > 0$, and consequently*

$$\operatorname{Re}\left(\frac{1}{1-z}\right) = \frac{1}{2}.$$

Proof. Since $z\bar{z} = 1$,

$$|u|^2 = (1-z)(1-\bar{z}) = 1 - z - \bar{z} + z\bar{z} = 2 - (z + \bar{z}) = 2\big(1 - \operatorname{Re} z\big) = 2\operatorname{Re}(u).$$

As $z \neq 1$ we have $u \neq 0$, so $\operatorname{Re}(u) = |u|^2/2 > 0$. Then

$$\operatorname{Re}\left(\frac{1}{u}\right) = \operatorname{Re}\left(\frac{\bar{u}}{|u|^2}\right) = \frac{\operatorname{Re}(u)}{|u|^2} = \frac{\operatorname{Re}(u)}{2\operatorname{Re}(u)} = \frac{1}{2}. \qquad \square$$

Now write $u_j = 1 - z_j$, so $u_j \neq 0$ for each j, and let

$$e_1, e_2, e_3, e_4 \quad \text{denote the elementary symmetric polynomials of } u_1, u_2, u_3, u_4.$$

Step 1: rewrite the expression in the shifted coordinate. We have $\sum_j z_j = \sum_j (1 - u_j) = 4 - e_1$ and

$$\prod_j z_j = \prod_j (1 - u_j) = 1 - e_1 + e_2 - e_3 + e_4.$$

Therefore, writing E for the expression in question,

$$E = 3 - \sum_j z_j + \prod_j z_j = 3 - (4 - e_1) + (1 - e_1 + e_2 - e_3 + e_4) = e_2 - e_3 + e_4.$$

Step 2: invert. Since every $u_j \neq 0$ we have $e_4 = u_1u_2u_3u_4 \neq 0$, so $E = 0$ if and only if $E/e_4 = 0$. Put $v_j = 1/u_j$ and let ε_k denote the elementary symmetric polynomials of v_1, v_2, v_3, v_4. Dividing a product of k of the u_j by $u_1u_2u_3u_4$ leaves the reciprocals of the complementary $4 - k$ of them, so

$$\frac{e_2}{e_4} = \sum_{i<j} \frac{1}{u_iu_j} = \varepsilon_2, \qquad \frac{e_3}{e_4} = \sum_j \frac{1}{u_j} = \varepsilon_1,$$

(in the first identity the complement of a pair is a pair, and every pair occurs exactly once). Hence

$$\frac{E}{e_4} = \varepsilon_2 - \varepsilon_1 + 1.$$

Step 3: use the line. By the Lemma, $\mathrm{Re}(v_j) = \frac{1}{2}$, so we may write

$$v_j = \tfrac{1}{2} + it_j \qquad (t_j \in \mathbb{R}).$$

Set $T = t_1 + t_2 + t_3 + t_4$ and $\sigma = \sum_{i<j} t_i t_j$. Then

$$\varepsilon_1 = 2 + iT,$$

and, since each of the six products contributes $\frac{1}{4} + \frac{i}{2}(t_i + t_j) - t_i t_j$ and each t_i occurs in exactly three pairs,

$$\varepsilon_2 = 6 \cdot \tfrac{1}{4} + \tfrac{i}{2}\,(3T) - \sigma = \tfrac{3}{2} + \tfrac{3i}{2}T - \sigma.$$

Therefore

$$\frac{E}{e_4} = \varepsilon_2 - \varepsilon_1 + 1 = \left(\tfrac{3}{2} - \sigma - 2 + 1\right) + i\left(\tfrac{3}{2}T - T\right) = \left(\tfrac{1}{2} - \sigma\right) + i\,\frac{T}{2}.$$

Step 4: conclude by a dichotomy. Suppose $E = 0$. Then both the real and the imaginary part above vanish, so $T = 0$ and $\sigma = \frac{1}{2}$. But for real numbers,

$$T^2 = \left(\sum_j t_j\right)^2 = \sum_j t_j^2 + 2\sigma,$$

so $T = 0$ forces $\sigma = -\frac{1}{2}\sum_j t_j^2 \le 0$, contradicting $\sigma = \frac{1}{2}$. Hence $E \neq 0$, that is,

$$\boxed{3 - z_1 - z_2 - z_3 - z_4 + z_1 z_2 z_3 z_4 \neq 0}$$

whenever $|z_j| = 1$ and $z_j \neq 1$ for $j = 1, 2, 3, 4$. ■

Sharpness. If $z_1 = z_2 = z_3 = 1$ then $E = 3 - 3 - z_4 + z_4 = 0$, so the hypothesis cannot be dropped entirely. If exactly two of them equal 1, say $z_1 = z_2 = 1$, then $E = 1 - z_3 - z_4 + z_3 z_4 = (1 - z_3)(1 - z_4) \neq 0$. If exactly one equals 1, the expression becomes $2 - z_2 - z_3 - z_4 + z_2 z_3 z_4$, which the same method handles in one line: shifting and inverting gives $E/e_3 = \varepsilon_1 - 1 = \frac{1}{2} + iT$, whose real part is $\frac{1}{2} \neq 0$. So the sharp hypothesis is that at most two of the z_j are equal to 1.

Layer 1: WHAT

Block 1. Evaluate the expression at $z_1 = \cdots = z_4 = 1$ and get zero; conclude that the forbidden point is a root and that $u_j = 1 - z_j$ is the right coordinate.

Block 2. Expand in u: the constant and linear terms cancel and $E = e_2 - e_3 + e_4$.

Block 3. Note $e_4 \neq 0$, divide by it, and use the reversal identity to get $E/e_4 = \varepsilon_2 - \varepsilon_1 + 1$ in terms of $v_j = 1/u_j$.

Block 4. Identify the constraint on v_j: the circle $|z_j| = 1$ punctured at $z = 1$ maps onto the line $\mathrm{Re}\, v = \frac{1}{2}$.

Block 5. Write $v_j = \frac{1}{2} + it_j$ and compute the real and imaginary parts of $\varepsilon_2 - \varepsilon_1 + 1$.

Block 6. Kill both parts at once: the imaginary part vanishes only if $T = 0$, and then the real part is $\frac{1}{2} + \frac{1}{2}\sum t_j^2 \ge \frac{1}{2}$.

The nominal route runs: "$|z_j| = 1$, so put $z_j = e^{i\theta_j}$ and split into real and imaginary parts." It produces two trigonometric equations in four angles and no way forward, and the honest continuation is Lagrange multipliers on a compact set, which is exactly the published first solution and about four times the length.

The bifurcation point is the treatment of the constraint. The zombie keeps $|z| = 1$ as a circle and parametrises by an angle. The ideal solver notices that the excluded point lies on the circle, applies the Mobius map that sends it to infinity, and thereby replaces four quadratic constraints by four real linear ones. Everything downstream is forced by that single decision.

Layer 2: HOW

Block 1. **Cue.** The constant is $3 = 4 - 1$, and the terms e_2 and e_3 in z are missing. **Why natural.** Test the symmetric point. It gives $3 - 4 + 1 = 0$, so the polynomial vanishes exactly where the hypothesis forbids you to be. A polynomial that vanishes at a forbidden point is telling you which coordinate system it was built in: the one centred at that point. So set $u_j = 1 - z_j$ and expect the expression to have no constant and no linear term in u. **Link to the next block.** In the u coordinate the expression should simplify, and it does.

Block 2. **Cue.** Both $\sum z_j$ and $\prod z_j$ expand cleanly under $z_j = 1 - u_j$. **Why natural.** $\prod(1 - u_j) = \sum_k (-1)^k e_k(u)$ is the generating identity for elementary symmetric polynomials, which you already know. The constant cancels against the 3 and the e_1 terms cancel against each other, leaving $e_2 - e_3 + e_4$. This is the payoff for choosing the right origin. **Link to the next block.** You now have a symmetric expression with no low order terms, which means it is divisible by nothing but is dominated by the top term e_4, and e_4 is nonzero.

Block 3. **Cue.** $e_4 = \prod u_j \neq 0$, and the expression $e_2 - e_3 + e_4$ has its indices running from the top downwards when divided by e_4. **Why natural.** Dividing a symmetric expression by the top elementary symmetric polynomial is the standard way to pass to the reciprocals of the variables, because reversing the coefficients of a monic polynomial inverts its roots. Concretely, $e_k(u)/e_n(u) = e_{n-k}(1/u)$. Do it, because the constraint you have not yet used is a statement about $1/u_j$, not about u_j. **Link to the next block.** Now everything is written in the variables $v_j = 1/(1 - z_j)$, which is where the hypothesis becomes usable.

Block 4. **Cue.** The hypothesis is a circle through the point you removed. **Why natural.** A Mobius transformation maps circles and lines to circles and lines, and a circle through the pole of the map goes to a line. The map with pole at $z = 1$ that appears naturally here is $z \mapsto 1/(1 - z)$, and the image line is found by evaluating at one point, say $z = -1$, giving $\frac{1}{2}$. The clean algebraic proof is the identity $|1 - z|^2 = 2\,\mathrm{Re}(1 - z)$, which is nothing more than $z\bar{z} = 1$ rearranged. A real linear constraint is the best possible hypothesis, because it leaves one real parameter per variable. **Link to the next block.** Substituting the parametrisation of the line reduces everything to real polynomials in four real numbers.

Block 5. **Cue.** Each v_j has a fixed real part and a free imaginary part. **Why natural.** Write $v_j = \frac{1}{2} + it_j$ and expand. The elementary symmetric polynomials of the v_j become polynomials in the elementary symmetric polynomials of the t_j with explicit binomial coefficients. Only ε_1 and ε_2 are needed, so only T and σ appear. Keep the real and imaginary parts separate from the start; do not combine them. **Link to the next block.** You now have one complex equation to contradict, which is two real equations.

Block 6. **Cue.** The imaginary part is $T/2$ and the real part is $\frac{1}{2} - \sigma$. **Why natural.** Two real conditions, and they are incompatible for a reason you already have installed: $(\sum t_j)^2 = \sum t_j^2 + 2\sigma$. If $T = 0$ then σ is a negative multiple of a sum of squares, hence at most zero, while the real part would demand $\sigma = \frac{1}{2}$. Do not try to prove the real part is positive in general (it is not, since σ can exceed $\frac{1}{2}$ when $T \neq 0$); prove the dichotomy instead. **Link to the next block.** The contradiction closes the proof.

Layer 3: WHY

The mathematical truth beneath this route is that a Mobius transformation converts one quadratic constraint into one linear constraint whenever the excluded point is on the constraint curve. The

unit circle is $\{|z| = 1\}$; the map $w = 1/(1-z)$ has its pole at $z = 1$, a point of the circle, so it sends the circle to a line and, more precisely, it sends the punctured circle bijectively onto the whole line $\operatorname{Re} w = \frac{1}{2}$. The inverse is $z = (w-1)/w$, and $|z| = 1$ says $|w-1| = |w|$, that is, w is equidistant from 0 and 1, which is exactly the perpendicular bisector $\operatorname{Re} w = \frac{1}{2}$. That last sentence is the geometric reason, and it is a sentence about a perpendicular bisector, not about complex analysis.

The second truth is the reversal identity for symmetric functions: if $p(x) = \prod_j (x - u_j) = \sum_k (-1)^k e_k x^{n-k}$, then $x^n p(1/x)/p(0)$ is the monic polynomial with roots $1/u_j$, and comparing coefficients gives $e_k(1/u) = e_{n-k}(u)/e_n(u)$. That is why dividing by e_4 converts an expression built from e_2, e_3, e_4 into one built from $\varepsilon_2, \varepsilon_1, \varepsilon_0$. Dividing by the top symmetric function *is* inverting the variables.

There is a third truth worth recording because it generalises. Because $\operatorname{Re} v_j = \frac{1}{2}$ we have $\bar{v}_j = 1 - v_j$, hence

$$\sum_{m=0}^{4} (-1)^m \varepsilon_m = \prod_j (1 - v_j) = \prod_j \bar{v}_j = \overline{\varepsilon_4}.$$

Using this, $E/e_4 = \varepsilon_2 - \varepsilon_1 + \varepsilon_0 = \overline{\varepsilon_4} + \varepsilon_3 - \varepsilon_4$, a compact reformulation of the problem. For a general number n of variables the same computation reduces the statement "$(n-1) - \sum z_j + \prod z_j \neq 0$" to "$\sum_{m=0}^{n-2} (-1)^m \varepsilon_m(v) \neq 0$ for points v_j on the line $\operatorname{Re} v = \frac{1}{2}$". For $n = 3$ and $n = 4$ this is settled above. For $n \geq 5$ I have not settled it, and a numerical search over the five variable case found no zero but is not a proof; the published remark connects the general question to the geometry of hypocycloids. Name that as a knowledge gap rather than pretending the argument extends.

The discernment mechanism is a three question filter. First: does the expression vanish at the excluded configuration? If yes, translate that configuration to the origin. Second: is the constraint a circle through the excluded point? If yes, invert, because circles through the pole become lines. Third: once the constraints are affine, can you finish by separating real and imaginary parts? If the answer to the first two is yes, the answer to the third is almost always yes, because affine constraints leave one real degree of freedom per variable and symmetric expressions in such variables have real and imaginary parts that are low degree real polynomials.

To train it, practise the identity $|1 - z|^2 = 2\operatorname{Re}(1 - z)$ for unimodular z until it is reflexive, and practise recognising which Mobius map to apply from where the excluded point sits. The excluded point always becomes the pole. That is the entire choice.

The examiners nudged you clearly. They wrote the constant as 3, one less than the number of variables, which forces the vanishing at the all ones point. They excluded exactly $z_j = 1$ and nothing else, which is the announcement that $z = 1$ is the special point of the circle. And they suppressed e_2 and e_3 from the expression, so that after the shift only the top three symmetric functions survive and the division by e_4 lands on $\varepsilon_2 - \varepsilon_1 + 1$, an expression short enough to expand by hand.

Revolving concepts

Unimodular identity. For $|z| = 1$, $|1 - z|^2 = 2 - (z + \bar{z}) = 2\operatorname{Re}(1 - z)$. Equivalently, $1 - z = 2\sin(\theta/2)\, e^{i(\theta - \pi)/2}$ for $z = e^{i\theta}$, so $1 - z$ traces the circle $|u - 1| = 1$ and $\operatorname{Re}(1 - z) = 1 - \cos\theta \geq 0$. Pointer: see Book 1, field: Complex numbers, topic: chords of the unit circle.

Mobius transformations and the circle to line principle. A Mobius map sends the family of circles and lines to itself; a circle is sent to a line exactly when it passes through the pole. The map $z \mapsto 1/(1-z)$ sends the unit circle minus $\{1\}$ bijectively onto $\{\operatorname{Re} w = \frac{1}{2}\}$, with inverse $w \mapsto (w-1)/w$; the image line is the perpendicular bisector of 0 and 1 because $|z| = 1$ reads $|w - 1| = |w|$. Pointer: see Book 1, field: Complex numbers, topic: Mobius transformations and the Cayley transform.

Elementary symmetric polynomials and reversal. $\prod_j (1 + x u_j) = \sum_k e_k(u) x^k$, and $\prod_j (1 - u_j) = \sum_k (-1)^k e_k(u)$. The recursion in the number of variables is $e_k^{(n)} = e_k^{(n-1)} + u_n e_{k-1}^{(n-1)}$ with $e_0 = 1$ and $e_k^{(n)} = 0$ for $k > n$; its general solution is the generating product above, and the same recursion

with u_n replaced by any weight is the standard method for every symmetric function computation. Reversal: $e_k(1/u_1, \dots, 1/u_n) = e_{n-k}(u)/e_n(u)$, which is the coefficient statement that reversing a polynomial inverts its roots. Pointer: see Book 1, field: Polynomials, topic: symmetric functions and Vieta.

Newton's identity in degree two. For real $t_1, \dots, t_n$, $\left(\sum t_j\right)^2 = \sum t_j^2 + 2e_2(t)$; hence $e_1(t) = 0$ implies $e_2(t) \le 0$ with equality only when all $t_j = 0$. Pointer: see Book 1, field: Inequalities, topic: power sums and symmetric means.

Conjugation on the critical line. If $\operatorname{Re} v = \frac{1}{2}$ then $\bar{v} = 1 - v$, so for any finite family on that line $\prod_j (1 - v_j) = \overline{\prod_j v_j}$, that is, $\sum_m (-1)^m \varepsilon_m = \overline{\varepsilon_n}$. Pointer: see Book 1, field: Complex numbers, topic: reflection identities on a vertical line.

Hypocycloid description of sums with fixed product. $S_n = \{w_1 + \cdots + w_n : |w_j| = 1,\ \prod_j w_j = 1\}$ is, for $n \le 4$, the closed interior of the hypocycloid $H_n = \{(n-1)w + w^{-(n-1)} : |w| = 1\}$; the boundary points are attained only by tuples taking at most two distinct values. This is the geometry underlying the published second solution and the natural home of the general n question. Pointer: see Book 1, field: Complex numbers, topic: hypocycloids and sums of roots of unity.

Circuit Bank

- **WHEN** an expression that vanishes exactly at the configuration the hypothesis forbids **DO** translate that configuration to the origin by substituting the differences as new variables **BECAUSE** the vanishing says the expression has no constant term in those coordinates, so the low order symmetric functions disappear and only the top ones survive
- **WHEN** a constraint $|z| = 1$ together with one excluded point on the circle **DO** apply the Mobius map whose pole is that excluded point **BECAUSE** a Mobius map sends circles through its pole to straight lines, so a quadratic constraint becomes a real linear one and each variable is left with a single real parameter
- **WHEN** a symmetric expression in $u_1, \dots, u_n$ that you wish to write in the reciprocals **DO** divide by $e_n(u)$ and read off $e_k(u)/e_n(u) = e_{n-k}(1/u)$ **BECAUSE** reversing the coefficients of a polynomial inverts its roots, so dividing by the top symmetric function is exactly the change of variable $u \mapsto 1/u$
- **WHEN** a complex quantity to be proved nonzero when each variable has a pinned real part **DO** separate real and imaginary parts and argue by dichotomy rather than bounding one part alone **BECAUSE** the two vanishing conditions are usually incompatible for a cheap reason, whereas neither one alone need be impossible
- **WHEN** the conditions $e_1(t) = 0$ and $e_2(t) = c > 0$ for real numbers t_j **DO** invoke $(\sum t_j)^2 = \sum t_j^2 + 2e_2(t)$ and stop **BECAUSE** a vanishing first symmetric function forces the second to be a nonpositive multiple of a sum of squares

Backtesting

1. First useful cue: the constant 3 equals the number of variables minus one, so the expression vanishes at the all ones point, which is exactly the excluded point.
2. Method the expression suggested: shift by $u_j = 1 - z_j$, invert by $v_j = 1/u_j$ so that the circle becomes the line $\operatorname{Re} v = \frac{1}{2}$, then separate real and imaginary parts.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at the inversion, which feels unmotivated if you only see it as algebra. Reinforce it geometrically: $|z| = 1$ becomes $|w - 1| = |w|$, the perpendicular bisector of 0 and 1. Once you can see the line, the substitution stops feeling like a manoeuvre and starts feeling like a change of chart.

4. The trap the question sets: the phrase $|z_j| = 1$ invites the parametrisation $z_j = e^{i\theta_j}$, which is the one coordinate system in which the constraint is free but every expression becomes trigonometric. The four angles then have to be handled by optimisation, which is the long published route.
5. What makes you recognise it faster next time: the pair "constraint is a circle" plus "one point of the circle is excluded". That pair means invert at the excluded point, always.
6. The reusable decision rule: choose coordinates in which the hypothesis is affine. A circle with a marked point becomes a line under inversion at that point, and affine hypotheses turn nonvanishing claims into real polynomial dichotomies.

Mirror drills

1. Prove that $2 - z_1 - z_2 - z_3 + z_1z_2z_3 \neq 0$ for $|z_j| = 1$, $z_j \neq 1$. *Answer:* the same shift gives $e_2(u) - e_3(u)$, and dividing by $e_3(u)$ gives $\varepsilon_1 - 1 = \frac{1}{2} + iT$, whose real part is $\frac{1}{2}$.
2. Show that $z \mapsto 1/(1-z)$ maps $\{|z| = 1\} \setminus \{1\}$ bijectively onto $\{\operatorname{Re} w = \frac{1}{2}\}$. *Answer:* the inverse is $z = (w-1)/w$, and $|z| = 1$ says $|w-1| = |w|$, the perpendicular bisector of 0 and 1.
3. Determine the sharp hypothesis: for how many of the z_j may we allow the value 1? *Answer:* at most two. Three forces the expression to vanish; two gives $(1 - z_3)(1 - z_4) \neq 0$; one reduces to the three variable statement of drill 1.

6.12 B6. Run lengths renormalise a Beatty sign sequence

Let n be a positive integer. Prove that

$$\sum_{k=1}^{n} (-1)^{\lfloor k(\sqrt{2}-1) \rfloor} \geq 0.$$

(As usual, $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x.)

Read it out loud

This is a partial sum positivity statement for a sign sequence generated by the parity of a Beatty sequence. The signs are $(-1)^{\lfloor k\alpha \rfloor}$ with $\alpha = \sqrt{2} - 1 \approx 0.4142$, and you are asked to show that the running total never dips below zero. The fields wired in are: Diophantine approximation of quadratic irrationals and their continued fractions; Beatty sequences and complementary sequences; self-similar (substitutive) sequences and their run length codings; and strong induction with a change of scale.

Write out the first few signs. Since $\lfloor k\alpha \rfloor$ for $k = 0, 1, 2, \ldots$ runs $0, 0, 0, 1, 1, 2, 2, 2, 3, 3, 4, 4, 4, \ldots$, the sign sequence for $k \geq 0$ is

$$+, +, +, -, -, +, +, +, -, -, +, +, +, \ldots$$

so it is a sequence of *runs*, alternating in sign, of lengths $3, 2, 3, 2, 3, \ldots$ Every run has length 2 or 3, because α lies strictly between $\frac{1}{3}$ and $\frac{1}{2}$. The partial sum can only be small at the end of a negative run, so the whole problem is the question: are the alternating partial sums of the run lengths always positive?

Now the crucial observation, and it is the only genuinely nonobvious idea in the problem. The irrational $\alpha = \sqrt{2} - 1$ was not chosen for flavour. Its reciprocal is

$$\beta = \frac{1}{\alpha} = \sqrt{2} + 1 = \alpha + 2,$$

so β and α differ by an *even* integer. Consequently, for every integer $j \geq 0$,

$$\lfloor j\beta \rfloor = \lfloor j\alpha \rfloor + 2j, \qquad \text{hence} \qquad (-1)^{\lfloor j\beta \rfloor} = (-1)^{\lfloor j\alpha \rfloor}.$$

That single line is the engine. The run lengths of the sign sequence are governed by the Beatty sequence of β, while the signs themselves are governed by α; because their parities agree, the run length data reproduces the original sign sequence at a compressed scale. The problem is therefore self-similar, and self-similar statements are proved by strong induction on the scale.

Within thirty seconds you should have written: $\alpha = \sqrt{2} - 1$, $1/\alpha = \alpha + 2$, "the reciprocal differs by an even integer"; the run lengths $3, 2, 3, 2, \ldots$; and the word "renormalise".

Approaches, ranked worst to best

1. **Compute a large partial sum and assert a pattern.** *Cue:* the statement is about all n, but the sequence looks regular. *Circuit:* numerical reconnaissance. You will confirm the inequality for millions of terms and prove nothing. Worse, the partial sums are unbounded, so no finite check tells you anything about the tail.

2. **Equidistribution of $\{k\alpha\}$.** *Cue:* α irrational, so the fractional parts equidistribute by Weyl. *Circuit:* averaging. Equidistribution says the density of each sign is $\frac{1}{2}$ and the partial sum is $o(n)$. That is a statement about the size of the partial sum, not its sign, and it is useless here: the partial sums genuinely reach arbitrarily large positive values, so no averaging argument can control the sign. Ranks low because it feels like progress and is not.

3. **Telescoping or a monotone quantity.** *Cue:* partial sum positivity often follows from a clever monotone functional. *Circuit:* invariants. There is no monotone functional: the partial sum decreases on every negative run. Any attempt to bound the sum below by a local quantity fails because the minima occur at unboundedly spaced places.

4. **Ostrowski representation via the continued fraction convergents.** *Cue:* α is a quadratic irrational, so its convergents obey a linear recursion. *Circuit:* expand k in the numeration system attached to α. With $q_0 = 0$, $q_1 = 1$, $q_j = 2q_{j-1} + q_{j-2}$ (the Pell numbers $0, 1, 2, 5, 12, 29, 70, \ldots$), one proves $f(n) = f(q_j) - f(n - q_j)$ for even j and $f(n) = f(n - q_j) + 1$ for odd j, then inducts. This is the published second solution and it is correct, but it requires the best approximation property of convergents, the parity fact $q_j \equiv j \pmod 2$, and a careful treatment of the exceptional indices $k \in \{q_j, 2q_j\}$ where $\lfloor k\alpha \rfloor$ and $\lfloor kq_{j-1}/q_j \rfloor$ disagree. Every one of those is a place to be wrong under time pressure.

5. **Run length renormalisation with an explicit self-similarity identity.** *Cue:* the signs form runs of lengths 2 and 3, and $1/\alpha - \alpha$ is an even integer. *Circuit:* code the sequence by its run lengths and observe that the coding reproduces the original sequence. Count how many k satisfy $\lfloor k\alpha \rfloor = i$, get $c_i \in \{2, 3\}$ with $c_i = 3$ exactly on the Beatty sequence $\lfloor j\beta \rfloor$, subtract 2 from every c_i and use the parity transfer to see that the alternating sum of the c_i is literally a partial sum of the original sequence, at an index compressed by the factor $\alpha < \frac{1}{2}$. Then strong induction on n, with the observation that the running total is minimised at the end of a negative run. No continued fraction theory beyond $\beta = \alpha + 2$, every constant explicit. This is the ideal route.

The ideal solution

Throughout, $\alpha = \sqrt{2} - 1$ and $\beta = \sqrt{2} + 1$, so that

$$\alpha\beta = 1, \qquad \beta = \alpha + 2, \qquad 0 < \alpha < \tfrac{1}{2} < 1 < \beta.$$

Set $a_k = \lfloor k\alpha \rfloor$ for $k \geq 0$ and

$$F(n) = \sum_{k=0}^{n} (-1)^{a_k} \qquad (n \geq 0).$$

Since $a_0 = 0$, the assertion to be proved is exactly $F(n) \geq 1$ for all $n \geq 0$.

Lemma 6.12.1 (parity transfer)**.** *For every integer $j \geq 0$, $\lfloor j\beta \rfloor = 2j + \lfloor j\alpha \rfloor$, and hence $(-1)^{\lfloor j\beta \rfloor} = (-1)^{a_j}$.*

Proof. $j\beta = 2j + j\alpha$ and $2j$ is an integer, so the floors differ by $2j$, which is even. □

Lemma 6.12.2 (run structure)**.** *For $i \geq 0$ let $c_i = \#\{k \geq 0 : a_k = i\}$ and $C_m = c_0 + c_1 + \cdots + c_m$. Then every c_i is finite and positive,*

$$C_m = \lfloor (m+1)\beta \rfloor + 1 = 2m + 3 + \lfloor (m+1)\alpha \rfloor \qquad (m \geq 0),$$

and $c_0 = 3$, while $c_i = 2 + \lfloor (i+1)\alpha \rfloor - \lfloor i\alpha \rfloor \in \{2, 3\}$ for $i \geq 1$.

Proof. Since $0 < \alpha < 1$, the sequence a_k is nondecreasing with $a_{k+1} - a_k \in \{0, 1\}$, starts at $a_0 = 0$ and tends to infinity, so it attains every nonnegative integer on a nonempty block of consecutive k. Next,

$$C_m = \#\{k \geq 0 : a_k \leq m\} = \#\{k \geq 0 : k\alpha < m + 1\} = \#\{k \geq 0 : k < (m+1)\beta\},$$

using $\alpha\beta = 1$. As β is irrational, $(m+1)\beta$ is not an integer, so this count is $\lfloor (m+1)\beta \rfloor + 1$. Substituting $\beta = \alpha + 2$ gives $C_m = 2(m+1) + \lfloor (m+1)\alpha \rfloor + 1$. Finally $c_0 = C_0 = 2 + \lfloor \alpha \rfloor + 1 = 3$ and, for $i \geq 1$, $c_i = C_i - C_{i-1} = 2 + \lfloor (i+1)\alpha \rfloor - \lfloor i\alpha \rfloor$, which lies in $\{2, 3\}$ because consecutive floors of $i\alpha$ differ by 0 or 1. □

Lemma 6.12.3 (long runs sit on a Beatty sequence)**.** *For every $i \geq 0$: $c_i = 3$ if and only if $i = \lfloor j\beta \rfloor$ for some integer $j \geq 0$; otherwise $c_i = 2$.*

Proof. For $i = 0$: $c_0 = 3$ and $0 = \lfloor 0 \cdot \beta \rfloor$, so the statement holds; and no $j \geq 1$ gives $\lfloor j\beta \rfloor = 0$, since $\lfloor \beta \rfloor = 2$ and $j \mapsto \lfloor j\beta \rfloor$ is increasing.

For $i \geq 1$, Lemma 2 gives $c_i = 3$ if and only if $\lfloor (i+1)\alpha \rfloor > \lfloor i\alpha \rfloor$, that is, if and only if there is an integer j with $i\alpha < j < (i+1)\alpha$; the inequalities are strict because $i\alpha$ and $(i+1)\alpha$ are irrational. Multiplying through by $\beta > 0$ and using $\alpha\beta = 1$, this says $i < j\beta < i+1$, that is, $i = \lfloor j\beta \rfloor$. Such a j satisfies $j > i\alpha > 0$, so $j \geq 1$. □

Lemma 6.12.4 (self-similarity)**.** *For $m \geq 0$ put $S_m = \sum_{i=0}^{m} (-1)^i c_i$ and $J_m = \lfloor (m+1)\alpha \rfloor$. Then*

$$S_m = F(J_m) + 2 \cdot [\, m \text{ even} \,].$$

Proof. Since $\sum_{i=0}^{m} (-1)^i$ equals 1 for even m and 0 for odd m,

$$S_m = \sum_{i=0}^{m} (-1)^i (c_i - 2) + 2 \cdot [\, m \text{ even} \,].$$

By Lemma 3, $c_i - 2$ is the indicator of the event $i \in \{\lfloor j\beta \rfloor : j \geq 0\}$, and $j \mapsto \lfloor j\beta \rfloor$ is strictly increasing, hence injective. Therefore

$$\sum_{i=0}^{m} (-1)^i (c_i - 2) = \sum_{\substack{j \geq 0 \\ \lfloor j\beta \rfloor \leq m}} (-1)^{\lfloor j\beta \rfloor} = \sum_{\substack{j \geq 0 \\ \lfloor j\beta \rfloor \leq m}} (-1)^{a_j},$$

the second equality by Lemma 1. Finally $\lfloor j\beta \rfloor \leq m$ if and only if $j\beta < m+1$, if and only if $j < (m+1)\alpha$; since $(m+1)\alpha$ is irrational, the admissible j are exactly $0, 1, \ldots, \lfloor (m+1)\alpha \rfloor = J_m$. So the displayed sum equals $F(J_m)$. □

Theorem 6.12.5. *$F(n) \geq 1$ for every integer $n \geq 0$; equivalently $\sum_{k=1}^{n} (-1)^{\lfloor k(\sqrt{2}-1) \rfloor} \geq 0$ for every $n \geq 1$.*

Proof. Strong induction on n. For $n = 0$, $F(0) = 1$. Let $n \geq 1$ and assume $F(n') \geq 1$ for all $0 \leq n' < n$.

Set $C_{-1} = 0$ and let $m \geq 0$ be the unique index with $C_{m-1} \leq n < C_m$, so that $k = n$ lies in the m-th run, on which every term equals $(-1)^m$. By Lemma 2, for $m \geq 1$,

$$C_{m-1} = 2(m-1) + 3 + \lfloor m\alpha \rfloor = 2m + 1 + \lfloor m\alpha \rfloor.$$

Also $F(C_m - 1) = S_m$ for every $m \geq 0$, since the first C_m terms of the sequence are exactly the runs $0, 1, \ldots, m$.

Case $m = 0$. Then $n < C_0 = 3$, and $a_0 = a_1 = a_2 = 0$, so $F(n) = n + 1 \geq 1$.

Case m odd. On the m-th run every term is -1, so F is decreasing there and

$$F(n) \geq F(C_m - 1) = S_m = F(J_m),$$

the last equality by Lemma 4 since m is odd. Now $J_m = \lfloor (m+1)\alpha \rfloor \leq (m+1)\alpha < \frac{m+1}{2} \leq m$ because $\alpha < \frac{1}{2}$ and $m \geq 1$, whereas $n \geq C_{m-1} \geq 2m + 1 > m$. Hence $J_m < n$ and the induction hypothesis gives $F(J_m) \geq 1$, so $F(n) \geq 1$.

Case m even, $m \geq 2$. On the m-th run every term is $+1$, so

$$F(n) = F(C_{m-1} - 1) + (n - C_{m-1} + 1) = S_{m-1} + (n - C_{m-1} + 1) \geq S_{m-1} + 1.$$

Since $m - 1$ is odd, Lemma 4 gives $S_{m-1} = F(J_{m-1})$ with $J_{m-1} = \lfloor m\alpha \rfloor$. And $n \geq C_{m-1} = 2m + 1 + \lfloor m\alpha \rfloor > \lfloor m\alpha \rfloor = J_{m-1}$, so the induction hypothesis gives $S_{m-1} \geq 1$ and therefore $F(n) \geq 2$.

In every case $F(n) \geq 1$, completing the induction. Subtracting the $k = 0$ term gives $\sum_{k=1}^{n} (-1)^{a_k} = F(n) - 1 \geq 0$. □

Layer 1: WHAT

Block 1. Include the term $k = 0$, which is $+1$, and restate the target as $F(n) \geq 1$. Working with a strictly positive quantity is easier to induct on than a nonnegative one.

Block 2. Record the parity transfer $(-1)^{\lfloor j\beta \rfloor} = (-1)^{\lfloor j\alpha \rfloor}$, which holds because $\beta - \alpha = 2$ is even.

Block 3. Decompose the sign sequence into alternating runs and count the run lengths: $c_i = \#\{k : \lfloor k\alpha \rfloor = i\}$, with the closed form $C_m = \lfloor (m+1)\beta \rfloor + 1$.

Block 4. Show $c_i \in \{2, 3\}$ and that $c_i = 3$ exactly when i lies in the Beatty sequence $\lfloor j\beta \rfloor$.

Block 5. Subtract 2 from every c_i and take the alternating sum. The parity transfer converts it into a partial sum of the original sequence, at index $J_m = \lfloor (m+1)\alpha \rfloor$, which is less than half of m.

Block 6. Observe that the running total is smallest at the end of a negative run, and run a strong induction: inside a negative run, bound $F(n)$ below by S_m; inside a positive run, bound it below by $S_{m-1} + 1$. Both bounds are values of F at strictly smaller indices.

The nominal route runs: "the fractional parts of $k\alpha$ equidistribute, so the signs are balanced on average; let me estimate the discrepancy". That gives $F(n) = o(n)$ and stops, because the true statement is about a sign, not a size, and the partial sums are unbounded.

The bifurcation point is Block 3. The zombie sums the sequence term by term and tries to control the running total directly. The ideal solver stops looking at terms and looks at *runs*, because a sequence whose values are ± 1 carries all of its information in its run lengths, and because the alternating sum of run lengths is precisely the quantity whose positivity is in question. Once you are summing runs rather than terms, the parity transfer does the rest.

Layer 2: HOW

Block 1. **Cue.** The sum starts at $k = 1$, but $\lfloor 0 \cdot \alpha \rfloor = 0$ contributes $+1$. **Why natural.** Adding the $k = 0$ term shifts the target from "≥ 0" to "≥ 1", and an induction that must prove strict positivity of an integer quantity is easier than one proving nonnegativity, because "positive integer" gives you a usable margin of one. Free moves that make the statement stronger and the bookkeeping cleaner should always be taken. **Link to the next block.** You now want the partial sums of the full sequence indexed from 0.

Block 2. **Cue.** The number $\sqrt{2} - 1$, not $\sqrt{2}$, not $\sqrt{3} - 1$. **Why natural.** Any competition problem naming a specific quadratic irrational is naming a specific continued fraction. Compute the reciprocal immediately: $1/(\sqrt{2} - 1) = \sqrt{2} + 1$, which exceeds α by exactly 2. An even difference means the floors of $j\alpha$ and $j\beta$ have equal parity. That is the only special property used anywhere in the proof, and it is the reason the problem is true. **Link to the next block.** The Beatty sequence of β is about to appear as the set of positions of long runs, and you will need its parities to match those of α.

Block 3. **Cue.** The sequence takes only two values, $+1$ and -1. **Why natural.** A two valued sequence is completely described by its run lengths. So change the object of study from the sequence to its run length sequence, exactly as you would replace a binary word by its coding. The run containing the value i is $\{k : i \leq k\alpha < i + 1\}$, and counting integers in an interval is the cheapest computation there is: C_m counts k with $k\alpha < m + 1$, that is, $k < (m+1)\beta$. **Link to the next block.** You now need to know which runs are long.

Block 4. **Cue.** c_i is 2 plus the increment of $\lfloor i\alpha \rfloor$, which is 0 or 1. **Why natural.** The increment of a floor is 1 exactly when an integer is crossed, that is, when some j lies in $(i\alpha, (i+1)\alpha)$. Dividing by α turns that into $j\beta \in (i, i+1)$, so the long runs sit precisely at the positions $\lfloor j\beta \rfloor$. This is a Beatty sequence, and it is the standard answer to "where are the exceptional indices?" in every problem

of this type. **Link to the next block.** The long run positions are indexed by j, and their signs are $(-1)^{\lfloor j\beta \rfloor}$, which by Block 2 equal $(-1)^{\lfloor j\alpha \rfloor}$.

Block 5. Cue. An alternating sum of numbers all equal to 2 or 3. **Why natural.** Subtract the common part. $\sum(-1)^i c_i = \sum(-1)^i(c_i - 2) + 2\sum(-1)^i$, and the second sum is 0 or 1 depending only on the parity of m. The first sum is an alternating sum over the long run positions, and by the parity transfer its terms are exactly the terms of the original sequence. So the run length data is the original sequence again, read at a scale compressed by α. Recognising the compressed copy is the moment the problem is solved. **Link to the next block.** Self-similarity plus a strict contraction of scale means induction.

Block 6. Cue. The partial sum decreases only during negative runs. **Why natural.** Therefore its minimum over any interval is attained at the end of a negative run, and those are exactly the values S_m with m odd, which Lemma 4 identifies with $F(J_m)$. You never need to control the sum inside a run; you only need to control it at the run boundaries. For a positive run you gain at least one over the previous boundary. The two cases are the only two cases. **Link to the next block.** The scale contracts by more than a factor of two each time, since $\alpha < \frac{1}{2}$, so the induction terminates at the base case $F(0) = 1$.

Layer 3: WHY

The mathematical truth beneath this route is that $(-1)^{\lfloor k\alpha \rfloor}$ is a Sturmian type sequence, and the run length coding of a Sturmian sequence is again Sturmian with a new parameter obtained by applying the Gauss map $T(x) = \{1/x\}$ to the old one. That is the general renormalisation theorem in this area: coding by runs on the sequence side corresponds to the continued fraction shift on the parameter side. It is why a run length argument ever closes.

For a generic irrational this renormalisation produces an infinite family of different parameters and no self-referential induction is possible. Here it does close, and for a specific reason: from $\beta = \alpha + 2$ we get $\alpha = 1/(2 + \alpha)$, so

$$\alpha = [0; 2, 2, 2, \dots],$$

a purely periodic continued fraction with all partial quotients equal to 2, and α is a fixed point of the Gauss map. The derived sequence is therefore the *same* sequence, not a relative of it. Two properties are needed and both hold: the partial quotients are all equal (so the renormalisation returns to the same parameter) and the difference $1/\alpha - \alpha = 2$ is even (so the parity survives the renormalisation). Drop either and the argument fails. For the golden section $\alpha = (\sqrt{5} - 1)/2$ the first property holds but $1/\alpha - \alpha = 1$ is odd, so the parity transfer picks up an extra factor $(-1)^j$ and the run length coding reproduces a *twisted* copy of the sequence rather than the sequence itself.

This also explains, without any extra work, the natural generalisation: for $\alpha_N = \sqrt{N^2 + 1} - N$ we have $1/\alpha_N = \sqrt{N^2 + 1} + N = \alpha_N + 2N$, an even difference, and $\alpha_N = [0; 2N, 2N, \dots]$, again a Gauss map fixed point. The identical proof gives $\sum_{k=1}^{n}(-1)^{\lfloor k\alpha_N \rfloor} \geq 0$ for every positive integer N, with run lengths in $\{2N, 2N + 1\}$ and the scale contracting by the factor $\alpha_N < 1/(2N)$.

The other structural fact worth naming is the complementary Beatty pair. With $\beta = 1 + \sqrt{2}$ and $\beta' = 1 + \frac{\sqrt{2}}{2}$ one has $1/\beta + 1/\beta' = (\sqrt{2} - 1) + (2 - \sqrt{2}) = 1$, so by Rayleigh's theorem the sequences $\lfloor j\beta \rfloor$ and $\lfloor j\beta' \rfloor$ partition the positive integers. In our language: the long runs sit at the β positions and the short runs at the β' positions. Knowing this in advance is what tells you that "which runs are long" has a clean answer at all.

The discernment mechanism is a short filter for any problem containing $\lfloor k\theta \rfloor$ with θ irrational. First, compute $1/\theta$ and its continued fraction. Second, ask what invariant of θ the question is testing: if the statement involves a parity, the relevant invariant is the parity of $1/\theta - \theta$ or of the partial quotients. Third, if the sequence takes finitely many values, code it by runs and see whether

the coding reproduces it. If the continued fraction is purely periodic with period one, the coding will reproduce the sequence exactly and an induction is available. Training this means doing the reciprocal computation reflexively and always writing out the first twenty terms, because the run structure is invisible in the formula and obvious in the list.

The examiners nudged you in three places. They chose $\sqrt{2}-1$ rather than $\sqrt{2}$ or $1/\sqrt{2}$, which is the choice that makes the reciprocal differ by an even integer. They asked for ≥ 0 rather than for an asymptotic, which signals that the mechanism is exact and structural rather than analytic. And they used a floor inside an exponent of -1, which is the standard packaging for "the parity of a Beatty sequence", a phrase you should be able to read off the notation instantly.

Revolving concepts

Beatty sequences and Rayleigh's theorem. For irrational $\beta > 1$, the Beatty sequence is $\lfloor j\beta \rfloor$, $j \geq 1$. If $\beta, \beta' > 1$ are irrational with $1/\beta + 1/\beta' = 1$, the two Beatty sequences partition the positive integers. Here $\beta = 1+\sqrt{2}$ and $\beta' = 1+\sqrt{2}/2$. Pointer: see Book 1, field: Number theory, topic: Beatty and complementary sequences.

Continued fractions of quadratic irrationals. Every quadratic irrational has an eventually periodic continued fraction (Lagrange), and it is purely periodic exactly when it is reduced. Here $\alpha = 1/(2+\alpha)$ gives $\alpha = [0;\overline{2}]$, and the convergent denominators satisfy the Pell recursion $q_j = 2q_{j-1} + q_{j-2}$ with $q_0 = 0$, $q_1 = 1$, giving $0, 1, 2, 5, 12, 29, 70, \ldots$ The general solution of that recursion is $q_j = (\beta^j - (-\alpha)^j)/(2\sqrt{2})$, and in particular $q_j \equiv j \pmod 2$. The same recursion with 2 replaced by $2N$ handles $\alpha_N = \sqrt{N^2+1} - N$. Pointer: see Book 1, field: Number theory, topic: continued fractions and Pell recursions.

The Gauss map and renormalisation. $T(x) = \{1/x\}$ acts as the shift on continued fraction expansions. Coding a Sturmian sequence by its run lengths corresponds to applying T to its parameter; a parameter fixed by T therefore has a run length coding equal to itself, which is exactly what makes a self-referential induction possible. Pointer: see Book 1, field: Number theory, topic: Gauss map and substitutive sequences.

Sturmian words and run length coding. A sequence $(\lfloor (k+1)\theta \rfloor - \lfloor k\theta \rfloor)$ with $0 < \theta < 1$ irrational takes two values and is Sturmian; equivalently, for $1/(L+1) < \theta < 1/L$ the level sets of $\lfloor k\theta \rfloor$ have sizes L or $L+1$, and the positions of the longer ones form a Beatty sequence. The recursion to install is: run lengths of the level sets of $\lfloor k\theta \rfloor$ equal $\lfloor 1/\theta \rfloor$ plus the increments of $\lfloor i\theta \rfloor$, and the increment indicator is again a Beatty indicator, so the construction can be iterated indefinitely. Pointer: see Book 1, field: Discrete structures and graphs, topic: Sturmian words.

Ostrowski representation. Every positive integer has a unique representation $n = \sum_j d_j q_j$ subject to digit constraints determined by the partial quotients; it is the numeration system in which $\lfloor n\theta \rfloor$ and quantities like $F(n)$ have clean recursive descriptions. This is the machinery behind the published second solution. Pointer: see Book 1, field: Number theory, topic: Ostrowski numeration.

Three distance theorem. For irrational θ and any N, the points $\{\theta\}, \{2\theta\}, \ldots, \{N\theta\}$ partition the circle into arcs of at most three distinct lengths. It is the geometric cousin of the run length statement and the reason such problems always have finitely many local patterns. Pointer: see Book 1, field: Number theory, topic: three distance theorem.

Circuit Bank

- **WHEN** an expression $(-1)^{\lfloor k\theta \rfloor}$ with θ irrational **DO** compute $1/\theta$ at once and check whether $1/\theta - \theta$ is an even integer **BECAUSE** $\lfloor j/\theta \rfloor$ and $\lfloor j\theta \rfloor$ differ by an integer multiple of j, so only an even difference preserves parity, and parity preservation is what lets the run length data reproduce the original sign sequence
- **WHEN** a sequence taking only two values whose partial sums you must control **DO** replace the sequence by its run length sequence and control the alternating sum of run lengths

BECAUSE a two valued sequence carries all of its information in its runs, and the partial sums are extremal exactly at run boundaries

- **WHEN** run lengths of the level sets of $\lfloor k\theta \rfloor$ **DO** write the cumulative count as $\lfloor (m+1)/\theta \rfloor + 1$ and take differences **BECAUSE** counting the k with $\lfloor k\theta \rfloor \leq m$ is counting integers below $(m+1)/\theta$, which is the cheapest possible computation and gives every run length at once
- **WHEN** run lengths lying in $\{L, L+1\}$ **DO** subtract L and read the remainder as the indicator of a Beatty sequence **BECAUSE** the difference of consecutive floors is exactly the indicator that an integer was crossed, and crossing points of $i\theta$ are the Beatty positions of $1/\theta$
- **WHEN** a statement about a self-similar sequence at all scales **DO** strong induction on the index, using that the renormalised index is smaller by a definite factor **BECAUSE** self-similarity gives an exact identity between a quantity at scale n and the same quantity at a strictly smaller scale, and a contraction factor bounded away from one guarantees the induction terminates
- **WHEN** a partial sum positivity claim **DO** prove positivity only at the local minima of the partial sum, and bound the rest by monotonicity **BECAUSE** the running total decreases only on negative runs, so its minimum over any range is at the end of one, and every other value is dominated by such a value

Backtesting

1. First useful cue: $\sqrt{2}-1$ is the reciprocal of $\sqrt{2}+1$ and differs from it by the even integer 2. Compute the reciprocal before doing anything else.
2. Method the expression suggested: code the two valued sign sequence by its run lengths, identify the long runs with the Beatty sequence of $1/\alpha$, and use the parity transfer to see the run length data as a compressed copy of the original sequence, then induct on the scale.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at the step where the alternating sum of $c_i - 2$ becomes a partial sum of the original sequence. Reinforce it by naming the two ingredients out loud: the long runs sit at $i = \lfloor j\beta \rfloor$, and $(-1)^{\lfloor j\beta \rfloor} = (-1)^{\lfloor j\alpha \rfloor}$. Without the second, the identity is about a different sequence and the induction has nothing to feed on.
4. The trap the question sets: the irrational invites you to reach for equidistribution and discrepancy bounds. Those are the right tools for questions about the *size* of $\sum(-1)^{\lfloor k\theta \rfloor}$ and completely useless for its sign, since the partial sums here are unbounded and the statement is exact.
5. What makes you recognise it faster next time: the triple "floor of k times a quadratic irrational", "inside a power of -1", "inequality valid for all n". That triple means Beatty parity plus renormalisation, and the very first computation is $1/\theta$.
6. The reusable decision rule: for a two valued sequence generated by $\lfloor k\theta \rfloor$, the run lengths are $\lfloor 1/\theta \rfloor$ or $\lfloor 1/\theta \rfloor + 1$, the long ones sit on the Beatty sequence of $1/\theta$, and if $1/\theta - \theta$ is an even integer the run length coding reproduces the original sign sequence at scale θ; then induct.

Mirror drills

1. Show that $(-1)^{\lfloor j(\sqrt{2}+1) \rfloor} = (-1)^{\lfloor j(\sqrt{2}-1) \rfloor}$ for every integer $j \geq 0$. *Answer:* $\sqrt{2}+1 = (\sqrt{2}-1)+2$, so the two floors differ by $2j$.
2. Prove that $\lfloor j(1+\sqrt{2}) \rfloor$ and $\lfloor j(1+\frac{\sqrt{2}}{2}) \rfloor$, for $j \geq 1$, partition the positive integers, and say what this means for the run lengths. *Answer:* $\frac{1}{1+\sqrt{2}} + \frac{1}{1+\sqrt{2}/2} = (\sqrt{2}-1) + (2-\sqrt{2}) = 1$, so

Rayleigh's theorem applies; the first sequence gives the positions of the runs of length 3 and the second the positions of the runs of length 2.

3. For which θ of the form $\sqrt{N^2+1}-N$ does the proof above go through unchanged, and why? *Answer:* all of them. Each satisfies $1/\theta = \theta + 2N$ with $2N$ even, and $\theta = [\,0;\overline{2N}\,]$ is a fixed point of the Gauss map, so the run lengths lie in $\{2N, 2N+1\}$ and the coding reproduces the same sequence.

Chapter 7

Putnam 2019

7.1 A1. Sum of cubes minus triple product, modulo nine

Determine all possible values of the expression

$$A^3 + B^3 + C^3 - 3ABC$$

where A, B, and C are nonnegative integers.

Read it out loud

This is the one symmetric cubic in three variables that every syllabus makes you factorise, evaluated on the nonnegative integer lattice, and the command term is "determine all possible values". Say it plainly: you are being asked for the image of a polynomial map $\mathbb{Z}^3_{\geq 0} \to \mathbb{Z}$. An image question is always two questions. First, an obstruction: every value has some property. Second, a construction: every integer with that property is hit. If you only do one of the two you have half the marks and you know it.

The fields wired in are polynomial identities (symmetric functions, the factorisation of $A^3 + B^3 + C^3 - 3ABC$) and elementary number theory (residues modulo 3 and 9). What is really being tested is whether you treat the expression as a structured object (a product of two symmetric factors, each with a known sign) or as three independent variables to be case-bashed.

The first useful cue is the expression itself: the circuit fires and you write $A^3 + B^3 + C^3 - 3ABC = (A + B + C)(A^2 + B^2 + C^2 - AB - BC - CA)$, and in the same breath the second factor as $\frac{1}{2}\big((A-B)^2 + (B-C)^2 + (C-A)^2\big)$. The second cue is "nonnegative": it is planted to make the first factor nonnegative, so the whole expression is nonnegative. Within thirty seconds you should also have a small table: $(1,0,0) \mapsto 1$, $(1,1,0) \mapsto 2$, $(2,1,1) \mapsto 4$, $(2,2,1) \mapsto 5$, $(3,2,2) \mapsto 7$, $(2,0,0) \mapsto 8$, $(2,1,0) \mapsto 9$, and 3 and 6 never appear. The conjecture "everything except 3 and 6 modulo 9" is on the page before any proof starts, and the rest of the problem is making the conjecture rigorous in both directions.

Approaches, ranked worst to best

1. **Tabulate and guess.** *Cue:* "determine all possible values". *Circuit:* small cases first. (Correct as a first move, worthless as a last one. The table gives the conjecture but says nothing about why 3, 6, 12, 15 are missing. It stalls exactly where the marks are.)

2. **Cubes modulo nine, case bash.** *Cue:* the cubes and the modulus 9 in the conjecture. *Circuit:* residues of cubes: $a^3 \equiv 0, \pm 1 \pmod 9$. (It works: if 3 divides the expression then $A + B + C \equiv 0 \pmod 3$, which leaves four residue patterns for (A, B, C) modulo 3, and each gives a multiple of 9. Four cases of consumption where one structural line would do; and it hides the reason.)

3. **Norm form on the circulant ring.** *Cue:* the expression is the determinant of the circulant matrix with rows $(A, B, C), (C, A, B), (B, C, A)$. *Circuit:* norm forms are multiplicative. (Conceptually the deepest: it explains why the value set is closed under multiplication and why 3 behaves like a ramified prime. It costs a whole ring's worth of theory for a ten-minute problem. Keep it for the WHY layer, not for the arena.)

4. **Factorise, then symmetric functions.** *Cue:* a symmetric polynomial asked modulo a prime power. *Circuit:* write everything in $s = A + B + C$ and $e = AB + BC + CA$. (Ideal. The obstruction is one line once $X = s^3 - 3se$ is on the page; the construction uses the smallest possible differences to make the second factor equal to 1 or 3.)

The ideal solution

Write $X = A^3 + B^3 + C^3 - 3ABC$, $s = A + B + C$ and $e = AB + BC + CA$. Expanding confirms the identity

$$X = (A+B+C)(A^2+B^2+C^2-AB-BC-CA) = \tfrac{1}{2}\, s\left((A-B)^2+(B-C)^2+(C-A)^2\right).$$

Since $A^2 + B^2 + C^2 = s^2 - 2e$, the second factor equals $s^2 - 3e$, so also $X = s^3 - 3se$.

Obstruction. Both factors are nonnegative when $A, B, C \geq 0$, so $X \geq 0$. Suppose $3 \mid X$. Then $3 \mid s^3 - 3se$ forces $3 \mid s^3$, and as 3 is prime, $3 \mid s$. But then $9 \mid s^3$ and $9 \mid 3se$, so $9 \mid X$. Thus no value of X is congruent to 3 or 6 modulo 9.

Construction. Let $n \geq 0$ be an integer. For $(A, B, C) = (n+1, n, n)$ the half sum of squares is $\frac{1}{2}(1+0+1) = 1$ and $s = 3n+1$, so $X = 3n+1$. For $(n+1, n+1, n)$ the half sum is $\frac{1}{2}(0+1+1) = 1$ and $s = 3n+2$, so $X = 3n+2$. For $n \geq 1$ and $(n+1, n, n-1)$ the half sum is $\frac{1}{2}(1+1+4) = 3$ and $s = 3n$, so $X = 9n$. Finally $(0, 0, 0)$ gives $X = 0$. As n runs over the nonnegative integers these families produce every nonnegative integer congruent to 1 or 2 modulo 3 and every nonnegative multiple of 9.

Combining the two parts, the set of values is

$$\boxed{\{\, m \in \mathbb{Z}_{\geq 0} :\ m \not\equiv 3, 6 \pmod 9 \,\}}$$

Layer 1: WHAT

Block 1: name the object. The expression is the standard cubic and factorises as s times a quadratic form. Block 2: read the quadratic factor as a half sum of squared differences, which fixes its sign and tells you its smallest positive values (1 when two variables agree and the third differs by one; 3 when the three variables are consecutive). Block 3: rewrite everything in the elementary symmetric functions, $X = s^3 - 3se$, so that divisibility by 3 and by 9 is read off s alone. Block 4: construct, using the two smallest quadratic factors and letting s sweep the residue classes. Block 5: box the answer as a set, with both directions stated.

Juxtapose with the nominal route: tabulate, conjecture "not 3 or 6 modulo 9", then try to prove it by listing the residues of A, B, C modulo 3 and checking cubes modulo 9 in each case. The bifurcation point is the moment you choose between three independent variables and two symmetric functions. The nominal solver keeps A, B, C separate because that is how they were handed over; the ideal solver notices the expression is symmetric and refuses to carry three letters when two symmetric quantities carry the same information with the prime already exposed.

Layer 2: HOW

Block 1. **Cue.** The literal shape $A^3 + B^3 + C^3 - 3ABC$. **Why natural.** This is a memorised identity, and the reason it is worth memorising is definitional: the expression vanishes whenever $A + B + C = 0$ (put $C = -A - B$ and expand), so $A + B + C$ is a factor by the factor theorem in three variables, and the cofactor is forced by degree and symmetry. **Link to the next block.** A product is easier to control than a sum: the sign, the size and the divisibility of X now come from two factors separately.

Block 2. **Cue.** The word "nonnegative" and the quadratic factor $A^2 + B^2 + C^2 - AB - BC - CA$. **Why natural.** A quadratic form that is symmetric and vanishes on the diagonal $A = B = C$ must be a combination of squared differences; here it is exactly half their sum. That gives nonnegativity for free and, more usefully, a catalogue of its small values in terms of how far apart the three variables are. **Link to the next block.** The obstruction needs the residues of X, and the construction needs the small values of the second factor; Block 2 supplies the second, Block 3 the first.

Block 3. **Cue.** A symmetric polynomial questioned modulo 3 and 9. **Why natural.** Symmetric polynomials are polynomials in s, e and ABC, and ABC has already cancelled. Writing $X = s^3 - 3se$ places the prime 3 in front of the only term that is not a pure power of s; divisibility questions then reduce to s. This is not a revelation, it is what "symmetric" means. **Link to the next block.** "$3 \mid X \Rightarrow 3 \mid s \Rightarrow 9 \mid X$" is the whole obstruction; the remaining question is whether every other residue occurs.

Block 4. **Cue.** The obstruction only mentions residues modulo 9, so the construction must be indexed by residue classes. **Why natural.** To hit a residue class you want one factor equal to a small constant and the other to sweep. The half sum of squares is 1 for $(n+1, n, n)$ type triples and 3 for consecutive triples, so X equals s or $3s$ with s controlled by n. **Link to the next block.** Two families cover residues 1 and 2 modulo 3; the consecutive family covers the multiples of 9; $(0, 0, 0)$ covers zero. That is every class the obstruction allowed.

Block 5. **Cue.** The command term "determine all". **Why natural.** A set answer needs both inclusions written explicitly; a grader looks for the sentence that says the construction exhausts the classes the obstruction permits. **Link to the next block.** Box the set and stop.

Layer 3: WHY

The mathematical truth beneath the route is Newton's identity for three variables: $A^3 + B^3 + C^3 = s^3 - 3se + 3ABC$, so $X = s^3 - 3se$ is not a coincidence you spotted, it is the power sum p_3 with the e_3 term deliberately subtracted. The examiners chose the coefficient -3 on ABC precisely so that e_3 disappears and the expression becomes a function of s and e only. Once you see it that way, the modulus 9 is not a mystery either: $X \equiv s^3 \equiv s \pmod 3$, and when $3 \mid s$ every term of $s^3 - 3se$ is divisible by 9.

A deeper reason the value set looks the way it does: X is the determinant of the circulant matrix with first row (A, B, C), equivalently the product $(A + B + C)(A + B\omega + C\omega^2)(A + B\omega^2 + C\omega)$ over a primitive cube root of unity ω. Products of circulants are circulants, so the value set is closed under multiplication, and $\{m \geq 0 : m \not\equiv 3, 6 \pmod 9\}$ is indeed multiplicatively closed. The prime 3 is special because $3 = -\omega^2(1-\omega)^2$ ramifies in the Eisenstein integers, and the three factors are congruent modulo $1 - \omega$; that is the ring-theoretic shadow of "if 3 divides X then it divides both rational factors".

The discernment mechanism: when a symmetric expression is questioned modulo a prime power, do not touch the individual variables. Move to the elementary symmetric functions, because the prime you care about will appear as a visible coefficient there. Train it by taking every symmetric identity you know ($p_2 = e_1^2 - 2e_2$, $p_3 = e_1^3 - 3e_1e_2 + 3e_3$) and reading off which primes they expose. The nudges here were three: the coefficient -3, the word "nonnegative", and the phrase "all possible values", which promises a clean description and therefore a modular one.

Revolving concepts

The cubic identity. $a^3 + b^3 + c^3 - 3abc = (a + b + c)(a^2 + b^2 + c^2 - ab - bc - ca)$, with the second factor $\frac{1}{2}\sum(a-b)^2 \geq 0$ and equal to zero only when $a = b = c$. This is also the AM-GM inequality for three terms in disguise: for $a, b, c \geq 0$ put $a = x^3$ and so on. See Book 1, field: Polynomials, topic: symmetric factorisations.

Newton's identities and their recursion. For three variables with $e_1 = s$, $e_2 = e$, $e_3 = ABC$, the power sums satisfy $p_1 = e_1$, $p_2 = e_1p_1 - 2e_2$, $p_3 = e_1p_2 - e_2p_1 + 3e_3$, and for $k > 3$ the pure recurrence $p_k = e_1p_{k-1} - e_2p_{k-2} + e_3p_{k-3}$. The generalised form for n variables is $p_k = e_1p_{k-1} - e_2p_{k-2} + \cdots + (-1)^{n-1}e_np_{k-n}$ for $k > n$, with the truncated versions (ending in $(-1)^{k-1}ke_k$) for $k \leq n$. Solution method: the power sums obey the same linear recurrence as the roots of

$t^n - e_1 t^{n-1} + \cdots$, so any question "power sums modulo a prime" becomes "coefficients modulo a prime". The same generalisation runs any problem where you are given $a+b+c$, $ab+bc+ca$, abc and asked about $a^k + b^k + c^k$. See Book 1, field: Polynomials, topic: Newton's identities.

Cubes modulo nine. $a^3 \equiv 0, 1$ or $-1 \pmod 9$ according as $a \equiv 0, 1, -1 \pmod 3$. See Book 1, field: Number theory, topic: residues of powers.

Circulant determinants and norm forms. The circulant with first row $(a_0, \ldots, a_{n-1})$ has eigenvectors $(1, \zeta^j, \zeta^{2j}, \ldots)$ and eigenvalues $\sum_k a_k \zeta^{jk}$, hence determinant $\prod_j \sum_k a_k \zeta^{jk}$. Multiplicativity of value sets (a^2+b^2, a^2+ab+b^2, a^2-db^2, and X here) always comes from a ring. See Book 1, field: Matrices and determinants, topic: circulants; and field: Number theory, topic: norm forms.

Circuit Bank

- **WHEN** you see $a^3+b^3+c^3-3abc$ **DO** write $(a+b+c)\cdot\frac{1}{2}\sum(a-b)^2$ and also s^3-3se **BECAUSE** because the first form gives sign and small values, the second exposes the prime 3
- **WHEN** you see "determine all possible values" of an integer expression **DO** write two headings, Obstruction and Construction, before anything else **BECAUSE** because an image is exactly the intersection of what is forbidden and what is achieved, and a grader marks the two halves separately
- **WHEN** you see a symmetric polynomial questioned modulo a prime power **DO** rewrite in elementary symmetric functions **BECAUSE** because Newton's identities put the prime in front of the mixed terms and the analysis collapses onto e_1
- **WHEN** you see the constraint "nonnegative integers" **DO** ask which factor it controls the sign of **BECAUSE** because a sign constraint is planted to kill one half of the value set and to hint that a factorisation exists

Backtesting

1. First useful cue: the literal expression $A^3+B^3+C^3-3ABC$; the factorisation must be written before reading the rest of the sentence.
2. Method the expression suggested: factor, sign from the half sum of squares, residues from s^3-3se, construction from triples with small spread.
3. Where you would hesitate, and the thought to reinforce: after conjecturing "not 3 or 6 modulo 9", the temptation is to case-bash residues of A, B, C. Reinforce: symmetric means work in s and e.
4. The trap the question sets: proving only the obstruction, or only the construction, and calling it done; also forgetting $X = 0$ and forgetting that $(n+1, n, n-1)$ needs $n \geq 1$.
5. What makes you recognise it faster next time: any coefficient -3 on abc is Newton's identity with e_3 cancelled; that is the signature of "this cubic factorises".
6. The reusable decision rule: image of a polynomial on the integers equals obstruction (modular) plus construction (parametrised families); both must appear on the page.

Mirror drills

1. (drill) Can $A^3+B^3+C^3-3ABC = 2019$ for nonnegative integers A, B, C? *Answer:* No; $2019 \equiv 3 \pmod 9$, and any value divisible by 3 is divisible by 9.

2. (drill) Show that if m and n are values of $A^3 + B^3 + C^3 - 3ABC$ on nonnegative integers then so is mn. *Answer:* X is the determinant of the circulant with first row (A, B, C); the product of two circulants with nonnegative entries is a circulant with nonnegative entries, and determinants multiply.
3. (drill) Show that $A^2 + B^2 + C^2 - AB - BC - CA$ is never congruent to 2 modulo 3. *Answer:* With $u = A - B$, $v = B - C$ it equals $u^2 + uv + v^2 \equiv (u - v)^2 \pmod 3$, a square modulo 3.

7.2 A2. Centroid and incentre parallel to a side

In the triangle $\triangle ABC$, let G be the centroid, and let I be the centre of the inscribed circle. Let α and β be the angles at the vertices A and B, respectively. Suppose that the segment IG is parallel to AB and that $\beta = 2\tan^{-1}(1/3)$. Find α.

Read it out loud

A triangle, two of its centres, a parallelism between the segment joining them and one side, and one angle handed to you not as β but as twice an arctangent. Read that last cue exactly as written: the examiners are telling you $\tan(\beta/2) = 1/3$. They could have written $\cos\beta = 4/5$ or "a 3, 4, 5 shape at B"; they chose the half angle. Half angles at the vertices live on the incircle, because the incentre sits on the bisectors and the touch points make right triangles with the vertices. That is the guiding light for the whole problem.

The fields wired in are triangle geometry (the centroid as an average of the vertices, the incentre through its tangent lengths and inradius), half-angle trigonometry, and the area bookkeeping $K = rs = \frac{1}{2}c\,h_c$. What is really being tested is a translation skill: can you convert each geometric condition into the parameter that defines the object named? G is defined by cutting each median in the ratio $2:1$, so its distance from AB is a third of the altitude from C. I is defined as the centre of a circle tangent to the three sides, so its distance from AB is r. "Parallel to AB" for two interior points is "same distance from AB". Nothing else is going on in the first half of the problem.

Within thirty seconds you should have written $r = h_c/3$, and next to it $\tan(\beta/2) = 1/3$, $\tan\beta = 3/4$, $\cos\beta = 4/5$, $\sin\beta = 3/5$. The rest is choosing which language to finish in. The ideal choice stays in the incircle's language, because the incircle already carries both hypotheses: r from the parallel condition and $\tan(\beta/2)$ from the given angle.

Approaches, ranked worst to best

1. **Sine rule with two unknown angles.** *Cue:* "find α" with β known. *Circuit:* angle sum plus sine rule. (After $a + b = 2c$ you get $\sin\alpha + \sin\beta = 2\sin(\alpha+\beta)$, which with $\sin\beta = 3/5$, $\cos\beta = 4/5$ reduces to $\sin\alpha + 2\cos\alpha = 1$. Squaring produces the spurious root $\cos\alpha = 4/5$, which must be rejected by a sign check. It works, but you paid for a squaring and a rejected root, and you still needed $a + b = 2c$ first.)

2. **Coordinates.** *Cue:* $\tan(\beta/2) = 1/3$ puts I at $(3r, r)$. *Circuit:* fix B at the origin, A on the x-axis. (With $B = (0,0)$, $I = (3r, r)$, C on the ray of slope $3/4$ at height $3r$, so $C = (4r, 3r)$; then $A = (t, 0)$ is fixed by $a + b = 2c$ or by tangency of CA to the incircle, giving $t = 4r$. Reliable, a page of ink, a quadratic in t to solve.)

3. **Law of cosines on $a + b = 2c$.** *Cue:* a linear relation between sides plus a known cosine. *Circuit:* eliminate one side. (With $b = 2c - a$ and $\cos\beta = 4/5$: $(2c-a)^2 = a^2 + c^2 - \frac{8}{5}ac$ gives $c = \frac{4}{5}a$, $b = \frac{3}{5}a$, so $a : b : c = 5 : 3 : 4$ and the right angle is opposite a, at A. Short and sound; it computes the whole triangle when only one angle was asked.)

4. **Synthetic: the altitude is tangent to the incircle.** *Cue:* $h_c = 3r$ and $\tan\beta = 3/4$ give the foot D of the altitude at $BD = 4r$, while the touch point F has $BF = 3r$. *Circuit:* a line at distance r from I is tangent. (Then $FD = r$, so the altitude line CD is tangent to the incircle; the tangents from C are CA and CB, and $CD \neq CB$, so $D = A$ and $\alpha = \pi/2$. Two lines, one delicate sentence about the two tangents from an external point.)

5. **Half-angle product identity.** *Cue:* a half angle given and the incentre present. *Circuit:* $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{s-c}{s}$. (Ideal. The parallel condition says $s = 3(s-c)$, the identity turns that into $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{1}{3}$, and the given $\tan\frac{\beta}{2} = \frac{1}{3}$ leaves $\tan\frac{\alpha}{2} = 1$. No squaring, no coordinates, nothing computed that was not asked.)

The ideal solution

Let $a = BC$, $b = CA$, $c = AB$, $s = \frac{1}{2}(a+b+c)$, let K be the area, r the inradius and h_c the altitude from C.

Step 1: the parallel condition means $a + b = 2c$. The centroid G lies on the median from C to the midpoint M of AB with $CG : GM = 2 : 1$, so the distance from G to the line AB is $\frac{1}{3}h_c$. The distance from I to AB is r, since the incircle is tangent to AB. Both points are interior, hence on the same side of AB, so $IG \parallel AB$ exactly when $r = \frac{1}{3}h_c$. (The segment is nondegenerate: $I = G$ would force an equilateral triangle, where $\tan\frac{\beta}{2} = \frac{1}{\sqrt{3}} \neq \frac{1}{3}$.) Now $2K = c\,h_c$, and $K = rs$ because the segments IA, IB, IC cut the triangle into three triangles of heights r on bases a, b, c. Hence $c\,h_c = 2rs$, and with $h_c = 3r$ this is $3c = a + b + c$, that is $a + b = 2c$. Equivalently $s = \frac{3}{2}c$ and $s - c = \frac{1}{2}c$, so

$$\frac{s-c}{s} = \frac{1}{3}.$$

Step 2: a lemma. In every triangle,

$$\tan\frac{\alpha}{2}\,\tan\frac{\beta}{2} = \frac{s-c}{s}.$$

Let the incircle touch BC, CA, AB at D, E, F. The two tangent segments from a point to a circle are equal, so with $AE = AF = x$, $BF = BD = y$, $CD = CE = z$ we have $x + y = c$, $y + z = a$, $z + x = b$, whence $x = s - a$, $y = s - b$, $z = s - c$. The incentre lies on the bisector at A and $IF \perp AB$ with $IF = r$, so in the right triangle AFI, $\tan\frac{\alpha}{2} = \frac{r}{s-a}$; similarly $\tan\frac{\beta}{2} = \frac{r}{s-b}$. Heron's formula $K^2 = s(s-a)(s-b)(s-c)$ together with $K = rs$ gives $r^2 = \frac{(s-a)(s-b)(s-c)}{s}$. Therefore

$$\tan\frac{\alpha}{2}\,\tan\frac{\beta}{2} = \frac{r^2}{(s-a)(s-b)} = \frac{s-c}{s}.$$

Step 3: conclude. By Steps 1 and 2 and the hypothesis $\tan\frac{\beta}{2} = \frac{1}{3}$,

$$\tan\frac{\alpha}{2}\cdot\frac{1}{3} = \frac{1}{3}, \qquad \text{so} \qquad \tan\frac{\alpha}{2} = 1.$$

Since $0 < \frac{\alpha}{2} < \frac{\pi}{2}$, this forces $\frac{\alpha}{2} = \frac{\pi}{4}$.

Step 4: the hypotheses are consistent. For $A = (0,0)$, $B = (4,0)$, $C = (0,3)$ we have $a = 5$, $b = 3$, $c = 4$, so $a + b = 2c$; $\tan\beta = \frac{3}{4}$ gives $\tan\frac{\beta}{2} = \frac{\sin\beta}{1+\cos\beta} = \frac{3/5}{9/5} = \frac{1}{3}$; and $r = K/s = 6/6 = 1$, so $I = (1,1)$ while $G = (\frac{4}{3}, 1)$, both on the line $y = 1$ parallel to AB.

$$\boxed{\alpha = \frac{\pi}{2}}$$

Layer 1: WHAT

Block 1: translate $IG \parallel AB$ into equal distances from AB, then through $K = rs = \frac{1}{2}ch_c$ into $a + b = 2c$, and store it as $(s-c)/s = 1/3$. Block 2: translate $\beta = 2\tan^{-1}(1/3)$ into $\tan(\beta/2) = 1/3$ and recognise a half angle as an incircle quantity. Block 3: prove the half-angle product lemma from tangent lengths and Heron. Block 4: solve $\tan(\alpha/2) = 1$ on the correct range, box, and exhibit the $3, 4, 5$ witness.

Juxtapose with the nominal route: coordinates for everything, or the sine rule with two unknown angles and a squaring. Both finish. The bifurcation point is the line after $a + b = 2c$: the nominal solver now computes the triangle (law of cosines, coordinates) because computing feels like progress; the ideal solver notices that the remaining hypothesis is a half angle, that half angles are ratios $r/(s-a)$, and that $a + b = 2c$ is itself a statement about s and $s - c$. Everything stays in one coordinate system, the incircle's.

Layer 2: HOW

Block 1. **Cue.** "The segment IG is parallel to AB." **Why natural.** Parallel to a line means equal signed distance from it. The distance of each named centre from a side is part of that centre's definition: G cuts the median in thirds, so its height is $h_c/3$; I is the centre of a circle tangent to AB, so its height is r. The identity $K = rs$ is then the natural bridge because it is the only standard formula containing r and the sides together, and $K = \frac{1}{2}ch_c$ is the only one containing h_c. **Link to the next block.** $a + b = 2c$ is the parallel condition in side language; rewriting it as $(s - c)/s = 1/3$ is what lets Block 3 use it.

Block 2. **Cue.** $\beta = 2\tan^{-1}(1/3)$. **Why natural.** The examiners chose to present the angle through its half. The half angle at B is the angle IBF in the right triangle with legs r and $BF = s - b$. So the given number is a ratio of incircle quantities, $r/(s - b) = 1/3$, not merely a trigonometric constant. **Link to the next block.** Both hypotheses are now statements about r, s, and the tangent lengths; the finishing identity must be one that relates exactly these.

Block 3. **Cue.** Two half-angle tangents and the ratio $(s - c)/s$ on the page. **Why natural.** $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{r^2}{(s-a)(s-b)}$ by definition of the tangent lengths, and Heron with $K = rs$ converts r^2 into $(s - a)(s - b)(s - c)/s$. The $(s - a)(s - b)$ cancels. There is no leap: you multiply the two things you have and simplify with the one formula that expresses r in terms of the sides. **Link to the next block.** The product equals $1/3$ by Block 1, and one factor equals $1/3$ by Block 2.

Block 4. **Cue.** $\tan(\alpha/2) = 1$ with α an angle of a triangle. **Why natural.** Inverting a tangent needs a range; $\alpha/2 \in (0, \pi/2)$ makes the inverse unique. The witness triangle is not decoration: a "find the angle" question with several hypotheses could be vacuous, and a $3, 4, 5$ check costs three lines and certifies the answer against a sign error. **Link to the next block.** Box and stop.

Layer 3: WHY

Three truths make this route ideal. First, both centres are weighted averages of the vertices: $G = \frac{A+B+C}{3}$ and $I = \frac{aA+bB+cC}{a+b+c}$. The signed distance of a weighted average from the line AB is the weight of C times h_c. So $IG \parallel AB$ is literally the equation of C-weights, $\frac{1}{3} = \frac{c}{a+b+c}$, which is $a + b = 2c$ without any area argument at all. That is the definitional reason the side relation is so clean, and it generalises instantly: $IG \parallel BC$ iff $b + c = 2a$.

Second, the natural coordinates of the incircle are the tangent lengths $(s - a, s - b, s - c)$ and r; every half angle is r over a tangent length, and Heron is the relation $r^2 s = (s - a)(s - b)(s - c)$ that closes the system. The identity $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{s-c}{s}$ is nothing but this closure. An equivalent closure is the angle-sum identity $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} + \tan\frac{\beta}{2}\tan\frac{\gamma}{2} + \tan\frac{\gamma}{2}\tan\frac{\alpha}{2} = 1$, from $\frac{\alpha+\beta+\gamma}{2} = \frac{\pi}{2}$; with $\tan\frac{\gamma}{2} = \frac{r}{s-c} = \frac{2r}{c}$ and $a + b = 2c$ it also yields $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{1}{3}$, so you can finish without Heron if you prefer.

Third, the geometry of the answer: $a + b = 2c$ places C on an ellipse with foci A and B, and fixing β restricts C to a ray from B, which meets the ellipse once. So the triangle is unique up to similarity and the answer had to be a specific number; the examiners chose $\tan\frac{\beta}{2} = \frac{1}{3}$ so that the number is $\pi/2$ and the triangle is the $3, 4, 5$ one. The discernment mechanism is: translate every hypothesis into the defining parameters of the objects named, then look for the one standard identity that links those parameters. Train it by writing, for each triangle centre you know, its barycentric weights and its distance to each side; after that, any "parallel to a side" or "on the bisector" condition is a one-line equation.

Revolving concepts

Centroid. $G = \frac{A+B+C}{3}$; it divides each median 2 : 1 from the vertex; its distance from side BC is $h_a/3$. See Book 1, field: Geometry, topic: triangle centres and barycentric coordinates.

Incentre and incircle. $I = \frac{aA+bB+cC}{a+b+c}$; $r = K/s$; the tangent lengths from A, B, C are $s-a, s-b, s-c$; $\tan\frac{\alpha}{2} = \frac{r}{s-a} = \sqrt{\frac{(s-b)(s-c)}{s(s-a)}}$; and the products $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{s-c}{s}$. The half-angle tangents satisfy $xy + yz + zx = 1$. See Book 1, field: Geometry, topic: incircle and tangent lengths.

Heron and the area identities. $K = \sqrt{s(s-a)(s-b)(s-c)} = rs = \frac{1}{2}ab\sin\gamma = \frac{1}{2}c\,h_c$. Each formula is a bridge between one family of parameters and another; you choose the pair of bridges that connect the hypotheses you were given. See Book 1, field: Geometry, topic: area formulae.

Parallelism in weighted-average language. If P and Q are weighted averages of A, B, C with normalised weights, then $PQ \parallel AB$ iff their C-weights agree, because the C-weight times h_c is the signed distance from AB. This is the general form of Step 1 and it handles the Nagel point, the symmedian point, and any other centre with known weights. See Book 1, field: Geometry, topic: barycentric coordinates.

Recursion. None in this problem; the only iteration is the protocol "translate, then link", applied once per hypothesis.

Circuit Bank

- **WHEN** you see "segment XY parallel to side AB" for two triangle centres **DO** write "distance of X from AB equals distance of Y from AB" in each centre's own parameters, or equate their C-weights **BECAUSE** because parallel means equal signed distance, and each centre's distance to a side is part of its definition
- **WHEN** you see an angle given as $2\tan^{-1}(t)$ **DO** write $\tan(\theta/2) = t$ and reach for the incircle: $\tan\frac{\alpha}{2} = r/(s-a)$ **BECAUSE** because half angles at the vertices are the acute angles of the right triangles formed by the incentre, a vertex and a touch point
- **WHEN** you see the incentre together with a relation among the sides **DO** write $s-a$, $s-b$, $s-c$, $K = rs$ and Heron **BECAUSE** because these are the incircle's coordinates and Heron is the single relation that closes them
- **WHEN** you see a "find the angle" question with several hypotheses **DO** after solving, exhibit one explicit triangle satisfying everything **BECAUSE** because a unique answer must be consistent and a witness catches a sign or squaring error at the cost of three lines

Backtesting

1. First useful cue: $\beta = 2\tan^{-1}(1/3)$, read as $\tan(\beta/2) = 1/3$, which points at the incircle before the parallel condition is even processed.
2. Method the expression suggested: convert the parallel condition to $r = h_c/3$ and then $a + b = 2c$; convert the half angle to $r/(s-b) = 1/3$; close with $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} = \frac{s-c}{s}$.
3. Where you would hesitate, and the thought to reinforce: after $a + b = 2c$, the pull toward the law of cosines or coordinates. Reinforce: the remaining datum is a half angle, so stay in tangent lengths; the finishing identity is one line.
4. The trap the question sets: solving $\sin\alpha + 2\cos\alpha = 1$ by squaring and keeping the spurious $\cos\alpha = 4/5$; or forgetting that $\tan\frac{\alpha}{2} = 1$ needs the range $\alpha/2 \in (0, \pi/2)$ to invert.
5. What makes you recognise it faster next time: any angle presented as a double arctangent is a half-angle cue, and half angles are incircle ratios; any centroid condition is a statement about thirds of an altitude.

6. The reusable decision rule: translate each hypothesis into the defining parameters of the objects it names, then finish with the one standard identity linking those parameters; compute the triangle only if the question asks for the triangle.

Mirror drills

1. (drill) In triangle ABC the segment IG is parallel to BC. Show that $b + c = 2a$. *Answer:* distances of G and I from BC are $h_a/3$ and r; $a\,h_a = 2K = 2rs$ with $h_a = 3r$ gives $3a = 2s$.
2. (drill) Prove $\tan\frac{\alpha}{2}\tan\frac{\beta}{2} + \tan\frac{\beta}{2}\tan\frac{\gamma}{2} + \tan\frac{\gamma}{2}\tan\frac{\alpha}{2} = 1$ in any triangle. *Answer:* $\frac{\alpha}{2} + \frac{\beta}{2} = \frac{\pi}{2} - \frac{\gamma}{2}$; take tangents, $\frac{x+y}{1-xy} = \frac{1}{z}$, and clear denominators.
3. (drill) In a triangle with $IG \parallel AB$, express $\tan\frac{\gamma}{2}$ through $x + y$ where $x = \tan\frac{\alpha}{2}$, $y = \tan\frac{\beta}{2}$. *Answer:* $xy = \frac{1}{3}$ by the lemma, so the identity of the previous drill gives $\tan\frac{\gamma}{2} = \frac{2}{3(x+y)}$.

7.3 A3. Mean root modulus by Vieta and AM-GM

Given real numbers $b_0, b_1, \ldots, b_{2019}$ with $b_{2019} \neq 0$, let $z_1, z_2, \ldots, z_{2019}$ be the roots in the complex plane of the polynomial

$$P(z) = \sum_{k=0}^{2019} b_k z^k.$$

Let $\mu = (|z_1| + \cdots + |z_{2019}|)/2019$ be the average of the distances from $z_1, z_2, \ldots, z_{2019}$ to the origin. Determine the largest constant M such that $\mu \geq M$ for all choices of $b_0, b_1, \ldots, b_{2019}$ that satisfy

$$1 \leq b_0 < b_1 < b_2 < \cdots < b_{2019} \leq 2019.$$

Read it out loud

A polynomial of degree 2019 whose real coefficients increase strictly from at least 1 to at most 2019, and you are asked about the arithmetic mean of the moduli of its roots. "The largest constant M such that $\mu \geq M$ for all choices" is the infimum of μ over the constraint set, so the command term is again double: prove a lower bound, then show the bound is attained (or at least approached). A bound without an extremal example is half an answer.

Now read the constraints as the examiners wrote them. Only two coefficients carry absolute bounds, $b_0 \geq 1$ and $b_{2019} \leq 2019$; all the middle coefficients are tied to each other by strict monotonicity alone. Which symmetric function of the roots depends on b_0 and b_{2019} and on nothing else? The product: $\prod z_i = (-1)^{2019} b_0/b_{2019}$. So the data control the product of the moduli, and the question asks about their arithmetic mean. The circuit that connects an arithmetic mean of nonnegative numbers to their product is AM-GM. Within thirty seconds you should have written $\prod |z_i| = b_0/b_{2019} \geq 1/2019$ and $\mu \geq (1/2019)^{1/2019}$.

The fields wired in are polynomials (the fundamental theorem of algebra, Vieta), inequalities (AM-GM with its equality case), and complex numbers (roots of unity, rescaling). What is really being tested is whether you attack a mean of moduli through the product rather than through the location of individual roots. One more planted cue: "real numbers", not integers. The equality case will need coefficients $2019^{k/2019}$, which are irrational; the examiners wrote "real" so that the extremal polynomial is admissible.

Approaches, ranked worst to best

1. **Root location theorems.** *Cue:* positive increasing coefficients. *Circuit:* Eneström and Kakeya: all roots lie in $|z| \leq 1$, refined to the annulus between min and max of b_k/b_{k+1}. (A theorem-name reflex that bounds every root separately. The lower radius $\min_k b_k/b_{k+1}$ can be as small as about $1/2019$ (take $b_0 = 1$, b_1 near 2019), so the mean is only shown to exceed roughly $1/2019$, far below the truth. Individual bounds are the wrong tool for a statement about a mean.)

2. **Triangle inequality on the sum of roots.** *Cue:* "average" suggests $\sum z_i = -b_{2018}/b_{2019}$. *Circuit:* mean of moduli is at least modulus of the mean. (Gives $\mu \geq b_{2018}/(2019\, b_{2019})$, which can be tiny, and in any case the triangle inequality can only certify a mean of moduli from below by a quantity that may be zero. Dead on arrival.)

3. **Small analogues first.** *Cue:* a 2019 that is clearly a placeholder. *Circuit:* replace 2019 by n and test $n = 2$. (For $1 \leq b_0 < b_1 < b_2 \leq 2$ the product of the moduli is $b_0/b_2 \geq \frac{1}{2}$ and $b_k = 2^{k/2}$ gives roots of modulus $2^{-1/2}$. This is not a proof but it tells you the answer is $2019^{-1/2019}$ and that a geometric progression of coefficients is the extremal shape. Worth sixty seconds.)

4. **Vieta, AM-GM, roots of unity.** *Cue:* bounds on b_0 and b_{2019} only; "average"; "largest constant". *Circuit:* mean $\geq$ geometric mean $=$ (product)$^{1/n}$, product from Vieta, equality forces all moduli equal, so rescale $1 + w + \cdots + w^{2019}$. (Ideal. Three lines for the bound, four for the example.)

The ideal solution

Since $b_{2019} \neq 0$, P has degree exactly 2019, so by the fundamental theorem of algebra it has 2019 roots $z_1, \ldots, z_{2019}$ counted with multiplicity and $P(z) = b_{2019} \prod_{i=1}^{2019} (z - z_i)$. Comparing constant terms, $b_0 = b_{2019}(-1)^{2019} \prod_i z_i$, hence

$$\prod_{i=1}^{2019} |z_i| = \frac{|b_0|}{|b_{2019}|} = \frac{b_0}{b_{2019}} \;\geq\; \frac{1}{2019},$$

using $b_0 \geq 1$ and $0 < b_{2019} \leq 2019$.

Lower bound. The moduli $|z_i|$ are nonnegative reals, so by the inequality of arithmetic and geometric means,

$$\mu = \frac{1}{2019} \sum_{i=1}^{2019} |z_i| \;\geq\; \Big(\prod_{i=1}^{2019} |z_i| \Big)^{1/2019} \;\geq\; 2019^{-1/2019}.$$

Sharpness. Let $\rho = 2019^{1/2019} > 1$ and put $b_k = \rho^k$ for $0 \leq k \leq 2019$. Then $b_0 = 1$, $b_{2019} = 2019$, and $b_0 < b_1 < \cdots < b_{2019}$ because $\rho > 1$; so this choice is admissible. For every z with $\rho z \neq 1$,

$$P(z) = \sum_{k=0}^{2019} (\rho z)^k = \frac{(\rho z)^{2020} - 1}{\rho z - 1},$$

and $P(1/\rho) = 2020 \neq 0$. So the roots of P are exactly the z with $(\rho z)^{2020} = 1$ and $\rho z \neq 1$, namely $z_j = \rho^{-1} e^{2\pi i j/2020}$ for $j = 1, \ldots, 2019$. These are 2019 distinct numbers, so they are all the roots, and each has modulus ρ^{-1}. Hence for this choice $\mu = \rho^{-1} = 2019^{-1/2019}$, and the lower bound is attained.

Therefore the largest constant is

$$\boxed{M = 2019^{-1/2019}}$$

Layer 1: WHAT

Block 1: parse "largest constant with $\mu \geq M$ for all choices" as an infimum, so the write-up has a Bound and a Sharpness paragraph. Block 2: identify the product of the moduli as the one symmetric function fixed by the two bounded coefficients, via Vieta. Block 3: pass from arithmetic mean to geometric mean by AM-GM. Block 4: read the equality conditions (all moduli equal, $b_0 = 1$, $b_{2019} = 2019$) as a specification for the extremal polynomial: coefficients in geometric progression with ratio $2019^{1/2019}$, which is the geometric series $\sum(\rho z)^k$, whose roots are rescaled roots of unity. Block 5: check admissibility and count the roots.

Juxtapose with the nominal route: try to locate the roots (Cauchy bounds, Eneström and Kakeya), or to bound the mean through the sum of the roots. The bifurcation point is the word "average": the nominal solver hears "each root is at least this far from the origin" and reaches for a location theorem; the ideal solver hears "arithmetic mean of nonnegative quantities" and asks which quantity the data control, finds the product, and writes AM-GM.

Layer 2: HOW

Block 1. Cue. "Determine the largest constant M such that $\mu \geq M$ for all choices." **Why natural.** This sentence defines M as an infimum. The largest valid lower bound is proved by two separate facts: the inequality, and an admissible choice where it is tight. Writing the two headings first prevents the classic half-answer. **Link to the next block.** The Bound paragraph needs a handle on μ; Block 2 finds it.

Block 2. **Cue.** Bounds on b_0 and b_{2019} only; the middle coefficients constrained merely by order. **Why natural.** Vieta says each coefficient ratio b_{2019-k}/b_{2019} is $(-1)^k$ times the k-th elementary symmetric function of the roots. The product e_{2019} is the only one that depends on the two bounded coefficients alone. The examiners bounded exactly those two; that is a planted cue, not a coincidence. Moduli enter because $|\prod z_i| = \prod |z_i|$. **Link to the next block.** You now hold $\prod |z_i| \geq 1/2019$ and want $\frac{1}{2019}\sum |z_i|$; an inequality between a mean and a product is required.

Block 3. **Cue.** An arithmetic mean of nonnegative reals to be bounded below, with their product known. **Why natural.** AM-GM is by definition the statement that the arithmetic mean dominates the geometric mean, and the geometric mean is the 2019-th root of the product. The moduli are nonnegative, which is exactly the hypothesis. One line. **Link to the next block.** The bound is $2019^{-1/2019}$; to show it is the largest constant you need equality, so read the equality conditions.

Block 4. **Cue.** Equality in AM-GM holds iff all $|z_i|$ are equal; equality in $b_0/b_{2019} \geq 1/2019$ iff $b_0 = 1$, $b_{2019} = 2019$. **Why natural.** A polynomial with all roots on one circle $|z| = 1/\rho$ is a rescaling $Q(\rho z)$ of a polynomial with all roots on the unit circle. Rescaling multiplies the k-th coefficient by ρ^k, so the coefficients of $Q(\rho z)$ form a geometric progression precisely when Q has constant coefficients, that is $Q(w) = 1 + w + \cdots + w^{2019}$, whose roots are the 2020-th roots of unity other than 1. Increasing coefficients means $\rho > 1$; the endpoint conditions fix $\rho^{2019} = 2019$. **Link to the next block.** The example must be verified, not merely proposed.

Block 5. **Cue.** A proposed extremal polynomial. **Why natural.** Three checks are owed to the grader: the coefficients satisfy the strict inequalities and the endpoint bounds; the roots are found completely (the geometric sum formula needs $\rho z \neq 1$, and $z = 1/\rho$ is not a root); and the number of distinct roots found equals the degree, so nothing is missing and μ is computed correctly. **Link to the next block.** Box M.

Layer 3: WHY

The route is ideal because of three truths. First, the elementary symmetric functions of the roots are the coefficients, and among them the product is the cheapest: it sees only the two ends of the coefficient list. Any constraint set that pins down b_0 and b_n is really a constraint on $\prod |z_i|$. Second, the arithmetic mean of nonnegative numbers is bounded below by their geometric mean and by nothing weaker that uses only the product; AM-GM is not one choice among many here, it is the exact inequality for the data you hold. Third, AM-GM is tight precisely on constant tuples, so the extremal configuration must have all roots on one circle, and the polynomial with all roots on a circle and monotone coefficients is a rescaled geometric series.

There is also a scaling truth worth repeating. If you replace z by z/λ the roots multiply by λ and the coefficients b_k become $b_k\lambda^{-k}$. So μ has the dimensions of a length and the constraints must fix a scale; they do it through b_0 and b_{2019}, whose ratio has dimension length^{2019}, which is why the answer is a 2019-th root. Dimensional reading alone predicts $M = (\text{ratio})^{1/2019}$ up to a constant.

The discernment mechanism: when asked to bound a mean of nonnegative quantities, ask which symmetric function the data control. Product: AM-GM. Sum of squares: QM-AM or Cauchy. Second symmetric function: Maclaurin. Train it by listing, for each Vieta coefficient, the inequality that links it to the arithmetic mean. The examiners nudged with three details: bounds on the two end coefficients only, the word "average", and "real numbers" so that the irrational geometric progression is admissible. The number 2019 is decoration; every step works for any degree n with bounds $L \leq b_0$, $b_n \leq U$, giving $M = (L/U)^{1/n}$.

Revolving concepts

Vieta's formulas. For $P(z) = b_n \prod(z - z_i)$, $e_k(z_1, \ldots, z_n) = (-1)^k b_{n-k}/b_n$; in particular $\prod z_i = (-1)^n b_0/b_n$ and $\sum z_i = -b_{n-1}/b_n$. The fundamental theorem of algebra guarantees the n roots with multiplicity. See Book 1, field: Polynomials, topic: Vieta and symmetric functions.

AM-GM and its relatives. For nonnegative x_i, $\frac{1}{n}\sum x_i \geq (\prod x_i)^{1/n}$ with equality iff all x_i are equal. The generalisation that orbits this problem is Maclaurin's chain: with $S_k = e_k/\binom{n}{k}$, $S_1 \geq S_2^{1/2} \geq \cdots \geq S_n^{1/n}$, so every symmetric mean of the moduli bounds the arithmetic mean from below. See Book 1, field: Inequalities, topic: means and Maclaurin.

Geometric sums, roots of unity, rescaling. $\sum_{k=0}^{n} w^k = \frac{w^{n+1}-1}{w-1}$ for $w \neq 1$, with roots the $(n+1)$-th roots of unity other than 1. The substitution $w = \rho z$ multiplies the coefficient of z^k by ρ^k and divides every root by ρ. The recursion beneath the geometric sum is $S_{m+1} = wS_m + 1$; its general solution $S_m = \frac{w^m-1}{w-1}$ is the template for every first-order linear recurrence with constant coefficients, $u_{m+1} = au_m + b$, solved by a fixed point plus a geometric term. See Book 1, field: Complex numbers, topic: roots of unity; and field: Sequences and recurrences, topic: first-order linear recurrences.

Eneström and Kakeya. If $0 < b_0 \leq b_1 \leq \cdots \leq b_n$ then every root of $\sum b_k z^k$ satisfies $|z| \leq 1$; more generally all roots lie in the annulus $\min_k(b_k/b_{k+1}) \leq |z| \leq \max_k(b_k/b_{k+1})$. Know it so that you know it bounds roots individually and is therefore the wrong instrument for a mean. See Book 1, field: Polynomials, topic: root location.

Circuit Bank

- **WHEN** you see a mean of nonnegative quantities to be bounded below, with their product controlled by the data **DO** write AM-GM: mean $\geq$ (product)$^{1/n}$, and immediately write its equality condition **BECAUSE** because AM-GM is the exact inequality between an arithmetic mean and a product, and the equality condition designs the extremal example
- **WHEN** you see absolute bounds on b_0 and b_n only **DO** write $\prod z_i = (-1)^n b_0/b_n$ **BECAUSE** because the product of the roots is the one Vieta quantity that depends on exactly those two coefficients
- **WHEN** you see "largest constant M such that $\ldots \geq M$ for all" **DO** write two headings, Bound and Sharpness, before any algebra **BECAUSE** because the largest constant is an infimum, and a bound without an admissible choice attaining it determines nothing
- **WHEN** you need a polynomial with all roots on a circle and monotone coefficients **DO** write $\sum_{k=0}^{n}(\rho z)^k$ with $\rho > 1$ **BECAUSE** because rescaling the geometric series puts the roots of unity on the circle $|z| = 1/\rho$ and makes the coefficients a geometric progression

Backtesting

1. First useful cue: the constraints bound b_0 from below and b_{2019} from above, and nothing else absolutely; that is Vieta pointing at the product of the roots.
2. Method the expression suggested: product of moduli from Vieta, arithmetic mean from AM-GM, extremal example from the equality case via a rescaled geometric series.
3. Where you would hesitate, and the thought to reinforce: after the bound, wondering whether $2019^{-1/2019}$ is really attained with strictly increasing coefficients. Reinforce: equality needs all moduli equal, and the geometric series with ratio $\rho > 1$ has exactly that and is strictly increasing.
4. The trap the question sets: reaching for a root-location theorem because the coefficients are positive and increasing; or presenting the example without verifying the strict inequalities, without excluding $\rho z = 1$, or without counting that 2019 distinct roots were found.

5. What makes you recognise it faster next time: "average of $|z_i|$" plus a polynomial is AM-GM on Vieta's product; "real coefficients" in a problem that could have said integers is a hint that the extremal example is irrational.
6. The reusable decision rule: bound a mean through the symmetric function the data control, then build the extremal case from the equality condition of the inequality you used.

Mirror drills

1. (drill) Real $b_0 < \cdots < b_n$ with $L \le b_0$ and $b_n \le U$ ($0 < L < U$). Find the largest M with $\mu \ge M$ always. *Answer:* $M = (L/U)^{1/n}$, by the same AM-GM bound and $b_k = L\,(U/L)^{k/n}$.
2. (drill) With $0 < b_0 \le b_1 \le \cdots \le b_n$, show every root of $\sum b_k z^k$ satisfies $|z| \le 1$. *Answer:* for $|z| > 1$, $|(1-z)P(z)| \ge b_n|z|^{n+1} - \left(b_0 + \sum_{k\ge 1}(b_k - b_{k-1})|z|^k\right) > b_n|z|^{n+1} - b_n|z|^n > 0$.
3. (drill) If $z^n + c_{n-1}z^{n-1} + \cdots + c_0$ has $|c_0| = 1$, show some root has modulus at least 1 and some root has modulus at most 1. *Answer:* the product of the moduli is $|c_0| = 1$, so they cannot all exceed 1 nor all fall below 1.

7.4 A4. Fixed radius spherical means have a kernel

Let f be a continuous real-valued function on $\mathbb{R}^3$. Suppose that for every sphere S of radius 1, the integral of $f(x,y,z)$ over the surface of S equals 0. Must $f(x,y,z)$ be identically 0?

Read it out loud

Say the structure out loud before you touch a single integral. This is an averaging question. There is one linear operator in the room: the map that sends a continuous function to its average over the sphere of radius 1 centred at a point. The hypothesis says that operator kills f. The question asks whether that operator is injective. So the entire problem is: does the fixed radius spherical mean have a nonzero kernel on $C(\mathbb{R}^3)$? The fields wired in are multivariable calculus (a genuine surface integral, with its Jacobian), a scrap of classical geometry (the area of a zone of a sphere), and, underneath everything, harmonic analysis (a translation invariant operator is a convolution, and a convolution is diagonalised by exponentials).

Now the cue that should have been circled within thirty seconds: *radius* 1. Not "every sphere". Not "spheres of every radius". One fixed radius. Ask immediately what the problem would look like with that restriction lifted: if the averages over spheres of *all* radii r vanished, then letting $r \to 0^+$ and using continuity of f would give $f(c) = 0$ at every centre c, and the answer would be a two line yes. The setter removed exactly that. When a "Must ...?" question carries a restriction whose removal would make the answer trivially yes, the restriction is the whole content of the problem and the answer is almost always no. You should have written "answer: no, build a counterexample" before doing any calculation.

Second observation, so that you do not waste time on a dimension count. The hypothesis is one scalar equation for each centre in $\mathbb{R}^3$, imposed on one unknown function on $\mathbb{R}^3$. The constraints and the unknowns are indexed by the same space, so no cheap counting argument decides the question either way. You have to look at the operator itself.

Third observation, and this is the one that shortens the search. The hypothesis is invariant under every translation and every rotation of $\mathbb{R}^3$. That means you are free to hunt for a counterexample inside a family that respects part of that symmetry, and the cheapest such family is functions of a single coordinate, $f(x,y,z) = g(z)$. Inside that family the surface integral collapses to a one dimensional integral, and a one dimensional integral condition is something you can solve exactly rather than estimate.

Approaches, ranked worst to best

1. **Shrink the spheres.** *Cue:* the phrase "for every sphere". *Circuit:* the mean value circuit: the average of a continuous function over a sphere of radius r about c tends to $f(c)$ as $r \to 0^+$. This is the reflex that fires if you read the question carelessly, and it is dead on arrival: you are given exactly one radius, so there is no limit to take. Cost: the entire time slot, and no marks, because you are trying to prove a false statement.

2. **Differentiate the identity in the centre.** *Cue:* the condition holds for every centre, so the centre is a free parameter. *Circuit:* differentiate an identity with respect to a free parameter to generate new identities. If f were smooth you would deduce that the spherical means of all its partial derivatives vanish too, and you would drift toward the Darboux equation for spherical means. Every statement is true and none of them closes the problem, and f is only assumed continuous, so you would also owe a smoothing argument. Costs a page, yields nothing.

3. **Hunt for a radial counterexample.** *Cue:* a sphere is the most rotationally symmetric object there is. *Circuit:* match the symmetry of the unknown to the symmetry of the constraint. Setting $f(v) = F(|v|)$ turns the condition into a one dimensional integral equation relating F

to a weighted average over an interval depending on $|c|$, of Volterra type. It is solvable, but the kernel is ugly and you will spend real ink before you see anything.

4. **Hunt for a counterexample depending on one coordinate.** *Cue:* a sphere meets each plane $z = \text{const}$ in a circle whose contribution to the area you can write down. *Circuit:* slice a surface integral along the axis of symmetry and integrate the slabs. Cheap and complete. The condition becomes a statement about the integral of a function of one variable over every window of length 2, which you can solve exactly. This is what you write up.

5. **Read the symbol of the operator.** *Cue:* "for every sphere of radius 1" is a convolution with the surface measure of the unit sphere. *Circuit:* the eigenfunction circuit: a translation invariant operator is diagonalised by $e^{i\xi\cdot v}$, so injectivity is exactly the statement that its symbol has no real zeros. This is the best route for understanding and it is how the example in route 4 is found rather than guessed. The average of $e^{i\xi\cdot v}$ over the unit sphere is $\sin|\xi|/|\xi|$, which vanishes whenever $|\xi|$ is a nonzero multiple of π. Take $|\xi| = \pi$ with ξ along the z axis and take the imaginary part: $\sin(\pi z)$. Write up route 4, but think in route 5.

The ideal solution

The answer is $\boxed{\text{no}}$: the hypothesis does not force f to vanish.

Take

$$f(x, y, z) = \sin(\pi z),$$

which is continuous on $\mathbb{R}^3$ and is not identically 0, since $f(0, 0, \frac{1}{2}) = 1$.

Let S be the sphere of radius 1 centred at an arbitrary point (x_0, y_0, z_0). Parametrise S by

$$S(\varphi, \theta) = (x_0 + \sin\varphi\cos\theta,\ y_0 + \sin\varphi\sin\theta,\ z_0 + \cos\varphi),$$

for $\varphi \in [0, \pi]$, $\theta \in [0, 2\pi]$. Then

$$\frac{\partial S}{\partial \varphi} = (\cos\varphi\cos\theta, \cos\varphi\sin\theta, -\sin\varphi), \qquad \frac{\partial S}{\partial \theta} = (-\sin\varphi\sin\theta, \sin\varphi\cos\theta, 0),$$

and a direct computation gives

$$\frac{\partial S}{\partial \varphi} \times \frac{\partial S}{\partial \theta} = -\sin\varphi\,\big(\sin\varphi\cos\theta,\ \sin\varphi\sin\theta,\ \cos\varphi\big),$$

whose length is $|\sin\varphi| = \sin\varphi$ on $[0, \pi]$. Hence $dS = \sin\varphi\, d\varphi\, d\theta$ and

$$\iint_S f\, dS = \int_0^{\pi}\int_0^{2\pi} \sin\big(\pi(z_0 + \cos\varphi)\big)\sin\varphi\, d\theta\, d\varphi = 2\pi\int_0^{\pi} \sin\big(\pi(z_0 + \cos\varphi)\big)\sin\varphi\, d\varphi.$$

Substituting $t = \cos\varphi$, so that $dt = -\sin\varphi\, d\varphi$ and the limits $\varphi = 0, \pi$ become $t = 1, -1$,

$$\iint_S f\, dS = 2\pi\int_{-1}^{1} \sin\big(\pi(z_0 + t)\big)\, dt = 2\pi\left[\frac{-\cos\big(\pi(z_0 + t)\big)}{\pi}\right]_{t=-1}^{t=1} = 2\big(\cos(\pi z_0 - \pi) - \cos(\pi z_0 + \pi)\big).$$

Both $\cos(\pi z_0 - \pi)$ and $\cos(\pi z_0 + \pi)$ equal $-\cos(\pi z_0)$, so the bracket is 0. Therefore

$$\iint_S f\, dS = 0$$

for every sphere S of radius 1, while $f \not\equiv 0$. The answer is no.

Remark 7.4.1. Nothing above is special to $\sin(\pi z)$ beyond two features: the integral over $t \in [-1, 1]$ of $g(z_0 + t)$ must vanish for every z_0, and $\sin(\pi t)$ has period 2 with mean zero over a period. In fact, writing G for an antiderivative of a continuous g, the condition $\int_{-1}^{1} g(z_0 + t)\, dt = 0$ for all z_0 says $G(z_0 + 1) = G(z_0 - 1)$ for all z_0, that is, G has period 2. So $f(x, y, z) = g(z)$ satisfies the hypothesis precisely when g is continuous, has period 2, and has mean zero over a period. That is an infinite dimensional kernel, not a lone counterexample.

Layer 1: WHAT

Block 1. Read "radius 1" as a deliberate restriction; check what happens if the radius is allowed to vary; conclude the answer is no and switch from proving to constructing.

Block 2. Use the translation and rotation invariance of the hypothesis to cut the search down: look for $f(x, y, z) = g(z)$, constant on planes perpendicular to one axis.

Block 3. Compute the surface measure on the unit sphere in spherical coordinates, $dS = \sin\varphi \, d\varphi \, d\theta$, and substitute $t = \cos\varphi$. The $\sin\varphi$ from the Jacobian is exactly the dt; that is Archimedes, and it is why this family is the cheap one.

Block 4. The condition becomes $\int_{-1}^{1} g(z_0 + t)\, dt = 0$ for every z_0, that is, every window of length 2 has zero integral.

Block 5. Solve that: an antiderivative of g has period 2. Pick the smallest example you can verify in one line, $g(t) = \sin(\pi t)$, and verify it.

The nominal route, the zombie route, is Block 1 misread: "for every sphere, the average is zero, so by the mean value property f is zero", followed by twenty minutes of trying to make a limit exist that does not exist, or by an attempt to expand f in spherical harmonics and prove injectivity coefficient by coefficient. The bifurcation point is the single sentence "what would change if the radius were free?". On the free radius side you get a proof; on the fixed radius side you get a counterexample. The problem is nothing but that sentence, and it is a sentence you must self-generate in the arena.

Layer 2: HOW

Block 1. Cue. The words "of radius 1", where the natural phrasing would have been "of every radius". **Why natural.** Restrictions in a "Must ...?" question are never decoration. Test the restriction by removing it: with all radii available, shrinking to a point recovers $f(c)$ by continuity, so the answer would be yes and the problem would be a warm up, not an A4. The setter therefore removed the only tool that would have worked, which tells you the answer is no. **Link to the next block.** Having decided to construct rather than prove, the next question is where to look, and the answer is dictated by the symmetry of the hypothesis.

Block 2. Cue. The hypothesis is stated for *every* sphere of radius 1, so it is invariant under all translations and rotations. **Why natural.** A condition invariant under a big group is cheapest to test on functions with a lot of symmetry of their own, because then the integral collapses. The most economical choice is a function of one coordinate: it is constant on each plane $z = \text{const}$, and a sphere meets that family of planes in a way you can integrate in closed form. Radial functions look more symmetric but produce a harder integral equation, because the intersection of a sphere with a sphere depends on two lengths rather than one. **Link to the next block.** Now you owe a computation of the surface measure.

Block 3. Cue. A surface integral over a parametrised surface: you need the Jacobian, and for a sphere it is standard. **Why natural.** On the unit sphere, $dS = \sin\varphi \, d\varphi \, d\theta$, and the substitution $t = \cos\varphi$ turns $\sin\varphi \, d\varphi$ into dt exactly. Geometrically this is Archimedes' theorem: the zone of a sphere of radius R cut off between two parallel planes at distance h apart has area $2\pi Rh$, independent of where the planes sit. In other words, the surface measure of the unit sphere pushes forward, under the height map, to 2π times Lebesgue measure on $[-1, 1]$. That is why a function of one coordinate is the family that makes the problem one dimensional with no distortion. **Link to the next block.** With the surface integral reduced, the hypothesis is now a statement about a function of one real variable.

Block 4. Cue. $\int_{-1}^{1} g(z_0 + t)\,dt = 0$ for every z_0: a sliding window with fixed width and zero integral. **Why natural.** Fixed width plus zero integral is a shift equation in disguise. Write G for an antiderivative; the condition is $G(z_0 + 1) = G(z_0 - 1)$, that is, $G(z + 2) = G(z)$. You have converted an integral condition into a periodicity condition, which has an explicit general solution rather than an estimate. **Link to the next block.** Choosing a periodic G with a continuous derivative is now a matter of writing down a sine.

Block 5. Cue. You need one example, and it must be verifiable in three lines by a grader. **Why natural.** $G(z) = -\cos(\pi z)/\pi$ has period 2, so $g = G' = \sin(\pi z)$ works. Do not present the general kernel as the answer; present the single function, verify it completely, and put the general description in a remark. A verified example scores; a described family with a half verification does not.

Layer 3: WHY

The mathematical truth underneath is one sentence: the fixed radius spherical mean is a convolution, and a convolution is injective exactly when its symbol has no real zeros. Precisely, let σ be the normalised surface measure on the unit sphere in $\mathbb{R}^3$ and let $(Mf)(c) = \int f(c + v)\,d\sigma(v)$. This is $f * \check{\sigma}$, so it commutes with translations, so it is diagonalised by the characters $v \mapsto e^{i\xi\cdot v}$. Feeding one in and taking $\xi = (0, 0, \rho)$ without loss of generality by rotational symmetry,

$$M\big(e^{i\xi\cdot v}\big)(c) = e^{i\xi\cdot c} \cdot \frac{1}{2}\int_{-1}^{1} e^{i\rho t}\,dt = e^{i\xi\cdot c}\,\frac{\sin\rho}{\rho},$$

using the same Archimedes pushforward as above. The symbol is $\operatorname{sinc}\rho = \sin\rho/\rho$, and it vanishes at $\rho = \pi, 2\pi, 3\pi, \ldots$. The operator therefore has an enormous kernel, and every element of it is built from frequencies of length exactly $k\pi$. Our $\sin(\pi z)$ is the imaginary part of the $\xi = (0, 0, \pi)$ character. There was no guessing.

The discernment mechanism, the thing you actually train, is this. Whenever a problem says "the average of f over every translate of a fixed shape is zero", you are being asked whether a specific Fourier multiplier has a real zero. Compute the multiplier, which is the Fourier transform of the shape. If it has a real zero, the answer is no and the counterexample is the corresponding plane wave, restricted to real values by taking a real or imaginary part. If it has none, the answer is yes and you must supply an inversion argument. The setter's nudge is the fixed size: fixed size means one value of the symbol's argument, and one value can easily be a zero, whereas all sizes at once means the whole symbol must vanish identically, which forces $f = 0$.

To train it, build a small table and know it cold: in $\mathbb{R}^3$ the unit sphere has symbol $\sin\rho/\rho$ with zeros at multiples of π; the unit ball has symbol proportional to $(\sin\rho - \rho\cos\rho)/\rho^3$ with zeros at the solutions of $\tan\rho = \rho$; in $\mathbb{R}$ the interval $[-1, 1]$ has symbol $\sin\rho/\rho$ again. In $\mathbb{R}^n$ the sphere's symbol is a Bessel function of order $\frac{n}{2} - 1$, which has infinitely many positive zeros for every n, so the answer to this problem is no in every dimension. Knowing that the answer is dimension independent is what lets you commit to "no" in the first thirty seconds without any computation at all.

Revolving concepts

Spherical mean. For f continuous on $\mathbb{R}^n$ and $r > 0$, the spherical mean is $(M_r f)(c) = \frac{1}{|S^{n-1}|}\int_{S^{n-1}} f(c + r\omega)\,d\omega$. Orbiting facts: $M_r f \to f$ uniformly on compacts as $r \to 0^+$ when f is continuous, which is the source of the "yes" answer when all radii are allowed; M_r commutes with translations and rotations; $u(c, r) = (M_r f)(c)$ satisfies the Darboux equation, so spherical means are the natural language for the wave equation in odd dimensions. See Book 1, field: Real analysis, topic: spherical means.

Archimedes' hat box theorem. The area of the zone of a sphere of radius R between two parallel planes a distance h apart (both meeting the ball) is $2\pi Rh$. Equivalently, the radial projection from the sphere onto the circumscribed cylinder preserves area, because it stretches in one direction by $1/\sin\varphi$ and compresses in the other by $\sin\varphi$. In our computation this is the identity $dS = \sin\varphi\, d\varphi\, d\theta \mapsto d\theta\, dt$. See Book 1, field: Geometry, topic: surfaces of revolution.

Fourier transform of a measure, and the symbol. For a finite measure μ on $\mathbb{R}^n$, $\widehat{\mu}(\xi) = \int e^{-i\xi\cdot v} d\mu(v)$, and the operator $f \mapsto f * \mu$ acts on the character $e^{i\xi\cdot v}$ by multiplication by $\widehat{\mu}(-\xi)$. Injectivity on nice function spaces is controlled by the zero set of $\widehat{\mu}$. For the normalised surface measure on the unit sphere in $\mathbb{R}^3$, $\widehat{\sigma}(\xi) = \sin|\xi|/|\xi|$; in $\mathbb{R}^n$ it is $\Gamma(\frac{n}{2})(2/|\xi|)^{n/2-1} J_{n/2-1}(|\xi|)$. See Book 1, field: Real analysis, topic: Fourier multipliers.

The Pompeiu problem and the two radius theorem. The general question "for which sets K does $\int_{K+c} f = 0$ for all c force $f \equiv 0$?" is the Pompeiu problem; balls and spheres fail it, as here. Delsarte's two radius theorem says that if the means over spheres of two radii r_1 and r_2 both vanish, then $f \equiv 0$ unless r_1/r_2 is a ratio of two zeros of the relevant Bessel function. That is the precise statement of "one radius is not enough, but almost two are". See Book 1, field: Real analysis, topic: Pompeiu problem.

The recursion hiding in the problem, generalised. The reduced condition $G(z+2) = G(z)$ is a linear shift equation, the continuous analogue of a recursion with step 2. The general statement worth installing: the equation $G(z+T) = G(z) + c$ for all real z has general solution $G(z) = cz/T + \Pi(z)$ with Π arbitrary of period T, and the corresponding condition on $g = G'$ is that g is periodic of period T with mean c/T over a period. Every "the integral over every window of fixed width T equals c" hypothesis, in any problem, collapses to exactly this. The discrete twin is $u_{n+T} = u_n + c$, whose general solution is $u_n = cn/T$ plus a T periodic sequence. See Book 1, field: Functional equations, topic: shift equations and periodicity.

Circuit Bank

- **WHEN** a hypothesis says the average of an unknown function over every translate of a *fixed* shape vanishes **DO** compute the Fourier transform of that shape and look for real zeros; if there is one, the answer is "no" and the counterexample is the matching plane wave **BECAUSE** convolution operators are diagonalised by characters, so the kernel is exactly the frequencies at which the symbol vanishes

- **WHEN** a "Must f be identically zero?" question carries a restriction whose removal would make the answer an easy yes **DO** stop trying to prove it and start constructing a counterexample **BECAUSE** the restriction is where the setter put all the content, and the removed hypothesis is precisely the tool your first instinct wanted to use

- **WHEN** you must integrate a function of one coordinate over a sphere **DO** slice by planes perpendicular to that coordinate and use that the unit sphere's surface measure pushes forward to $2\pi\, dt$ on $[-1,1]$ **BECAUSE** Archimedes: the area of a spherical zone depends only on its height, so the Jacobian $\sin\varphi$ is exactly consumed by the substitution $t = \cos\varphi$

- **WHEN** you are searching for a counterexample and the hypothesis is invariant under a group **DO** restrict the search to functions with a matching partial symmetry, choosing the one that makes the integral one dimensional **BECAUSE** the constraint then becomes a solvable equation in one variable instead of an estimate in three

Backtesting

1. First useful cue: the radius is pinned at 1 instead of being quantified over all positive values.

2. Method the expression suggested: treat the hypothesis as "a fixed convolution operator kills f", so test characters, find the symbol $\sin\rho/\rho$, and take the frequency $\rho = \pi$; write it up as a slicing computation with $f = \sin(\pi z)$.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at the start, wanting to believe that averaging over spheres determines a continuous function. Reinforce this: averaging at *all* scales determines it; averaging at *one* scale is a single Fourier multiplier and multipliers vanish somewhere all the time.
4. The trap the question sets: the phrase "must f be identically 0" invites you to spend the session proving a false statement, and the mean value property is the plausible looking tool that does not apply.
5. What makes you recognise it faster next time: the pattern "fixed size averaging" should now be a single unit in your head, tagged with the symbol table (sphere in $\mathbb{R}^3$: $\sin\rho/\rho$; ball in $\mathbb{R}^3$: $(\sin\rho - \rho\cos\rho)/\rho^3$; interval in $\mathbb{R}$: $\sin\rho/\rho$).
6. The reusable decision rule: when a hypothesis is a convolution against a fixed shape, the answer to "is f forced to be zero" is exactly "does the shape's Fourier transform avoid zero on the reals".

Mirror drills

1. Find a nonzero continuous f on $\mathbb{R}^3$ whose integral over every *ball* of radius 1 is 0. *Answer:* $f(x,y,z) = \cos(\rho z)$ where $\rho > 0$ satisfies $\tan\rho = \rho$ (the first such ρ is about 4.4934). Slicing gives $\int_{-1}^{1} \pi(1-t^2)\cos(\rho(z_0+t))\,dt$, the odd part vanishes, and $\int_{-1}^{1}(1-t^2)\cos(\rho t)\,dt = 4(\sin\rho - \rho\cos\rho)/\rho^3 = 0$.
2. Show that if the surface integral of a continuous f over every sphere of every radius vanishes, then $f \equiv 0$. *Answer:* the average over the sphere of radius r about c is 0 for all $r > 0$; by continuity of f at c it tends to $f(c)$ as $r \to 0^+$, so $f(c) = 0$.
3. Determine all continuous $g : \mathbb{R} \to \mathbb{R}$ with $\int_{x-1}^{x+1} g(t)\,dt = 0$ for every x. *Answer:* exactly the continuous g of period 2 with mean zero over a period; take an antiderivative G, note the condition is $G(x+1) = G(x-1)$, so G has period 2, hence so does $g = G'$, and $\int_0^2 g = G(2) - G(0) = 0$.

7.5 A5. Euler operator reads the vanishing order

Let p be an odd prime number, and let $\mathbb{F}_p$ denote the field of integers modulo p. Let $\mathbb{F}_p[x]$ be the ring of polynomials over $\mathbb{F}_p$, and let $q(x) \in \mathbb{F}_p[x]$ be given by

$$q(x) = \sum_{k=1}^{p-1} a_k x^k,$$

where

$$a_k = k^{(p-1)/2} \bmod p.$$

Find the greatest nonnegative integer n such that $(x-1)^n$ divides $q(x)$ in $\mathbb{F}_p[x]$.

Read it out loud

Read the structure, not the symbols. There is a polynomial over a finite field whose coefficient of x^k is a fixed power of the index k, and you are asked for the multiplicity of the root $x = 1$. Three fields are wired in: number theory (the exponent $(p-1)/2$ is Euler's criterion, so a_k is the Legendre symbol of k), polynomial algebra in characteristic p (the root 1, the formal derivative, the fact that $x^p - 1$ is a perfect pth power), and generating functions (a coefficient that is a function of the index is a signal, not a coincidence).

The first useful cue is the shape of the coefficient. When the coefficient of x^k is k^m, you are never looking at an arbitrary polynomial: you are looking at the Euler operator $D = x\frac{d}{dx}$ applied m times to something simpler, because D acts on the monomial basis diagonally, $D(x^k) = kx^k$. That is the entire content of "sum with k^m in front". Within thirty seconds you should have written

$$D = x\frac{d}{dx}, \qquad D^m\Big(\sum_k c_k x^k\Big) = \sum_k k^m c_k x^k,$$

and then asked which simple polynomial has all coefficients equal to 1 for $k = 0, \ldots, p-1$. The answer is $s(x) = 1 + x + \cdots + x^{p-1}$, so $q = D^{(p-1)/2}s$, with the $k = 0$ term dying because $0^{(p-1)/2} = 0$.

The second cue is the modulus you are asked about: $(x-1)^n$. In $\mathbb{F}_p[x]$ the polynomial $x^p - 1$ is not a product of p distinct cyclotomic factors, it is $(x-1)^p$, because the Frobenius map makes $(x-1)^p = x^p - 1$. So $s(x) = (x^p - 1)/(x-1) = (x-1)^{p-1}$, and the object you are differentiating is already a pure power of exactly the quantity the question is asking about. That is not luck; the question was designed around it.

So the problem is: start at $(x-1)^{p-1}$, apply D a total of $(p-1)/2$ times, and track the order of vanishing at 1. If each application of D costs exactly one factor of $(x-1)$, the answer is $p - 1 - \frac{p-1}{2} = \frac{p-1}{2}$. Verify that on the smallest case before you commit: for $p = 3$, $q = x + 2x^2 = 2x(x-1)$, so $n = 1$; for $p = 5$, $q = x + 4x^2 + 4x^3 + x^4$, and $q(1) = 0$, $q'(1) = 0$, $q''(1) = 4 \neq 0$, so $n = 2$. Both match $(p-1)/2$.

Approaches, ranked worst to best

1. **Compute small cases and guess.** *Cue:* "find the greatest nonnegative integer n" wants a formula in p. *Circuit:* tabulate, guess, then prove. As a solution this is worth nothing, because a Putnam grader pays for the proof and not the pattern. As reconnaissance it is compulsory: two minutes on $p = 3, 5, 7$ fixes the target $(p-1)/2$ and stops you from proving the wrong thing. Rank it last as a solution, first as a habit.

2. **Split into quadratic residues and nonresidues.** *Cue:* a_k is the Legendre symbol by Euler's criterion. *Circuit:* when you see a Legendre symbol, reach for Gauss sums and character

theory. You would write $q(x) = \sum_{\text{QR}} x^k - \sum_{\text{NQR}} x^k$ and then try to evaluate its behaviour near $x = 1$. Gauss sums live in $\mathbb{Z}[\zeta_p]$ or in an extension of $\mathbb{F}_p$, and they compute values at roots of unity, not orders of vanishing at 1. In $\mathbb{F}_p$ all pth roots of unity have already collapsed to 1, so the entire machinery you are reaching for has degenerated. It stalls, and expensively.

3. **Substitute $x = 1 + y$ and expand blindly.** *Cue:* "divides by $(x-1)^n$" means "lowest surviving power of y". *Circuit:* translate the root to the origin so that divisibility becomes a statement about low order coefficients. Correct instinct, correct first move, but if you expand $(1+y)^k$ and stare at $\sum_k a_k \binom{k}{j}$ without a tool for evaluating power sums modulo p, you stall. Costs a page of binomial identities.

4. **Substitute $x = 1 + y$ and finish with the power sum lemma.** *Cue:* the same cue, plus the sum $\sum_{k=1}^{p-1} k^t$ appearing after you expand $\binom{k}{j}$ as a polynomial in k. *Circuit:* reduce every coefficient to power sums modulo p, which are all known. This is a complete and honest solution. The coefficient of y^j is $\sum_{k=1}^{p-1} a_k \binom{k}{j}$; since $\binom{k}{j}$ is a polynomial in k of degree j with leading coefficient $1/j!$ (invertible for $j \le p-1$), that coefficient is a combination of $S_t = \sum_{k=1}^{p-1} k^t$ for $t = \frac{p-1}{2}, \dots, \frac{p-1}{2} + j$. Since $S_t = -1$ when $(p-1) \mid t$ and 0 otherwise, everything vanishes until $j = \frac{p-1}{2}$, where the leading term contributes $S_{p-1}/j! = -1/j! \neq 0$. It works. It just costs three times the ink of the next route.

5. **Euler operator on the collapsed cyclotomic polynomial.** *Cue:* the coefficient is $k^{(p-1)/2}$, a power of the index, and the modulus is a power of $(x-1)$. *Circuit:* coefficients depending on the index means an operator applied to a simpler series; divisibility by $(x-1)^n$ in characteristic p means use Frobenius to rewrite everything as powers of $(x-1)$. Identify $q = D^{(p-1)/2} s$ with $s = (x-1)^{p-1}$, prove one short lemma saying D lowers the order of vanishing at 1 by exactly one, and iterate. Shortest, and it explains why the answer is $(p-1)/2$ rather than merely confirming it. This is the write up.

The ideal solution

The answer is $\boxed{n = \dfrac{p-1}{2}}$.

Throughout, D denotes the Euler operator $Du = xu'$ on $\mathbb{F}_p[x]$, where u' is the formal derivative. For a nonzero $u \in \mathbb{F}_p[x]$ write $\operatorname{ord}(u)$ for the largest n with $(x-1)^n \mid u$.

Lemma 7.5.1. *Let $u \in \mathbb{F}_p[x]$ be nonzero with $\operatorname{ord}(u) = n$, and suppose $1 \le n \le p-1$. Then $Du \neq 0$ and $\operatorname{ord}(Du) = n - 1$.*

Proof. Write $u = (x-1)^n v$ with $v \in \mathbb{F}_p[x]$ and $v(1) \neq 0$. The product rule holds for formal derivatives over any commutative ring, so

$$u' = n(x-1)^{n-1} v + (x-1)^n v', \qquad Du = xu' = (x-1)^{n-1} \big[\, n\, x\, v(x) + x(x-1) v'(x) \,\big].$$

Call the bracket $w(x)$. Then $w(1) = n\, v(1)$. Since $1 \le n \le p-1$, the residue n is nonzero in $\mathbb{F}_p$, and $v(1) \neq 0$, so $w(1) \neq 0$. Hence $Du = (x-1)^{n-1} w$ with $w(1) \neq 0$, which is exactly the statement that Du is nonzero of order $n-1$. □

Lemma 7.5.2. *In $\mathbb{F}_p[x]$, $s(x) := 1 + x + \cdots + x^{p-1} = (x-1)^{p-1}$.*

Proof. The Frobenius endomorphism of $\mathbb{F}_p[x]$ gives $(x-1)^p = x^p + (-1)^p = x^p - 1$, using that p is odd. Also $(x-1)s(x) = x^p - 1$ by telescoping. Therefore $(x-1)s(x) = (x-1)^p$, and since $\mathbb{F}_p[x]$ is an integral domain and $x - 1 \neq 0$, we may cancel to get $s(x) = (x-1)^{p-1}$. □

Now identify q. Since $D(x^k) = kx^k$, we have $D^m(x^k) = k^m x^k$ for every $m \geq 0$, and therefore

$$D^m s(x) = \sum_{k=0}^{p-1} k^m x^k.$$

Take $m = \frac{p-1}{2}$, which is at least 1 because p is an odd prime. The $k = 0$ term is $0^{(p-1)/2} = 0$, so it drops out and

$$D^{(p-1)/2} s(x) = \sum_{k=1}^{p-1} k^{(p-1)/2} x^k = q(x).$$

Finally, iterate the first lemma along the chain $s,\ Ds,\ D^2 s,\ \ldots,\ D^{(p-1)/2} s$. By the second lemma $\operatorname{ord}(s) = p - 1$. Suppose that for some j with $0 \leq j \leq \frac{p-1}{2} - 1$ we have $\operatorname{ord}(D^j s) = p - 1 - j$. Then

$$\frac{p-1}{2} + 1 \ \leq\ p - 1 - j \ \leq\ p - 1,$$

so the order lies in the range $1 \leq n \leq p - 1$ required by the first lemma, which therefore applies and gives $\operatorname{ord}(D^{j+1} s) = p - 2 - j$. By induction on j,

$$\operatorname{ord}\big(D^{(p-1)/2} s\big) = p - 1 - \frac{p-1}{2} = \frac{p-1}{2}.$$

That is, $\operatorname{ord}(q) = \frac{p-1}{2}$: the polynomial $(x-1)^{(p-1)/2}$ divides q, and $(x-1)^{(p-1)/2+1}$ does not. Note that the chain also shows $q \neq 0$, so the order is well defined.

Remark 7.5.3. The hypothesis $1 \leq n \leq p - 1$ in the first lemma is not decoration. If $p \mid n$ the conclusion is false: $u = (x-1)^p = x^p - 1$ has $u' = px^{p-1} = 0$, so $Du = 0$ and the order jumps to infinity rather than dropping by one. This is the standard failure of the derivative test for multiple roots in characteristic p, and it is exactly why the chain above was checked to stay inside the range 1 to $p - 1$ at every step.

Layer 1: WHAT

Block 1. Recognise $a_k = k^{(p-1)/2}$ as the Legendre symbol by Euler's criterion, note $a_k = \pm 1$ for every k in range, and do the two minute reconnaissance on $p = 3$ and $p = 5$ to fix the target answer $(p-1)/2$.

Block 2. Read the coefficient shape: coefficient $= k^m$ means $q = D^m$(something with all coefficients 1), with $D = x\frac{d}{dx}$.

Block 3. Simplify the something: $s = 1 + x + \cdots + x^{p-1} = (x-1)^{p-1}$ in $\mathbb{F}_p[x]$, by Frobenius. Now both the object and the question are expressed in powers of $(x-1)$.

Block 4. Prove the single lemma that carries all the weight: D lowers the order of vanishing at 1 by exactly one, provided the current order is between 1 and $p - 1$.

Block 5. Iterate $(p-1)/2$ times, checking at each step that the order stays in range, and read off $(p-1)/2$.

The nominal route splits from the ideal route at Block 2. The zombie reads $k^{(p-1)/2}$ as an arithmetic object, says "Legendre symbol", and reaches for quadratic residues, Gauss sums, and character sums. The ideal route reads $k^{(p-1)/2}$ as an *operator eigenvalue* and reaches for a generating function. That is the bifurcation point, and it is worth naming precisely: the question is not asking anything about which k are squares, it is asking about the $(x-1)$ adic order of a polynomial. The arithmetic content of a_k is a decoy; only the fact that a_k is a power of k matters.

Layer 2: HOW

Block 1. **Cue.** The exponent $(p-1)/2$ attached to k modulo an odd prime. **Why natural.** Euler's criterion is atomic: $k^{(p-1)/2} \equiv \pm 1$ according to whether k is a quadratic residue. Recognising it tells you the coefficients are units, hence $\deg q = p-1$ exactly and $q \neq 0$, which you need before you can speak of an order of vanishing at all. It also tells you $q(1) = \sum_k a_k = 0$, since there are equally many residues and nonresidues, so $n \geq 1$ before you have done any work. **Link to the next block.** You now know the answer is at least 1 and you have small cases; the question is what structure produces the exact value.

Block 2. **Cue.** The coefficient of x^k is a function of k alone, namely a power of k. **Why natural.** This is the generating function circuit in its purest form. On monomials, $x\frac{d}{dx}$ is diagonal with eigenvalue k, so multiplying the kth coefficient by k is not a manipulation of the coefficients at all, it is the application of a single operator to the whole polynomial. Whenever the coefficient is $P(k)$ for a polynomial P, the series is $P(D)$ applied to the plain series. Formally, this is the same circuit as evaluating $\sum k^2 z^k$ by differentiating the geometric series twice and multiplying by z each time; the only new content is that it is legal over $\mathbb{F}_p$ because the formal derivative is defined over any ring. **Link to the next block.** The operator has isolated a simpler polynomial, namely the all ones polynomial of degree $p-1$; simplify it next.

Block 3. **Cue.** A sum of all powers x^0 to x^{p-1} over $\mathbb{F}_p$, and a question about divisibility by $(x-1)^n$. **Why natural.** Two facts collide productively. First, $(x-1)s(x) = x^p - 1$ by telescoping. Second, in characteristic p the Frobenius map is a ring homomorphism, so $(x-1)^p = x^p - 1$. Hence $s = (x-1)^{p-1}$. The structural reading: over $\mathbb{Q}$ the polynomial $x^p - 1$ has p distinct roots and s is the irreducible pth cyclotomic polynomial, but over $\mathbb{F}_p$ all p roots of unity have collapsed onto the single point 1, so s is totally degenerate there. Any question about $\sum_k c_k x^k$ modulo powers of $(x-1)$ in $\mathbb{F}_p[x]$ is therefore a question about a single filtration, and every tool that separates roots of unity is useless. **Link to the next block.** Everything is now a power of $(x-1)$ acted on by D; you need to know what D does to that power.

Block 4. **Cue.** You are applying one operator repeatedly to an object measured by one integer, the order of vanishing. **Why natural.** The natural thing to compute is the effect of one application on that integer. Write $u = (x-1)^n v$ with $v(1) \neq 0$, apply the product rule, and factor out $(x-1)^{n-1}$: the leftover evaluated at 1 is $n\,v(1)$. It is nonzero precisely when n is invertible modulo p. This is the honest statement of the derivative test for multiplicity in characteristic p, and the caveat is real, not pedantic: $x^p - 1$ has zero derivative. **Link to the next block.** One step is now understood; the chain is an induction, and all that remains is checking the hypothesis holds at every step.

Block 5. **Cue.** A repeated application with a side condition means an induction with a range check. **Why natural.** Starting at order $p-1$ and dropping by one each time, the orders visited are $p-1, p-2, \ldots, \frac{p-1}{2}$, all strictly between 0 and p, so the lemma's hypothesis holds at every step. Do not skip the range check in the write up: it is the only place the argument could break, and a grader looking for the characteristic p subtlety will look exactly there.

Layer 3: WHY

Two mathematical truths make this route the ideal one.

The first is that in $\mathbb{F}_p[x]$ the ideal $(x-1)$ carries a discrete valuation, the order of vanishing at 1, and everything relevant in this problem is a statement in that valuation. Frobenius says $x^p - 1 = (x-1)^p$, so the all ones polynomial, the object that is spiritually "sum over the pth roots of unity", is a pure power of the uniformiser. In other words the arithmetic of pth roots of unity, which

over $\mathbb{Q}$ is a rich cyclotomic story, has degenerated in characteristic p to a single ramified point. Once you know that, the choice of tool is forced: use the valuation, not the roots.

The second is that $D = x\frac{d}{dx}$ is the operator whose eigenbasis is the monomials, with eigenvalue equal to the exponent. This makes D the exact translation between two languages: "do arithmetic to the index k" and "apply an operator to the polynomial". The coefficient sequence $k \mapsto k^{(p-1)/2}$ was constructed by the setter to be an eigenvalue sequence, which is precisely why the problem has a clean answer. The interaction between the two truths is the lemma: D is a derivation multiplied by x, and x is a unit at the point $x = 1$, so D shifts the valuation down by exactly one, with the single degeneracy at orders divisible by p.

The discernment mechanism is a two question filter you should run on any problem of this shape. Question one: does the coefficient of x^k depend on k through a polynomial or an exponential? If yes, write the series as an operator applied to a simpler series ($P(D)$ for polynomial dependence, a substitution $x \mapsto \lambda x$ for exponential dependence). Question two: does the question ask about divisibility by a power of a linear factor? If yes, rewrite every object as a power of that linear factor first, using Frobenius if you are in characteristic p. If both answers are yes, the solution is the composition, and it is short. Here both are yes, which is why the ideal solution fits on half a page.

To train the discernment, do the reverse exercise. Take the answer $(p-1)/2$ and ask what the setter had to do to build a problem with that answer. They needed an object of known order $p-1$ (the collapsed cyclotomic polynomial), an operator that peels off exactly one order per application (D), and a number of applications equal to $(p-1)/2$ (the Euler criterion exponent, which also disguises the operator as arithmetic). Reading a problem backwards from its construction is how you learn to see the construction forwards, in the arena, under time.

Revolving concepts

Euler's criterion and the Legendre symbol. For odd prime p and k not divisible by p, $k^{(p-1)/2} \equiv \left(\frac{k}{p}\right) \pmod p$, which is $+1$ if k is a nonzero quadratic residue and -1 otherwise. Orbiting facts: the residues form the index two subgroup of $\mathbb{F}_p^\times$, so there are exactly $(p-1)/2$ of each and the sum of all Legendre symbols is 0; the symbol is completely multiplicative. See Book 1, field: Number theory, topic: quadratic residues.

Frobenius and the freshman's dream. In any commutative ring of characteristic p, the map $a \mapsto a^p$ is a ring homomorphism, so $(a+b)^p = a^p + b^p$. Consequences used here: $(x-1)^p = x^p - 1$ for odd p, and more generally $x^{p^e} - c^{p^e} = (x-c)^{p^e}$. See Book 1, field: Abstract algebra, topic: characteristic p and Frobenius.

Formal derivative and multiplicity in characteristic p. The formal derivative is defined over any commutative ring and satisfies linearity and the product rule. Over a field of characteristic 0, c is a root of u of multiplicity exactly n if and only if $u(c) = \cdots = u^{(n-1)}(c) = 0 \neq u^{(n)}(c)$. In characteristic p this equivalence survives only for $n \leq p-1$, because $u^{(n)}$ carries a factor $n!$; the standard counterexample is $x^p - 1$, whose derivative is identically zero. See Book 1, field: Polynomials, topic: repeated roots and the derivative test.

Order of vanishing as a valuation. For a field K and $c \in K$, ord_c on $K[x]$ satisfies $\mathrm{ord}_c(uv) = \mathrm{ord}_c(u) + \mathrm{ord}_c(v)$ and $\mathrm{ord}_c(u+v) \geq \min$. Divisibility questions by $(x-c)^n$ are valuation questions, and the right move is always to find one operation whose effect on the valuation you can compute exactly. See Book 1, field: Polynomials, topic: valuations and factorisation.

Power sums modulo p. For $t \geq 0$, $\sum_{k=1}^{p-1} k^t \equiv -1 \pmod p$ if $(p-1) \mid t$, and $\equiv 0$ otherwise. Proof: let g be a primitive root, so the sum is the geometric series $\sum_{i=0}^{p-2}(g^t)^i$; if $g^t = 1$ each term is 1 and the sum is $p-1$, otherwise the series sums to $(g^{t(p-1)} - 1)/(g^t - 1) = 0$. This is the engine of the alternative solution and appears constantly in finite field problems. See Book 1, field: Number theory, topic: power sums and primitive roots.

The Euler operator, and the recursion generalised. $D = x\frac{d}{dx}$ satisfies $D(x^k) = kx^k$, hence

$P(D)\sum_k c_k x^k = \sum_k P(k)c_k x^k$ for any polynomial P. Over $\mathbb{R}$ this is how you evaluate $\sum_k k^2 z^k$ from the geometric series. The recursion in this problem is the sequence of orders $n_j = \operatorname{ord}(D^j s)$, which satisfies the linear recursion $n_{j+1} = n_j - 1$ with $n_0 = p-1$, valid exactly while $p \nmid n_j$. Its solution is $n_j = p - 1 - j$. The general form worth installing: if an operator T shifts a valuation by a fixed amount d away from a set of degenerate values Δ, then $\operatorname{ord}(T^j u) = \operatorname{ord}(u) + jd$ for as long as the trajectory avoids Δ, and the entire proof is the induction that checks that avoidance. The same pattern governs orders of vanishing under repeated differentiation, degrees under repeated multiplication, and p adic valuations under repeated division. See Book 1, field: Generating functions, topic: operator calculus on power series.

Circuit Bank

- **WHEN** the coefficient of x^k is $P(k)$ for some polynomial P (here $k^{(p-1)/2}$) **DO** write the series as $P(D)$ applied to the plain series, where $D = x\frac{d}{dx}$ **BECAUSE** D is diagonal on monomials with eigenvalue k, so index arithmetic and operator application are the same thing
- **WHEN** you are in $\mathbb{F}_p[x]$ and the question concerns divisibility by a power of $(x-c)$ **DO** rewrite every object as a power of $(x-c)$ first, using $x^p - c^p = (x-c)^p$ **BECAUSE** Frobenius collapses the pth roots of unity onto one point, so the whole problem lives in a single valuation and no root separating tool can help
- **WHEN** you want to use "least m with the mth derivative nonzero" to find a multiplicity in characteristic p **DO** first check the multiplicity is at most $p-1$, and state that check in the write up **BECAUSE** the mth derivative carries a factor $m!$, which vanishes for $m \geq p$; $x^p - 1$ has zero derivative and infinite apparent order
- **WHEN** a sum $\sum_{k=1}^{p-1} k^t$ appears modulo p **DO** replace it by -1 if $(p-1) \mid t$ and by 0 otherwise **BECAUSE** multiplying the sum by g^t for a primitive root g permutes the terms, so the sum is fixed by multiplication by g^t and must vanish unless $g^t = 1$

Backtesting

1. First useful cue: the coefficient of x^k is a power of the index k, which is the signature of the Euler operator $D = x\frac{d}{dx}$.
2. Method the expression suggested: write $q = D^{(p-1)/2}s$ with $s = 1 + x + \cdots + x^{p-1}$, collapse s to $(x-1)^{p-1}$ by Frobenius, and track the order of vanishing under each application of D.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at "is the derivative test even legal here?". Reinforce the exact statement: in characteristic p, one application of D drops the order by one if and only if the current order is not divisible by p, and here the orders run from $p-1$ down to $(p-1)/2$, all safely in range.
4. The trap the question sets: dressing the coefficient as a Legendre symbol so that you spend the session on quadratic residues and Gauss sums. The arithmetic meaning of a_k is never used; only the fact that it is k raised to a power.
5. What makes you recognise it faster next time: pair the two cues as one unit. "Coefficient is a power of the index" plus "divisibility by a power of $(x-1)$ over $\mathbb{F}_p$" equals "apply D^m to $(x-1)^{p-1}$ and count".
6. The reusable decision rule: to find how many times a linear factor divides a polynomial built by an operator, compute the operator's effect on the order of vanishing once, then iterate with a range check.

Mirror drills

1. In $\mathbb{F}_p[x]$, find the greatest n with $(x-1)^n$ dividing $\sum_{k=1}^{p-1} kx^k$. *Answer:* $n = p-2$. That polynomial is $Ds = x\big((x-1)^{p-1}\big)' = (p-1)x(x-1)^{p-2}$, and $p-1 \neq 0$ in $\mathbb{F}_p$.
2. Show that the test "c is a repeated root of u if and only if $u(c) = u'(c) = 0$" can fail in $\mathbb{F}_p[x]$, and say exactly which direction fails. *Answer:* the direction that detects multiplicity fails only through degeneracy of the derivative: $u = x^p - 1$ has $u' = px^{p-1} = 0$, so u' vanishes everywhere and gives no information; here $u = (x-1)^p$ really does have a repeated root, but the same computation would say nothing for a polynomial such as $x^p - x$ whose roots are all simple.
3. Prove that $\sum_{k=1}^{p-1} k^t \equiv -1 \pmod{p}$ when $(p-1) \mid t$ and $\equiv 0$ otherwise. *Answer:* let g be a primitive root modulo p; the sum equals $\sum_{i=0}^{p-2} (g^t)^i$, which is $p-1 \equiv -1$ if $g^t = 1$, and otherwise is $\big((g^t)^{p-1} - 1\big)/(g^t - 1) = 0$ by Fermat.

7.6 A6. Taylor windows and Landau's interpolation inequality

Let g be a real-valued function that is continuous on the closed interval $[0,1]$ and twice differentiable on the open interval $(0,1)$. Suppose that for some real number $r > 1$,

$$\lim_{x\to 0^+} \frac{g(x)}{x^r} = 0.$$

Prove that either

$$\lim_{x\to 0^+} g'(x) = 0 \qquad \textit{or} \qquad \limsup_{x\to 0^+} x^r |g''(x)| = \infty.$$

Read it out loud

Read the logical shape first, because it dictates everything. The conclusion is a disjunction, "P or Q". You never prove a disjunction by choosing a side; you prove it by assuming the negation of one alternative and deriving the other. Which one do you negate? The one whose negation is a usable object. Negating P gives you a sequence along which g' stays away from 0, which is weak. Negating Q gives you an inequality: $\limsup_{x\to 0^+} x^r|g''(x)| < \infty$ means there are $M > 0$ and $\delta_0 > 0$ with

$$|g''(x)| \le Mx^{-r} \quad \text{for all } x \in (0, \delta_0).$$

That is a hard quantitative bound, and it is the only quantitative object the problem is willing to hand you. So the shape of the work is fixed within the first minute: assume the curvature bound, prove $g' \to 0$.

Now read the analytic content. You are given control of g near 0 (it is $o(x^r)$), you have just assumed control of g'' near 0 (it is $O(x^{-r})$), and you must control g'. A hypothesis on a function and on its second derivative, with the conclusion about the first derivative, is the exact signature of an *interpolation inequality*. The relevant one is Landau's: on a half line, $\|f'\|_\infty^2 \le 4\|f\|_\infty\|f''\|_\infty$. The first derivative is always controlled by the geometric mean of size and curvature, for a reason you can state in one sentence: if f' is large at a point and f'' is bounded, then f' stays large over a whole window, so f must move a lot over that window, and a large movement contradicts f being small.

The third thing to read is the exponent. The same r appears in both hypotheses: g is small like x^r, and g'' is allowed to blow up like x^{-r}. Their product is $x^r \cdot x^{-r} = 1$: scale free. That is not decoration, it is the design of the problem. Since Landau's bound is the geometric mean $\sqrt{\|f\|\,\|f''\|}$, the geometric mean of εx^r and Mx^{-r} is $\sqrt{\varepsilon M}$, with the x dependence cancelling exactly. So the final bound on $|g'(x)|$ will be independent of x and proportional to $\sqrt{\varepsilon}$, and since ε is arbitrary that is precisely the statement $g' \to 0$. Within thirty seconds of reading, you should have written down "negate the limsup", "$|g''| \le Mx^{-r}$", and "Landau: $|g'| \lesssim \sqrt{\varepsilon M}$, scale free".

What remains is the one genuinely technical decision: over which window do you run Taylor's theorem? That is where the hypothesis $r > 1$ lives, and it is the only place it is used.

Approaches, ranked worst to best

1. **Prove the first alternative from the hypothesis alone.** *Cue:* $\lim g(x)/x^r = 0$ looks strong, so surely g' must be small too. *Circuit:* differentiate an asymptotic estimate. This is false and you must know it is false before you start, or you will burn the session. Take $g(x) = x^{r+1}\sin(x^{-N})$ with $N > r$, extended by $g(0) = 0$. Then $g(x)/x^r = x\sin(x^{-N}) \to 0$, but $g'(x) = (r+1)x^r\sin(x^{-N}) - Nx^{r-N}\cos(x^{-N})$, which is unbounded near 0. Asymptotic control of a function never transfers to its derivative without curvature control. Ranked worst.

2. **Prove the two alternatives are mutually exclusive.** *Cue:* the word "either ... or" reads as exclusive in ordinary English. *Circuit:* case analysis. Also false, and cheap to test. With

$g(x) = x^5 \sin(x^{-3})$ and $r = 2$ both alternatives hold simultaneously. Two minutes spent checking that costs nothing; twenty minutes spent trying to prove exclusivity costs the problem.

3. **Mean value theorem on the dyadic window $[x, 2x]$.** *Cue:* you need to connect g, g' and g'', and $[x, 2x]$ is the standard window near a singularity. *Circuit:* compare a derivative at a point with a difference quotient, then correct using the second derivative. The mean value theorem gives $\xi \in (x, 2x)$ with $|g'(\xi)| \le (2^r + 1)\varepsilon x^{r-1}$, which does tend to 0 since $r > 1$. But transporting that bound from ξ to x costs $\int_x^{2x} |g''| \le Mx^{-r} \cdot x = Mx^{1-r}$, which blows up. The route fails, and the failure is diagnostic: the window is too long by a factor of x^{r-1}. Ranked here because the idea is right and only the scale is wrong; fixing the scale is the whole solution.

4. **Contradiction with a sequence and a fixed window.** *Cue:* negating both alternatives gives a sequence $x_n \to 0$ with $|g'(x_n)| > \varepsilon_0$, and a bound on g''. *Circuit:* double negation, then quantitative contradiction. Assume both alternatives fail. On the window $[x_n, x_n + \varepsilon_0 x_n^r/(2M)]$ the curvature bound keeps $|g'| > \varepsilon_0/2$, so by the mean value theorem g moves by more than $\varepsilon_0^2 x_n^r/(4M)$, while the hypothesis says g is at most $\varepsilon(2^r + 1)x_n^r$ in size there. Letting $\varepsilon \to 0$ contradicts $\varepsilon_0 > 0$. This is complete and correct. It costs more ink than it needs to, because of the sequence bookkeeping and the two layers of negation.

5. **Taylor on a window of unspecified length, then optimise.** *Cue:* g and g'' controlled, g' wanted: interpolation. *Circuit:* leave the window length ℓ as a free parameter, bound $|g'(x)| \le A/\ell + B\ell/2$, and minimise over ℓ. Direct, quantitative, no contradiction, and it produces an explicit bound $|g'(x)| \le \sqrt{2(2^r + 1)M\varepsilon}$ valid for all small x. It is Landau's inequality localised at the point x. Shortest and clearest; this is the write up.

The ideal solution

Since the conclusion is a disjunction, it suffices to prove that if the second alternative fails then the first holds. So assume

$$\limsup_{x \to 0^+} x^r |g''(x)| < \infty.$$

By the definition of lim sup there exist a constant $M > 0$ and a $\delta_0 \in (0, 1)$ such that

$$|g''(x)| \le Mx^{-r} \qquad \text{for all } x \in (0, \delta_0). \tag{1}$$

We prove that $\lim_{x \to 0^+} g'(x) = 0$.

Let $\varepsilon > 0$ be arbitrary. Since $g(x)/x^r \to 0$ as $x \to 0^+$, there is $\delta_1 \in (0, \delta_0)$ with

$$|g(x)| \le \varepsilon x^r \qquad \text{for all } x \in (0, \delta_1). \tag{2}$$

Put

$$c = \sqrt{\frac{2(2^r + 1)\varepsilon}{M}} > 0, \qquad \delta = \min\Big\{ \frac{\delta_1}{2},\ c^{-1/(r-1)} \Big\} > 0,$$

where the second entry is finite and positive precisely because $r > 1$. Fix any $x \in (0, \delta)$ and set

$$\ell = c\, x^r.$$

Because $x \le c^{-1/(r-1)}$ we have $cx^{r-1} \le 1$, hence $\ell \le x$; and because $x < \delta_1/2$ we get

$$[x,\ x + \ell] \subseteq [x, 2x] \subset (0, \delta_1) \subset (0, \delta_0) \subset (0, 1).$$

In particular g is differentiable on an open interval containing $[x, x + \ell]$ and g'' exists on $(x, x + \ell)$, so Taylor's theorem with Lagrange remainder applies: there is $\xi \in (x, x + \ell)$ with

$$g(x + \ell) = g(x) + g'(x)\,\ell + \tfrac{1}{2} g''(\xi)\,\ell^2.$$

Solving for $g'(x)$ and taking absolute values,

$$|g'(x)| \;\le\; \frac{|g(x+\ell)|+|g(x)|}{\ell} \;+\; \tfrac{1}{2}|g''(\xi)|\,\ell. \tag{3}$$

Now bound the three pieces. Since $x+\ell \le 2x < \delta_1$, estimate (2) gives $|g(x+\ell)| \le \varepsilon(x+\ell)^r \le \varepsilon(2x)^r = 2^r\varepsilon x^r$ and $|g(x)| \le \varepsilon x^r$. Since $\xi > x$ and $t \mapsto t^{-r}$ is decreasing on $(0,\infty)$, estimate (1) gives $|g''(\xi)| \le M\xi^{-r} \le Mx^{-r}$. Substituting into (3),

$$|g'(x)| \;\le\; \frac{(2^r+1)\varepsilon x^r}{\ell} + \frac{Mx^{-r}\ell}{2} \;=\; \frac{(2^r+1)\varepsilon}{c} + \frac{Mc}{2},$$

where the last equality is exactly the substitution $\ell = cx^r$, and this is the point of choosing ℓ proportional to x^r: both terms become independent of x. With $c = \sqrt{2(2^r+1)\varepsilon/M}$ the two terms are equal, and

$$\frac{(2^r+1)\varepsilon}{c} + \frac{Mc}{2} = 2\sqrt{\frac{(2^r+1)M\varepsilon}{2}} = \sqrt{2(2^r+1)M\varepsilon}.$$

Therefore

$$|g'(x)| \le \sqrt{2(2^r+1)M\varepsilon} \qquad \text{for every } x \in (0,\delta).$$

Here M and r are fixed while $\varepsilon > 0$ was arbitrary, and $\sqrt{2(2^r+1)M\varepsilon} \to 0$ as $\varepsilon \to 0^+$. Hence for every $\eta > 0$ we may choose $\varepsilon = \eta^2/\big(2(2^r+1)M\big)$ and obtain $\delta > 0$ with $|g'(x)| \le \eta$ on $(0,\delta)$. That is $\lim_{x\to 0^+} g'(x) = 0$, which is the first alternative. ■

Remark 7.6.1. The hypothesis $r > 1$ was used exactly once, to guarantee that a window of length proportional to x^r fits inside $[x, 2x]$, so that both estimates (1) and (2) are valid throughout the window with the constants they have at the scale x. Everything else in the argument is Landau's interpolation inequality applied on that window.

Remark 7.6.2. The "or" is inclusive and you should not spend a second trying to make it exclusive. Take $g(x) = x^5\sin(x^{-3})$ on $(0,1]$ with $g(0) = 0$, and $r = 2$. Then $g(x)/x^2 = x^3\sin(x^{-3}) \to 0$; also $g'(x) = 5x^4\sin(x^{-3}) - 3x\cos(x^{-3}) \to 0$, so the first alternative holds; and $g''(x) = (20x^3 - 9x^{-3})\sin(x^{-3}) - 18\cos(x^{-3})$, so $x^2g''(x) = (20x^5 - 9x^{-1})\sin(x^{-3}) - 18x^2\cos(x^{-3})$ is unbounded along $x^{-3} = \frac{\pi}{2} + 2k\pi$, so the second alternative holds as well.

Layer 1: WHAT

Block 1. Convert the disjunction into an implication by negating the alternative whose negation is quantitative: assume $\limsup x^r|g''(x)| < \infty$, extract $|g''(x)| \le Mx^{-r}$ near 0, and aim at $g' \to 0$.

Block 2. Check that this is necessary, that is, that the hypothesis on g alone does not force $g' \to 0$; the example $x^{r+1}\sin(x^{-N})$ settles it. Now you know the curvature bound must appear in the proof, and you know roughly where.

Block 3. Recognise the pattern "control on f and f'', conclusion about f'" as interpolation, and write the only identity that links all three at once: Taylor to second order over a window $[x, x+\ell]$, solved for $g'(x)$.

Block 4. Do not choose ℓ. Leave it free, obtain $|g'(x)| \le A/\ell + B\ell/2$ with $A = (2^r+1)\varepsilon x^r$ and $B = Mx^{-r}$, and only then minimise, which gives $\ell \propto x^r$ and the scale free bound $\sqrt{2AB} = \sqrt{2(2^r+1)M\varepsilon}$.

Block 5. Check the window fits ($\ell \le x$, which is where $r > 1$ is spent), and let $\varepsilon \to 0$.

The nominal route and the ideal route split at Block 4. The zombie picks a window because it looks natural, almost always $[x, 2x]$ or $[0, x]$, computes with it, gets Mx^{1-r} blowing up, and either gives up or starts patching with sequences. The ideal route refuses to name the window until the algebra tells it which window to name. *That* is the bifurcation point: introducing ℓ as an unknown rather than guessing it. It is the same discipline as not choosing a substitution until you see which one clears the denominator, and it is worth stating as a law: when the estimate contains a free scale, carry the scale as a symbol and optimise at the end.

Layer 2: HOW

Block 1. **Cue.** The conclusion is "either P or Q", and Q is a statement that something is infinite. **Why natural.** A disjunction is proved as "not Q implies P", and you pick the side to negate by asking which negation is usable. "$\limsup = \infty$" is a statement of unboundedness; its negation is boundedness, which is an inequality with a constant, and inequalities with constants are the raw material of analysis. Negating P instead gives you only a sequence along which $|g'|$ stays above some ε_0, which is a much thinner object to work with. **Link to the next block.** You now hold two bounds near 0, one on g and one on g'', and a target concerning g'.

Block 2. **Cue.** The problem offers you two hypotheses; before using both, check whether one suffices. **Why natural.** If the hypothesis on g alone forced $g' \to 0$, the second alternative would be redundant and the problem would be misstated. It does not: $g(x) = x^{r+1}\sin(x^{-N})$ with $N > r$ is $o(x^r)$ with an unbounded derivative. Running this check costs two minutes and buys certainty about what the proof must contain, namely an essential use of the curvature bound. On a six mark boundary problem, knowing what the proof must contain is worth more than any single estimate. **Link to the next block.** The proof must link g, g' and g'' simultaneously, so you need an identity containing all three.

Block 3. **Cue.** Control on a function and on its second derivative, with the first derivative as the target. **Why natural.** There is exactly one elementary identity that contains a function, its derivative at a point, and its second derivative: Taylor's theorem to second order with Lagrange remainder, $g(x+\ell) = g(x) + g'(x)\ell + \frac{1}{2}g''(\xi)\ell^2$. Solving it for $g'(x)$ is not a manipulation, it is the definition of what interpolation means here: the derivative equals a difference quotient (controlled by the size of g) minus a curvature correction (controlled by g''). Landau's inequality is nothing but this identity plus an optimisation. **Link to the next block.** The two error terms move in opposite directions as the window grows, so the window length is a genuine degree of freedom.

Block 4. **Cue.** An upper bound of the form $A/\ell + B\ell/2$ with ℓ at your disposal. **Why natural.** Short windows make the difference quotient blow up; long windows make the curvature term blow up. Any bound of the shape $A/\ell + B\ell/2$ is minimised where the two terms are equal, at $\ell = \sqrt{2A/B}$, with minimum value $\sqrt{2AB}$: that is the arithmetic mean to geometric mean inequality, nothing more. Here $A = (2^r+1)\varepsilon x^r$ and $B = Mx^{-r}$, so $AB = (2^r+1)\varepsilon M$ is free of x, which is the design of the problem showing itself. Note the shape of the answer before you compute it: a constant times $\sqrt{\varepsilon}$, uniform in x. **Link to the next block.** Optimising has told you the window must have length proportional to x^r, so now you must check such a window is available.

Block 5. **Cue.** A window $[x, x+\ell]$ with $\ell = cx^r$, next to the singular point 0. **Why natural.** Both of your estimates were stated with constants valid at the scale x: $|g''| \le Mx^{-r}$ needs ξ to be comparable to x, and $|g(x+\ell)| \le \varepsilon(x+\ell)^r$ needs $x+\ell$ still inside the region where (2) holds. Requiring $\ell \le x$ delivers both at once, and $\ell = cx^r \le x$ is exactly $cx^{r-1} \le 1$, which is available for small x if and only if $r > 1$. Say this explicitly in the write up; it is the only line where the hypothesis $r > 1$ is consumed, and a grader will look for it.

Layer 3: WHY

The mathematical truth beneath the route is Landau's interpolation inequality. In its cleanest form: if f is twice differentiable on $[a, \infty)$ with $|f| \le A$ and $|f''| \le B$ throughout, then $|f'| \le 2\sqrt{AB}$ throughout. The proof is one line of Taylor plus one line of optimisation, exactly as above. The content is a statement about incompatible behaviours: a function cannot be simultaneously small, gently curved, and steep. Steepness at a point plus bounded curvature forces steepness on a whole

window of length proportional to $|f'|/B$; steepness over that window forces a displacement of order $|f'|^2/B$; and smallness caps the displacement at $2A$. Rearranging gives $|f'|^2 \le 2AB$ up to constants. Every proof of this problem, including the official ones, is that sentence in some packaging.

The second truth is scale invariance, and it is the discernment mechanism. The hypotheses are $|g| \le \varepsilon x^r$ and $|g''| \le Mx^{-r}$; their geometric mean is $\sqrt{\varepsilon M}$, free of x. When a problem's two hypotheses carry exponents whose product is scale free, the sharp mechanism is a geometric mean inequality and the conclusion will be uniform in the scale. That observation tells you the answer before you compute: $|g'|$ is bounded by a constant multiple of $\sqrt{\varepsilon}$ on a whole punctured neighbourhood, not merely along a sequence, and since ε is arbitrary the limit is 0. Conversely, if you compute and find an x dependent bound, you have chosen the wrong window and should go back to Block 4.

How do you train this? Two exercises, both cheap. First, dimensional analysis: give x the dimension of length and ask what length the problem provides. Here the only lengths available near the point x are x itself and x^r (through the ratio $\sqrt{A/B} = x^r\sqrt{\varepsilon/M}$), and since $r > 1$ these are genuinely different scales, so the problem is a two scale problem and the small scale is the one you must work at. Second, always leave the scale symbolic. Students lose this problem not because they cannot do Taylor but because they commit to a window in the first line and never revisit it.

What were the examiners nudging toward? Three nudges, all deliberate. The lim sup rather than lim in the second alternative is a nudge to negate that side, since only the lim sup has a clean finite negation. The identical exponent r on both sides is the scale invariance nudge. And the condition $r > 1$, which looks like a technical hypothesis, is the nudge that the window must be shorter than the distance to the singularity. Three cues, three decisions, one page of proof.

Revolving concepts

Limit superior. For h defined near 0^+, $\limsup_{x\to 0^+} h(x) = \lim_{\delta\to 0^+} \sup_{0<x<\delta} h(x)$, a limit of a nonincreasing function of δ, hence always defined in $[-\infty, +\infty]$. The negation used here: $\limsup_{x\to 0^+} h(x) < \infty$ if and only if there are M and $\delta_0 > 0$ with $h(x) \le M$ on $(0, \delta_0)$. Also $\limsup h = \infty$ if and only if there is a sequence $x_n \to 0^+$ with $h(x_n) \to \infty$. See Book 1, field: Real analysis, topic: limsup and liminf.

Taylor's theorem with Lagrange remainder. If f is n times differentiable on an open interval containing $[a, b]$, with $f^{(n-1)}$ continuous on $[a, b]$ and $f^{(n)}$ existing on (a, b), then there is $\xi \in (a, b)$ with $f(b) = \sum_{k=0}^{n-1} \frac{f^{(k)}(a)}{k!}(b-a)^k + \frac{f^{(n)}(\xi)}{n!}(b-a)^n$. The case $n = 1$ is the mean value theorem. The hypotheses matter here: we only know g is twice differentiable on the open interval $(0, 1)$, which is why the window is taken strictly inside. See Book 1, field: Real analysis, topic: Taylor's theorem.

Landau and Kolmogorov interpolation inequalities. Landau: on a half line, $\|f'\|_\infty^2 \le 4\|f\|_\infty\|f''\|_\infty$; on the whole line the constant improves to 2. Kolmogorov's general form: $\|f^{(k)}\|_\infty \le C_{k,n}\|f\|_\infty^{1-k/n}\|f^{(n)}\|_\infty^{k/n}$ for $0 < k < n$. The intermediate derivative is always controlled by a geometric mean of the two extremes, with weights given by the exponents. See Book 1, field: Real analysis, topic: interpolation inequalities.

The recursion hiding in interpolation, generalised. Apply Landau's inequality to $f^{(k-1)}$ instead of f: it gives $\|f^{(k)}\|_\infty^2 \le 4\|f^{(k-1)}\|_\infty\|f^{(k+1)}\|_\infty$ for every $k \ge 1$. Writing $L_k = \log\|f^{(k)}\|_\infty$, this is the recursion $2L_k \le L_{k-1} + L_{k+1} + \log 4$, that is, the sequence of logarithms of derivative norms is midpoint convex up to an additive constant. Iterating a midpoint convexity recursion of the form $2L_k \le L_{k-1} + L_{k+1} + c$ downward from the endpoints L_0 and L_n yields $L_k \le (1 - \frac{k}{n})L_0 + \frac{k}{n}L_n + C_{k,n}$, which is precisely Kolmogorov's inequality. The general lesson to install: any two term inequality that relates consecutive members of a family is a convexity recursion, and convexity recursions are solved by interpolating between the endpoints. See Book 1, field: Sequences and recurrences, topic: convexity recursions.

Two scale arguments and optimised windows. The general template: to estimate a quantity at a point x near a singularity, introduce a window of unspecified length ℓ, obtain a bound of the form

$A(x)/\ell^a + B(x)\ell^b$ with $a, b > 0$, and minimise over ℓ, giving $\ell \propto (aA/(bB))^{1/(a+b)}$ and a bound proportional to $A^{b/(a+b)}B^{a/(a+b)}$. Then check the optimal ℓ lies in the admissible range; that check is where the problem's structural hypotheses are consumed. See Book 1, field: Inequalities, topic: optimisation of a free parameter.

Circuit Bank

- **WHEN** the conclusion is a disjunction "P or Q" and Q asserts that something is infinite **DO** assume Q fails, convert that into an explicit inequality with a constant, and prove P **BECAUSE** the negation of an unboundedness statement is a usable quantitative bound, whereas the negation of a limit statement gives only a sequence
- **WHEN** you are given control of f and of f'' and asked about f' **DO** write Taylor to second order over a window and solve for f' **BECAUSE** this is the only elementary identity containing all three objects at once, and it is exactly Landau's interpolation inequality before optimisation
- **WHEN** an estimate contains a window length, a cutoff, or any other free scale **DO** carry it as a symbol and optimise at the very end, never guess it at the start **BECAUSE** the two error terms move in opposite directions, so the correct scale is the one that balances them, and guessing usually picks the scale that makes one term diverge
- **WHEN** two hypotheses carry exponents whose product is scale free (here x^r and x^{-r}) **DO** expect a geometric mean inequality and expect the final bound to be independent of position **BECAUSE** scale invariance of the data forces scale invariance of the sharp conclusion, so an x dependent answer is a signal that you optimised badly
- **WHEN** you are asked to prove a dichotomy **DO** first test whether either branch is a theorem on its own, using the cheapest oscillating counterexample you know, such as $x^a \sin(x^{-N})$ **BECAUSE** if one branch were always true the problem would be misstated, and the counterexample tells you which hypothesis the proof is obliged to use

Backtesting

1. First useful cue: the conclusion is a disjunction with a lim sup $= \infty$ on one side, so negate that side and read off $|g''(x)| \le Mx^{-r}$.
2. Method the expression suggested: interpolation. Control of g and g'' with g' as the target is Landau's inequality, realised as Taylor to second order over a window whose length you leave free and then optimise.
3. Where you would hesitate, and the thought to reinforce: the hesitation is choosing the window. Reinforce this sentence: do not choose it, carry ℓ as a symbol, and let the arithmetic mean to geometric mean balance point choose it for you; then check the chosen ℓ is admissible.
4. The trap the question sets: reading "either ... or" as exclusive, or believing that $g = o(x^r)$ already forces $g' \to 0$. Both cost the whole session, and both are refuted in two minutes by $x^a \sin(x^{-N})$ examples.
5. What makes you recognise it faster next time: the pair of exponents x^r and x^{-r} in the two hypotheses. Whenever the data are scale free in that way, the answer is a geometric mean and the bound is uniform in x.
6. The reusable decision rule: to bound an intermediate derivative, interpolate between the function and a higher derivative on an optimised window, and spend the structural hypothesis (here $r > 1$) on checking the window fits.

Mirror drills

1. Prove Landau's inequality: if f is twice differentiable on $[0,\infty)$ with $|f| \le A$ and $|f''| \le B$ everywhere, then $|f'(x)| \le 2\sqrt{AB}$ for all $x \ge 0$. *Answer:* Taylor gives $f(x+\ell) = f(x) + f'(x)\ell + \tfrac{1}{2}f''(\xi)\ell^2$, so $|f'(x)| \le 2A/\ell + B\ell/2$; minimise at $\ell = 2\sqrt{A/B}$ to get $2\sqrt{AB}$.

2. Show that $\lim_{x\to 0^+} g(x)/x^r = 0$ alone does not force $\lim_{x\to 0^+} g'(x) = 0$. *Answer:* take $g(x) = x^{r+1}\sin(x^{-N})$ with $N > r$ and $g(0) = 0$; then $g(x)/x^r = x\sin(x^{-N}) \to 0$, while along $x^{-N} = 2k\pi$ we get $g'(x) = -Nx^{r-N}$, which is unbounded.

3. Show the two alternatives in the problem can hold simultaneously. *Answer:* $g(x) = x^5\sin(x^{-3})$ with $r = 2$; then $g'(x) = 5x^4\sin(x^{-3}) - 3x\cos(x^{-3}) \to 0$, while $x^2 g''(x) = (20x^5 - 9x^{-1})\sin(x^{-3}) - 18x^2\cos(x^{-3})$ is unbounded along $x^{-3} = \tfrac{\pi}{2} + 2k\pi$.

7.7 B1. Counting squares in a Gaussian power of two lattice

Denote by $\mathbb{Z}^2$ the set of all points (x, y) in the plane with integer coordinates. For each integer $n \geq 0$, let P_n be the subset of $\mathbb{Z}^2$ consisting of the point $(0, 0)$ together with all points (x, y) such that $x^2 + y^2 = 2^k$ for some integer $k \leq n$. Determine, as a function of n, the number of four-point subsets of P_n whose elements are the vertices of a square.

Read it out loud

This is a lattice counting question whose real content is a factorisation question. Read the defining condition aloud: $x^2 + y^2 = 2^k$. You are not being asked about a circle. You are being asked which Gaussian integers have norm a power of two, because $x^2 + y^2$ is by definition the norm of $x + iy$ in $\mathbb{Z}[i]$. That is the whole cue. When a Diophantine condition on a pair (x, y) is exactly a norm form, you stop treating x and y as two unknowns and start treating $z = x + iy$ as one object, because norms are multiplicative and two unknowns are not.

The second structural feature is the range $k \leq n$. The exponent runs over an interval, so P_n is not a random point set: it is a geometric family, one level per exponent. Whatever the level m points are, the level $m + 1$ points ought to be their image under something of norm 2. Multiplication by $1 + i$ is a rotation by $\pi/4$ composed with a scaling by $\sqrt{2}$, and P_n is self-similar under it. Self-similarity plus a finite range is the signature of a descent argument, so expect the final classification to be proved by dividing by $1 + i$ and repeating.

The third feature is the word square. Squares are rigid: the four vertices of a square are $c + i^a w$ for $a = 0, 1, 2, 3$, where c is the centre and w a half diagonal. Rotation by a right angle is multiplication by i. So the geometric condition is also multiplicative, in the same ring as the arithmetic condition. Three multiplicative readings in one question is not a coincidence: the examiners built this problem on $\mathbb{Z}[i]$ and are waiting to see whether you notice.

Fields embedded: Gaussian integers and unique factorisation (number theory), plane geometry expressed through complex multiplication, and a small amount of careful counting. Within thirty seconds you should have written three lines: $N(z) = x^2 + y^2$; $2 = -i(1 + i)^2$ so $N(1 + i) = 2$; therefore $N(z) = 2^k$ forces $z = i^a(1 + i)^k$. Also note the one piece of small print: k is allowed to be negative, but 2^k is then not an integer, so negative k contributes nothing and the exponent range is really $0 \leq k \leq n$. State that, do not silently assume it.

Approaches, ranked worst to best

1. **Full coordinate case bash.** *Cue:* the set is finite and explicit for each n. *Circuit:* list the points, take all pairs as diagonals, impose integrality. You can list the $4n + 5$ points, take each of the $\binom{4n+5}{2}$ pairs as a diagonal, compute the other two vertices and test membership. It terminates, but it is $O(n^2)$ pairs with sign cases inside each, all done symbolically in n. It costs pages of ink and it teaches you nothing, because at the end you still cannot say why the answer is linear in n.

2. **Modulo four descent on coordinates.** *Cue:* $x^2 + y^2 \equiv 0 \pmod 4$ and squares are 0 or 1 mod 4. *Circuit:* parity descent to identify P_n, then a geometric case split. For $k \geq 2$ the relation $x^2 + y^2 \equiv 0 \pmod 4$ forces x and y both even, so $(x/2, y/2)$ lies two levels down, and induction identifies P_n as axis points at even levels and diagonal points at odd levels. This is honest and it is the standard route, but the description that comes out is split into two families, and the square classification then needs four separate cases (both vertices axial, both diagonal, one of each, and one at the origin). It works, it is slower, and the two families are an artefact of coordinates rather than a fact about the set.

3. **Convex hull recursion.** *Cue:* P_n is P_{n-1} plus exactly four new extreme points. *Circuit:* count the increment, not the total. The convex hull of P_n is a square whose vertices are the four

new points, and one shows each step adds exactly five squares. This is a good idea and gives $5n+1$ quickly, but it still needs the structure of P_n as an input and the hull argument requires care about which squares can touch an extreme point.

4. **Work in $\mathbb{Z}[i]$.** *Cue:* x^2+y^2 is the norm form; rotation by a right angle is multiplication by i. *Circuit:* unique factorisation plus a valuation descent. One lemma describes the entire set as $\{0\}\cup\{i^a(1+i)^m\}$. The vertex formula for a square becomes a formula about dividing by $1+i$. The classification collapses to a single dichotomy, whether the origin is a vertex or the centre, and the descent is two lines. This is best because every step is forced by the definition of the objects rather than by coordinates.

The ideal solution

Identify $\mathbb{Z}^2$ with the ring of Gaussian integers $\mathbb{Z}[i]$, writing $z=x+iy$ and $N(z)=z\bar{z}=x^2+y^2$. Set $u=1+i$, so $N(u)=2$ and $u^2=2i$.

Lemma 7.7.1. *For every $n\geq 0$,*

$$P_n=\{0\}\cup\{\, i^a u^m \;:\; a\in\{0,1,2,3\},\ 0\leq m\leq n\,\},$$

and the $4(n+1)$ listed nonzero points are pairwise distinct.

Proof. If $x^2+y^2=2^k$ with x,y integers, then 2^k is a nonnegative integer, so $k\geq 0$; negative k gives no points at all. Since $N(u)=2$ is a rational prime, u is prime in $\mathbb{Z}[i]$, and $\mathbb{Z}[i]$ is Euclidean, hence a unique factorisation domain whose units are $1, i, -1, -i$. From $N(z)=2^k$ and $2=-iu^2$ we get $z\bar{z}=(-i)^k u^{2k}$, so the only prime dividing z is u, and comparing norms gives $z=i^a u^k$ for some a. Conversely $N(i^a u^m)=2^m$. Distinctness is immediate: $|i^a u^m|=2^{m/2}$ separates the levels, and the four unit multiples at a fixed level are distinct. □

Call m the *level* of the nonzero point $i^a u^m$; it is the u-adic valuation $v(z)$, and $|P_n|=4n+5$.

Lemma 7.7.2. *If z and w are opposite vertices of a square, the other two vertices are*

$$\frac{iz+w}{u} \qquad \text{and} \qquad \frac{z+iw}{u}.$$

Proof. The centre is $c=(z+w)/2$ and the remaining vertices are $c\pm i(w-z)/2$, since a half diagonal rotates into the other half diagonal. Expanding, $c+i(w-z)/2=\frac{1}{2}\big((1-i)z+(1+i)w\big)$ and $c-i(w-z)/2=\frac{1}{2}\big((1+i)z+(1-i)w\big)$. Now $\frac{1-i}{2}=\frac{1}{u}$ and $\frac{1+i}{2}=\frac{i}{u}$, which gives the two stated expressions. □

Now count. Every square with vertices in P_n either has 0 as a vertex or does not.

Case A: the origin is a vertex. Let $w=i^a u^m$ be the opposite vertex. By Lemma 2 the other two vertices are $w/u=i^a u^{m-1}$ and $iw/u=i^{a+1}u^{m-1}$. If $m\geq 1$ these are level $m-1$ points of P_n, so the four points do lie in P_n and they do form a square. If $m=0$ then $w/u=w\bar{u}/2$ has norm $1/2$ and is not a Gaussian integer, so no square arises. Distinct choices of w give distinct squares, since w is recovered from the square as the vertex opposite 0. The count is therefore 4 choices of a times n choices of $m\in\{1,\dots,n\}$, that is $4n$ squares.

Case B: the origin is not a vertex. Then all four vertices are nonzero, so each has a level. I claim any such square is $\{i^a u^m : a=0,1,2,3\}$ for a single m, the square centred at the origin with its four vertices at one level.

First, a parity dichotomy. Reduction modulo the prime u gives $\mathbb{Z}[i]/(u)\cong\mathbb{F}_2$, in which every unit is congruent to 1 and every element of positive level is congruent to 0. If z,w are opposite vertices, Lemma 2 says $(iz+w)/u$ is a vertex, hence a Gaussian integer, so $iz+w\equiv 0 \pmod{u}$.

Since $i \equiv 1$, this reads $[v(z) = 0] + [v(w) = 0] \equiv 0 \pmod 2$: two opposite vertices have level zero together, or positive level together.

Second, suppose some pair of opposite vertices are both units, $z = i^a$ and $w = i^b$ with $b \neq a$. If $b = a + 1$ then $(z + iw)/u = i^a(1 + i^2)/u = 0$ and if $b = a + 3$ then $(iz + w)/u = i^{a+1}(1 + i^2)/u = 0$, both excluded because the origin is not a vertex. So $b = a + 2$, and then Lemma 2 gives the remaining vertices i^{a+1} and i^{a+3}: the square is $\{1, i, -1, -i\}$, the level zero square centred at the origin.

Third, descend. If some vertex has level 0, its opposite vertex has level 0 by the dichotomy, so by the previous paragraph the square is $\{1, i, -1, -i\}$. Otherwise all four vertices have level at least 1, hence all are divisible by u; dividing all four by u is a spiral similarity, so it produces another square, again with four nonzero vertices in $\{i^a u^m : m \geq 0\}$, and every level has dropped by exactly one. The minimum level is a nonnegative integer that strictly decreases, so after finitely many divisions some vertex has level 0 and the reduced square is $\{1, i, -1, -i\}$. Multiplying back by u^m, where m was the original minimum level, the original square is $\{i^a u^m : a = 0, 1, 2, 3\}$, as claimed.

Conversely each set $\{i^a u^m : a = 0, 1, 2, 3\}$ is a square, since multiplication by i rotates it by a right angle about 0, and it lies in P_n exactly when $0 \leq m \leq n$. That is $n + 1$ squares.

Adding the two cases, the number of four-point subsets of P_n forming a square is

$$4n + (n + 1) = \boxed{5n + 1}.$$

As a check, $n = 0$ gives 1 (only $\{1, i, -1, -i\}$, since P_0 has five points) and $n = 1$ gives 6 (the two centred squares plus the four unit squares with a vertex at the origin), both correct by direct inspection.

Layer 1: WHAT

Block 1. Recognise $x^2 + y^2 = 2^k$ as $N(z) = 2^k$ and move into $\mathbb{Z}[i]$.

Block 2. Factor: $2 = -iu^2$ with $u = 1 + i$ prime, so unique factorisation forces $z = i^a u^k$. The whole set is one geometric orbit, $P_n = \{0\} \cup \{i^a u^m : 0 \leq m \leq n\}$.

Block 3. Express the square condition multiplicatively: from opposite vertices z, w the other two are $(iz + w)/u$ and $(z + iw)/u$.

Block 4. Split on whether the origin is a vertex, because the origin is the only point of P_n that is not of the form $i^a u^m$ and the only one fixed by multiplication by i.

Block 5. Case A is a one line computation: the partner of w at level m lives at level $m - 1$, so the count is $4n$.

Block 6. Case B: reduce modulo u to force opposite vertices to share the level zero or not, then divide by u and descend. Only the centred squares survive: $n + 1$ of them.

Block 7. Add and sanity check small n.

Juxtaposition. The nominal route writes P_n as two coordinate families (axis points at even k, diagonal points at odd k), then takes all pairs of points as candidate diagonals and grinds through four configuration cases with sign analysis and triangle inequalities. It reaches the same $5n + 1$ after several pages. The bifurcation point is exactly the line where you first write the condition down. Either you write $x^2 + y^2 = 2^k$ as an equation in two integer unknowns and start reducing modulo 4, or you write $N(z) = 2^k$ and start factoring. Everything after that is downstream of that single choice, including whether you ever have to mention the words axis and diagonal at all.

Layer 2: HOW

Block 1. Cue. The expression $x^2 + y^2$ set equal to a prime power, with x, y ranging over integers. **Why natural.** $x^2 + y^2$ is not a random quadratic form: it is the norm form of $\mathbb{Z}[i]$, and the defining property of a norm is multiplicativity. A Diophantine condition that says a norm equals a prime power is a condition about factorisation, and factorisation questions are answered in the ring, not in coordinates. This is the same reflex that converts $a^2 + b^2 = c^2$ into Gaussian integers and $a^2 + 3b^2$

into Eisenstein integers. **Link to the next block.** Once you are in $\mathbb{Z}[i]$ you need the factorisation of the right hand side, so the next question asks itself: how does 2 factor in $\mathbb{Z}[i]$?

Block 2. **Cue.** The right hand side is a power of the single rational prime 2. **Why natural.** In $\mathbb{Z}[i]$ the prime 2 ramifies: $2 = -i(1+i)^2$, and $1+i$ has prime norm so it is prime. A number whose norm is 2^k can therefore have no prime factor other than $1+i$, and counting norms pins the exponent exactly. The four units $\pm 1, \pm i$ are the only remaining freedom, which is why there are exactly four points per level. **Link to the next block.** You now have a complete, uniform description of the point set with one integer parameter, the level. A one-parameter set invites a one-parameter answer, so you should expect the final count to be linear in n.

Block 3. **Cue.** The word square, and the fact that you have to decide membership of two unknown vertices. **Why natural.** A square is generated from any vertex by rotating about the centre by right angles, and rotation by a right angle in $\mathbb{C}$ is multiplication by i. Writing the two remaining vertices in terms of two opposite ones therefore produces an expression built from $1 \pm i$, and dividing by 2 turns into dividing by u because $\frac{1-i}{2} = \frac{1}{u}$. This is why the geometry and the arithmetic speak the same language: both are statements about divisibility by u. **Link to the next block.** The formula contains a division by u, and divisibility by u is exactly the level being positive. So the next move is to track levels.

Block 4. **Cue.** The set has one special point, 0, which has no level and is the fixed point of every rotation about it. **Why natural.** Any classification of configurations in a set with one distinguished point splits on whether the configuration uses that point. Here the split is genuinely structural rather than cosmetic, because the origin is the centre of symmetry of P_n: squares through it and squares around it are different animals. **Link to the next block.** Case A is finite and immediate; Case B needs an argument, so do Case A first to bank the easy count.

Block 5. **Cue.** A square with a vertex at 0 and opposite vertex w; Lemma 2 with $z = 0$ collapses to w/u and iw/u. **Why natural.** Dividing by u drops the level by one. So a square with a vertex at the origin exists exactly when the opposite vertex has a level to spare, that is $m \geq 1$, and the failure at $m = 0$ is not an accident but the statement that u does not divide a unit. **Link to the next block.** Case A used only divisibility by u. Case B should use the same tool, so reduce modulo u.

Block 6. **Cue.** Four unknown vertices, all nonzero, with a formula forcing $u \mid iz + w$. **Why natural.** The residue field $\mathbb{Z}[i]/(u)$ has two elements, so reduction modulo u can only see whether a point is a unit multiple of a power of u with exponent zero or positive. That is precisely the piece of information the formula constrains, and it yields the parity dichotomy for free. Once every vertex has positive level, every vertex is divisible by u, and dividing by u is a similarity of the plane, so it preserves squares. Any self-similar set forces this style of argument: strip a factor, repeat, terminate because the valuation is a nonnegative integer. **Link to the next block.** Two cases counted, so add them and test the answer on the smallest n you can enumerate by hand.

Block 7. **Cue.** A closed form in n that you have not yet stress tested. **Why natural.** $n = 0$ and $n = 1$ are cheap and independent of the argument, so they catch an off by one in the level ranges, which is where this problem punishes people. The two ranges to check are $1 \leq m \leq n$ in Case A and $0 \leq m \leq n$ in Case B, and mixing them up changes the constant.

Layer 3: WHY

The mathematical truth beneath the route is that P_n is not a set of points, it is a truncated orbit of a single multiplicative map. Once you know $P_n \setminus \{0\}$ is the set of associates of u^m for $0 \le m \le n$, the set is closed under multiplication by i and almost closed under multiplication by u: multiplying by u is a bijection from level m to level $m+1$. Any question about configurations inside such a set is a question about the group of similarities generated by these maps, and the answer must be an arithmetic progression in n because each new level contributes a fixed number of new configurations. The count $5n+1$ is linear for that reason, not by accident: four new squares through the origin, one new centred square, per level.

The discernment mechanism, stated as a rule: when the constraint defining your point set is a norm form of a quadratic ring, the coordinates are the wrong variables. Sums of two squares belong to $\mathbb{Z}[i]$; x^2+xy+y^2 and x^2+3y^2 belong to the Eisenstein integers $\mathbb{Z}[\omega]$; x^2-2y^2 belongs to $\mathbb{Z}[\sqrt{2}]$ with its unit group and Pell structure. The tell is that the form is multiplicative, so a product of two representable numbers is representable. If you catch yourself doing case analysis modulo 4 on x and y separately, you have already lost the multiplicative structure that the problem was built from.

The second truth is that descent by valuation is the correct tool for self-similar sets. The general shape is: a configuration all of whose members are divisible by a fixed prime π can be divided by π, producing a similar configuration in the same set with smaller valuation; therefore it is enough to classify configurations in which some member is a unit. That single sentence replaced an entire case analysis here, and it is the same sentence used in infinite descent proofs of the irrationality of $\sqrt{2}$, in the classification of Pythagorean triples, and in Fermat style arguments. Train it by asking, of any set you are given, whether it is invariant under some scaling, and if so by what.

The examiners nudged in three ways at once. They wrote x^2+y^2 rather than a general quadratic condition; they wrote a power of two rather than a general integer, which is the one case where the Gaussian prime is ramified and the answer is a single orbit; and they asked for the answer as a function of n, which tells you in advance that the structure is uniform across levels. Reading a question for what it has been made easy is part of reading it at all.

Revolving concepts

Gaussian integers. $\mathbb{Z}[i] = \{x+iy : x,y \in \mathbb{Z}\}$ with norm $N(x+iy) = x^2+y^2$, which is multiplicative: $N(zw) = N(z)N(w)$. It is Euclidean with respect to N, hence a principal ideal domain and a unique factorisation domain. Units are exactly the elements of norm 1, namely $1, i, -1, -i$. A rational prime p splits as a product of two conjugate Gaussian primes when $p \equiv 1 \pmod 4$, stays prime when $p \equiv 3 \pmod 4$, and ramifies for $p=2$ with $2 = -i(1+i)^2$. Any element of prime norm is prime. See Book 1, field: Number theory, topic: Gaussian integers.

Norm forms and sums of two squares. An integer m is a sum of two squares if and only if every prime $\equiv 3 \pmod 4$ occurs to an even power in m; the number of representations is $r_2(m) = 4(d_1(m) - d_3(m))$ where d_j counts divisors congruent to j modulo 4. For $m = 2^k$ this gives $r_2 = 4$, which is the count you proved by hand here. See Book 1, field: Number theory, topic: Sums of two squares.

Valuations and descent. For a prime π of a unique factorisation domain, $v_\pi(z)$ is the exponent of π in z, with $v_\pi(zw) = v_\pi(z) + v_\pi(w)$ and $v_\pi(z+w) \ge \min(v_\pi(z), v_\pi(w))$, with equality when the two valuations differ. Reduction modulo π lands in the residue field, here $\mathbb{Z}[i]/(1+i) \cong \mathbb{F}_2$. See Book 1, field: Number theory, topic: Valuations.

The recursion, generalised. The set here satisfies $P_n \setminus \{0\} = U \cup u\,(P_{n-1} \setminus \{0\})$ with $U = \{1, i, -1, -i\}$, a self-similar recursion with ratio u. The general statement worth installing: if a set X satisfies $X = F \cup \alpha X$ for a similarity α with $|\alpha| > 1$ and a finite base F, then any finite configuration $C \subseteq X$ which is closed under nothing in particular can be normalised by repeatedly dividing by α until $C \cap F \neq \emptyset$; hence configurations in X are exactly the α^m images of configurations

meeting F. Counting configurations then gives a count of the form (configurations meeting F) $\times$ (number of levels) plus boundary corrections, which is why linear answers in n appear. The same normalise then descend pattern solves questions about geometric sequences in lattices, about binary expansions, and about p-adic digit configurations. See Book 1, field: Sequences and recurrences, topic: Self-similar sets.

Squares via complex rotation. Four points z_1, z_2, z_3, z_4 form a square in this cyclic order if and only if $z_2 - z_1 = i(z_1 - z_4)$ and the quadrilateral closes, equivalently $z_1 + z_3 = z_2 + z_4$ and $z_3 - z_1 = -i(z_4 - z_2)$. Given opposite vertices, the other pair is $(iz + w)/(1 + i)$ and $(z + iw)/(1 + i)$. See Book 1, field: Complex numbers, topic: Rotations and similar figures.

Circuit Bank

- **WHEN** you see $x^2 + y^2$ constrained to a specific set of values **DO** rewrite it as $N(z)$ with $z = x + iy$ and factor in $\mathbb{Z}[i]$ **BECAUSE** $x^2 + y^2$ is the norm form of $\mathbb{Z}[i]$, and norms are multiplicative, so the condition is about factorisation, which coordinates hide
- **WHEN** you see a point set that is invariant under a scaling or spiral similarity **DO** classify configurations by minimal valuation, divide by the scaling factor, descend **BECAUSE** the valuation is a nonnegative integer, so the descent terminates and only the base level needs a case analysis
- **WHEN** you see squares or right angles with lattice constraints **DO** move to $\mathbb{C}$ and write rotation by a right angle as multiplication by i **BECAUSE** the rigid motion group of the plane is generated by complex multiplication and translation, so geometric conditions become algebraic identities with no coordinates
- **WHEN** you see a set with exactly one distinguished point, such as the centre of symmetry **DO** split the count on whether the configuration uses that point **BECAUSE** the distinguished point is the only one whose stabiliser differs, so the two families obey different formulas and merging them creates errors
- **WHEN** you finish with a closed form in n **DO** test $n = 0$ and $n = 1$ by hand **BECAUSE** the only realistic failure mode left is an off by one in an exponent range, and small cases expose it in seconds

Backtesting

1. First useful cue: $x^2 + y^2 = 2^k$ read as $N(z) = 2^k$, so the point set is the set of associates of powers of $1 + i$.
2. Method the expression suggested: unique factorisation in $\mathbb{Z}[i]$ to describe P_n, then complex rotation to describe squares, then a valuation descent to classify them.
3. Where you would hesitate, and the thought to reinforce: at Case B, when four unknown vertices appear at once. The thought is that you never need all four levels; you only need whether each is zero or positive, and reduction modulo $1 + i$ answers exactly that.
4. The trap the question sets: the condition is $k \leq n$, not $0 \leq k \leq n$, and the two exponent ranges in the two cases differ by one. A candidate who writes $4(n + 1)$ in Case A or n in Case B loses the problem while believing the argument is complete.
5. What makes you recognise it faster next time: the pairing of a norm form with a prime power right hand side. That pairing has exactly one meaning, and it appears every few years on the Putnam.
6. The reusable decision rule: norm form condition means move into the quadratic ring; self-similar set means descend by valuation; square means multiply by i.

Mirror drills

1. How many points $(x, y) \in \mathbb{Z}^2$ satisfy $x^2 + y^2 = 5^k$ for a fixed $k \geq 0$? *Answer:* $4(k+1)$. Since $5 = (2+i)(2-i)$ with the two factors non-associate primes, $z = i^a(2+i)^j(2-i)^{k-j}$ with $0 \leq j \leq k$.
2. Let $T_n = \{0\} \cup \{(x, y) : x^2 + y^2 = 3^k,\ k \leq n\}$. How many four-point subsets of T_n form a square? *Answer:* $\lfloor n/2 \rfloor + 1$. Here 3 is inert, so $N(z) = 3^k$ forces k even and $z = i^a 3^{k/2}$; every nonzero point has odd norm, so $1 + i$ divides no vertex and no square can use the origin. Only the centred squares survive.
3. Show that if all four vertices of a square lie in $\mathbb{Z}^2$, then its area is an integer. *Hint:* the side vector v satisfies $v \in \mathbb{Z}[i]$, and the area is $N(v)$.

7.8 B2. Telescoping a trigonometric sum of cubic growth

For all $n \geq 1$, let

$$a_n = \sum_{k=1}^{n-1} \frac{\sin\left(\frac{(2k-1)\pi}{2n}\right)}{\cos^2\left(\frac{(k-1)\pi}{2n}\right)\cos^2\left(\frac{k\pi}{2n}\right)}.$$

Determine

$$\lim_{n\to\infty} \frac{a_n}{n^3}.$$

Read it out loud

Say the summand aloud in words, not symbols. The denominator is the product of the squares of the cosines at two consecutive angles, $(k-1)\theta$ and $k\theta$, where $\theta = \pi/(2n)$. The numerator is the sine at the angle exactly halfway between the doubled angles, since $\frac{1}{2}\big(2(k-1)\theta + 2k\theta\big) = (2k-1)\theta$. A denominator built from consecutive values of one sequence, and a numerator built from the mean of the corresponding doubled angles, is not a random arrangement. It is the fingerprint of a difference. The examiners built this summand backwards from a telescoping identity, and your job is to run the construction in reverse.

The second reading is about size. There is a limit being asked for, and the normalisation is n^3, not n or n^2. A sum of $n-1$ terms that grows like n^3 cannot be an ordinary Riemann sum, because a Riemann sum grows like n times a bounded average. So either the individual terms are of size n^2, or a few terms near one end are enormous. Look at the last term: $\cos\big((n-1)\pi/(2n)\big) = \sin(\pi/(2n))$, which is of size $1/n$. Its square sits in the denominator alongside another cosine of size $1/n$, so the denominator is of size $1/n^4$ while the numerator is of size $1/n$, making that single term of size n^3 by itself. So the last few terms already carry the whole answer. That is a warning: the naive continuum approximation has a non integrable singularity at the right endpoint, and anyone who writes down $\int_0^1$ has already lost the problem.

Fields embedded: trigonometric identities (double angle, sum to product), telescoping sums, and one elementary limit. There is no analysis here beyond $\sin x/x \to 1$. What is really being tested is whether you look for exact structure before you reach for approximation. The Putnam rewards the candidate who computes a_n in closed form and then takes a limit in one line, and punishes the candidate who starts estimating.

Within thirty seconds you should have written $\theta = \pi/(2n)$, $c_k = \cos(k\theta)$, and the question: what is $\dfrac{1}{c_k^2} - \dfrac{1}{c_{k-1}^2}$?

Approaches, ranked worst to best

1. **Riemann sum in $t = k/n$.** *Cue:* a sum of $n-1$ terms with argument $k\pi/(2n)$. *Circuit:* recognise $\frac{1}{n}\sum f(k/n)$ as $\int_0^1 f$. Substituting $t = k/n$ turns the summand into roughly $2\sin(\pi t/2)/\cos^3(\pi t/2)$, so you would write $a_n/n^3 \approx \frac{1}{n}\int_0^1$ of something. It fails twice over. The integrand blows up like $(1-t)^{-3}$ at $t = 1$, so the integral diverges, and the powers of n do not even balance. This route is not merely inefficient, it produces a wrong answer with a confident face on it.

2. **Endpoint asymptotics after reversing the index.** *Cue:* the terms near $k = n$ dominate. *Circuit:* substitute $j = n-k$ and expand each term for fixed j. Putting $j = n-k$ turns the summand into $\sin((2j+1)\theta)/\big(\sin^2((j+1)\theta)\sin^2(j\theta)\big)$, which for fixed j and small θ is $\frac{2j+1}{j^2(j+1)^2}\theta^{-3}(1+O(\theta^2))$. Summing gives $\theta^{-3}\sum_{j\geq 1}\frac{2j+1}{j^2(j+1)^2} = \theta^{-3}$, and $\theta^{-3} = (2n/\pi)^3$. Correct, and it is the honest analytic route, but full marks demand a uniform error bound over the whole range of j, not just fixed j, which costs a paragraph of estimates. Notice too that the

final series telescopes, which is the problem telling you a second time that you should have telescoped from the start.

3. **Compute small cases, guess the closed form, prove by induction.** *Cue:* the answer is a clean constant, so a_n is probably clean. *Circuit:* conjecture and verify. From $a_2 = \sqrt{2}$ and a_3, a_4 you can guess $a_n = \cos^2\theta/\sin^3\theta$ and then prove it by induction on the partial sums. It gets there, but induction on a sum whose upper limit and whose angle both depend on n is awkward, because θ changes with n and the partial sums of different n are unrelated.
4. **Telescope exactly.** *Cue:* the denominator is a product of consecutive values of $\cos^2$. *Circuit:* write the summand as a difference of $1/\cos^2$ at consecutive indices. Compute $\cos^2((k-1)\theta) - \cos^2(k\theta)$ by double angle and sum to product; it equals $\sin((2k-1)\theta)\sin\theta$, which is the numerator times $\sin\theta$. The sum collapses to two boundary terms, a_n is exact, and the limit is immediate. Best because it uses the exact structure the setter built in, it needs no estimates at all, and it produces more than was asked, namely a closed form valid for every n.

The ideal solution

Write $\theta = \dfrac{\pi}{2n}$ and $c_k = \cos(k\theta)$, so that the summand is $\dfrac{\sin((2k-1)\theta)}{c_{k-1}^2 c_k^2}$.

By the double angle formula $\cos^2 x = \frac{1}{2}(1+\cos 2x)$,

$$c_{k-1}^2 - c_k^2 = \tfrac{1}{2}\Big(\cos\big(2(k-1)\theta\big) - \cos\big(2k\theta\big)\Big),$$

and by the sum to product identity $\cos A - \cos B = 2\sin\frac{A+B}{2}\sin\frac{B-A}{2}$ with $A = 2(k-1)\theta$, $B = 2k\theta$,

$$c_{k-1}^2 - c_k^2 = \sin\big((2k-1)\theta\big)\sin\theta.$$

For $1 \le k \le n-1$ every angle $k\theta$ lies in $[0, \pi/2)$, so $c_k \neq 0$ and we may divide by $c_{k-1}^2 c_k^2$:

$$\frac{1}{c_k^2} - \frac{1}{c_{k-1}^2} = \frac{c_{k-1}^2 - c_k^2}{c_{k-1}^2 c_k^2} = \sin\theta \cdot \frac{\sin((2k-1)\theta)}{c_{k-1}^2 c_k^2}.$$

Therefore each summand equals $\dfrac{1}{\sin\theta}\left(\dfrac{1}{c_k^2} - \dfrac{1}{c_{k-1}^2}\right)$, and the sum telescopes:

$$a_n = \frac{1}{\sin\theta}\left(\frac{1}{c_{n-1}^2} - \frac{1}{c_0^2}\right).$$

Now $c_0 = 1$ and $c_{n-1} = \cos\big(\frac{\pi}{2} - \theta\big) = \sin\theta$, so for every $n \ge 2$,

$$a_n = \frac{1}{\sin\theta}\left(\frac{1}{\sin^2\theta} - 1\right) = \frac{\cos^2\theta}{\sin^3\theta}, \qquad \theta = \frac{\pi}{2n}.$$

(For $n = 1$ the sum is empty and $a_1 = 0$, which does not affect a limit.)

Finally, $\lim\limits_{x\to 0}\dfrac{\sin x}{x} = 1$ gives $n\sin\theta = \frac{\pi}{2}\cdot\frac{\sin\theta}{\theta} \to \frac{\pi}{2}$, and $\cos\theta \to 1$. Hence

$$\frac{a_n}{n^3} = \cos^2\theta \cdot \frac{1}{(n\sin\theta)^3} \longrightarrow 1 \cdot \left(\frac{2}{\pi}\right)^3,$$

so

$$\lim_{n\to\infty}\frac{a_n}{n^3} = \boxed{\frac{8}{\pi^3}}.$$

As a check, $n = 2$ gives $\theta = \pi/4$ and $\cos^2\theta/\sin^3\theta = \frac{1}{2}/(\frac{\sqrt{2}}{2})^3 = \sqrt{2}$, which is the single term $\sin(\pi/4)/(1\cdot\cos^2(\pi/4)) = \sqrt{2}$.

Layer 1: WHAT

Block 1. Name the angle: $\theta = \pi/(2n)$, $c_k = \cos(k\theta)$. The summand is $\sin((2k-1)\theta)/(c_{k-1}^2 c_k^2)$.

Block 2. Read the denominator as a product of consecutive values of one sequence, which is the standing cue for telescoping, and therefore compute $c_{k-1}^2 - c_k^2$.

Block 3. Double angle turns squares of cosines into cosines of doubled angles; sum to product turns that difference into $\sin((2k-1)\theta)\sin\theta$, which is exactly the numerator times a constant.

Block 4. Divide: the summand is $\frac{1}{\sin\theta}\left(c_k^{-2} - c_{k-1}^{-2}\right)$. Check no denominator vanishes on the range.

Block 5. Telescope to the two endpoints, and use $c_{n-1} = \sin\theta$, which is the complementary angle identity and the reason the answer is clean.

Block 6. Closed form $a_n = \cos^2\theta/\sin^3\theta$, then divide by n^3 and use $n\sin\theta \to \pi/2$.

Juxtaposition. The nominal route is: this is a sum with k/n inside trigonometric functions, so it must be a Riemann sum, so approximate and integrate. That route dies on the singularity at the right endpoint. A slightly better zombie route notices the domination by the last terms, reverses the index and estimates term by term, arriving at the answer after a page of O notation. The bifurcation point is the very first decision after reading the summand: do you look at the denominator as a shape, or at the summand as a function of k/n? Looking at the shape gives an exact identity, looking at the function gives approximations. On a problem whose answer is a clean constant, the shape wins.

Layer 2: HOW

Block 1. Cue. Every angle in sight is an integer multiple of $\pi/(2n)$. **Why natural.** Naming θ costs one line and removes the n from the interior of every trigonometric function, so the summand becomes a statement about a sequence indexed by k with a parameter θ. Identities are found on sequences, not on cluttered fractions. Notation that hides the parameter is notation that lets you see the structure. **Link to the next block.** With the notation clean, the shape of the denominator becomes visible: $c_{k-1}^2 c_k^2$.

Block 2. Cue. The denominator is $g(k-1)g(k)$ for $g(k) = c_k^2$, a product of consecutive values of one sequence. **Why natural.** This is the exact algebraic signature of $\frac{1}{g(k-1)} - \frac{1}{g(k)} = \frac{g(k)-g(k-1)}{g(k-1)g(k)}$. Whenever a summand has consecutive values of one sequence multiplied in the denominator, the summand is a difference in disguise provided the numerator matches $g(k) - g(k-1)$. So you do not guess: you compute $g(k-1) - g(k)$ and compare. The same reading turns $\frac{1}{k(k+1)}$ into $\frac{1}{k} - \frac{1}{k+1}$ and $\frac{2k+1}{k^2(k+1)^2}$ into $\frac{1}{k^2} - \frac{1}{(k+1)^2}$. **Link to the next block.** The comparison requires simplifying a difference of squares of cosines, so reach for the identities that linearise squares.

Block 3. Cue. A difference of $\cos^2$ at two angles. **Why natural.** $\cos^2 x = \frac{1}{2}(1+\cos 2x)$ is the only identity that converts a square into something additive, and the additive constants cancel in a difference. Once you have a difference of two cosines, the sum to product identity is forced, since it is the identity that converts additive differences of cosines into products of sines. The numerator that emerges, sin of the mean of the two doubled angles, is precisely the numerator printed in the question. That is the confirmation that the setter constructed the summand this way. **Link to the next block.** The identity now reads: numerator times $\sin\theta$ equals the difference of the denominators' factors. Divide through.

Block 4. Cue. A proposed identity involving division by $c_{k-1}^2 c_k^2$. **Why natural.** Division needs the divisor to be nonzero, and here k runs only to $n-1$, so $k\theta \le (n-1)\pi/(2n) < \pi/2$ and every cosine is strictly positive. The upper limit $n-1$ rather than n is not decoration: at $k = n$ the cosine would vanish and the sum would not exist. Saying this in one clause is what separates a full-marks

write-up from a nearly right one. **Link to the next block.** With the identity established termwise, sum it.

Block 5. Cue. A sum of consecutive differences. **Why natural.** Telescoping is the discrete fundamental theorem of calculus: $\sum_{k=1}^{m} \big(G(k) - G(k-1)\big) = G(m) - G(0)$. The only work left is evaluating the two boundary values. Here $c_0 = 1$ is trivial and $c_{n-1} = \cos(\pi/2 - \theta) = \sin\theta$ by the complementary angle identity, which is why the closed form involves only θ. **Link to the next block.** You now have an exact formula, so the limit is a routine calculation rather than an estimate.

Block 6. Cue. A closed form with $\theta = \pi/(2n) \to 0$ and a division by n^3. **Why natural.** The only asymptotic input is $\sin x \sim x$. Group the expression so that each factor has an obvious limit: $\cos^2\theta \to 1$, and $n\sin\theta \to \pi/2$ because $n\theta = \pi/2$ exactly. Grouping before taking limits is what stops you from writing ∞/∞ and losing control of the constant.

Layer 3: WHY

The truth underneath is that this sum is a telescoping sum by construction, and telescoping is the discrete antiderivative. In the continuous world, $\int f$ is easy when $f = G'$. In the discrete world, $\sum_k t_k$ is easy when $t_k = G(k) - G(k-1)$. Everything in the ideal solution is the search for G. The setter chose $G(k) = \sec^2(k\theta)$, computed $G(k) - G(k-1)$, tidied it with two standard identities, and printed the result as the question. Reading a summand as the output of that process is a trainable skill, and the trigger is structural: consecutive indices multiplied in the denominator.

The discernment mechanism for choosing telescoping over approximation is a two part test. First, is the requested answer exact or asymptotic in a clean form? A clean constant such as $8/\pi^3$ suggests an exact closed form is available. Second, is the sum dominated by a bounded number of terms, or spread evenly? If a handful of terms carry the whole sum, a Riemann sum is meaningless and you should be looking for exact cancellation instead. Both tests point the same way here. Train the reflex by taking any sum you meet and asking the single question: is the denominator a product of consecutive values of one sequence? If yes, compute the difference of those values before doing anything else.

There is a second, deeper reason the problem is clean. The function $x \mapsto \sec^2 x$ is the derivative of $\tan x$, and the identity used is the discrete shadow of $\frac{d}{dx}\tan x = \sec^2 x$. In fact the general identity behind the whole problem is

$$\sum_{k=1}^{m} \frac{\sin((2k-1)x)}{\cos^2((k-1)x)\cos^2(kx)} = \frac{\tan^2(mx)}{\sin x},$$

valid whenever no cosine vanishes, and the problem is the case $m = n-1$, $x = \pi/(2n)$, chosen so that mx is the complement of x and the answer becomes a function of x alone. The examiners hid a one parameter identity by tying the two parameters together. Recognising a specialised identity, and asking what the free version looks like, is exactly how you convert one Putnam problem into a piece of permanent equipment.

The nudges are visible once you know they are there: the upper limit $n-1$ instead of n (which keeps the cosines nonzero and makes the last term the dominant one), the argument $(2k-1)\pi/(2n)$ in the numerator (which is sin of the mean of the doubled denominator angles, the sum to product signature), and the normalisation n^3 (which announces that the sum is not an average but is driven by a small number of large terms).

Revolving concepts

Telescoping. If $t_k = G(k) - G(k-1)$ then $\sum_{k=a}^{b} t_k = G(b) - G(a-1)$. This is the discrete fundamental theorem of calculus and the discrete counterpart of an exact differential. The search for G is

the discrete analogue of antidifferentiation, and it is decidable for hypergeometric summands by Gosper's algorithm, whose creative telescoping extension (Zeilberger) settles most binomial identities mechanically. Knowing that such an algorithm exists tells you that hunting for G is a legitimate first move rather than a hopeful one. See Book 1, field: Sequences and recurrences, topic: Telescoping and summation by parts.

The generalised recursion. Set $G(k) = \sec^2(k\theta)$ and the problem is the first order recursion $G(k) = G(k-1) + \sin\theta \cdot t_k$, whose solution is $G(m) = G(0) + \sin\theta \sum_{k\le m} t_k$. Every telescoping sum is exactly this: a first order linear recursion with unit coefficient, solved by summing. The general version worth carrying is: for $t_k = \dfrac{g(k) - g(k-1)}{g(k-1)g(k)}$ the partial sums are $\dfrac{1}{g(a-1)} - \dfrac{1}{g(b)}$, so any summand whose denominator is a product of consecutive values of g and whose numerator is the corresponding difference collapses. Special cases include $\sum \frac{1}{k(k+1)}$, $\sum \frac{2k+1}{k^2(k+1)^2}$, and $\sum \frac{1}{\sin(kx)\sin((k+1)x)}$ via $\sin x = \sin((k+1)x - kx)$.

Sum to product and product to sum. $\cos A - \cos B = 2\sin\frac{A+B}{2}\sin\frac{B-A}{2}$, $\sin A - \sin B = 2\cos\frac{A+B}{2}\sin\frac{A-B}{2}$, and the double angle forms $\cos^2 x = \frac{1+\cos 2x}{2}$, $\sin^2 x = \frac{1-\cos 2x}{2}$. These are the standard toolkit for converting between additive and multiplicative trigonometric structure, and telescoping needs the additive direction. See Book 1, field: Calculus and integrals, topic: Trigonometric identities.

Small angle asymptotics. $\sin x = x + O(x^3)$, $\cos x = 1 + O(x^2)$, and $\lim_{x\to 0}\frac{\sin x}{x} = 1$. For sums where the answer is only asymptotic, the uniform version $\frac{\sin x}{x} = 1 + O(x^2)$ on $[0, \pi]$ is what licenses termwise replacement inside a sum, and the error must be summed, not waved at. See Book 1, field: Real analysis, topic: Asymptotic expansions.

Domination by boundary terms. A positive sum $\sum_{k=1}^{n-1} f(k/n)$ where f has a pole of order $p > 1$ at the endpoint grows like n^p times a convergent series, not like n times an integral. The correct normalisation is read off from the order of the pole, and the constant comes from the convergent series in the reversed index. See Book 1, field: Real analysis, topic: Sums with endpoint singularities.

Circuit Bank

- **WHEN** you see a summand whose denominator is a product of consecutive values of one sequence **DO** compute the difference of those consecutive values and compare it with the numerator **BECAUSE** such a summand is $\frac{g(k)-g(k-1)}{g(k-1)g(k)} = \frac{1}{g(k-1)} - \frac{1}{g(k)}$, so the sum telescopes and no estimation is needed

- **WHEN** you see a difference of squares of sines or cosines **DO** apply the double angle formula first, then sum to product **BECAUSE** squares are multiplicative objects and differences are additive, so you must linearise before you can factor a difference

- **WHEN** you are asked for an asymptotic limit with a clean closed constant **DO** look for an exact closed form of the finite sum before estimating **BECAUSE** clean constants are the residue of exact identities, and an exact formula converts the limit into one line with no error control

- **WHEN** you see a sum whose terms blow up at one end of the range **DO** reverse the index so the blow up is at $j = 1$ and check the resulting series converges **BECAUSE** the growth rate in n is then read off the order of the pole, and no Riemann sum is available since the limiting integral diverges

- **WHEN** you divide by a trigonometric expression inside a sum **DO** state the range of angles that keeps it nonzero **BECAUSE** the upper limit $n - 1$ was chosen precisely to avoid a zero cosine, and a grader reads that clause as the difference between an identity and an assertion

Backtesting

1. First useful cue: the denominator is $\cos^2((k-1)\theta)\cos^2(k\theta)$, consecutive values of one sequence multiplied together.
2. Method the expression suggested: build the difference $\sec^2(k\theta) - \sec^2((k-1)\theta)$, confirm by double angle and sum to product that it reproduces the numerator times $\sin\theta$, and telescope.
3. Where you would hesitate, and the thought to reinforce: the temptation to treat the sum as a Riemann sum because $k\pi/(2n)$ looks like a sampled variable. The thought is that a Riemann sum grows like n, and you were handed n^3, so the analogy is dead on arrival.
4. The trap the question sets: the n^3 normalisation invites approximation, and the approximation that first comes to hand diverges. The problem is built so that the lazy analytic route is not merely long but wrong.
5. What makes you recognise it faster next time: the numerator being the sine of the mean of the doubled denominator angles. That is the sum to product identity read backwards, and it can only have been put there deliberately.
6. The reusable decision rule: consecutive values in a denominator means telescoping; verify by computing the difference of those values, never by hoping.

Mirror drills

1. Evaluate $\displaystyle\sum_{k=1}^{m} \frac{\sin((2k-1)x)}{\cos^2((k-1)x)\cos^2(kx)}$ for $0 < x < \frac{\pi}{2m}$. *Answer:* $\dfrac{\tan^2(mx)}{\sin x}$, by the same telescoping; the problem is the case $m = n-1$, $x = \pi/(2n)$.
2. Evaluate $\displaystyle\sum_{k=1}^{m} \frac{1}{\sin(kx)\sin((k+1)x)}$. *Answer:* $\dfrac{\cot x - \cot((m+1)x)}{\sin x}$, since $\sin x = \sin\big((k+1)x - kx\big)$ expands to $\sin((k+1)x)\cos(kx) - \cos((k+1)x)\sin(kx)$, giving $\cot(kx) - \cot((k+1)x)$ after dividing.
3. Show $\displaystyle\sum_{k=1}^{n-1} \frac{2k+1}{k^2(k+1)^2} = 1 - \frac{1}{n^2}$, and say in one sentence which feature of the summand told you so. *Answer:* the numerator is $(k+1)^2 - k^2$ and the denominator is the product of the corresponding consecutive squares, so the summand is $\frac{1}{k^2} - \frac{1}{(k+1)^2}$. This is the algebraic shadow of the trigonometric identity in the problem.

7.9 B3. A reflection restores the eigenvalue one

Let Q be an n-by-n real orthogonal matrix, and let $u \in \mathbb{R}^n$ be a unit column vector (that is, $u^T u = 1$). Let $P = I - 2uu^T$, where I is the n-by-n identity matrix. Show that if 1 is not an eigenvalue of Q, then 1 is an eigenvalue of PQ.

Read it out loud

Three objects, each of which must be translated from matrix language into geometric language before you write anything. Q is orthogonal, so it preserves lengths and inner products. The hypothesis that 1 is not an eigenvalue of Q says Q has no nonzero fixed vector, and equivalently that $Q - I$ is invertible. Say the equivalent form out loud, because an invertible matrix is something you can divide by, and a hypothesis you can divide by is a hypothesis you should use as a division. $P = I - 2uu^T$ is a rank one update of the identity, and rank one updates are never mysterious: P fixes everything orthogonal to u and sends u to $-u$, so P is the reflection in the hyperplane $u^\perp$, with $P^2 = I$ and $\det P = -1$. The conclusion, 1 is an eigenvalue of PQ, says PQ has a nonzero fixed vector.

Now read the whole statement as one sentence of geometry: composing a fixed-point-free orthogonal map with a reflection always produces a map with a fixed direction. That is a statement about parity, and it should smell like one, because reflections flip orientation and fixed-point-free orthogonal maps are exactly those of a definite determinant sign. Two proofs are therefore available in principle: a parity proof by determinants, and a constructive proof that produces the fixed vector. On the Putnam the constructive proof is worth more per line, because a produced vector cannot be wrong.

The decisive structural fact is that P differs from I by rank one. If $PQw = w$ then, multiplying by P and using $P^2 = I$, we get $Qw = Pw$, and Pw differs from w by a multiple of u. So the equation to solve is $(Q - I)w \in \operatorname{span}(u)$. Since $Q - I$ is invertible, that determines w up to scale immediately: there is exactly one candidate line, spanned by $(Q - I)^{-1}u$. There is no search here and no cleverness required. The rank one structure leaves exactly one place for the answer to be.

Fields embedded: the real orthogonal group, eigenvalues as fixed vectors, rank one perturbations, and (in the alternative route) the spectrum of real orthogonal matrices and the Cayley transform. Within thirty seconds you should have written $P^2 = I$, $Pw = w - 2(u^T w)u$, and $w = (Q - I)^{-1}u$.

Approaches, ranked worst to best

1. **Expand the characteristic polynomial of PQ.** *Cue:* eigenvalue questions are about characteristic polynomials. *Circuit:* compute $\det(PQ - \lambda I)$ in coordinates. There is nothing to expand: Q is an arbitrary orthogonal matrix of arbitrary size. Any attempt to write entries produces symbols with no relations you can exploit. This is zombie maths in its purest form, doing the thing the words suggest with no reading of the structure.

2. **Test $n = 2$ and $n = 3$.** *Cue:* small cases build intuition. *Circuit:* specialise and look for the mechanism. For $n = 2$, Q is a rotation by $\theta \neq 0$ and P is a reflection, so PQ is a reflection and every reflection of the plane fixes a line. That is genuinely informative, because it shows the mechanism is orientation and not size. It is not a proof, and generalising the two-dimensional picture directly does not work, because in higher dimensions PQ is not a reflection.

3. **Determinant parity.** *Cue:* orthogonal, fixed-point-free, and a reflection all point at det. *Circuit:* spectrum of a real orthogonal matrix, then a contradiction. The eigenvalues of a real orthogonal matrix lie on the unit circle and the non-real ones occur in conjugate pairs of product 1. If 1 is not an eigenvalue, then $\det Q = (-1)^k$ where k is the multiplicity of -1, and $k \equiv n$

(mod 2) since the remaining eigenvalues pair up. Hence $\det Q = (-1)^n$. If PQ also had no eigenvalue 1, the same applies to the orthogonal matrix PQ, giving $\det(PQ) = (-1)^n$, while $\det(PQ) = \det P \det Q = -(-1)^n$. Contradiction. This is a complete and memorable proof; it costs you the spectral lemma and it yields no eigenvector.

4. **Matrix determinant lemma.** *Cue:* P is a rank one update, and rank one updates have a determinant formula. *Circuit:* $\det(A + vw^T) = \det A\,(1 + w^T A^{-1} v)$. Since $P^{-1} = P$, $\det(PQ - I) = \det P \det(Q - P) = -\det\big((Q - I) + 2uu^T\big) = -\det(Q - I)\big(1 + 2u^T(Q - I)^{-1}u\big)$, and the bracket vanishes because $u^T(Q - I)^{-1}u = -\frac{1}{2}$. Short and structural, but it imports a lemma and proves only that something is singular.

5. **Construct the fixed vector.** *Cue:* $Pw - w$ is always a multiple of u, and $Q - I$ is invertible. *Circuit:* solve $(Q - I)w \in \operatorname{span}(u)$, then normalise. The candidate $w = (Q - I)^{-1}u$ is forced, not guessed, and the single scalar that must come out right, $u^T w$, is pinned by the one property of Q you have not yet used: it preserves norms. Best because it is four lines, uses no lemma beyond the definition of orthogonality, produces the eigenvector explicitly, and even shows the eigenspace is exactly one dimensional.

The ideal solution

Since $u^T u = 1$,

$$P^2 = (I - 2uu^T)^2 = I - 4uu^T + 4u(u^T u)u^T = I,$$

so P is invertible with $P^{-1} = P$, and for every w,

$$Pw = w - 2(u^T w)\,u.$$

Because 1 is not an eigenvalue of Q, the matrix $Q - I$ is invertible. Define

$$w = (Q - I)^{-1}u.$$

Then $w \neq 0$ (as $u \neq 0$), and by construction

$$Qw = w + u.$$

Now use the only property of Q not yet used: Q preserves the Euclidean norm, so $\|Qw\| = \|w\|$. Squaring $Qw = w + u$,

$$\|w\|^2 = \|Qw\|^2 = \|w + u\|^2 = \|w\|^2 + 2u^T w + \|u\|^2 = \|w\|^2 + 2u^T w + 1,$$

and therefore

$$u^T w = -\tfrac{1}{2}.$$

Consequently

$$Pw = w - 2\left(-\tfrac{1}{2}\right)u = w + u = Qw,$$

and applying P to both sides,

$$PQw = P(Pw) = P^2 w = w.$$

Since $w \neq 0$, the number 1 is an eigenvalue of PQ, with eigenvector $(Q - I)^{-1}u$. ■

Remark 7.9.1. The eigenvector is unique up to scale. Indeed $PQw = w$ is equivalent to $Qw = Pw$, that is to $(Q - I)w = -2(u^T w)u \in \operatorname{span}(u)$, so any such w lies in $(Q - I)^{-1}\operatorname{span}(u)$, a one dimensional space. Hence the eigenvalue 1 of PQ has geometric multiplicity exactly 1.

Remark 7.9.2. The scalar identity proved above is worth isolating: for every real orthogonal Q without eigenvalue 1 and every unit vector u,

$$u^T(Q - I)^{-1}u = -\tfrac{1}{2}.$$

Since this holds for every unit u, the symmetric part of $(Q - I)^{-1}$ is exactly $-\frac{1}{2}I$, that is, $(Q - I)^{-1} + \frac{1}{2}I$ is skew-symmetric. Equivalently $(Q + I)(Q - I)^{-1}$ is skew-symmetric, which is the Cayley transform.

Layer 1: WHAT

Block 1. Translate: P is the reflection in $u^\perp$; $P^2 = I$; $Pw = w - 2(u^T w)u$; 1 is an eigenvalue of M means M has a nonzero fixed vector.

Block 2. Rewrite the target: $PQw = w$ if and only if $Qw = Pw$, because P is its own inverse.

Block 3. Read the right hand side: Pw differs from w by a multiple of u, so the equation says $(Q - I)w \in \text{span}(u)$.

Block 4. Use the hypothesis as a division: $Q - I$ is invertible, so the candidate is forced to be $w = (Q - I)^{-1}u$ up to scale.

Block 5. Test the candidate. It satisfies $Qw = w + u$ by construction, so everything reduces to the single scalar $u^T w$.

Block 6. Pin the scalar with the unused hypothesis: $\|Qw\| = \|w\|$ gives $u^T w = -\frac{1}{2}$, exactly the value that makes $Pw = w + u$.

Block 7. Conclude, and note the eigenspace is one dimensional.

Juxtaposition. The nominal route treats the conclusion as a determinant statement and goes hunting for $\det(PQ - I) = 0$: either by expanding, or by the spectral parity argument, or by importing the matrix determinant lemma. All three are proofs by contradiction or by formula, and all three tell you only that a fixed vector exists. The bifurcation point is the sentence you write after reading the conclusion. If you write $\det(PQ - I) = 0$, you are committed to determinants and to whichever lemma you can remember. If you write $PQw = w$ for some $w \neq 0$, then the next line is $Qw = Pw$, and from there the rank one structure of P does the rest. Eigenvalue equals fixed vector is the more useful translation whenever the matrix has any explicit structure at all.

Layer 2: HOW

Block 1. **Cue.** $P = I - 2uu^T$ with u a unit vector, a rank one update of the identity. **Why natural.** A rank one update acts as the identity on a hyperplane, so its whole content is what it does to one direction. Computing $Pu = u - 2u(u^T u) = -u$ and $Pv = v$ for $v \perp u$ identifies P instantly as a reflection, and $P^2 = I$ follows either from that description or from one line of algebra. Knowing $P^{-1} = P$ is what lets you move P across an equation later without introducing an inverse you would have to analyse. **Link to the next block.** Since P is its own inverse, the fixed vector equation for PQ can be rearranged into a comparison of Q and P directly.

Block 2. **Cue.** The conclusion names an eigenvalue, and the eigenvalue is 1. **Why natural.** Eigenvalue 1 has a coordinate-free meaning: a nonzero fixed vector. That reading is superior here because you are being asked to prove existence, and existence proofs are cheapest when you can exhibit the object. Rewriting $PQw = w$ as $Qw = Pw$ removes the composition and turns the problem into a question about where the two maps agree. **Link to the next block.** You now have two explicit expressions to compare, and one of them, Pw, has a very restricted form.

Block 3. **Cue.** $Pw = w - 2(u^T w)u$, so $Pw - w$ always lies on the fixed line $\text{span}(u)$. **Why natural.** This is the definition-level consequence of rank one. Whatever w is, P can only move it parallel to u. So the equation $Qw = Pw$ can only hold if Q also moves w parallel to u, that is $(Q - I)w \in \text{span}(u)$. The unknown vector has been reduced from n degrees of freedom to a one dimensional condition without any computation. **Link to the next block.** A membership condition of the form $(Q - I)w \in \text{span}(u)$ is solved by inverting $Q - I$, and the hypothesis has just handed you that inverse.

Block 4. **Cue.** The hypothesis says 1 is not an eigenvalue of Q, which is the same as $\det(Q - I) \neq 0$. **Why natural.** A negative hypothesis is useless until you turn it into a positive tool. Not an eigenvalue means invertible means you may divide. So the hypothesis is not a condition to check at the end, it

is the engine that produces the candidate. Applying $(Q-I)^{-1}$ to u gives the unique line on which the equation can possibly hold, so the candidate is forced rather than guessed. **Link to the next block.** You have a candidate. All that remains is a consistency check on one scalar.

Block 5. Cue. The candidate satisfies $Qw = w + u$ exactly, and the requirement is $Qw = w - 2(u^T w)u$. **Why natural.** Comparing the two gives the single scalar equation $-2(u^T w) = 1$. Everything vectorial has already matched. Reducing an n dimensional problem to one scalar identity is the normal outcome of a rank one structure, and the moment you see it you should ask which hypothesis has not yet been spent, because that hypothesis must be what makes the scalar come out right. **Link to the next block.** The unused hypothesis is that Q is orthogonal, in the sense of preserving norms.

Block 6. Cue. $Qw = w + u$ and $\|Qw\| = \|w\|$. **Why natural.** Orthogonality is most usable in its metric form. Squaring an equation whose two sides have equal norms kills the $\|w\|^2$ terms and leaves a linear relation between $u^T w$ and $\|u\|^2$. Geometrically this is the statement that $0, w, Qw$ is an isosceles triangle, so the chord $Qw - w$ is perpendicular to the mean direction $Qw + w$; the value $-\frac{1}{2}$ is the projection of that isosceles condition onto the chord. Note that it is exactly here that u being a unit vector is used, and if u had norm c the same computation would give $u^T w = -c^2/2$, matching the rescaled reflection. **Link to the next block.** The scalar is correct, so substitute back and finish.

Block 7. Cue. A verified identity $Qw = Pw$ with $w \neq 0$. **Why natural.** Applying P to both sides and using $P^2 = I$ converts the agreement into the required eigenvector equation with no further work. Recording that the solution set was forced also gives the multiplicity for free, which costs one sentence and shows the grader you understood why the vector was unique.

Layer 3: WHY

The route is ideal because it uses each hypothesis exactly once, in its most usable form: P rank one gives the direction, $Q - I$ invertible gives the vector, Q orthogonal gives the scalar. When a proof consumes each hypothesis exactly once and in a different way, you can be confident it is the intended route and that nothing has been wasted.

Underneath, the identity $u^T(Q-I)^{-1}u = -\frac{1}{2}$ is doing all the work, and it is not an accident about u. It says the symmetric part of $(Q-I)^{-1}$ equals $-\frac{1}{2}I$ for every fixed-point-free orthogonal Q. Equivalently $A = (Q+I)(Q-I)^{-1}$ is skew-symmetric, which is the Cayley transform: the map $A \mapsto (I-A)(I+A)^{-1}$ is a bijection from skew-symmetric matrices onto orthogonal matrices without eigenvalue -1, and it is the rational parametrisation of the orthogonal group, the matrix analogue of the tangent half angle substitution. Recognising that the scalar identity is a shadow of the Cayley transform is what upgrades this problem from a one-off to a piece of equipment.

The parity explanation is the other half of the truth, and you should hold both. A real orthogonal matrix has all eigenvalues on the unit circle, with non-real ones in conjugate pairs whose product is 1; if 1 is not among them then the multiplicity of -1 has the same parity as n, so $\det Q = (-1)^n$. Fixed-point-free orthogonal maps therefore have a forced determinant, and multiplying by a reflection breaks that forcing, so the product must acquire a fixed vector. In the language of topology, an orthogonal map with no eigenvalue 1 has no fixed point on the sphere S^{n-1}, and a fixed-point-free self map of the sphere is homotopic to the antipodal map, whose degree is $(-1)^n$. The determinant is the degree. That is why the statement is really about orientation.

Discernment mechanism: when a conclusion is of the form some specific number is an eigenvalue, ask whether the matrix has structure that makes the eigenvector constructible. Rank one updates, projections, reflections, permutation-like structure, and circulants all do. If yes, construct.

If not, fall back on determinants, traces or parity. Train it by rewriting every eigenvalue statement you meet as a statement about a vector, and every non-eigenvalue hypothesis as an invertibility statement you are allowed to divide by.

The examiners nudged clearly: they wrote P in the explicit form $I - 2uu^T$ rather than saying let P be a reflection, which is an invitation to use the algebraic rank one form; and they gave the hypothesis as a negative statement about an eigenvalue, which is an invitation to convert it into invertibility.

Revolving concepts

Householder reflections. For a unit u, $P = I - 2uu^T$ satisfies $P = P^T$, $P^2 = I$, $Pu = -u$, $Pv = v$ for $v \perp u$, and $\det P = -1$. Every orthogonal matrix is a product of at most n Householder reflections; the minimum number needed is $n - \dim\ker(Q - I)$, so Q needs exactly n precisely when it has no eigenvalue 1. This gives a third proof: if neither Q nor PQ had eigenvalue 1, each would be a product of exactly n reflections and so both would have determinant $(-1)^n$, which contradicts $\det(PQ) = -\det Q$. See Book 1, field: Linear algebra, topic: Householder reflections and QR.

Spectrum of a real orthogonal matrix. All eigenvalues have modulus 1; non-real eigenvalues come in conjugate pairs; the real canonical form is a block diagonal matrix of 2×2 rotation blocks together with ± 1 entries. Consequences: $\det Q = \pm 1$; if n is odd and $\det Q = 1$ then 1 is an eigenvalue (Euler's rotation theorem); if 1 is not an eigenvalue then $\det Q = (-1)^n$. See Book 1, field: Linear algebra, topic: Orthogonal group and normal forms.

Rank one updates. $\det(A + vw^T) = \det A\,(1 + w^T A^{-1} v)$ (matrix determinant lemma) and $(A + vw^T)^{-1} = A^{-1} - \frac{A^{-1} v w^T A^{-1}}{1 + w^T A^{-1} v}$ (the Sherman and Morrison formula). Both say the same thing: a rank one perturbation changes everything through a single scalar $w^T A^{-1} v$. Recognising that scalar in a problem is usually the whole problem. See Book 1, field: Matrices and determinants, topic: Rank one updates.

Cayley transform. If A is skew-symmetric then $I + A$ is invertible and $Q = (I - A)(I + A)^{-1}$ is orthogonal without eigenvalue -1; conversely if Q is orthogonal without eigenvalue -1 then $A = (I - Q)(I + Q)^{-1}$ is skew. The version used here is the reflection of that statement about the eigenvalue 1: for Q orthogonal without eigenvalue 1, $(Q + I)(Q - I)^{-1}$ is skew, equivalently $(Q - I)^{-1} + \frac{1}{2} I$ is skew. See Book 1, field: Linear algebra, topic: Cayley transform.

Fixed vectors and iteration. There is no recursion in this problem, but the associated generalisation is worth recording: for any invertible M and rank one update $P = I - 2uu^T$, the fixed vectors of PM are exactly the solutions of $(M - I)w \in \operatorname{span}(u)$ together with the normalisation $u^T w = -\frac{1}{2} \cdot$ (a factor determined by M). The general principle is that composing with a rank one update reduces an n dimensional existence question to a single scalar equation, and the scalar is $u^T (M - I)^{-1} u$. See Book 1, field: Linear algebra, topic: Fixed vectors of perturbed maps.

Circuit Bank

- **WHEN** you see the hypothesis that some number is not an eigenvalue **DO** immediately rewrite it as the corresponding matrix being invertible, and plan to divide by it **BECAUSE** a negative spectral statement is unusable, an invertibility statement is a tool that constructs vectors
- **WHEN** you are asked to show that 1 is an eigenvalue **DO** try to exhibit a nonzero fixed vector rather than to show a determinant vanishes **BECAUSE** a produced vector is a complete proof with no lemma imported, and it also delivers the eigenspace dimension
- **WHEN** you see $I + vw^T$ or $I - 2uu^T$ **DO** note that the map moves every vector only in the direction v, and reduce the problem to one scalar **BECAUSE** rank one means the entire perturbation is carried by a single scalar, usually $w^T A^{-1} v$

- **WHEN** you have used all the structure except orthogonality and one scalar remains **DO** square the vector identity and use $\|Qx\| = \|x\|$ **BECAUSE** norm preservation is the metric form of orthogonality and is the cheapest way to extract a scalar relation from a vector equation
- **WHEN** you see a reflection composed with something **DO** think determinant and orientation as a sanity check on whether the claim can be true **BECAUSE** reflections have determinant -1, so any claim about a composition with a reflection is usually a parity statement in disguise

Backtesting

1. First useful cue: P is a rank one update, so $Pw - w$ is always parallel to u.
2. Method the expression suggested: rewrite $PQw = w$ as $Qw = Pw$, deduce $(Q - I)w \in \text{span}(u)$, invert $Q - I$, and check the one scalar with norm preservation.
3. Where you would hesitate, and the thought to reinforce: after writing $w = (Q - I)^{-1}u$, whether the normalisation can be arranged by scaling w. It cannot, because scaling w scales both sides of $(Q - I)w = -2(u^T w)u$ equally; the scalar identity is a genuine constraint and is exactly what orthogonality supplies.
4. The trap the question sets: pushing you toward determinants. The determinant route needs the spectral structure of real orthogonal matrices, which is fine if you know it and fatal if you half remember it, and it never produces the vector.
5. What makes you recognise it faster next time: the pairing of a rank one update with an invertibility hypothesis. That pairing always means the answer is a specific vector obtained by inverting the invertible part and applying it to the rank one direction.
6. The reusable decision rule: eigenvalue statements become fixed vector statements; not an eigenvalue becomes invertible; rank one becomes one scalar.

Mirror drills

1. Prove Euler's rotation theorem: if n is odd and Q is orthogonal with $\det Q = 1$, then 1 is an eigenvalue of Q. *Answer:* $\det(Q - I) = \det(Q - QQ^T) = \det Q \det(I - Q^T) = \det Q(-1)^n \det(Q - I)$, so for n odd and $\det Q = 1$ we get $\det(Q - I) = -\det(Q - I) = 0$.
2. Show that for real orthogonal Q without eigenvalue 1, the matrix $(Q + I)(Q - I)^{-1}$ is skew-symmetric. *Hint:* the problem proved $u^T(Q - I)^{-1}u = -\frac{1}{2}$ for every unit u, so the symmetric part of $(Q - I)^{-1}$ is $-\frac{1}{2}I$; now write $(Q + I)(Q - I)^{-1} = I + 2(Q - I)^{-1}$.
3. Let M be any real matrix with 1 not an eigenvalue, let u be a unit vector, and set $P = I - 2uu^T$. Show that 1 is an eigenvalue of PM if and only if $u^T(M - I)^{-1}u = -\frac{1}{2}$. *Answer:* the fixed vectors of PM are the nonzero w with $(M - I)w \in \text{span}(u)$, hence $w = c(M - I)^{-1}u$, and substituting into $Mw = w - 2(u^T w)u$ gives the stated scalar condition. Orthogonality of M is exactly what guarantees it.

7.10 B4. Manufacturing a mixed partial with Euler operators

Let $\mathcal{F}$ be the set of functions $f(x,y)$ that are twice continuously differentiable for $x \geq 1$, $y \geq 1$ and that satisfy the following two equations (where subscripts denote partial derivatives):

$$xf_x + yf_y = xy\ln(xy),$$
$$x^2 f_{xx} + y^2 f_{yy} = xy.$$

For each $f \in \mathcal{F}$, let

$$m(f) = \min_{s \geq 1} \left(f(s+1,s+1) - f(s+1,s) - f(s,s+1) + f(s,s)\right).$$

Determine $m(f)$, and show that it is independent of the choice of f.

Read it out loud

Read the quantity being minimised first, because it is the only thing you actually need. The expression $f(s+1,s+1) - f(s+1,s) - f(s,s+1) + f(s,s)$ is the mixed second difference of f over the unit square $[s,s+1]^2$. For a twice continuously differentiable f it equals $\iint_{[s,s+1]^2} f_{xy}$, by applying the fundamental theorem of calculus once in each variable. So the entire problem is: find f_{xy}. Nothing else about f matters, and in particular the mixed second difference annihilates every function of the form $u(x) + v(y)$, which is exactly why the answer can be independent of f.

Now read the two given equations and notice what is missing. You are given f_x, f_y combined with coefficients x and y, and f_{xx}, f_{yy} combined with coefficients x^2 and y^2. You are not given anything mixed. So the problem has a single mechanical task at its heart: manufacture f_{xy} out of pure derivatives. There is only one way to make a mixed derivative from first order data, namely to differentiate the first equation again, once in x and once in y, and to arrange the algebra so that the pure second derivatives cancel against the second equation.

The coefficients tell you which derivative to use. A first derivative multiplied by x and a second derivative multiplied by x^2 is the signature of the Euler operator $D = x\frac{\partial}{\partial x}$, which in the variable $u = \ln x$ is simply $\partial/\partial u$. That is the same signature you meet in the Cauchy and Euler differential equation $ax^2y'' + bxy' + cy = 0$, and the response is identical: pass to logarithms and the operators become plain derivatives with constant coefficients. Once you are in that language, the algebra to perform is the completed square $2DE = (D+E)^2 - D^2 - E^2$, and the two given equations supply precisely $D + E$ and (after correction) $D^2 + E^2$.

Fields embedded: multivariable calculus (mixed partials, Clairaut, iterated integrals), linear differential operators and their algebra, and a one line optimisation at the end. Within thirty seconds you should have written: target $= \iint f_{xy}$; $D = x\partial_x$, $E = y\partial_y$; $x^2 f_{xx} = D(D-1)f$; and the plan, apply $D + E$ to the first equation.

Approaches, ranked worst to best

1. **Solve the system for all f.** *Cue:* two partial differential equations and a family $\mathcal{F}$. *Circuit:* find the general solution, then compute. You can do it: the general element is $\frac{1}{2}(xy\ln(xy) - xy) + a\ln(x/y) + b$. But finding the general solution of a coupled system is a much bigger job than the question asks, and if you fail to prove you have found all of them you cannot conclude independence of f. Maximum ink, maximum risk.

2. **Guess a particular solution, then prove independence separately.** *Cue:* both equations are symmetric in x and y and involve xy. *Circuit:* posit $f = h(xy)$, solve for h, then classify the difference of two solutions. Substituting $f = h(xy)$ into the first equation gives $2xyh'(xy) = xy\ln(xy)$, hence $h(t) = \frac{1}{2}(t\ln t - t)$, which does satisfy both equations. Then show any two members of $\mathcal{F}$ differ by something with vanishing mixed second difference. This is a

genuinely good route and it has the bonus of proving $\mathcal{F}$ is nonempty, but it is two jobs instead of one.

3. **The magic linear combination.** *Cue:* you need f_{xy}, so differentiate. *Circuit:* form $x\partial_x(1) + y\partial_y(1) - (1) - (2)$ and watch the pure terms cancel. This is the shortest route in raw lines, and it works, but presented as a combination pulled out of the air it is unrepeatable. If you cannot say why those four terms with those signs, you have memorised an answer rather than installed a circuit.

4. **Euler operators, then complete the square.** *Cue:* coefficients x with f_x and x^2 with f_{xx}. *Circuit:* $D = x\partial_x$, $x^2\partial_{xx} = D(D-1)$, and $2DE = (D+E)^2 - D^2 - E^2$. The two given equations become $(D+E)f$ known and $D^2f + E^2f$ known. Applying $D+E$ to the first gives $(D+E)^2 f$, and subtracting delivers $2DEf = 2xyf_{xy}$. Every step is forced by the algebra of operators, and the same manipulation is exactly why the combination in the previous approach works. Best because it is short, motivated, and reusable.

The ideal solution

Write $D = x\dfrac{\partial}{\partial x}$ and $E = y\dfrac{\partial}{\partial y}$, and label the given equations

$$(1)\ \ Df + Ef = xy\ln(xy), \qquad (2)\ \ x^2 f_{xx} + y^2 f_{yy} = xy.$$

Since $D(Df) = x\partial_x(xf_x) = xf_x + x^2 f_{xx}$, we have the operator identity

$$x^2 f_{xx} = D^2 f - Df, \qquad \text{and similarly} \qquad y^2 f_{yy} = E^2 f - Ef.$$

Also $DEf = x\partial_x\big(yf_y\big) = xy\, f_{yx}$, and $EDf = xy\, f_{xy}$; since f is twice continuously differentiable, Clairaut's theorem gives $f_{xy} = f_{yx}$, so $DEf = EDf = xyf_{xy}$. All the expressions written here involve only first and second derivatives of f, so every step below is legitimate for $f \in \mathcal{F}$; no third derivative is ever required.

Substituting the operator identity into (2) and then using (1),

$$D^2 f + E^2 f = xy + \big(Df + Ef\big) = xy + xy\ln(xy). \tag{2′}$$

Now apply the operator $D+E$ to both sides of (1). On the left, $(D+E)(Df+Ef) = D^2f + E^2f + 2DEf$. On the right,

$$D\big(xy\ln(xy)\big) = x\big(y\ln(xy) + y\big) = xy\ln(xy) + xy,$$

and by the symmetry of $xy\ln(xy)$ in x and y the operator E gives the same, so

$$D^2 f + E^2 f + 2DEf = 2xy\ln(xy) + 2xy. \tag{1′}$$

Subtracting (2′) from (1′),

$$2DEf = xy\ln(xy) + xy, \qquad \text{that is} \qquad 2xy\, f_{xy} = xy\big(\ln(xy) + 1\big).$$

Since $x, y \geq 1$ we may divide by xy:

$$f_{xy}(x,y) = \tfrac{1}{2}\big(\ln x + \ln y + 1\big) \qquad \text{for all } x, y \geq 1.$$

This is the same for every $f \in \mathcal{F}$, which is where the independence of f comes from.

Now evaluate the mixed second difference. Fix $s \geq 1$ and put

$$\Delta(s) = f(s+1, s+1) - f(s+1, s) - f(s, s+1) + f(s, s).$$

For each fixed x, the fundamental theorem of calculus applied to $y \mapsto f(x,y)$ gives $f(x,s+1) - f(x,s) = \int_s^{s+1} f_y(x,y)\,dy$. Applying it again to $x \mapsto f_y(x,y)$ and using continuity of f_{xy} to interchange the (proper, continuous) iterated integrals,

$$\Delta(s) = \int_s^{s+1}\int_s^{s+1} f_{xy}(x,y)\,dx\,dy = \frac{1}{2}\int_s^{s+1}\int_s^{s+1} \big(\ln x + \ln y + 1\big)\,dx\,dy.$$

The square has unit area, and by symmetry the two logarithm terms contribute equally, so

$$\Delta(s) = \frac{1}{2}\left(2\int_s^{s+1} \ln x\,dx + 1\right) = \frac{1}{2} + \int_s^{s+1} \ln x\,dx.$$

Differentiating, $\Delta'(s) = \ln(s+1) - \ln s > 0$ for $s \geq 1$, so Δ is strictly increasing on $[1,\infty)$ and its minimum there is attained at $s = 1$. Since $\int_1^2 \ln x\,dx = \big[x\ln x - x\big]_1^2 = 2\ln 2 - 1$,

$$m(f) = \frac{1}{2} + 2\ln 2 - 1 = \boxed{2\ln 2 - \tfrac{1}{2}}$$

for every $f \in \mathcal{F}$, independently of f. (Numerically $m(f) \approx 0.8863$.)

Remark 7.10.1. $\mathcal{F}$ is nonempty: $f_0(x,y) = \frac{1}{2}\big(xy\ln(xy) - xy\big)$ satisfies $xf_{0,x} = yf_{0,y} = \frac{1}{2}xy\ln(xy)$ and $x^2 f_{0,xx} = y^2 f_{0,yy} = \frac{1}{2}xy$. Moreover $\mathcal{F} = \{f_0 + a\ln(x/y) + b : a,b \in \mathbb{R}\}$: if $f_0 + g \in \mathcal{F}$ then $g_{xy} = 0$, so $g = u(x) + v(y)$, and $xg_x + yg_y = 0$ says $xu'(x) = -yv'(y)$, a function of x equal to a function of y, hence a constant a. Every such g is a sum of a function of x and a function of y, so its mixed second difference vanishes, which is a second and completely explicit reason for the independence of f.

Layer 1: WHAT

Block 1. Read the target as a mixed second difference, hence as $\iint f_{xy}$ over the unit square. Conclude that only f_{xy} matters.

Block 2. Notice the data are pure derivatives with Euler coefficients, so introduce $D = x\partial_x$, $E = y\partial_y$ and record $x^2\partial_{xx} = D(D-1)$, $DE = xy\partial_{xy}$.

Block 3. Rewrite equation (2) as $D^2 f + E^2 f = xy + xy\ln(xy)$, using equation (1) to absorb the first order correction.

Block 4. Apply $D + E$ to equation (1) to obtain $(D+E)^2 f$, which is $D^2 f + E^2 f + 2DEf$.

Block 5. Subtract. The pure squares cancel by construction and $2xyf_{xy} = xy(\ln(xy) + 1)$ falls out. Independence of f is now proved, since f_{xy} is determined.

Block 6. Integrate over the unit square: $\Delta(s) = \frac{1}{2} + \int_s^{s+1} \ln x\,dx$.

Block 7. $\Delta' = \ln(s+1) - \ln s > 0$, so minimise at $s = 1$ and evaluate.

Juxtaposition. The nominal route reads the words let $\mathcal{F}$ be the set of functions satisfying, and starts solving the system for f. That is the consumption reflex: the question mentioned a family, so describe the family. It works, eventually, and it is three times the length. The bifurcation point is the first line you write after reading the definition of $m(f)$. If you write out the four values of f and start hunting for f itself, you are solving the wrong problem. If you write $\iint f_{xy}$, you have converted the entire question into a single unknown function and the rest is forced.

Layer 2: HOW

Block 1. **Cue.** Four values of f at the corners of a unit square with alternating signs. **Why natural.** That pattern is the definition of a second mixed difference, and the second mixed difference is the discrete analogue of $\partial^2/\partial x\partial y$. Applying the fundamental theorem of calculus in each variable turns it into $\iint f_{xy}$ exactly, not approximately. The immediate structural payoff is that any additively separable part of f contributes nothing, which is the mechanism that will make $m(f)$ independent of f. **Link to the next block.** You now need f_{xy} and nothing else, so look at the data and ask how to produce a mixed derivative.

Block 2. **Cue.** First derivatives carrying a factor x, second derivatives carrying a factor x^2. **Why natural.** That is the fingerprint of the Euler operator $D = x\partial_x$, whose powers are $D^2 = x^2\partial_{xx} + x\partial_x$. Equivalently, in the variable $u = \ln x$ the operator D is ∂_u: the substitution that makes a Cauchy and Euler differential equation constant coefficient is the same substitution that makes this system translation invariant. Introducing D and E costs one line and converts the two equations into statements about a commutative algebra of operators, where you can do ordinary algebra. **Link to the next block.** In that algebra you want DE, and you have $D + E$ and, after a correction, $D^2 + E^2$.

Block 3. **Cue.** Equation (2) contains $D^2 - D$ and $E^2 - E$, while equation (1) contains $D + E$. **Why natural.** The correction terms $-Df - Ef$ are already known from equation (1), so you can eliminate them and get a clean expression for $D^2f + E^2f$. Always normalise your data into the same basis before combining them; here the basis is $\{D + E, D^2 + E^2, DE\}$. **Link to the next block.** With $D + E$ and $D^2 + E^2$ both known, the identity $2DE = (D + E)^2 - D^2 - E^2$ tells you exactly what to do next.

Block 4. **Cue.** You possess $(D + E)f$ as an explicit function and you want $(D + E)^2 f$. **Why natural.** Applying $D + E$ to a known function is just differentiation of the right hand side, which is elementary. This is the standard move for extracting higher information from a first order equation: differentiate the equation, not the unknown. Note that only second derivatives of f appear on the left, so the C^2 hypothesis is enough and no regularity is being assumed that you were not given. **Link to the next block.** The two expressions now share the terms $D^2f + E^2f$, so subtract.

Block 5. **Cue.** Two equations whose difference isolates $2DEf$. **Why natural.** This is completing the square in operator form, the same identity as $2ab = (a + b)^2 - a^2 - b^2$. The reason the setter's system is solvable at all is that they gave you the two symmetric functions of D and E that determine the product. Dividing by xy is legal because the domain is $x, y \geq 1$, and saying so is worth a clause. **Link to the next block.** You have f_{xy} explicitly; feed it into the integral from Block 1.

Block 6. **Cue.** $\iint$ of $\frac{1}{2}(\ln x + \ln y + 1)$ over a unit square. **Why natural.** The integrand is separable, so the double integral splits into single integrals, and the square has unit side so the constant term contributes $\frac{1}{2}$ and each logarithm contributes $\frac{1}{2}\int_s^{s+1} \ln$. Symmetry in x and y means you compute one integral, not two. **Link to the next block.** You have a closed form in s; minimise it.

Block 7. **Cue.** A minimum over the closed unbounded interval $s \geq 1$. **Why natural.** Differentiating under the integral sign is unnecessary here since $\frac{d}{ds}\int_s^{s+1} \ln = \ln(s + 1) - \ln s$ by the fundamental theorem, and this is positive because ln is strictly increasing. A strictly increasing continuous function on $[1, \infty)$ attains its minimum at the left endpoint, so the minimum exists and equals $\Delta(1)$. Say that the minimum is attained; a question that writes min rather than inf expects you to notice.

Layer 3: WHY

The route is ideal because it identifies the single invariant of the family. The problem hands you a set $\mathcal{F}$ and asks for a quantity that turns out not to depend on which member you take. Whenever that happens, the honest question is: which functional of f is pinned down by the data? Here the data pin down f_{xy} exactly, and the quantity asked for depends only on f_{xy}. Once those two facts are in place the independence is not a surprise to be verified at the end but the reason the question was posed. Contrast with the route that solves for all of $\mathcal{F}$: it reaches the same conclusion by describing a two parameter family and then observing that the two parameters happen to cancel. Same answer, far less understanding.

Underneath sits the algebra of the operators $D = x\partial_x$ and $E = y\partial_y$. They commute, they generate scalings, and under $x = e^u$, $y = e^v$ they become ∂_u and ∂_v. In those coordinates the

problem reads: given $F_u + F_v$ and $F_{uu} + F_{vv}$, find F_{uv}; and the answer is the identity $2F_{uv} = (\partial_u + \partial_v)(F_u + F_v) - (F_{uu} + F_{vv})$. That is why the problem is well posed with exactly two equations: the two symmetric functions $p = D + E$ and $q = D^2 + E^2$ determine the elementary symmetric product DE through $2DE = p^2 - q$, exactly as $2ab = (a+b)^2 - (a^2 + b^2)$ recovers a product from a sum and a sum of squares. A setter who wants to hide f_{xy} behind two innocent equations has essentially no other choice.

The second truth is the exactness of the discrete to continuous translation. The second mixed difference over $[s, s+1]^2$ is not approximately f_{xy} at some point, it is the exact double integral of f_{xy} over the square, by two applications of the fundamental theorem of calculus. That exactness is what lets you compute $m(f)$ in closed form rather than estimate it. The same reflex applies whenever you see an alternating sum of function values at the corners of a box: in one dimension a difference is $\int f'$, in two a mixed second difference is $\iint f_{xy}$, in n dimensions an alternating sum over 2^n corners is the integral of $\partial^n f / \partial x_1 \cdots \partial x_n$.

Discernment mechanism: read the objective before the data. The objective told you that only f_{xy} matters; the data then had exactly one job. Candidates who read the data first spend their time on the family $\mathcal{F}$ because that is what the first half of the question talks about. Train this by making it a rule: on any problem that defines a set and then asks for a number, find out which feature of the set members the number depends on before you study the set.

The examiners nudged with the coefficients. They could have written the first equation as $f_x + f_y = \ldots$, and they did not; the factors x, y, x^2, y^2 are there to tell you to use D and E, and the perfect symmetry of both equations in x and y is there to tell you that only the symmetric combinations matter.

Revolving concepts

Mixed second differences. For f of class C^2, $f(a+h, b+k) - f(a+h, b) - f(a, b+k) + f(a, b) = \int_a^{a+h} \int_b^{b+k} f_{xy}$. Consequences: the mixed second difference vanishes identically if and only if f is additively separable, $f = u(x) + v(y)$; it is the model case of the inclusion and exclusion pattern over the 2^n vertices of a box, which for n variables computes $\int \partial^n f / \partial x_1 \cdots \partial x_n$. See Book 1, field: Calculus and integrals, topic: Mixed partials and differences.

Clairaut's theorem. If f_{xy} and f_{yx} exist and are continuous near a point then they are equal there. This is what licenses $DEf = EDf$ and the interchange of the order of integration in the iterated integral (both integrands being continuous on a compact square, so Fubini for continuous functions applies). See Book 1, field: Real analysis, topic: Equality of mixed partials.

Euler operators and the Cauchy and Euler equation. With $D = x\frac{d}{dx}$ one has $x^k \frac{d^k}{dx^k} = D(D-1)\cdots(D-k+1)$, and under $x = e^u$, $D = \frac{d}{du}$. Hence any equation whose kth derivative carries the coefficient x^k becomes constant coefficient in logarithmic variables. The archetype is $ax^2y'' + bxy' + cy = 0$, solved by $y = x^r$ with $ar(r-1) + br + c = 0$, which is the indicial equation. Recognising the pattern is the trigger for the substitution. See Book 1, field: Sequences and recurrences, topic: Euler operators and equidimensional equations.

Symmetric functions of commuting operators. For commuting D, E one has $2DE = (D+E)^2 - D^2 - E^2$, and more generally the power sums $p_k = D^k + E^k$ determine the elementary symmetric functions through Newton's identities. The generalised statement worth carrying: if a problem gives you p_1 and p_2 for two commuting objects, it has given you their product, and if it gives you $p_1, \ldots, p_k$ for k objects it has given you their entire characteristic polynomial. See Book 1, field: Polynomials, topic: Newton's identities.

Recursion note. There is no sequence here, but the operator relation is a recursion in disguise: define $T_k = D^k f + E^k f$; then applying $D + E$ gives $T_{k+1} = (D+E)T_k - DE\,T_{k-1}$, the standard second order recursion satisfied by power sums of two quantities with known sum and product. In this problem you run it once, in reverse, to extract DE from T_1 and T_2. The general solution method

is Newton's identities, and the same recursion is what makes power sums of roots computable from a polynomial's coefficients.

Monotone functions defined by moving integrals. $s \mapsto \int_s^{s+1} g$ has derivative $g(s+1) - g(s)$, so it is increasing exactly when g is increasing on the relevant range. This is the cleanest way to compare sliding window averages and it removes any need to differentiate under the integral sign. See Book 1, field: Calculus and integrals, topic: Fundamental theorem with variable limits.

Circuit Bank

- **WHEN** you see an alternating sum of values at the corners of a rectangle **DO** write it as the double integral of the mixed partial over that rectangle **BECAUSE** the mixed second difference is exactly $\iint f_{xy}$, so the problem reduces to one function and all additively separable parts vanish
- **WHEN** you see first derivatives with coefficient x and second derivatives with coefficient x^2 **DO** introduce $D = x\partial_x$ and recall $x^2\partial_{xx} = D(D-1)$, equivalently substitute $x = e^u$ **BECAUSE** these are the equidimensional (Euler) operators, which become constant coefficient derivatives in logarithmic variables, so the equations become algebra
- **WHEN** you need a mixed derivative but are given only pure ones **DO** differentiate the first order equation again with the same operator and use $2DE = (D+E)^2 - D^2 - E^2$ **BECAUSE** a sum and a sum of squares determine the product, so two symmetric equations are exactly enough to isolate the cross term
- **WHEN** you are asked to show a quantity is independent of the choice of object **DO** identify which functional of the object the quantity depends on, and prove the data determine that functional **BECAUSE** independence is a statement about invariants, and finding the invariant is shorter and more informative than classifying all objects
- **WHEN** you must minimise a sliding integral $\int_s^{s+1} g$ **DO** differentiate to get $g(s+1) - g(s)$ and read off monotonicity **BECAUSE** the fundamental theorem with variable limits makes this immediate, and monotonicity moves the minimum to an endpoint where it is attained

Backtesting

1. First useful cue: the four corner values with alternating signs, which is a mixed second difference and therefore $\iint f_{xy}$.
2. Method the expression suggested: introduce the Euler operators, rewrite both equations in terms of $D+E$ and D^2+E^2, and extract $2DE$ by completing the square.
3. Where you would hesitate, and the thought to reinforce: worrying that applying $D+E$ to equation (1) needs third derivatives. It does not: $(D+E)(Df+Ef)$ involves only second derivatives of f, and Clairaut makes the two mixed ones equal. State that in the write-up, because a grader will look for it.
4. The trap the question sets: the definition of $\mathcal{F}$ invites you to solve for f. The word min invites you to worry about attainment, which is real but cheap; the real trap is spending the time on the family.
5. What makes you recognise it faster next time: the coefficient pattern x with f_x and x^2 with f_{xx}. That pattern has exactly one meaning, and it is logarithmic variables.
6. The reusable decision rule: corner alternating sum means mixed partial; equidimensional coefficients mean Euler operators; sum plus sum of squares means product.

Mirror drills

1. Solve $x^2y'' + xy' - y = 0$ for $x > 0$. *Answer:* with $x = e^t$ the equation becomes $\ddot{y} - y = 0$, so $y = Ae^t + Be^{-t} = Ax + B/x$.
2. Let g be C^2 on $\mathbb{R}^2$ with $g_x + g_y = e^{x+y}(x+y)$ and $g_{xx} + g_{yy} = e^{x+y}(x+y+1)$. Find g_{xy}. *Answer:* $\frac{1}{2}e^{x+y}(x+y+1)$; this is the present problem written in logarithmic coordinates.
3. Show that a C^2 function on a rectangle has vanishing mixed second difference over every subrectangle if and only if it is of the form $u(x) + v(y)$. *Hint:* vanishing for all subrectangles forces $\iint f_{xy} = 0$ over every subrectangle, hence $f_{xy} \equiv 0$ by continuity; then integrate.

7.11 B5. Interpolating Fibonacci and the missing top difference

Let F_m be the mth Fibonacci number, defined by $F_1 = F_2 = 1$ and $F_m = F_{m-1} + F_{m-2}$ for all $m \geq 3$. Let $p(x)$ be the polynomial of degree 1008 such that $p(2n+1) = F_{2n+1}$ for $n = 0, 1, 2, \ldots, 1008$. Find integers j and k such that $p(2019) = F_j - F_k$.

Read it out loud

A polynomial of degree N is being pinned down by its values at $N+1$ equally spaced points, and then evaluated at the next point of the same arithmetic progression. Say that aloud and the whole problem is already classified. Equally spaced nodes plus one step beyond the last node is the exact setting of the finite difference calculus, not of Lagrange interpolation. The reason is the single fact that a polynomial of degree N has vanishing $(N+1)$st difference, so its value at the next node is a fixed integer combination of the known values with binomial coefficients.

The nodes are the odd numbers $1, 3, 5, \ldots, 2017$ and the target is 2019, so substitute $x = 2t+1$ once and forget the odd numbers: with $P(t) = p(2t+1)$ you have a polynomial of degree $N = 1008$ with $P(n) = F_{2n+1}$ for $n = 0, \ldots, N$, and you want $P(N+1)$. Uniform notation is not cosmetic; it is what lets you see that the answer will be a statement about N.

Now look at the data through the difference operator. The Fibonacci recurrence in the form $F_{m+2} - F_m = F_{m+1}$ says exactly that differencing the sequence $n \mapsto F_{2n+c}$ produces the sequence $n \mapsto F_{2n+c+1}$. So every row of the difference table of this data is the Fibonacci sequence again, shifted by one index. That is the observation the problem is built on. The difference table of the data never terminates, but the difference table of a degree N polynomial must, and the answer records precisely the discrepancy between the two.

Fields embedded: finite differences and interpolation (the discrete calculus), linear recurrences and Binet's formula, and binomial identities. Within thirty seconds you should have written $P(t) = p(2t+1)$, the difference identity $F_{m+2} - F_m = F_{m+1}$, and the plan: build the difference table, use Newton's forward formula at $t = N+1$.

Approaches, ranked worst to best

1. **Find p explicitly.** *Cue:* a polynomial is its coefficients. *Circuit:* solve the linear system for 1009 coefficients. There is no version of this that a human completes. It is the purest form of doing what the words say without reading the structure.

2. **Lagrange interpolation with Binet.** *Cue:* 1009 points and a value, so use the interpolation formula. *Circuit:* write $p(2019) = \sum p(x_i) \prod \frac{2019 - x_j}{x_i - x_j}$, then split F by Binet. The products collapse to $(-1)^{1008-n}\binom{1009}{n}$, and applying the binomial theorem to $\sum_n \gamma^{2n+1}(-1)^{1008-n}\binom{1009}{n}$ with $\gamma^2 = \gamma + 1$ gives $\gamma^{2019} - \gamma^{1010}$. It is a real solution, but it carries a sign error in every second line if you are not careful, and it uses interpolation machinery where difference machinery is native to equally spaced nodes.

3. **Difference table with edge climbing.** *Cue:* equally spaced nodes and a recursive data set. *Circuit:* build $\Delta^k P$, note the top row is constant, then climb back along the boundary. Since $\Delta^k P(n) = F_{2n+k+1}$ and $\Delta^N P$ is a constant equal to F_{N+1}, you can compute the first unknown entry in each row from the row below, giving $P(N+1) = \sum_{j=N+1}^{2N+1} F_j$, and then $\sum_{j=1}^{m} F_j = F_{m+2} - 1$ finishes it. Fully elementary, no Binet, no binomial identity. It costs two inductions and a telescoping sum, and it hides the reason the answer is a difference of two Fibonacci numbers.

4. **Newton's forward difference formula.** *Cue:* a polynomial of degree N evaluated one node past the data. *Circuit:* $P(N+1) = \sum_k \binom{N+1}{k} \Delta^k P(0)$, with $\Delta^k P(0) = F_{k+1}$. The sum is the

complete binomial Fibonacci sum $\sum_{k=0}^{N+1}\binom{N+1}{k}F_{k+1} = F_{2N+3}$ with exactly one term missing, the top one, F_{N+2}. Three lines, and it explains the shape of the answer: the polynomial reproduces the Fibonacci continuation except for the single term its degree forbids. Best.

The ideal solution

Write $N = 1008$ and $P(t) = p(2t+1)$, a polynomial in t of degree N with

$$P(n) = F_{2n+1} \qquad (n = 0, 1, \ldots, N),$$

and note $p(2019) = P(N+1)$ since $2(N+1)+1 = 2019$. Let Δ denote the forward difference, $(\Delta G)(t) = G(t+1) - G(t)$.

Lemma 7.11.1. *For $0 \le k \le N$ and $0 \le n \le N-k$ we have $(\Delta^k P)(n) = F_{2n+k+1}$.*

Proof. Induction on k. For $k = 0$ this is the hypothesis. Assume it for $k-1$. For $0 \le n \le N-k$ both n and $n+1$ lie in $[0, N-(k-1)]$, so

$$(\Delta^k P)(n) = (\Delta^{k-1}P)(n+1) - (\Delta^{k-1}P)(n) = F_{2n+k+2} - F_{2n+k},$$

and the Fibonacci recurrence in the form $F_{m+2} - F_m = F_{m+1}$ (valid for $m \ge 1$, and here $m = 2n+k \ge 1$ unless $n = k = 0$, where $F_3 - F_1 = F_2$ holds directly) gives F_{2n+k+1}. □

In particular $(\Delta^k P)(0) = F_{k+1}$ for $0 \le k \le N$, and since $\deg P = N$ we have $\Delta^k P \equiv 0$ for $k > N$.

Lemma 7.11.2 (Newton's forward difference formula)**.** *For any function G on the integers and any integer $t \ge 0$, $G(t) = \sum_{k\ge 0}\binom{t}{k}(\Delta^k G)(0)$, the sum being finite when G is a polynomial.*

Proof. Let E be the shift, $(EG)(t) = G(t+1)$. Then $E = I + \Delta$ as operators, and I and Δ commute, so the binomial theorem for commuting operators gives $E^t = (I+\Delta)^t = \sum_{k=0}^{t}\binom{t}{k}\Delta^k$ for every nonnegative integer t. Evaluating both sides at 0 gives the formula. If G is a polynomial of degree N then $\Delta^k G \equiv 0$ for $k > N$ and the sum terminates. □

Applying Lemma 7.11.2 at $t = N+1$ and inserting Lemma 7.11.1,

$$P(N+1) = \sum_{k=0}^{N}\binom{N+1}{k}(\Delta^k P)(0) = \sum_{k=0}^{N}\binom{N+1}{k}F_{k+1}. \tag{$*$}$$

Lemma 7.11.3. *For every integer $M \ge 0$, $\displaystyle\sum_{k=0}^{M}\binom{M}{k}F_{k+1} = F_{2M+1}$.*

Proof. Let $\alpha = \frac{1+\sqrt{5}}{2}$ and $\beta = \frac{1-\sqrt{5}}{2}$, the roots of $z^2 = z+1$, so that $F_m = \frac{\alpha^m - \beta^m}{\sqrt{5}}$ (Binet). For $\gamma \in \{\alpha, \beta\}$ we have $1 + \gamma = \gamma^2$, hence by the binomial theorem

$$\sum_{k=0}^{M}\binom{M}{k}\gamma^{k+1} = \gamma(1+\gamma)^M = \gamma\cdot\gamma^{2M} = \gamma^{2M+1}.$$

Subtracting the two cases and dividing by $\sqrt{5}$ gives the claim. □

Now apply Lemma 7.11.3 with $M = N+1$:

$$F_{2N+3} = \sum_{k=0}^{N+1}\binom{N+1}{k}F_{k+1} = \left(\sum_{k=0}^{N}\binom{N+1}{k}F_{k+1}\right) + \binom{N+1}{N+1}F_{N+2},$$

so by $(*)$,

$$P(N+1) = F_{2N+3} - F_{N+2}.$$

With $N = 1008$ this reads

$$p(2019) = F_{2019} - F_{1010}, \qquad \text{so} \qquad \boxed{(j,k) = (2019,\ 1010)}.$$

Remark 7.11.4. Sanity check on small N. For $N = 1$, p is the line through $(1, F_1) = (1, 1)$ and $(3, F_3) = (3, 2)$, so $p(5) = 3 = 5 - 2 = F_5 - F_3$, matching $F_{2N+3} - F_{N+2}$. For $N = 2$, $P(0), P(1), P(2) = 1, 2, 5$ has second difference 2, so $P(3) = 10 = 13 - 3 = F_7 - F_4$, again matching.

Remark 7.11.5. The pair is unique in positive integers. Write $V = F_{2019} - F_{1010} > 0$ and suppose $F_j - F_k = V$. Since $V > 0$ we may take $k < j$, so $F_j - F_k \geq F_j - F_{j-1} = F_{j-2}$. If $j \geq 2021$ then $F_j - F_k \geq F_{2019} > V$, and if $j \leq 2018$ then $F_j - F_k \leq F_{2018} < F_{2019} - F_{1010} = V$ because $F_{2019} - F_{2018} = F_{2017} > F_{1010}$. So $j \in \{2019, 2020\}$. For $j = 2019$ we need $F_k = F_{1010}$, giving $k = 1010$. For $j = 2020$ we would need $F_k = F_{2020} - F_{2019} + F_{1010} = F_{2018} + F_{1010}$, which lies strictly between F_{2018} and $F_{2018} + F_{2017} = F_{2019}$ and so is not a Fibonacci number.

Layer 1: WHAT

Block 1. Normalise the nodes: $P(t) = p(2t + 1)$ has degree $N = 1008$, matches F_{2n+1} at $n = 0, \ldots, N$, and the target is $P(N + 1)$.

Block 2. Read the Fibonacci recurrence as a difference statement: $F_{m+2} - F_m = F_{m+1}$, so differencing the data shifts the Fibonacci index by one.

Block 3. Hence $(\Delta^k P)(0) = F_{k+1}$: the leading column of the difference table is the Fibonacci sequence itself.

Block 4. Use $E = I + \Delta$ to extrapolate one node: $P(N + 1) = \sum_k \binom{N+1}{k} F_{k+1}$, the sum truncated at $k = N$ because $\deg P = N$.

Block 5. Recognise the untruncated sum: $\sum_{k=0}^{M} \binom{M}{k} F_{k+1} = F_{2M+1}$, which is $1 + \gamma = \gamma^2$ under Binet.

Block 6. The truncation removes exactly the top term F_{N+2}, so $P(N + 1) = F_{2N+3} - F_{N+2}$; substitute $N = 1008$.

Juxtaposition. The nominal route sees 1009 interpolation conditions and reaches for Lagrange, then fights the product $\prod_{j \neq n} \frac{2019-(2j+1)}{(2n+1)-(2j+1)}$ into a binomial coefficient with an alternating sign, then applies Binet, then applies the binomial theorem. It works. The bifurcation point is the moment you notice the nodes are equally spaced. Equally spaced nodes make the difference operator the native tool, and the difference operator converts the Fibonacci recurrence into a triviality, whereas Lagrange interpolation ignores the spacing and forces you to rebuild the binomial coefficients by hand from a ratio of products.

Layer 2: HOW

Block 1. Cue. Nodes $1, 3, 5, \ldots$ and a target 2019 in the same progression. **Why natural.** Any statement about a polynomial on an arithmetic progression is a statement about the polynomial composed with a linear map, and the composition is again a polynomial of the same degree. Substituting once removes an irrelevant factor of two from every subsequent line. It also reveals the real index set $n = 0, \ldots, N$ and the real target $n = N + 1$, which is the shape that finite difference theory is written in. **Link to the next block.** With integer nodes $0, 1, \ldots, N$ you are in the discrete calculus, so ask what the difference operator does to the data.

Block 2. Cue. Data given by a linear recurrence, and a difference operator waiting to act on it. **Why natural.** The Fibonacci recurrence can be written $F_{m+2} = F_{m+1} + F_m$ or, rearranged, $F_{m+2} - F_m = F_{m+1}$. The second form is the one adapted to your data because your data steps the index by two, and the difference of two data points is a difference of Fibonacci numbers two apart. Rearranging a recurrence into the form the problem needs is not a manipulation, it is reading. In the Binet picture, Δ multiplies α^{2n} by $\alpha^2 - 1 = \alpha$, so Δ acts exactly as multiplication by the golden ratio. **Link to the next block.** Iterating gives every row of the difference table, so compute the leading column.

Block 3. **Cue.** The difference table of a polynomial has finitely many nonzero rows, and its leading column is what an extrapolation formula consumes. **Why natural.** Newton's forward formula uses only $(\Delta^k P)(0)$. Lemma 7.11.1 at $n = 0$ gives F_{k+1}, so the leading column is $1, 1, 2, 3, 5, \ldots$: the entire difference table is the Fibonacci sequence, which is the strongest possible sign that you are on the intended route. Note carefully the range restriction $n \leq N - k$: values of $\Delta^k P$ outside that range are not given by the Fibonacci pattern, and confusing the two is the way to lose this problem. **Link to the next block.** You have the leading column; now you need the formula that turns it into a value at a new point.

Block 4. **Cue.** A polynomial of known degree, values on a full block of consecutive integers, and a request for the next value. **Why natural.** The shift operator factors as $E = I + \Delta$, and since I and Δ commute, $E^t = (I + \Delta)^t$ expands by the binomial theorem. This is the discrete analogue of a Taylor expansion, with $\binom{t}{k}$ in place of $t^k/k!$ and Δ^k in place of the kth derivative. The degree hypothesis enters exactly once, to truncate the sum at $k = N$, and that truncation is the entire content of the answer. **Link to the next block.** You have a binomial sum of Fibonacci numbers; identify it.

Block 5. **Cue.** $\sum \binom{M}{k} F_{k+1}$, a binomial transform of the Fibonacci sequence. **Why natural.** Any identity of this shape is proved in one line by Binet, because F is a linear combination of two geometric sequences and the binomial theorem loves geometric sequences. The specific miracle, $\sum \binom{M}{k} \gamma^{k+1} = \gamma(1 + \gamma)^M$, becomes γ^{2M+1} only because $1 + \gamma = \gamma^2$, which is the defining equation of the golden ratio. So the identity is not a coincidence about Fibonacci numbers, it is the defining property of α in binomial clothing. **Link to the next block.** The identity is for the full range $0 \leq k \leq M$; your sum stops one short, so compare and subtract.

Block 6. **Cue.** A complete binomial identity and an incomplete sum differing by one term. **Why natural.** The missing term is $\binom{N+1}{N+1} F_{N+2} = F_{N+2}$. That is why the answer is a difference of two Fibonacci numbers rather than one: the polynomial faithfully reproduces the Fibonacci continuation to the extent that its difference table allows, and the single row it is forbidden to have (the $(N+1)$st difference, which must vanish for a degree N polynomial but equals F_{N+2} for the true Fibonacci data) is subtracted off. The answer is the error term of a truncation, and knowing that, you could have predicted the shape before computing.

Layer 3: WHY

The mathematical truth here is that the finite difference operator acts on Fibonacci data as multiplication by the golden ratio. Under Binet, the data are $P(n) = \frac{\alpha^{2n+1} - \beta^{2n+1}}{\sqrt{5}}$, and Δ multiplies γ^{2n} by $\gamma^2 - 1 = \gamma$. So on the two dimensional space spanned by the sequences (α^{2n}) and (β^{2n}), the operator Δ is diagonal with eigenvalues α and β. Everything in this problem is that one sentence. The difference table is the Fibonacci sequence because applying Δ shifts the index; Newton's formula computes $(I + \Delta)^{N+1}$, which on the α eigenvector is $(1 + \alpha)^{N+1} = \alpha^{2N+2}$, which is the Fibonacci continuation F_{2N+3} after multiplying by the remaining α; and the truncation at $k = N$ removes exactly the top eigen-contribution F_{N+2}.

That gives the discernment mechanism. The polynomial p is a degree N approximation to a sequence that satisfies a linear recurrence with eigenvalues α, β. Polynomials cannot satisfy that recurrence, so they must eventually disagree with it, and the first disagreement occurs at the first node beyond the data, in a controlled way: the disagreement is precisely the $(N+1)$st difference the polynomial is not allowed to have. Whenever you meet a polynomial interpolating a recurrent sequence and are asked for the next value, expect the answer to be the true next value minus the $(N+1)$st difference of the true sequence, and expect that difference to be another term of the same sequence when the recurrence is nice.

The training instruction is concrete. Equally spaced nodes means finite differences, never Lagrange. Build the operator dictionary once and keep it: $E = I + \Delta$; Δ^{N+1} kills degree N; $\Delta^{N+1}P(0) = \sum_i (-1)^{N+1-i}\binom{N+1}{i}P(i)$; Newton forward formula $P(t) = \sum \binom{t}{k}\Delta^k P(0)$; and the discrete Taylor analogy that makes all of them memorable. The single most common Putnam use of this dictionary is exactly this problem's shape, a degree N interpolant of a nice sequence evaluated one step out.

The examiners nudged three times. They chose odd nodes so that the data would obey $F_{m+2} - F_m = F_{m+1}$ rather than the raw recurrence, which is what makes the difference table clean. They chose the degree to be exactly one less than the number of nodes, so the interpolant is unique and the difference table terminates exactly at the right place. And they asked for the answer in the form $F_j - F_k$ rather than as a number, which is a broad hint that the answer is a truncation error and not a fresh quantity.

Revolving concepts

Finite difference calculus. $(\Delta G)(t) = G(t+1) - G(t)$; $E = I + \Delta$; Δ lowers the degree of a polynomial by exactly one, so $\Delta^{N+1}P \equiv 0$ for $\deg P \le N$; $\Delta^k G(0) = \sum_{i=0}^{k}(-1)^{k-i}\binom{k}{i}G(i)$; Newton's forward formula $G(t) = \sum_k \binom{t}{k}\Delta^k G(0)$; and the basis of binomial polynomials $\binom{t}{k}$ with $\Delta\binom{t}{k} = \binom{t}{k-1}$, which is the discrete analogue of differentiating $t^k/k!$. See Book 1, field: Polynomials, topic: Finite differences and interpolation.

Linear recurrences and Binet. A sequence satisfying $G_m = aG_{m-1} + bG_{m-2}$ has $G_m = A\lambda_1^m + B\lambda_2^m$ when the characteristic roots λ_i of $z^2 = az + b$ are distinct, and $(A + Bm)\lambda^m$ when they coincide. For Fibonacci, $\alpha, \beta = \frac{1\pm\sqrt{5}}{2}$ and $F_m = \frac{\alpha^m - \beta^m}{\sqrt{5}}$. The identity used repeatedly here is $1 + \gamma = \gamma^2$, the characteristic equation itself. See Book 1, field: Sequences and recurrences, topic: Linear recurrences with constant coefficients.

The generalised recursion, and its generalisation. The data satisfy the second order recursion $G_{n+1} = 3G_n - G_{n-1}$ where $G_n = F_{2n+1}$, since α^2 and β^2 are the roots of $z^2 - 3z + 1 = 0$. More generally, if a sequence satisfies a linear recurrence with characteristic roots λ_i, then Δ acts on the corresponding geometric components with eigenvalues $\lambda_i - 1$, so the difference table is again a combination of geometric sequences and every row is computable in closed form. The general solution method for problems of this type: diagonalise Δ on the span of the geometric components, apply $(I + \Delta)^t$, and truncate at the degree. This handles $\sum \binom{M}{k}G_k$ for any linear recurrent G, giving $\sum_i A_i\lambda_i(1 + \lambda_i)^M$. See Book 1, field: Sequences and recurrences, topic: Binomial transforms.

Binomial transforms of Fibonacci. $\sum_k \binom{M}{k}F_k = F_{2M}$ and $\sum_k \binom{M}{k}F_{k+1} = F_{2M+1}$, both instances of $1 + \gamma = \gamma^2$. Combinatorially, F_{n+1} counts tilings of a $1 \times n$ strip by squares and dominoes, and the identity has a direct tiling proof by splitting on the number of squares among the first M tiles. See Book 1, field: Combinatorics, topic: Fibonacci identities.

Telescoping Fibonacci sums. $\sum_{j=1}^{m} F_j = F_{m+2} - 1$, from $F_j = F_{j+2} - F_{j+1}$. This is what the elementary difference table route needs at the end, and it is worth having alongside $\sum F_j^2 = F_m F_{m+1}$ and $\sum F_{2j-1} = F_{2m}$. See Book 1, field: Sequences and recurrences, topic: Standard Fibonacci sums.

Circuit Bank

- **WHEN** you see a polynomial of degree N prescribed at $N + 1$ equally spaced points and evaluated at another such point **DO** build the difference table and use $P(t) = \sum_k \binom{t}{k}\Delta^k P(0)$ **BECAUSE** $E = I + \Delta$ makes extrapolation a binomial sum, and the degree hypothesis enters only as the truncation of that sum, which is where the answer comes from

- **WHEN** you see data sampled at every other index of a recurrent sequence **DO** rewrite the recurrence in the form matching your sampling, here $F_{m+2} - F_m = F_{m+1}$ **BECAUSE** a

recurrence is a relation between whichever indices you need, and rearranging it to match the difference operator is what makes the difference table computable

- **WHEN** you see $\sum_k \binom{M}{k} G_k$ with G satisfying a linear recurrence **DO** split G into geometric components by Binet and use the binomial theorem **BECAUSE** a binomial transform of a geometric sequence is another geometric sequence with ratio $1+\lambda$, so the whole sum collapses without any combinatorial argument
- **WHEN** you are asked for an answer in the form (something) minus (something) **DO** expect a truncation or an error term, and look for the complete identity that your sum falls short of **BECAUSE** answers presented as differences are usually a clean object minus the piece the hypothesis forbids, and naming the clean object first makes the computation short
- **WHEN** you have a general parameter hidden behind a specific number such as 1008 **DO** solve for general N and substitute at the end **BECAUSE** the general computation is no harder, it protects you from arithmetic slips, and it lets you sanity check on $N = 1$ and $N = 2$ by hand

Backtesting

1. First useful cue: the interpolation nodes are equally spaced and the target is the next node, which is the native setting of finite differences.
2. Method the expression suggested: substitute $t = (x-1)/2$, build the difference table using $F_{m+2} - F_m = F_{m+1}$, apply Newton's forward formula, and recognise the resulting binomial sum.
3. Where you would hesitate, and the thought to reinforce: whether $\Delta^k P(n) = F_{2n+k+1}$ holds for all n. It does not; it holds only for $n \le N - k$, because beyond that range the polynomial and the Fibonacci sequence part company. That parting is the answer, so keep the range explicit.
4. The trap the question sets: the number 1008 invites arithmetic panic and the phrase polynomial of degree 1008 invites Lagrange. Both push you toward a longer route.
5. What makes you recognise it faster next time: any polynomial interpolating a linear recurrent sequence on consecutive integers, evaluated one step beyond, has value equal to the true next term minus the $(N+1)$st difference of the true sequence.
6. The reusable decision rule: equally spaced nodes means difference table, never Lagrange; and a requested answer in the form of a difference means look for a truncation.

Mirror drills

1. Let q be the polynomial of degree N with $q(n) = 2^n$ for $n = 0, \dots, N$. Find $q(N+1)$. *Answer:* $2^{N+1} - 1$, since every $\Delta^k q(0) = 1$ and $\sum_{k=0}^{N} \binom{N+1}{k} = 2^{N+1} - 1$: the same missing top term structure.
2. Prove $\sum_{k=0}^{M} \binom{M}{k} F_k = F_{2M}$. *Answer:* by Binet, $\sum_k \binom{M}{k} \gamma^k = (1+\gamma)^M = \gamma^{2M}$ for $\gamma \in \{\alpha, \beta\}$; subtract and divide by $\sqrt{5}$.
3. Show that $G_n = F_{2n+1}$ satisfies $G_{n+1} = 3G_n - G_{n-1}$, and use it to compute the first three rows of the difference table directly. *Answer:* α^2, β^2 are the roots of $z^2 - 3z + 1 = 0$ since $\alpha^2 + \beta^2 = 3$ and $\alpha^2\beta^2 = 1$; the rows are $F_{2n+1}, F_{2n+2}, F_{2n+3}$.

7.12 B6. Perfect codes tile the lattice in every dimension

Let $\mathbb{Z}^n$ be the integer lattice in $\mathbb{R}^n$. Two points in $\mathbb{Z}^n$ are called neighbors *if they differ by exactly 1 in one coordinate and are equal in all other coordinates. For which integers $n \geq 1$ does there exist a set of points $S \subset \mathbb{Z}^n$ satisfying the following two conditions?*

(1) If p is in S, then none of the neighbors of p is in S.

(2) If $p \in \mathbb{Z}^n$ is not in S, then exactly one of the neighbors of p is in S.

Read it out loud

Two conditions, and they are begging to be merged. Condition (1) says S is an independent set in the lattice graph. Condition (2) says every point outside S is dominated exactly once. Put them together and say the merged statement aloud: every point of $\mathbb{Z}^n$ lies in exactly one closed ball of radius 1 centred at a point of S. A point of S lies in its own ball and, by (1), in nobody else's; a point outside S lies in exactly one by (2). So the problem is a tiling problem. The closed balls of radius 1, which are the crosses consisting of a centre and its $2n$ neighbours, must partition $\mathbb{Z}^n$.

That reformulation immediately hands you the only number in the problem: each cross has $2n+1$ points, so S must have density exactly $\frac{1}{2n+1}$. You have not constructed anything yet and the modulus $2n+1$ is already forced. This matters enormously for how you feel about the eventual answer. The function $x \mapsto x_1 + 2x_2 + \cdots + nx_n \bmod (2n+1)$ looks like a rabbit pulled from a hat if you meet it cold; it looks inevitable if you have already deduced that whatever S is, it has one point in every $2n+1$, and that the cheapest sets of density $\frac{1}{m}$ in $\mathbb{Z}^n$ are the kernels of homomorphisms onto a group of order m.

The third reading is the phrase for which integers n. That phrasing is bait. It suggests an obstruction exists for some n, and it invites you to spend forty minutes hunting for a parity or divisibility argument that rules out, say, even n. There is no obstruction: the answer is every n. Learn to read the interrogative form of a Putnam question as neutral, because setters use for which n precisely when the honest answer is all of them and they want you to doubt it.

Fields embedded: lattices and subgroups of $\mathbb{Z}^n$, finite abelian groups, tilings, and coding theory (this is a perfect code in the Lee metric, or an efficient dominating set of the lattice graph). What is really being tested is whether you convert two local combinatorial conditions into one global structural condition, and whether you then reach for algebra rather than for drawing. Within thirty seconds you should have written: balls of size $2n+1$ tile; density $\frac{1}{2n+1}$; try $S = \ker \varphi$ for $\varphi : \mathbb{Z}^n \to \mathbb{Z}/(2n+1)\mathbb{Z}$.

Approaches, ranked worst to best

1. **Draw and hope.** *Cue:* small cases are easy to sketch. *Circuit:* find a pattern in $\mathbb{Z}^2$, then guess. For $n = 2$ you will, with patience, find the diagonal pattern of period five, and for $n = 3$ you may find period seven. Drawing three dimensional lattices is slow and error prone, and even after you find the patterns you must still explain what they are before you can generalise. This is the route that consumes an hour and produces a conjecture rather than a proof.

2. **Induct on n.** *Cue:* a set in $\mathbb{Z}^n$ ought to build a set in $\mathbb{Z}^{n+1}$. *Circuit:* lift or take products. There is no natural lift. The ball sizes for consecutive dimensions are $2n+1$ and $2n+3$, which are coprime, so no product or layering construction can turn a tiling of one into a tiling of the other; the densities are incompatible. Recognising quickly that an induction cannot work is itself worth training, because the alternative is losing the problem to a route that never closes.

3. **Colour the lattice.** *Cue:* a partition into $2n+1$ classes, one per ball position. *Circuit:* assign to each point a colour in $\{0, 1, \ldots, 2n\}$ so that every ball is rainbow. This is the right idea in disguise. You want a colouring $c : \mathbb{Z}^n \to \mathbb{Z}/(2n+1)\mathbb{Z}$ such that each closed ball receives all

$2n+1$ colours exactly once, and then take $S = c^{-1}(0)$. Stated this way you still have to invent the colouring, but at least you know exactly what you are looking for.

4. **Force the modulus, then look for a subgroup.** *Cue:* density $\frac{1}{2n+1}$ and an ambient group $\mathbb{Z}^n$. *Circuit:* take S to be the kernel of a homomorphism onto a group of order $2n+1$, and translate the two conditions into a condition on the images of the basis vectors. A homomorphism φ is determined by $c_i = \varphi(e_i)$. The two conditions become: the $2n$ elements $\pm c_1, \ldots, \pm c_n$ are distinct and nonzero, hence exhaust the nonzero residues. Modulo $2n+1$ the choice $c_i = i$ does exactly that. Best, because every step is forced by the previous one and the verification is four lines.

The ideal solution

The answer: such a set S exists for every integer $n \geq 1$.

Fix $n \geq 1$, write $m = 2n+1$, and let $e_1, \ldots, e_n$ be the standard basis of $\mathbb{Z}^n$. Define the group homomorphism

$$\varphi : \mathbb{Z}^n \longrightarrow \mathbb{Z}/m\mathbb{Z}, \qquad \varphi(x_1, \ldots, x_n) = x_1 + 2x_2 + \cdots + nx_n \pmod{m},$$

and set

$$S = \varphi^{-1}(0) = \{\, p \in \mathbb{Z}^n : \varphi(p) \equiv 0 \pmod{m} \,\}.$$

Lemma 7.12.1. *The $2n$ residues $\pm 1, \pm 2, \ldots, \pm n$ are distinct and nonzero modulo $m = 2n+1$, and therefore they are exactly the $2n$ nonzero residues modulo m, each occurring once.*

Proof. Each i with $1 \leq i \leq n$ satisfies $0 < i < m$, so $i \not\equiv 0$ and likewise $-i \not\equiv 0$. If $i \equiv j$ with $1 \leq i, j \leq n$ then $m \mid i - j$ and $|i - j| \leq n - 1 < m$, so $i = j$. If $i \equiv -j$ then $m \mid i + j$ and $2 \leq i + j \leq 2n < m$, which is impossible. So the $2n$ listed residues are pairwise distinct and nonzero; since there are exactly $2n$ nonzero residues modulo m, they exhaust them. □

Every neighbour of a point p has the form $q = p + \delta e_i$ with $i \in \{1, \ldots, n\}$ and $\delta \in \{+1, -1\}$, and distinct pairs (i, δ) give distinct neighbours. Since φ is a homomorphism with $\varphi(e_i) = i$,

$$\varphi(q) = \varphi(p) + \delta i \pmod{m}. \tag{$*$}$$

Condition (1). Let $p \in S$, so $\varphi(p) \equiv 0$, and let $q = p + \delta e_i$ be a neighbour. By $(*)$, $\varphi(q) \equiv \delta i$, which is nonzero by Lemma 7.12.1. Hence $q \notin S$, and no neighbour of a point of S lies in S.

Condition (2). Let $p \notin S$ and put $c = \varphi(p) \not\equiv 0$. By $(*)$, a neighbour $q = p + \delta e_i$ lies in S if and only if $\delta i \equiv -c \pmod{m}$. Since $c \not\equiv 0$, also $-c \not\equiv 0$, so by Lemma 7.12.1 there is exactly one pair (i, δ) with $i \in \{1, \ldots, n\}$, $\delta \in \{\pm 1\}$ and $\delta i \equiv -c$. Distinct pairs give distinct neighbours, so exactly one neighbour of p lies in S.

Both conditions hold for every $n \geq 1$, so the answer is: all n. ■

Remark 7.12.2. Equivalently, the closed balls of radius 1 (the crosses of $2n+1$ points) centred at the points of S partition $\mathbb{Z}^n$: conditions (1) and (2) together say exactly that each point lies in exactly one such ball. The construction is therefore a lattice tiling of $\mathbb{Z}^n$ by crosses, and S is a perfect single error correcting code in the Lee metric.

Remark 7.12.3. Nothing forces the group to be cyclic. Let G be any abelian group of order $2n+1$. Since $|G|$ is odd, no nonzero element satisfies $g = -g$, so the nonzero elements split into exactly n pairs $\{g, -g\}$; choose one representative c_i from each pair and define $\varphi(x) = \sum x_i c_i$. The proof above applies verbatim, since $\{\pm c_1, \ldots, \pm c_n\} = G \setminus \{0\}$ with no repetitions. Also, every translate of S works, so each of the $2n+1$ fibres $\varphi^{-1}(c)$ is itself a valid set.

Layer 1: WHAT

Block 1. Merge conditions (1) and (2) into one: the closed unit balls centred at points of S partition $\mathbb{Z}^n$.

Block 2. Count: each ball has $2n+1$ points, so S has density $\frac{1}{2n+1}$ and the modulus $2n+1$ is forced.

Block 3. Decide to look for a structured S: subgroups of index $2n+1$, that is kernels of homomorphisms $\varphi : \mathbb{Z}^n \to \mathbb{Z}/(2n+1)\mathbb{Z}$.

Block 4. A homomorphism on $\mathbb{Z}^n$ is free: it is determined by arbitrary images $c_i = \varphi(e_i)$. So the search space is n residues.

Block 5. Translate the tiling condition: moving to a neighbour changes φ by $\pm c_i$, so the condition is that $\pm c_1, \ldots, \pm c_n$ hit each nonzero residue exactly once.

Block 6. Solve it: $c_i = i$ modulo $2n+1$, because $\pm 1, \ldots, \pm n$ are exactly the nonzero residues.

Block 7. Verify both original conditions directly from $\varphi(p + \delta e_i) = \varphi(p) + \delta i$, and conclude the answer is all n.

Juxtaposition. The nominal route treats the question as a search: draw $\mathbb{Z}^2$, find a pattern, draw $\mathbb{Z}^3$, find a pattern, guess a general rule, then try to verify it. Some strong candidates get there, because the $n = 2$ pattern really is the diagonal one of period five, but they arrive without a mechanism and with no clean way to write the general case. The bifurcation point is the sentence you write after reading condition (2). If you write examples of S in low dimensions, you are searching. If you write the balls tile, so the density is $\frac{1}{2n+1}$, you have converted a search into an equation, and the equation has an obvious solution.

Layer 2: HOW

Block 1. **Cue.** Two conditions, one about points inside S and one about points outside S, both about neighbours. **Why natural.** A pair of conditions that split on membership in S is almost always one condition in disguise, and the way to find it is to ask what both say about a fixed point p: how many points of S lie in the closed ball around p? Condition (1) says one when $p \in S$ (namely p itself), condition (2) says one when $p \notin S$. So the count is uniformly one, and by symmetry of the neighbour relation that is the same as saying the balls centred at points of S cover every point exactly once. Merging the cases removes the case analysis from every later step. **Link to the next block.** A partition is a counting statement, so count.

Block 2. **Cue.** A partition of an infinite set into congruent finite pieces of size $2n+1$. **Why natural.** Density is the only invariant of such a partition, and it is forced: in any large box, the number of points of S is asymptotically $\frac{1}{2n+1}$ of the number of lattice points. The number $2n+1$ is now attached to the problem permanently, and any construction you propose must produce it. This is the step that removes all mystery from the final formula; if you later find yourself writing modulo $2n+1$, it is because you deduced it here, not because you guessed it. **Link to the next block.** You want a set of density $\frac{1}{2n+1}$ inside $\mathbb{Z}^n$; the most economical such sets are subgroups of that index.

Block 3. **Cue.** The ambient object $\mathbb{Z}^n$ is a group, and the neighbour relation is translation invariant. **Why natural.** The whole problem is invariant under translation by any lattice vector, so it is natural to look for an S that is itself translation invariant under a subgroup, and the cleanest possibility is that S is a subgroup, or a coset of one, of index $2n+1$. Subgroups of index m in $\mathbb{Z}^n$ are exactly the kernels of surjections onto groups of order m. Choosing $\mathbb{Z}/(2n+1)\mathbb{Z}$ is the simplest such group. This is not a guess; it is the standard way to manufacture a set of prescribed density in a group, and it is the same reflex that produces arithmetic progressions when the ambient group is $\mathbb{Z}$. **Link to the next block.** A homomorphism out of a free abelian group is determined by its values on a basis, so parametrise.

Block 4. **Cue.** $\mathbb{Z}^n$ is free abelian with basis $e_1, \ldots, e_n$. **Why natural.** Freeness means there are no relations to respect: any choice whatsoever of $c_1, \ldots, c_n$ in $\mathbb{Z}/m\mathbb{Z}$ defines a homomorphism. So the search space collapses from all subsets of $\mathbb{Z}^n$ to n residues, which is a finite and tiny space. Reducing an infinite search to a finite parametrisation is the single most valuable move in this problem. **Link to the next block.** Now push the conditions through the parametrisation.

Block 5. **Cue.** φ is a homomorphism and neighbours differ by $\pm e_i$. **Why natural.** Therefore $\varphi(q) - \varphi(p) = \pm c_i$: the value of φ is a coordinate that changes in a completely controlled way along every lattice edge. Condition (1) becomes: no $\pm c_i$ is zero. Condition (2) becomes: for every nonzero c there is exactly one (i, δ) with $\delta c_i = -c$. Together: the list $\pm c_1, \ldots, \pm c_n$ has $2n$ entries and must hit each of the $2n$ nonzero residues exactly once, which is a bijection condition, and bijections between sets of equal finite size need only injectivity. That is why the verification is short. **Link to the next block.** You need n residues whose $2n$ signed versions are all distinct and nonzero, so write down the obvious candidate.

Block 6. **Cue.** Modulo an odd number $2n + 1$, the nonzero residues split into n pairs $\{c, -c\}$. **Why natural.** Oddness is exactly what makes the pairing work: in $\mathbb{Z}/m\mathbb{Z}$ with m odd, $c \equiv -c$ forces $2c \equiv 0$ hence $c \equiv 0$, so no nonzero element is its own negative. Therefore a set of n pair representatives exists, and the residues $1, 2, \ldots, n$ are the most obvious choice. If the ball size had been even, no such choice could exist, which is a good way to remember why the centre point of the ball is essential to the whole construction. **Link to the next block.** Everything is determined; write the verification.

Block 7. **Cue.** A concrete φ and two conditions to check. **Why natural.** The verification must be done in the original language, because that is what the question asked. Each check is a single application of $\varphi(p + \delta e_i) = \varphi(p) + \delta i$ together with the lemma. Also record explicitly that distinct pairs (i, δ) give distinct neighbours, since condition (2) says exactly one neighbour, not exactly one solution of a congruence.

Layer 3: WHY

The mathematical truth is that this problem is the existence of a perfect single error correcting code in the Lee metric on $\mathbb{Z}^n$, and perfect codes exist exactly when a sphere packing bound is met with equality by an algebraic construction. The sphere packing (or Hamming) bound says that if the closed balls of radius r are disjoint then the density of centres is at most the reciprocal of the ball size; a code is perfect when equality holds, that is when the balls also cover. So conditions (1) and (2) are packing and covering respectively, and asking for both is asking for a perfect code. Once you see the problem that way, you know what the possible answers look like: perfect codes are rare and almost always come from algebra, specifically as kernels of homomorphisms, exactly as Hamming codes are kernels of parity check matrices over $\mathbb{F}_2$.

The mechanism that makes radius 1 work in every dimension is a coincidence of shape and group. The ball is a cross of $2n + 1$ points, and its $2n$ non-central points come in n opposite pairs $\pm e_i$. A group of odd order $2n + 1$ has its $2n$ nonzero elements in n opposite pairs $\pm g$. The construction is nothing more than matching those two pairings with each other. That is the whole idea, and it is why the natural generalisation is the second remark above: any abelian group of order $2n + 1$ works, since the pairing is all that was used.

That also tells you where the difficulty lies for larger radius, and this is worth knowing because it stops you from over-generalising. For radius r the ball in $\mathbb{Z}^n$ has $\sum_{k \geq 0} 2^k \binom{n}{k} \binom{r}{k}$ points, and the analogous requirement is that the displacement vectors of the ball map to distinct residues, which is a perfect difference set condition, far stronger than a pairing. For $n = 2$ such codes exist for

every r (for example $\varphi(x,y) = x + 5y$ modulo 13 works for $r = 2$: the thirteen displacements with $|a| + |b| \leq 2$ map to the thirteen distinct residues). For $n \geq 3$ and $r \geq 2$, Golomb and Welch conjectured in 1970 that no perfect Lee code exists, and the conjecture is proved in various cases but remains open in general. So the case $r = 1$ in this problem is not a sample of a general phenomenon; it is the one radius where the shape and the group agree, and that is why the answer is all n rather than some n.

The discernment mechanism is this. When a problem imposes local conditions on a subset of an infinite homogeneous structure, first merge the conditions into a single statement about a partition or covering; then extract the density; then look for the subgroup or homomorphism that realises that density. Homogeneous ambient structure means the answer is likely to be homogeneous too, and homogeneous subsets of a group are cosets of subgroups. Train this by rewriting every combinatorial condition you meet in terms of the group action that preserves it.

Knowledge gaps, stated plainly. Before this problem is even approachable you need four things installed, none of which is in A-level or Further Mathematics. First, homomorphisms of abelian groups and the fact that a homomorphism out of $\mathbb{Z}^n$ is freely determined by the images of the basis vectors; without it the parametrisation step is invisible. Second, the notion of index and density of a subgroup, and the ability to convert a counting statement about an infinite set into a modular arithmetic statement. Third, the language of packings and coverings, so that you can recognise conditions (1) and (2) as packing and covering and merge them; the sphere packing bound with equality is the definition of a perfect code. Fourth, the Lee metric on $\mathbb{Z}^n$ and the fact that its balls are crosses, together with enough of the Golomb and Welch story to know that the radius 1 case is special and that you should not expect the same trick to survive at larger radius. If you meet this problem without those four, you will find the $n = 2$ picture by hand and stall.

Revolving concepts

Perfect codes and the sphere packing bound. A code S in a metric space with all balls of radius r of the same size V is perfect if the balls of radius r about the codewords partition the space, equivalently if it attains the sphere packing bound density $1/V$ with equality. The classical examples are the Hamming codes over $\mathbb{F}_q$, the binary and ternary Golay codes, and (in the Lee metric) the constructions of this problem. Perfect codes are almost always kernels of linear maps, because a partition into congruent pieces with a transitive symmetry group is usually a coset decomposition. See Book 1, field: Discrete structures and graphs, topic: Perfect codes and coverings.

Efficient domination. In graph language, S is an efficient dominating set (or perfect code) of a graph G if the closed neighbourhoods $N[s]$ for $s \in S$ partition $V(G)$. Conditions (1) and (2) are precisely that for the lattice graph $\mathbb{Z}^n$ with edges between points differing by one in a single coordinate. This is a standard reformulation to keep, because the merged statement is easier to reason about than the pair of conditions. See Book 1, field: Discrete structures and graphs, topic: Domination in graphs.

Subgroups of $\mathbb{Z}^n$ and index. $\mathbb{Z}^n$ is free abelian of rank n; a homomorphism to any abelian group is uniquely determined by arbitrary images of $e_1, \ldots, e_n$. Subgroups of index m are kernels of surjections onto groups of order m, and the quotient of $\mathbb{Z}^n$ by such a subgroup has exactly m cosets, each of density $1/m$. This is the standard machine for producing periodic subsets of the lattice of prescribed density. See Book 1, field: Abstract algebra, topic: Free abelian groups and quotients.

Odd order and the involution free pairing. In an abelian group G of odd order, $g = -g$ implies $2g = 0$ implies $g = 0$ (since 2 is invertible modulo an odd number), so $G \setminus \{0\}$ splits into $\frac{|G|-1}{2}$ pairs $\{g, -g\}$. This is a special case of the standard involution counting argument: if an involution acts on a finite set, the number of fixed points has the same parity as the size of the set. That argument also proves Fermat's two squares theorem in Zagier's one sentence form, and it is worth carrying as a general tool. See Book 1, field: Abstract algebra, topic: Involutions and pairing arguments.

The Lee metric and cross tilings. The Lee distance on $\mathbb{Z}^n$ is $\sum_i |x_i - y_i|$; the ball of radius r has $\sum_{k\geq 0} 2^k \binom{n}{k}\binom{r}{k}$ points, which is $2n+1$ for $r=1$. A perfect Lee code is a lattice or non-lattice tiling of $\mathbb{Z}^n$ by these crosses. Golomb and Welch (1970) proved existence for $r=1$ in all dimensions and for all r when $n=2$, and conjectured non-existence for $n\geq 3$, $r\geq 2$; the conjecture is still open in general. Related is Minkowski's conjecture on lattice tilings by cubes, proved by Hajós. See Book 1, field: Discrete structures and graphs, topic: Lattice tilings.

Recursion note and its generalisation. There is no recursion in the construction, and it is worth saying why, because the instinct to induct on n is strong and wastes time. The obstruction is arithmetic: a tiling of $\mathbb{Z}^n$ by crosses has density $\frac{1}{2n+1}$ and a tiling of $\mathbb{Z}^{n+1}$ has density $\frac{1}{2n+3}$, and $\gcd(2n+1, 2n+3) = 1$, so no construction that layers copies of an n dimensional solution can produce an $(n+1)$ dimensional one. The generalised statement to keep is: when a family of problems indexed by n has an invariant (here density) whose values for consecutive n are incompatible, look for a uniform construction rather than a recursive one. The uniform construction here is the group homomorphism, whose definition is the same sentence for every n. See Book 1, field: Combinatorics, topic: Uniform versus recursive constructions.

Circuit Bank

- **WHEN** you see one condition for points in a set and another for points outside it, both about neighbours **DO** merge them into a single statement counting how many points of the set lie in each closed ball **BECAUSE** a pair of complementary local conditions is usually one covering or partition statement, and the merged form removes every later case split
- **WHEN** you have a partition of an infinite homogeneous set into congruent pieces of size V **DO** write down the density $\frac{1}{V}$ immediately and treat V as a forced modulus **BECAUSE** density is the only invariant of such a partition, so any construction must realise it, and the modulus you will eventually work with is determined before you construct anything
- **WHEN** you need a subset of a group of prescribed density $\frac{1}{m}$ **DO** take the kernel or a coset of a homomorphism onto a group of order m **BECAUSE** a homomorphism out of a free abelian group is determined by arbitrary images of the basis vectors, so an infinite search collapses to finitely many parameters
- **WHEN** you have parametrised by images c_i of basis vectors and a condition about moving to neighbours **DO** write the condition as a statement about the multiset $\pm c_1, \ldots, \pm c_n$ and count **BECAUSE** moving by $\pm e_i$ changes the homomorphism by $\pm c_i$, so all local conditions become statements about that finite list, and equal cardinalities turn surjectivity into injectivity
- **WHEN** you read for which n does there exist **DO** do not assume an obstruction exists; test the construction instinct first and only then look for an invariant **BECAUSE** setters use the interrogative form both when the answer is a restricted set and when it is everything, so treating it as a hint costs you the problem when the answer is all n

Backtesting

1. First useful cue: conditions (1) and (2) merge into the statement that the closed unit balls about points of S partition $\mathbb{Z}^n$.
2. Method the expression suggested: count the ball size to get the modulus $2n+1$, look for S as the kernel of $\varphi : \mathbb{Z}^n \to \mathbb{Z}/(2n+1)\mathbb{Z}$, translate the conditions into a statement about $\pm\varphi(e_i)$, and take $\varphi(e_i) = i$.
3. Where you would hesitate, and the thought to reinforce: after the density calculation, whether a subgroup really can do the job. The thought is that the problem is translation invariant, so it

is reasonable that a translation invariant solution exists, and translation invariant subsets of a group of density $\frac{1}{m}$ are cosets of index m subgroups.

4. The trap the question sets: the phrase for which integers n, which invites a hunt for an obstruction, and the low dimensional pictures, which invite an hour of drawing. Both consume time without producing a general argument.

5. What makes you recognise it faster next time: the shape of the ball. A cross whose non-central points come in opposite pairs $\pm e_i$ matches the pairing of the nonzero elements of any odd order abelian group; the moment you see that matching you have the construction.

6. The reusable decision rule: merge local conditions into a covering statement, read the density, then realise the density as a subgroup and push the conditions through the homomorphism.

Mirror drills

1. Show that conditions (1) and (2) together are equivalent to: the closed balls of radius 1 centred at points of S partition $\mathbb{Z}^n$. *Answer:* for $p \in S$, the balls containing p are those centred at p and at neighbours of p in S, so (1) says exactly one; for $p \notin S$, the balls containing p are those centred at neighbours of p in S, so (2) says exactly one. The converse reverses both readings.

2. Classify all valid S for $n = 1$. *Answer:* the balls are the blocks $\{p-1, p, p+1\}$ and they must tile $\mathbb{Z}$, so consecutive elements of S differ by exactly 3 and $S = 3\mathbb{Z} + a$ for some $a \in \{0, 1, 2\}$; these are exactly the fibres of $\varphi(x) = x \bmod 3$.

3. Exhibit a perfect code of radius 2 in $\mathbb{Z}^2$, that is a set S such that the thirteen point Lee balls of radius 2 centred at its points tile $\mathbb{Z}^2$. *Answer:* $S = \{(x, y) : x + 5y \equiv 0 \pmod{13}\}$; the thirteen displacement vectors (a, b) with $|a| + |b| \leq 2$ have values $a + 5b$ running over all thirteen residues modulo 13 exactly once.

Chapter 8

Putnam 2018

8.1 A1. Unit fractions collapse into a factor pair

Find all ordered pairs (a, b) of positive integers for which

$$\frac{1}{a} + \frac{1}{b} = \frac{3}{2018}.$$

Read it out loud

Two positive unit fractions add to a third fraction that is not itself a unit fraction. Say the equation out loud and the shape matters more than the numbers: $1/a + 1/b = k/n$ is the Egyptian fraction template, and every question in that template wants the same treatment. Clear denominators and you get a relation with exactly one ab term and two linear terms, nothing else. An expression of that shape, one bilinear term, two linear terms, no pure squares, is one scalar multiple away from $(pa - q)(pb - q) =$ constant, because expanding that product is precisely a bilinear term plus two linear terms plus a constant term. You do not need to search for this factorisation. You need to recognise that the number of terms and their degrees have already told you it exists, and multiply by whatever scalar makes the cross term's coefficient match the square of the linear coefficient.

The second thing to read off in the first thirty seconds is 2018 itself. A Putnam problem set in a specific year routes its arithmetic through that year's own factorisation far more often than coincidence would suggest, so you check it immediately: $2018 = 2 \times 1009$, and you confirm 1009 is prime by trial division up to $\lfloor \sqrt{1009} \rfloor = 31$ rather than assuming it. The fields wired into this problem are pure number theory: completing a product, counting divisors of a number given in prime factorised form, and a congruence condition that falls out of the algebra automatically rather than being imposed as an extra restriction. There is no calculus, no geometry, nothing combinatorial beyond counting divisors. Within thirty seconds you should have $2018a + 2018b = 3ab$ on your page and your pen already moving to multiply by 3.

Approaches, ranked worst to best

1. **Search from the forced lower bound.** *Cue:* the phrase "positive integers" with no algebra attempted first. *Circuit:* brute force enumeration. (Positivity of $1/a$ and $1/b$ forces each to be smaller than $3/2018$, so $a, b \geq 673$. Nothing in that bound stops a from being enormous, and indeed the true maximum is $a = 1358114$. You would only discover that ceiling by accident, and you could never write a complete proof of termination without secretly reconstructing the factorisation below. Unbounded ink for an answer you cannot certify complete.)

2. **Isolate b and chase divisibility.** *Cue:* solving $b = \dfrac{2018a}{3a - 2018}$ and asking when the quotient is an integer. *Circuit:* direct divisibility chasing. (This does reach $3a - 2018 \mid 2018^2$, but only after noticing that $3a - 2018$ must divide $3(2018a) = 2018(3a - 2018) + 2018^2$ and subtracting off the multiple of itself. It works, yet it destroys the symmetry between a and b that the factored equation displays for free, so you effectively do the casework once for a and would need to redo it for b to be sure you have not lost a pair.)

3. **Complete the product.** *Cue:* a single ab term, two linear terms, no squares. *Circuit:* force a fixed product by matching the scalar to the cross term. (Multiplying by the right constant produces $(3a - 2018)(3b - 2018) = 2018^2$ in one line, manifestly symmetric, and reduces the entire problem to divisor enumeration of a single fully factorised number. This is the route that costs the least ink and proves completeness as a side effect of the method rather than as an afterthought.)

The ideal solution

Multiply $\frac{1}{a}+\frac{1}{b}=\frac{3}{2018}$ through by $2018ab$:

$$2018b+2018a=3ab, \qquad \text{that is,} \qquad 3ab-2018a-2018b=0.$$

Multiply by 3 so that the coefficient of each linear term becomes the product of 3 (the coefficient of ab) and the constant 2018, which is exactly what completing the product needs:

$$9ab-3(2018)a-3(2018)b=0 \implies (3a-2018)(3b-2018)=2018^2,$$

since expanding the left side gives $9ab-3(2018)a-3(2018)b+2018^2$, and adding 2018^2 to both sides of the previous line produces exactly that.

Because $2018=3(672)+2$, we have $2018\equiv 2 \pmod 3$, so for any positive integer a, $3a-2018\equiv -2\equiv 1 \pmod 3$. This holds automatically, for every positive integer a, not just for the ones that end up solving the equation; the same is true of $3b-2018$. So whatever factor pair of 2018^2 we find, both factors are forced to lie in the residue class 1 mod 3.

Factorise $2018=2\times 1009$; 1009 is prime, since none of $2,3,5,7,11,13,17,19,23,29,31$ divides it and $31^2<1009<32^2$. So $2018^2=2^2\times 1009^2$ has $(2+1)(2+1)=9$ positive divisors, one for each choice of the exponent of 2 in $\{0,1,2\}$ and of 1009 in $\{0,1,2\}$. Since $1009=3(336)+1$, every power of 1009 is $\equiv 1 \pmod 3$, while $2\equiv 2$ and $2^2=4\equiv 1 \pmod 3$. A divisor $2^i\cdot 1009^j$ is therefore $\equiv 1 \pmod 3$ exactly when $i\in\{0,2\}$, regardless of j. Writing $d=3a-2018$, we need d to be one of these six divisors:

$$1,\ 4,\ 1009,\ 4036,\ 1009^2,\ 4\cdot 1009^2,$$

and the complementary divisor $2018^2/d=2^{2-i}\cdot 1009^{2-j}$ automatically has its own power of 2 equal to 2^{2-i}, which lies in $\{0,2\}$ exactly when i does; so the cofactor is independently guaranteed to be $\equiv 1 \pmod 3$ too, for each of the six choices, with no separate check required.

For each such d, solve $a=(d+2018)/3$ and $b=(2018^2/d+2018)/3$:

d	a	b
1	673	1358114
4	674	340033
1009	1009	2018
4036	2018	1009
1009^2	340033	674
$4\cdot 1009^2$	1358114	673

all six of which are positive integers, giving the ordered pairs $(673,1358114)$, $(674,340033)$, $(1009,2018)$, and their three reversals.

It remains to rule out negative divisors. If $d<0$ then $2018^2/d<0$ too, and $a=(d+2018)/3$ is a positive integer only if $d>-2018$; among the negative divisors of 2018^2 that are $\equiv 1 \pmod 3$ by the same rule, namely -2, -2018, and $-2\cdot 1009^2$, only $d=-2$ satisfies $d>-2018$, giving $a=672$. But then $2018^2/d=-1009^2\cdot 2\cdot 1009/$ more directly $2018^2/(-2)=-2036162$, so $b=(-2036162+2018)/3=-678048<0$, which is not a positive integer. The other two negative divisors give $a\le 0$ outright. No negative-divisor case survives.

$$\boxed{(a,b)\in\{(673,1358114),\ (674,340033),\ (1009,2018),\ (2018,1009),\ (340033,674),\ (1358114,673)\}}$$

Layer 1: WHAT

Block 1: read the equation's shape (one bilinear term, two linear terms) and recognise the completed-product target. Block 2: choose the scalar multiplier that makes the cross term's coefficient equal the

product of the two linear coefficients, and complete the product. Block 3: factorise 2018 and count divisors of 2018^2 by the exponent grid. Block 4: observe that the linear shift forces every valid factor into the residue class 1 mod 3, and use this as a filter rather than extra casework. Block 5: solve for (a, b) across the six admissible positive divisors, then explicitly dispose of the negative-divisor branch by positivity.

The nominal route enumerates a from 673 upward (or chases divisibility on $b = 2018a/(3a - 2018)$ one variable at a time) with no certified stopping point and no symmetry to halve the work. The bifurcation point is the exact line after $3ab - 2018a - 2018b = 0$: the ideal route multiplies by 3 and completes a product; the nominal route instead tries to isolate a variable and search. Everything before that line is identical for both; everything after diverges completely, and the nominal route never regains the symmetry it threw away.

Layer 2: HOW

Block 1. **Cue.** The equation, once denominators are cleared, has exactly one ab term and two linear terms with no pure squares. **Why natural.** A degree-two expression with that exact term structure is the expansion of a product of two linear binomials minus a constant, so the target factorisation is not found by search, it is read off the shape. **Link to the next block.** The next step is arithmetic: find the scalar that turns the shape into that exact product.

Block 2. **Cue.** The linear coefficients are both -2018 and the ab coefficient is 3. **Why natural.** Multiplying the whole equation by 3 makes each linear coefficient equal to 3×2018, which is precisely the cross term you get from expanding $(3a - 2018)(3b - 2018)$, so the completion requires no guessing, only matching coefficients. **Link to the next block.** Once the equation reads $(3a - 2018)(3b - 2018) = 2018^2$, the problem has become a divisor-pair count.

Block 3. **Cue.** 2018 is a specific number handed to you by the year of the competition. **Why natural.** Any fixed integer factorises into a finite, checkable list of primes, and once you have that list, divisor counting is mechanical rather than exploratory. **Link to the next block.** Knowing the prime factorisation lets you predict, before listing anything, how many divisors there are and which residues they occupy.

Block 4. **Cue.** Both factors are linear expressions of the form $3(\text{integer}) - 2018$. **Why natural.** Every integer of that form lies in a single residue class mod 3 purely because 2018 itself has a fixed residue mod 3; this is true before you know which divisor you will use, so it is a filter you get for free rather than a condition you must verify case by case. **Link to the next block.** This filter cuts the nine divisors of 2018^2 down to the six that can possibly arise, before you do a single division.

Block 5. **Cue.** Six candidate values of d remain, plus the logical possibility of negative factor pairs. **Why natural.** A full-marks solution to "find all pairs" must dispose of every branch explicitly; the negative branch is not automatically empty just because the positive branch works, since a negative times a negative is also positive. **Link to the next block.** Checking positivity directly on each negative candidate closes the proof without leaving an unjustified gap.

Layer 3: WHY

The truth underneath this problem is that the equation $pxy = q(x + y)$, for any fixed nonzero p and q, always reduces to $(px - q)(py - q) = q^2$ by multiplying through by p; this is not special to 2018 or to unit fractions, it is an identity in p, q, x, y that holds before any numbers are substituted. The discernment mechanism is counting terms and degrees: whenever a relation between two variables has total degree two, contains exactly one mixed term, and contains no pure squares, the

completed-product form exists and is unique up to the choice of scalar. You train this by drilling the recognition step in isolation, away from any specific numbers, until seeing "xy term plus two linear terms, no squares" fires the same reflex every time, regardless of how the problem dresses up the coefficients as fractions, denominators, or a competition year.

The examiners nudged you toward this by choosing a right-hand side, $3/2018$, whose numerator is not 1. Had the equation been $1/a + 1/b = 1/n$, the completed form $(a-n)(b-n) = n^2$ is even more familiar and the problem would test only divisor counting; forcing a numerator of 3 makes you notice that the scalar you multiply by must match that numerator, and it is precisely this extra step, multiplying by 3 rather than by 1, that separates a candidate who has memorised one instance of the trick from one who has understood the identity behind it.

Revolving concepts

Unit fraction (Egyptian fraction) equations: the general family $\frac{1}{x} + \frac{1}{y} = \frac{k}{n}$ always reduces to $(kx - n)(ky - n) = n^2$; the number of ordered solution pairs equals the number of divisors of n^2 that lie in the residue class forced by n mod k, which is a clean generalisation of the case $k = 1$ that most students meet first. See Book 1, field: Number theory, topic: Diophantine equations from cleared denominators.

Divisor counting from a prime factorisation: if $N = \prod p_i^{e_i}$ then N has $\prod(e_i + 1)$ positive divisors, and this generalises immediately to counting divisors in a fixed residue class mod a small prime m, by tracking which combinations of exponents give which residue, exactly as done above for $m = 3$. See Book 1, field: Number theory, topic: multiplicative functions and the divisor function.

Residues under multiplication: the nonzero residues mod 3 form a cyclic group of order 2 under multiplication, $\{1, 2\}$, so the residue of a product is determined by the parity of how many factors are $\equiv 2$; this is why tracking only the exponent of the single prime that is $\equiv 2 \pmod 3$ (here, the prime 2) was enough to classify every divisor, while the prime $1009 \equiv 1 \pmod 3$ contributed nothing to the residue. See Book 1, field: Number theory, topic: multiplicative structure of $(\mathbb{Z}/m\mathbb{Z})^\times$.

Circuit Bank

- **WHEN** a relation has one bilinear term and two linear terms in two variables, no pure squares **DO** multiply by the scalar that makes the cross-term coefficient equal the product of the two linear coefficients, then complete the product **BECAUSE** this shape is always the expansion of a two-factor product minus a constant, so the factorisation is guaranteed to exist, not merely hoped for

- **WHEN** a fixed integer must be factorised to enumerate divisors **DO** write its full prime factorisation first and read off the exponent grid, rather than testing candidate divisors one at a time **BECAUSE** the number and residue class of every divisor is determined instantly from the exponents, which turns an open-ended search into a finite, checkable list

- **WHEN** a family of candidate values is generated by a linear shift of an unconstrained integer **DO** check what residue class the shift forces before doing any casework **BECAUSE** the residue is fixed by the shift alone and applies to every candidate simultaneously, so it filters the whole list in one step instead of one check per candidate

- **WHEN** a factorisation $XY = k$ with $k > 0$ is used to solve for positive unknowns **DO** explicitly substitute the negative divisors and check positivity, rather than assuming they fail **BECAUSE** negative times negative is also positive, so the sign condition on X, Y individually is not implied by the sign of their product and must be checked

Backtesting

1. First useful cue: two unit fractions summing to a fixed rational number with numerator greater than 1.
2. Method the expression suggested: clear denominators, then complete the product by multiplying by the numerator of the target fraction.
3. Where you would hesitate, and the thought to reinforce: whether the scalar to multiply by is the numerator or the denominator; reinforce by matching the cross-term coefficient to the product of the two linear coefficients algebraically, not by memory.
4. The trap the question sets: assuming the negative-divisor branch is impossible without checking it, when in fact it fails only because of an explicit positivity computation.
5. What makes you recognise it faster next time: seeing exactly one ab term and two linear terms with no squares, regardless of how the coefficients are dressed up as fractions.
6. The reusable decision rule: for $\frac{1}{a}+\frac{1}{b}=\frac{k}{n}$, always reduce to $(ka-n)(kb-n)=n^2$, then filter the divisors of n^2 by the residue that $-n$ forces mod k.

Mirror drills

1. Find all positive integer pairs with $\frac{1}{a}+\frac{1}{b}=\frac{1}{12}$. *Answer:* reduces to $(a-12)(b-12)=144=2^4\cdot 3^2$, which has 15 positive divisors, so there are 15 ordered pairs.
2. Find all positive integer pairs with $\frac{2}{a}+\frac{5}{b}=\frac{1}{3}$. *Answer:* clear denominators to $6b+15a=ab$, then complete the product as $(a-6)(b-15)=90$, and check which of the resulting pairs keep a,b positive.
3. For $\frac{1}{a}+\frac{1}{b}=\frac{3}{N}$ with N an odd prime greater than 3, explain why the number of solution pairs depends on N mod 3 rather than on N itself. *Answer:* because N^2 has exactly the divisors $1, N, N^2$, and which of $3N-N=2N$-type shifts (that is, which of $1, N, N^2$) lie in the forced residue class is controlled entirely by whether $N\equiv 1$ or $2 \pmod 3$.

8.2 A2. Determinant of the subset intersection matrix

Let $S_1, S_2, \ldots, S_{2^n-1}$ be the nonempty subsets of $\{1, 2, \ldots, n\}$ in some order, and let M be the $(2^n - 1) \times (2^n - 1)$ matrix whose (i, j) entry is

$$m_{ij} = \begin{cases} 0 & \text{if } S_i \cap S_j = \emptyset; \\ 1 & \text{otherwise.} \end{cases}$$

Calculate the determinant of M.

Read it out loud

The size of the matrix is $2^n - 1$, and that number is not a free choice, it is exactly the count of nonempty subsets of $\{1, \ldots, n\}$. Whenever a construction is indexed by all subsets (or all nonempty subsets) of $\{1, \ldots, n\}$, and the entries only depend on set relations between the indices, you should reach for induction on n by splitting on whether the largest element n belongs to a given subset. This is the one recursive structure that the power set of $\{1, \ldots, n\}$ actually has for free: the subsets not containing n are exactly the subsets of $\{1, \ldots, n-1\}$, and the subsets containing n are exactly those same subsets of $\{1, \ldots, n-1\}$ with n adjoined, so the whole object at size n is built out of two copies of itself at size $n - 1$.

The phrase "in some order" is a second cue, and a generous one: it is telling you the answer cannot depend on the labelling, which means you are free to choose whichever ordering makes the recursive split visible as an actual block structure in the matrix, rather than being stuck with whatever order the subsets happened to be listed in. Before using that freedom you must earn it, by checking that relabelling really does leave the determinant alone; this is a short argument and it belongs at the very start, not as an afterthought.

The field here is linear algebra sitting on top of combinatorics: the entries of M are not numbers you compute from a formula, they are the answer to a set-theoretic question (is the intersection empty), so the tools you reach for are the ones that behave well when a whole block of an matrix is constant, namely elementary row and column operations and block-triangular determinant formulas, rather than eigenvalues or a closed-form entrywise pattern. Within the first minute you should have written $2^n - 1 = (2^{n-1} - 1) + 2^{n-1}$ and be looking for the block matrix that decomposition produces.

Approaches, ranked worst to best

1. **Compute small cases by hand and guess.** *Cue:* the temptation to just write out the 3×3 and 7×7 matrices for $n = 2, 3$. *Circuit:* pattern spotting. (This gets you $\det M_2 = -1$ quickly and $\det M_3 = -1$ with real effort, but a guessed pattern from two data points is not a proof, and the 7×7 hand computation is already painful enough that $n = 4$ is not realistic by this route at all.)

2. **Attack the eigenvalues directly.** *Cue:* a $0/1$ matrix defined by a Boolean condition on the indices. *Circuit:* spectral methods. (The matrix is symmetric, so it has real eigenvalues, but the entries are not linear or bilinear in any convenient basis of indicator vectors, since $m_{ij} = 1$ means $S_i \cap S_j \neq \emptyset$, a threshold condition, not $S_i \cap S_j$ itself. There is no clean way to write M as a low rank perturbation of a matrix with known spectrum, so this route stalls almost immediately.)

3. **Recursive block decomposition on n.** *Cue:* index set size exactly $2^n - 1$, and an explicit invitation ("in some order") to relabel freely. *Circuit:* split on membership of the top element, then reduce the resulting block matrix. (This is the route that works: choosing the ordering that respects the split "does not contain n, contains n but is not $\{n\}$, equals $\{n\}$" turns M_n

into an explicit block matrix built from M_{n-1}, and elementary operations finish it in a short recursion.)

The ideal solution

Write M_n for the matrix built from the nonempty subsets of $\{1, \dots, n\}$, in whatever order.

Order independence. Every permutation of a finite list is a product of transpositions, so it suffices to check invariance under swapping the labels of two subsets, say S_p and S_q. Relabelling swaps rows p, q of M and, since the entries also depend on the column index through the same subsets, it separately swaps columns p, q too. Swapping two rows multiplies the determinant by -1, and swapping two columns multiplies it by -1 again, so the net effect on $\det M$ is multiplication by $(-1)^2 = 1$. Hence $\det M_n$ does not depend on the chosen ordering, and we are free to choose the most convenient one.

The recursive ordering. For $n \geq 2$, let $S'_1, \dots, S'_{2^{n-1}-1}$ be the nonempty subsets of $\{1, \dots, n-1\}$ in any fixed order, with matrix M_{n-1}, entries m'_{ij}. Order the $2^n - 1$ nonempty subsets of $\{1, \dots, n\}$ as: $S_i = S'_i$ for $1 \leq i \leq 2^{n-1} - 1$; $S_i = S'_{i-2^{n-1}+1} \cup \{n\}$ for $2^{n-1} \leq i \leq 2^n - 2$; and $S_{2^n-1} = \{n\}$. This lists every subset not containing n, then every subset containing n other than $\{n\}$ itself, then $\{n\}$; the counts are $(2^{n-1} - 1) + (2^{n-1} - 1) + 1 = 2^n - 1$, correct.

The block structure. For $i, j \leq 2^{n-1} - 1$, neither S_i nor S_j contains n, so $m_{ij} = m'_{ij}$: the top-left block is M_{n-1}. For $i \leq 2^{n-1} - 1$ and j in the middle range, $S_i \cap S_j = S_i \cap (S'_j \cup \{n\}) = S_i \cap S'_j$ since $n \notin S_i$, so this entry also equals the corresponding entry of M_{n-1}; by the symmetry of intersection the same holds with rows and columns exchanged, so both off-diagonal blocks equal M_{n-1}. For i, j both in the middle range, $S_i \cap S_j \supseteq \{n\}$ always, since both sets contain n; that whole block is entrywise 1. The last row and column (index $2^n - 1$, the set $\{n\}$) meet the first block in 0 (since $\{n\}$ is disjoint from any subset of $\{1, \dots, n-1\}$), meet the middle block in 1 (since every set there contains n), and meet itself in 1. Altogether,

$$M_n = \left(\begin{array}{ccc|ccc|c} & & & & & & 0 \\ & M_{n-1} & & & M_{n-1} & & \vdots \\ & & & & & & 0 \\ \hline & & & 1 & \cdots & 1 & 1 \\ & M_{n-1} & & \vdots & \ddots & \vdots & \vdots \\ & & & 1 & \cdots & 1 & 1 \\ \hline 0 & \cdots & 0 & 1 & \cdots & 1 & 1 \end{array}\right).$$

Clearing the constant block. Subtract the final row from each of the middle rows. Each middle row currently reads $(M_{n-1}\text{-row}, \mathbf{1}^T, 1)$ and the final row reads $(\mathbf{0}, \mathbf{1}^T, 1)$, so after subtraction every middle row becomes $(M_{n-1}\text{-row}, \mathbf{0}^T, 0)$; the M_{n-1} entries survive because the final row is 0 there. This changes nothing about the determinant. Now, using the current matrix, subtract the final column from each middle column. The final column currently reads 0 in the top rows, 0 in the (already updated) middle rows, and 1 in the final row, so this operation only changes the final row: its middle entries, currently 1 each, become $1 - 1 = 0$. The result is

$$\left(\begin{array}{ccc|ccc|c} & & & & & & 0 \\ & M_{n-1} & & & M_{n-1} & & \vdots \\ & & & & & & 0 \\ \hline & & & 0 & \cdots & 0 & 0 \\ & M_{n-1} & & \vdots & \ddots & \vdots & \vdots \\ & & & 0 & \cdots & 0 & 0 \\ \hline 0 & \cdots & 0 & 0 & \cdots & 0 & 1 \end{array}\right),$$

whose final row and final column now both vanish except for the corner entry 1. Expanding along that row (a row with a single nonzero entry, in the bottom right corner, so the cofactor sign is $(-1)^{(2^n-1)+(2^n-1)} = +1$) gives

$$\det M_n = \det \begin{pmatrix} M_{n-1} & M_{n-1} \\ M_{n-1} & 0 \end{pmatrix}.$$

Collapsing the block matrix. Let $k = 2^{n-1} - 1$, the common block size. Swap the top k rows with the bottom k rows. Moving a block of k rows past another block of k rows is a permutation with sign $(-1)^{k^2}$ (it is realised by k^2 transpositions, one for each pair consisting of one row from each block, each such pair needing to cross exactly once). For $n \geq 2$, 2^{n-1} is even, so $k = 2^{n-1} - 1$ is odd, and k^2 is odd; the sign is -1. After the swap the matrix reads $\begin{pmatrix} M_{n-1} & 0 \\ M_{n-1} & M_{n-1} \end{pmatrix}$, which is block lower triangular, so its determinant is the product of the diagonal blocks' determinants, $(\det M_{n-1})^2$. Undoing the sign of the swap,

$$\det M_n = -(\det M_{n-1})^2 \qquad (n \geq 2).$$

Base case and conclusion. For $n = 1$ the only nonempty subset of $\{1\}$ is $\{1\}$ itself, $M_1 = (1)$, and $\det M_1 = 1$. For $n \geq 2$, whatever the sign of $\det M_{n-1}$, its square is 1, so $\det M_n = -1$. In particular $\det M_2 = -(\det M_1)^2 = -1$, and every subsequent value stays at -1 once it lands there.

$$\boxed{\det M = 1 \text{ if } n = 1, \qquad \det M = -1 \text{ if } n \geq 2.}$$

Layer 1: WHAT

Block 1: prove the answer is independent of the ordering, using transpositions, so a convenient ordering may be chosen freely. Block 2: choose the ordering that splits subsets by membership of the top element n, and set up the resulting index ranges. Block 3: derive the four block types of M_n from first principles (two copies of M_{n-1}, one all-ones block, and a corner row and column), using only the definitions of intersection and union. Block 4: use elementary row and column operations, subtracting the constant last row and last column, to zero out the all-ones contamination and shrink to a genuine block matrix. Block 5: swap the two remaining row blocks, track the sign of that swap exactly, and read off a block lower triangular determinant. Block 6: solve the resulting recursion $\det M_n = -(\det M_{n-1})^2$ using the base case $n = 1$.

The nominal route computes $n = 2, 3$ by hand and conjectures -1, or hunts for eigenvalues of a matrix with no clean spectral structure. The bifurcation point is the decision, immediately after proving order-independence, to spend effort choosing a specific recursive ordering rather than accepting whatever ordering the problem handed you; every subsequent block operation exists only because that specific ordering was chosen, and none of it is visible under an arbitrary ordering.

Layer 2: HOW

Block 1. Cue. The problem explicitly says "in some order," which is an unusual thing to say unless it matters. **Why natural.** If a stated quantity is independent of an arbitrary choice, that independence is normally worth proving first, because it licenses you to make the choice yourself later, and a swap of two labels is the smallest possible instance of "the choice changes," so checking it settles every relabelling by composition. **Link to the next block.** Once relabelling is free, the natural next move is to choose the ordering that makes the object's recursive structure visible.

Block 2. Cue. The index count is $2^n - 1$, the count of nonempty subsets of an n-element set. **Why natural.** This number has exactly one natural recursive decomposition, splitting on whether the

largest element belongs to the subset, because the subsets of $\{1, \ldots, n\}$ are in exact bijection with (subsets of $\{1, \ldots, n-1\}$) union (subsets of $\{1, \ldots, n-1\}$, each with n attached). **Link to the next block.** This bijection becomes a literal block structure in the matrix once the ordering respects it.

Block 3. **Cue.** Every entry of M is determined purely by whether an intersection is empty. **Why natural.** Once you know which of the two groups ("contains n" or not) each index belongs to, the intersection condition simplifies immediately: dropping n from a computation costs nothing when the other set does not contain n either, and containing n on both sides is enough on its own to guarantee nonempty intersection, regardless of anything else in the sets. **Link to the next block.** This is why two full blocks collapse to M_{n-1} exactly and one full block collapses to the constant 1, with no computation left to do inside those blocks.

Block 4. **Cue.** One whole block of the matrix is entirely 1s, and it sits next to a corner row and column that are mostly 0 or mostly 1. **Why natural.** A block that is entirely one constant value is exactly what elementary row or column operations are built to erase: subtracting a row that matches the constant block everywhere zeroes that block out without touching anything the operation does not act on. **Link to the next block.** Once erased, the last row and column have only one nonzero entry each, so a single cofactor expansion strips them off entirely.

Block 5. **Cue.** What remains is a $2k \times 2k$ matrix built from two copies of M_{n-1} and one zero block, not yet triangular. **Why natural.** Two identical off-diagonal-style blocks separated by a zero block become block triangular as soon as the two row blocks are swapped, and a block triangular determinant is just the product down the diagonal, so the swap is the one operation standing between the current shape and a shape you already know how to evaluate. **Link to the next block.** The only remaining task is bookkeeping: the swap has a sign, and that sign must be computed, not assumed.

Block 6. **Cue.** The identity $\det M_n = -(\det M_{n-1})^2$ holds for every $n \geq 2$, and the base value $\det M_1 = 1$ has magnitude one. **Why natural.** Squaring a real number always destroys its sign, so from $n = 2$ onward the recursion produces a negative number no matter what sign $\det M_{n-1}$ carried, and a magnitude of 1 squared stays 1 forever, so nothing ever grows or shrinks. **Link to the next block.** The recursion is therefore solved completely by one step of induction, needing no closed form beyond -1.

Layer 3: WHY

The truth underneath this problem is that the power set of $\{1, \ldots, n\}$ is the one combinatorial object on n elements with a perfectly free recursive decomposition: removing the top element splits it into two copies of the power set on $n-1$ elements, with no leftover cases beyond the single set $\{n\}$ itself, and the same split works for any construction whose entries depend only on intersection, union or containment between the indexing subsets, because those three operations interact with "does this set contain n" in a fixed, computable way. The discernment mechanism is recognising the count $2^n - 1$ (or 2^n) as the fingerprint of an object indexed by all subsets, distinct from counts like n, n^2 or $\binom{n}{k}$ which suggest entirely different recursions (on a single element, on a pair, or on choosing a subset of fixed size). You train it by asking, every time you see a matrix, sequence or sum indexed by all subsets of a finite set, "what happens to the top element," before trying anything else.

The examiners nudged you toward this by writing "in some order," a phrase that does no logical work unless the reader is meant to notice that the answer is ordering-independent and therefore free to be computed under a self-chosen, structure-revealing order. A candidate who fixes an arbitrary order (say, binary or lexicographic) and tries to compute directly is fighting the problem; a candidate who reads the phrase as permission is working with it.

Revolving concepts

Recursive self-similarity of the subset lattice: nonempty subsets of $\{1,\dots,n\}$ split into nonempty subsets of $\{1,\dots,n-1\}$ together with all subsets of $\{1,\dots,n-1\}$ (including empty) unioned with $\{n\}$; the same split, applied to unions, symmetric differences, or containment relations instead of intersections, is the standard first move for any subset-indexed recursive construction. See Book 1, field: Combinatorics, topic: induction on the power set by top-element removal.

Block triangular and block-swap determinant rules: for square diagonal blocks A, D, $\det \begin{pmatrix} A & B \\ 0 & D \end{pmatrix} = \det A \cdot \det D$, and swapping two contiguous row blocks of size k each multiplies the determinant by $(-1)^{k^2}$, since the swap is realised by exactly k^2 pairwise crossings between the two blocks' rows. See Book 1, field: Matrices and determinants, topic: block matrix determinants.

Power recursions and sign absorption: for a real recursion $x_n = c\, x_{n-1}^k$ with k even, $x_{n-1}^k \geq 0$ for every real x_{n-1}, so x_n takes the sign of c from the very next term onward regardless of the sign of x_{n-1}; the magnitude separately follows $|x_n| = |c|\,|x_{n-1}|^k$, which becomes linear after taking logarithms, $\log|x_n| = \log|c| + k\log|x_{n-1}|$, and unrolls to $\log|x_n| = \log|c| \cdot \frac{k^{n-1}-1}{k-1} + k^{n-1}\log|x_1|$ for $k \neq 1$. In this problem $k = 2$, $c = -1$, and $|x_1| = |\det M_1| = 1$, so the magnitude stays fixed at 1 forever and the sign is absorbed into -1 after a single step; this is the general reason the recursion resolves in one line rather than needing a closed form in n. See Book 1, field: Sequences and recurrences, topic: nonlinear recursions solved by separating sign and magnitude.

Circuit Bank

- **WHEN** an object is indexed by all subsets (or all nonempty subsets) of $\{1,\dots,n\}$, with entries depending only on set relations **DO** split the index set by membership of the largest element and look for a block matrix built from the same object at size $n-1$ **BECAUSE** the power set decomposes exactly and losslessly under this split, with no leftover cases beyond the single top-element set
- **WHEN** a problem states that an answer does not depend on an arbitrary ordering or labelling **DO** prove that independence first, via invariance under a single transposition, before choosing a convenient ordering to exploit **BECAUSE** every permutation is a product of transpositions, so checking one transposition certifies every relabelling at once
- **WHEN** a block of a matrix is entirely one constant value **DO** subtract a row or column that already equals that constant everywhere in the block **BECAUSE** elementary row and column operations do not change the determinant, and a row equal to the block's constant value zeroes it out in a single step
- **WHEN** two identical blocks sit off the diagonal of a matrix separated by a zero block **DO** swap the corresponding row (or column) blocks and track the sign of the swap from the parity of the block size **BECAUSE** the swap turns the matrix block triangular, whose determinant is just the product of the diagonal blocks, but the sign of a block swap must be computed, not assumed

Backtesting

1. First useful cue: an index set of size $2^n - 1$, the exact count of nonempty subsets of an n-element set.
2. Method the expression suggested: split on membership of the top element n and look for a block matrix expressed in terms of the same construction at size $n-1$.

3. Where you would hesitate, and the thought to reinforce: whether the answer depends on the chosen ordering of the subsets; reinforce by proving invariance under a single transposition before doing anything else.
4. The trap the question sets: being satisfied with a guessed pattern from $n = 2, 3$ computed by hand, instead of the structural recursion a full proof requires.
5. What makes you recognise it faster next time: any matrix, sum or sequence indexed by all subsets of a finite set, with entries depending only on intersection, union or containment.
6. The reusable decision rule: prove order-independence by transposition, split by the top element, and turn the resulting block matrix into a recursion by zeroing constant blocks with elementary row and column operations.

Mirror drills

1. Redo the construction with $m_{ij} = 1$ if $S_i \cup S_j = \{1, \ldots, n\}$, else 0. *Hint:* split on the top element again; now the middle block (both sets containing n) is not automatically all 1s, but the block where neither set contains n is automatically all 0s, so the contamination to clear sits in a different corner.
2. For $m_{ij} = |S_i \cap S_j| \bmod 2$, compute $\det M$ for $n = 2$ and $n = 3$ directly, then check whether the same top-element split still produces two off-diagonal copies of M_{n-1}. *Hint:* it does not, because dropping n from a parity count is not always cost-free the way dropping it from a Boolean "empty or not" test is.
3. Solve $x_n = -3x_{n-1}^2$ with $x_1 = -1$ for x_4. *Answer:* the sign is fixed at -1 for every $n \geq 2$ (since $x_{n-1}^2 \geq 0$ always), and the magnitude follows $|x_n| = 3|x_{n-1}|^2$, giving $|x_2| = 3$, $|x_3| = 27$, $|x_4| = 3 \cdot 27^2 = 2187$, so $x_4 = -2187$.

8.3 A3. Smoothing cubes to maximise a cosine sum

Determine the greatest possible value of $\sum_{i=1}^{10} \cos(3x_i)$ *for real numbers* $x_1, x_2, \ldots, x_{10}$ *satisfying* $\sum_{i=1}^{10} \cos(x_i) = 0$.

Read it out loud

The constraint is stated purely in $\cos(x_i)$, and the target is $\cos(3x_i)$, also purely in terms of the same angle, with no sin anywhere and no information about the actual value of any x_i beyond its cosine. That is the whole cue: the ten angles themselves are irrelevant, only $y_i = \cos x_i \in [-1, 1]$ matters, and $\cos(3\theta)$ expands as a polynomial in $\cos\theta$ through the triple angle identity $\cos 3\theta = 4\cos^3\theta - 3\cos\theta$. Within the first thirty seconds you should have rewritten the whole problem as a polynomial optimisation on a box: maximise $4\sum y_i^3 - 3\sum y_i$ subject to $y_i \in [-1, 1]$ and $\sum y_i = 0$, and notice immediately that the constraint kills the linear term, leaving a bare sum of cubes to maximise.

The second thing to notice is what kind of optimisation this is. It is not a textbook convexity problem, because $y \mapsto y^3$ is concave on $[-1, 0]$ and convex on $[0, 1]$; it changes its mind in the middle of the domain. Jensen's inequality and other single global convexity arguments do not apply directly to a function that is not convex (or concave) across the whole feasible range, and recognising that failure fast is itself useful information: it tells you the extremiser is not going to be a single clean formula, it is going to be a boundary and case structure, found by looking at how the objective responds to small local trades between coordinates rather than by one global inequality.

The field is real analysis dressed as trigonometry: existence of a maximum by compactness, then an exchange argument (fix the sum of a pair of coordinates, see how the objective moves as you shift value between them) to pin down exactly what an optimal configuration must look like.

Approaches, ranked worst to best

1. **Lagrange multipliers on the interior only.** *Cue:* a sum to maximise subject to a sum constraint. *Circuit:* unconstrained critical point equations. (Differentiating $4\sum y_i^3 - \mu\sum y_i$ gives $12y_i^2 = \mu$ for every i, so every y_i is one of two values $\pm\sqrt{\mu/12}$; this finds interior critical points, but the true maximum can and does sit partly on the boundary of the box $y_i = 1$, which Lagrange multipliers alone do not see, so this route needs an entire second argument for the boundary before it can claim anything about a global maximum.)

2. **Global convexity or Jensen on $\sum y_i^3$.** *Cue:* a sum of the same power applied to each variable. *Circuit:* a single inequality covering all cases at once. (This looks like the natural tool for a sum of identical single-variable functions, but y^3 is not convex on all of $[-1, 1]$, only on $[0, 1]$; applying Jensen without splitting the domain produces a bound that is simply not valid for negative y_i, so this route is not merely inefficient, it is not correct as stated.)

3. **Pairwise smoothing at fixed sum.** *Cue:* a box constraint together with a non-globally-convex separable objective. *Circuit:* fix the sum of two coordinates and study the resulting one-variable function's second derivative. (This is the correct route: it needs no global convexity, only the sign of a constant second derivative on a one-dimensional slice, and it converts the whole ten-dimensional problem into a short, finite classification of what an optimal pair can look like.)

The ideal solution

Since $\cos 3\theta = 4\cos^3\theta - 3\cos\theta$ for every real θ, writing $y_i = \cos x_i \in [-1,1]$ turns the target into

$$\sum_{i=1}^{10} \cos(3x_i) = 4\sum_{i=1}^{10} y_i^3 - 3\sum_{i=1}^{10} y_i = 4\sum_{i=1}^{10} y_i^3,$$

the linear term vanishing by the given constraint. So the task is to maximise $4\sum_{i=1}^{10} y_i^3$ over $y \in [-1,1]^{10}$ with $\sum y_i = 0$. This feasible set is a box intersected with a hyperplane, hence closed and bounded, hence compact, and the objective is continuous, so a maximiser exists and may be reasoned about directly.

For an even integer $0 \le n \le 10$, let $V(n)$ be the maximum of $4\sum_{i=1}^{n} y_i^3$ over $y \in [-1,1]^n$ with $\sum y_i = 0$; we compute $V(10)$ by first understanding every even $V(n)$.

Pairwise smoothing. At a maximiser of $V(n)$, fix indices i, j and let $s = y_i + y_j$. Replacing (y_i, y_j) by any $(y, s-y)$ staying in $[-1,1]^2$ changes neither the sum nor the box constraints, so $y = y_i$ must maximise $g(y) = y^3 + (s-y)^3$ on the resulting interval. Since $g''(y) = 6s$ is constant, g is strictly convex when $s > 0$, strictly concave when $s < 0$, and identically zero when $s = 0$ (as $(s-y)^3 = -y^3$ there). A strictly convex function on a closed interval is maximised only at an endpoint; the feasible interval for y is $[s-1, 1]$ when $s \in (0,2]$, and both its endpoints force one of y_i, y_j to equal 1 exactly. A strictly concave function is maximised at its unique critical point, here $y = s/2$, which lies inside the feasible interval $[-1, s+1]$ for every $s \in [-2, 0)$, forcing $y_i = y_j = s/2 \in [-1, 0)$. When $s = 0$, the pair contributes 0 regardless of its values, so deleting it (dropping to $n-2$ coordinates, still summing to 0) changes nothing.

Consequence. At a maximiser of $V(n)$ with $n \ge 2$: either some pair is deletable at no cost, giving $V(n) = V(n-2)$, or no pair is deletable, in which case any two coordinates that are not equal to 1 must, by the case analysis above, be equal and negative (two equal values have $s \ge 0$, and $s > 0$ would force one of them to be 1, so two equal non-1 values force $s < 0$). So every coordinate is either one of $a \ge 0$ copies of 1, or one of $b = n - a$ copies of a single common value $y \le 0$. If $a = 0$ then $ny = 0$ forces $y = 0$ and the objective is 0. If $a \ge 1$, the sum condition $a + by = 0$ gives $y = -a/b$, and $y \ge -1$ forces $a \le b$; the objective becomes

$$4\Big(a \cdot 1^3 + by^3\Big) = 4a\left(1 - \frac{a^2}{b^2}\right) =: W(a,b), \qquad a + b = n,\ 1 \le a \le b.$$

For fixed a, $W(a,b)$ is strictly increasing in b, since $\partial W/\partial b = 8a^3/b^3 > 0$.

Closing the induction. Write $W(n)$ for the largest value of $W(a,b)$ over $a + b = n$, $1 \le a \le b$ (and $W(0) = 0$). The argument above shows $V(n) = \max\big(V(n-2), W(n)\big)$ for $n \ge 2$, with $V(0) = 0 = W(0)$. Given $a + b = n - 2$ with $a \le b$, the pair $(a, b+2)$ has $a + (b+2) = n$ and still $a \le b + 2$, and by the strict monotonicity above $W(a, b+2) > W(a,b)$; taking the best such (a,b) shows $W(n) > W(n-2)$ for every $n \ge 2$. By induction, $V(n) = W(n)$ for every even n, since $W(n) > W(n-2) = V(n-2)$ makes the maximum in $V(n) = \max(V(n-2), W(n))$ always the second term.

It remains to compute $W(10)$: the admissible pairs are $(1,9), (2,8), (3,7), (4,6), (5,5)$, giving

$$W(1,9) = \tfrac{320}{81}, \quad W(2,8) = \tfrac{15}{2}, \quad W(3,7) = \tfrac{480}{49}, \quad W(4,6) = \tfrac{80}{9}, \quad W(5,5) = 0,$$

of which $480/49$ is the largest, attained at $a = 3, b = 7$: three of the x_i with $\cos x_i = 1$ and seven with $\cos x_i = -3/7$.

$$\boxed{\max \sum_{i=1}^{10} \cos(3x_i) = \frac{480}{49}}$$

Layer 1: WHAT

Block 1: expand $\cos 3\theta$ via the triple angle identity and substitute $y_i = \cos x_i$, reducing to a polynomial optimisation on $[-1,1]^{10}$. Block 2: note compactness and continuity to secure that a maximiser genuinely exists. Block 3: derive the pairwise smoothing lemma by holding a pair's sum fixed and reading the constant sign of the resulting one-variable function's second derivative. Block 4: classify every pair at a maximiser into "deletable", "one coordinate equals 1", or "both equal and negative", and note that a deletable pair reduces the problem to a smaller size with no loss. Block 5: strengthen the target claim to every even size n, so the reduction in Block 4 becomes an honest induction, and show the best non-deletable configuration is strictly increasing in size. Block 6: evaluate the finite list of integer configurations at $n = 10$ and identify the maximum.

The nominal route reaches for Lagrange multipliers on the open box or for a direct convexity bound on $\sum y_i^3$, both of which either miss the boundary or are simply invalid on the part of the domain where y^3 is concave. The bifurcation point is the moment you notice $y \mapsto y^3$ changes convexity inside $[-1,1]$: before that realisation the zombie route tries to apply one global inequality; after it, the only honest option is a local, pairwise argument, and everything that follows is a direct consequence of that one observation.

Layer 2: HOW

Block 1. **Cue.** The constraint and the target are both stated purely through $\cos(x_i)$, with no other trigonometric information anywhere in the problem. **Why natural.** Nothing about the actual angles is used or needed, so the substitution $y_i = \cos x_i$ loses no information, and the triple angle identity is exactly the tool that rewrites $\cos(3x_i)$ as a polynomial in that same variable. **Link to the next block.** Once the problem is a polynomial optimisation on a box, existence of a maximum becomes the next thing worth securing before manipulating it.

Block 2. **Cue.** The feasible set is a closed box intersected with a hyperplane, and the objective is a polynomial. **Why natural.** Closed, bounded sets in $\mathbb{R}^{10}$ are compact, and continuous functions on compact sets attain their extrema, so it is legitimate to talk about "the maximiser" as an actual point rather than a value that might only be approached. **Link to the next block.** With a genuine maximiser to reason about, the next question is what local property it must satisfy.

Block 3. **Cue.** $y \mapsto y^3$ is not convex across all of $[-1,1]$, so no single global inequality can be trusted. **Why natural.** Fixing the sum of any two coordinates turns a piece of the high-dimensional problem into a one-variable function whose second derivative is a constant depending only on that fixed sum, so its convexity type (and hence where its maximum sits) is known exactly, without needing convexity everywhere else. **Link to the next block.** The sign of that constant sorts every pair into one of three mutually exclusive behaviours.

Block 4. **Cue.** Three behaviours emerge: strictly convex slice (sum positive), strictly concave slice (sum negative), or a slice with zero curvature (sum exactly zero). **Why natural.** A strictly convex function's maximum on a closed interval is always at an endpoint, a strictly concave function's maximum is at its unique interior critical point when that point is feasible, and a flat function is indifferent to its value entirely, and each of these three facts is a one-line calculus statement, not an assumption. **Link to the next block.** The flat case is the useful one: it means a pair can be discarded without cost, which is precisely the move that lets the argument induct on the number of variables.

Block 5. **Cue.** Discarding a pair changes the problem into the same problem with two fewer variables. **Why natural.** An induction only closes if the statement being proved is stated for every relevant size, not just the size in the original question, so the claim must be generalised to every

even n up to 10 before the discard move can be used at all. **Link to the next block.** Comparing the best achievable value across sizes then tells you which size to actually use.

Block 6. Cue. For a fixed number of coordinates pinned to 1, the achievable value strictly increases as more coordinates are added at the shared negative value. **Why natural.** More coordinates means more total negative mass available to balance the fixed positive coordinates while keeping each individual value further from the boundary -1, and the algebraic derivative of the target with respect to that count is manifestly positive. **Link to the next block.** So the largest available size, $n = 10$, is guaranteed to beat every smaller size, and only a short finite check across integer splits of 10 remains.

Layer 3: WHY

The truth underneath this problem is that optimising a sum of identical copies of a function that is convex on part of its domain and concave on the rest, subject to a linear constraint, is never solved by a single global inequality; it is solved by asking what a two-coordinate trade can do at fixed total, because that question only ever depends on the sign of one number, the constant second derivative of the resulting one-variable slice. This is the same mechanism behind exchange arguments across competition mathematics: whenever a configuration is claimed to be optimal, the cheapest way to test it is to ask whether transferring a small amount between two coordinates, at fixed total, could do better, and the answer is governed entirely by local curvature, never by the values of the other coordinates.

The discernment mechanism is spotting the phrase "sum of y_i^k subject to $\sum y_i$ fixed, y_i boxed" and checking, before attempting any bound, whether $y \mapsto y^k$ is convex across the whole box. If it is, the entire problem collapses to a boundary/all-equal comparison in one step; if it changes sign, as it does here on $[-1, 1]$ for odd $k \geq 3$, the pairwise argument is not a fallback, it is the only correct route, and recognising that in the first minute saves you from chasing an inequality that does not exist.

The examiners nudged you toward this specifically by choosing the target as $\cos(3x_i)$, a cube of the cosine, rather than $\cos(2x_i)$, a square: a square is globally convex on $[-1, 1]$ and the whole problem would collapse to ordinary convexity in one line. The cube is the smallest odd power for which the sign-changing behaviour, and hence the need for the pairwise argument, appears at all.

Revolving concepts

Multiple angle formulas as Chebyshev polynomials: $\cos(n\theta) = T_n(\cos\theta)$ for the Chebyshev polynomial T_n, defined by $T_0(y) = 1$, $T_1(y) = y$, and the three-term recursion $T_{n+1}(y) = 2y\,T_n(y) - T_{n-1}(y)$; this recursion generates $T_2(y) = 2y^2 - 1$ and $T_3(y) = 4y^3 - 3y$ directly, and generalises to computing $\cos(k\theta)$ for any k purely algebraically once $\cos\theta$ is known, without ever needing the angle θ itself. See Book 1, field: Polynomials, topic: Chebyshev polynomials and multiple angle formulas.

Existence of extrema by compactness: a continuous real function on a nonempty closed, bounded subset of $\mathbb{R}^n$ attains both a maximum and a minimum (the extreme value theorem); this single fact is what licenses speaking of "the maximiser" throughout the argument, rather than a supremum that might not be attained. See Book 1, field: Real analysis, topic: compactness and the extreme value theorem.

Pairwise smoothing and exchange arguments: for an objective that is a sum of a single-variable function applied to each coordinate, subject to a linear constraint and box bounds, fixing the sum of any two coordinates reduces the local optimality question to the sign of a constant second derivative on that pair's one-dimensional slice; this generalises to any function that is convex on part of the

domain and concave on the rest, and is the standard tool whenever a global convexity or Jensen argument fails outright. See Book 1, field: Inequalities, topic: exchange and smoothing arguments.

Strengthening a claim to enable induction: when an argument naturally produces a smaller instance of the same problem (here, deleting a pair drops n by two), the claim being proved must be restated for every relevant size of the parameter, not only the size asked about, before the reduction step can be used as an inductive step; this is a general proof-shape, not specific to this problem, and appears whenever an extremal argument allows removing part of a configuration without loss. See Book 1, field: Inequalities, topic: induction via a reduction or removal move.

Circuit Bank

- **WHEN** a constraint and a target are both stated purely through $\cos(x_i)$ for the same angles, with no other trigonometric data given **DO** substitute $y_i = \cos x_i$ and expand the target's multiple angle as the corresponding Chebyshev polynomial in y_i **BECAUSE** the actual angles carry no information the problem uses, so the substitution is lossless and turns a trigonometric question into a polynomial one on a box
- **WHEN** an objective is a sum of the same power y^k (k odd, $k \geq 3$) over variables in a box, subject to a linear constraint **DO** fix the sum of any two coordinates and read the sign of the constant second derivative of the resulting one-variable slice **BECAUSE** y^k changes convexity inside the box for odd $k \geq 3$, so no single global inequality can be trusted, and the pairwise sign test is the finest tool that still applies everywhere
- **WHEN** a pairwise optimality test produces a case where a pair contributes nothing to the objective regardless of its values **DO** delete that pair and restate the claim for every smaller relevant size of the parameter, so the deletion becomes a genuine inductive step **BECAUSE** the removable case is exactly what an induction on the size parameter needs, and proving the general-size claim costs no more than proving the single specific size asked about
- **WHEN** several candidate optimal configurations differ only in how much of a fixed total resource (here, the number of variables) each uses **DO** check monotonicity in that resource directly from the formula before assuming more is better **BECAUSE** the derivative of the achievable value with respect to the resource is a specific algebraic expression that must actually be checked, not assumed positive

Backtesting

1. First useful cue: $\cos(3x_i)$ in the target against a constraint stated purely in $\cos(x_i)$, no other trigonometric information anywhere.
2. Method the expression suggested: substitute $y_i = \cos x_i$, expand via the triple angle (Chebyshev) identity, and reduce to a polynomial optimisation on $[-1, 1]^{10}$.
3. Where you would hesitate, and the thought to reinforce: reaching for Jensen's inequality on $\sum y_i^3$ directly; reinforce by checking that y^3 changes convexity inside $[-1, 1]$, which rules out any single global inequality.
4. The trap the question sets: solving only the interior Lagrange multiplier equations and missing that part of the true maximiser sits exactly on the box boundary $y_i = 1$.
5. What makes you recognise it faster next time: a sum of the same odd power of each variable, subject to a linear constraint and box bounds, whenever that power is at least three.
6. The reusable decision rule: fix the sum of any two coordinates, compute the constant second derivative of the resulting one-variable slice, and classify the pair as boundary-extremal, common-interior-value, or costlessly removable.

Mirror drills

1. Maximise $\sum_{i=1}^{6} \cos(3x_i)$ given $\sum_{i=1}^{6} \cos(x_i) = 0$. *Hint:* same reduction, now check $(a, b) \in \{(1,5), (2,4), (3,3)\}$ with $a + b = 6$; the largest is at $(2,4)$.
2. Maximise $\sum_{i=1}^{n} y_i^3$ over $y_i \in [-1, 1]$ with $\sum y_i = 0$, for n odd. *Hint:* the removable pairs still work two at a time, so the argument reduces an odd n down to $n = 1$, where the only feasible point is $y_1 = 0$, giving objective 0; check whether any odd n can beat this using the same a, b structure with $a + b = n$.
3. Find the greatest possible value of $\sum_{i=1}^{4} \cos(2x_i)$ given $\sum_{i=1}^{4} \cos(x_i) = 1$. *Hint:* $\cos 2\theta = 2\cos^2\theta - 1$ has constant, globally positive second derivative in $\cos\theta$, so ordinary convexity (Jensen, or pushing every variable to a box extreme) applies directly with no pairwise casework needed, a useful contrast to A3 itself.

8.4 A4. A Christoffel word forces two generators to commute

Let m and n be positive integers with $\gcd(m,n) = 1$, and let

$$a_k = \left\lfloor \frac{mk}{n} \right\rfloor - \left\lfloor \frac{m(k-1)}{n} \right\rfloor$$

for $k = 1, 2, \ldots, n$. Suppose that g and h are elements in a group G and that

$$gh^{a_1}gh^{a_2}\cdots gh^{a_n} = e,$$

where e is the identity element. Show that $gh = hg$. (As usual, $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x.)

Read it out loud

This is a group theory question wearing a number theory coat. You are handed one relation in a completely unknown group G and asked to produce a second one. Since nothing at all is known about G, nothing about G may be used: every line of the proof has to be a formal consequence of the single word $W = gh^{a_1}gh^{a_2}\cdots gh^{a_n}$ being trivial. Say that constraint out loud and it hands you the shape of the answer before you have done any arithmetic. From $W = e$ there are exactly two families of consequences available at zero cost: conjugates $uWu^{-1} = e$, and cyclic rotations of W, which are the conjugates by a prefix. So if a short proof exists it consists of writing $ghg^{-1}h^{-1}$ as a rotation of W. Everything else in the problem is bookkeeping that tells you which rotation.

Now the coat. The quantity a_k counts the integers lying in the interval $(m(k-1)/n,\ mk/n]$, the partial sums telescope to $a_1 + \cdots + a_k = \lfloor mk/n \rfloor$, and so the total is $a_1 + \cdots + a_n = m$. Thus W contains exactly n letters g and exactly m letters h. The defining formula makes sense for every integer k, not only for k between 1 and n, and it is periodic in k with period n, because replacing k by $k+n$ adds m to both floors. Each a_k lies in $\{\lfloor m/n \rfloor,\ \lfloor m/n \rfloor + 1\}$. Geometrically W records the staircase lattice path from $(0,0)$ to (n,m) that hugs the line $y = mx/n$ from below, reading g for a step to the right and h for a step up. That object has a name, the Christoffel word of slope m/n, and knowing the name is worth a great deal here, because it tells you that the whole family of such words is one cyclic orbit.

The hypothesis $\gcd(m,n) = 1$ is doing two jobs and you must decide which one you want to spend. Job one: the line $y = mx/n$ meets no lattice point strictly between its endpoints, so the staircase is unambiguous and every a_k is well defined by the geometry. Job two: m is invertible modulo n. The second job is the one that pays, because the exponent sequence is n periodic, so the cyclic group of index shifts acts on it, and coprimality hands you a distinguished shift for free.

Within thirty seconds of reading this you should have on paper: $\sum_k a_k = m$; $a_{k+n} = a_k$; W has n letters g and m letters h; the target is $ghg^{-1}h^{-1} = e$; and the note "only rotations and conjugates are free".

Approaches, ranked worst to best

1. **Small cases in a free group.** *Cue:* a short word equation in two letters. *Circuit:* write out $n = 2, 3, 4$ and hope a cancellation appears. (Worst. G is arbitrary, so there is nothing to cancel against: the only letters you are permitted to move are the ones the relation itself moves. You will fill a page, learn the answer for $n = 3$, and learn nothing that scales.)
2. **Tietze transformations on the presentation.** *Cue:* a one relator group $\langle g, h \mid W \rangle$. *Circuit:* simplify a presentation by adding and deleting generators. (Correct in principle. By von Dyck's theorem it is enough to prove $gh = hg$ inside the universal group $\langle g, h \mid W \rangle$, since any g, h in any G satisfying the relation give a homomorphism out of it. But with no idea which substitution to make you are searching blind through a large space.)

3. **Fundamental group of a punctured torus.** *Cue:* a word in two letters attached to a lattice path. *Circuit:* read a word as a loop and let topology do the algebra. (This works, and it is the most illuminating account of why the result is true: the staircase and the straight segment from $(0,0)$ to (n,m) are homotopic in the once punctured torus, an element of $GL_2(\mathbb{Z})$ carries that segment to a single loop, and the presented group collapses to $\mathbb{Z}$. It ranks low only because it needs fundamental groups, covering space intuition and a homotopy argument, none of which the problem actually requires.)
4. **Euclidean descent on the pair (m,n).** *Cue:* $\gcd(m,n)=1$ sitting next to floor functions. *Circuit:* run the Euclidean algorithm and induct on $m+n$. (Strong, and the natural first thought. When $m>n$ the increments for (m,n) are those for $(m-n,n)$ each raised by one, and $h^{a_k+1} = h\,h^{a_k}$ lets you absorb the extra letter into a new generator $\tilde{g}=gh$, so that case descends in two lines. When $m<n$ you must exchange the roles of m and n, which amounts to reflecting the staircase in the line $y=x$ and traversing it backwards. That step is true, but it is a picture argument, and turning it into airtight prose costs more ink than the whole of the next method.)
5. **Index shift by the inverse of m modulo n.** *Cue:* $\gcd(m,n)=1$ together with an n periodic exponent sequence. *Circuit:* when a periodic sequence is built from m and has period n, shift the index by m^{-1} modulo n. (Best. One identity about floors, one substitution, and the commutator turns out to be a cyclic rotation of the relator. No induction, no pictures, and every step is checkable in the arena.)

The ideal solution

Throughout, extend the definition to every integer k by the same formula, $a_k = \lfloor mk/n \rfloor - \lfloor m(k-1)/n \rfloor$. For a statement P write $[P]$ for 1 if P holds and 0 otherwise.

Lemma 8.4.1. *For every integer k we have $a_{k+n} = a_k$, and $a_1 + a_2 + \cdots + a_n = m$.*

Proof. Replacing k by $k+n$ increases each of mk/n and $m(k-1)/n$ by the integer m, hence increases both floors by m and leaves their difference unchanged. The sum telescopes to $\lfloor mn/n \rfloor - \lfloor 0 \rfloor = m$. □

Lemma 8.4.2. *Let $x \in \{1,\dots,n\}$ satisfy $mx \equiv 1 \pmod{n}$. Then for every integer k,*

$$a_{k-x} \;=\; a_k + [\,n \mid k-1\,] - [\,n \mid k\,].$$

Proof. First, for any integer N,

$$\left\lfloor \frac{N-1}{n} \right\rfloor - \left\lfloor \frac{N}{n} \right\rfloor = -\,[\,n \mid N\,].$$

Indeed, write $N = qn + r$ with $0 \le r < n$. If $r \ge 1$ then $N-1 = qn + (r-1)$ with $0 \le r-1 < n$, so both floors equal q; if $r = 0$ then $N-1 = (q-1)n + (n-1)$, so the left floor is $q-1$ and the right floor is q.

Since $\gcd(m,n)=1$ there are integers x, y with $mx + ny = 1$, and we may take the representative $x \in \{1,\dots,n\}$. From $-mx = ny - 1$ we get $m(k-x) = mk + ny - 1$ and $m(k-x-1) = m(k-1) + ny - 1$. Pulling the integer y out of each floor and cancelling the two copies,

$$a_{k-x} = \left\lfloor \frac{mk-1}{n} \right\rfloor - \left\lfloor \frac{m(k-1)-1}{n} \right\rfloor.$$

Subtract a_k and apply the displayed floor identity twice, once with $N = mk$ and once with $N = m(k-1)$:

$$a_{k-x} - a_k = -\,[\,n \mid mk\,] + [\,n \mid m(k-1)\,].$$

Because $\gcd(m,n)=1$, n divides mk if and only if n divides k, and likewise with $k-1$ in place of k. □

Proof that $gh=hg$. Write $W=gh^{a_1}gh^{a_2}\cdots gh^{a_n}$, so that $W=e$, and fix x as in Lemma 2.

Since $g\cdot(h^{a_1}gh^{a_2}\cdots gh^{a_n})=W=e$ and G is a group, $g^{-1}=h^{a_1}gh^{a_2}\cdots gh^{a_n}$. Substituting this into the commutator,

$$ghg^{-1}h^{-1}=gh\,(h^{a_1}gh^{a_2}\cdots gh^{a_n})\,h^{-1}=gh^{a_1+1}gh^{a_2}\cdots gh^{a_{n-1}}gh^{a_n-1}.$$

The exponents on the right are a_k for every k, with 1 added when $k=1$ and 1 subtracted when $k=n$; that is, they are $a_k+[\,n\mid k-1\,]-[\,n\mid k\,]$ for $k=1,\dots,n$. (When $n=1$ both corrections land on the same exponent and cancel, and this description is still correct.) By Lemma 2 that list is precisely a_{1-x}, a_{2-x}, …, a_{n-x}, so

$$ghg^{-1}h^{-1}=gh^{a_{1-x}}gh^{a_{2-x}}\cdots gh^{a_{n-x}}.$$

Now use periodicity. For $1\le k\le x$ we have $a_{k-x}=a_{k-x+n}$, and as k runs over $1,\dots,x$ the index $k-x+n$ runs over $n+1-x,\dots,n$; for $x<k\le n$ the index $k-x$ runs over $1,\dots,n-x$. Hence, setting

$$P=gh^{a_{n+1-x}}gh^{a_{n+2-x}}\cdots gh^{a_n},\qquad Q=gh^{a_1}gh^{a_2}\cdots gh^{a_{n-x}}$$

(with Q the empty product when $x=n$, which happens only for $n=1$), we have

$$ghg^{-1}h^{-1}=PQ,\qquad QP=gh^{a_1}\cdots gh^{a_n}=W=e.$$

Finally, in a group a one sided inverse is a two sided inverse: from $QP=e$ we get $P=Q^{-1}(QP)=Q^{-1}$, and therefore $PQ=Q^{-1}Q=e$. Thus $ghg^{-1}h^{-1}=e$, that is, $gh=hg$. □

Corollary 8.4.3. *Under the same hypotheses there is an element $c\in G$ with $g=c^m$ and $h=c^{-n}$; in particular the subgroup $\langle g,h\rangle$ is cyclic.*

Proof. Since g and h commute, W collapses to $g^nh^{a_1+\cdots+a_n}=g^nh^m$ by Lemma 1, so $g^nh^m=e$. Choose integers u,v with $um+vn=1$ and put $c=g^vh^{-u}$. From $g^nh^m=e$ we get $h^{-m}=g^n$, hence $h^{-um}=g^{un}$ and $c^m=g^{vm}h^{-um}=g^{vm+un}=g$. From the same relation $g^{-n}=h^m$, hence $g^{-vn}=h^{vm}$ and $c^{-n}=g^{-vn}h^{un}=h^{vm+un}=h$. □

Layer 1: WHAT

Block 1. Note what is free. From $W=e$ you may conjugate and you may rotate cyclically; nothing else is available because G is arbitrary.

Block 2. Turn the target into a word. Read g^{-1} off the relation and substitute it into $ghg^{-1}h^{-1}$.

Block 3. Read the result. It is the same word with 1 added to the first exponent and 1 subtracted from the last.

Block 4. Ask which operation on the index sequence does that. The sequence is n periodic, so shifts are available; coprimality supplies the shift by $x=m^{-1}$ modulo n.

Block 5. Prove the shift identity from the definition, using Bezout and one floor identity.

Block 6. Recognise the shifted word as PQ where $QP=W=e$, and close with the one sided inverse fact.

Juxtaposition. The nominal route is either the small case expansion (which learns nothing) or Euclidean induction on $m+n$ (which is genuinely good but needs the reflection step for $m<n$). The bifurcation point is Block 3. The zombie stares at $gh^{a_1+1}gh^{a_2}\cdots gh^{a_n-1}$ and asks "what can I cancel here?" The ideal solver asks "is this word again a member of the same family, and if so which member?" One question leads to a page of nothing, the other to two lines.

Layer 2: HOW

Block 1. **Cue.** The group G is unnamed and unconstrained, and exactly one relation is given. **Why natural.** When a hypothesis is a single equation in an unknown group, the deductive closure you are allowed to use is the normal closure of that relator. The cheapest members of it are conjugates uWu^{-1} and, as a special case, the cyclic rotations obtained by taking u to be a prefix of W. So before touching the arithmetic, you decide that the target is going to be manufactured as one of those. **Link to the next block.** That decision means you need the target written as a word in g and h with no inverses of g left loose, which is exactly what the relation lets you arrange.

Block 2. **Cue.** The relation begins with the letter g, and the target begins with gh and then wants g^{-1}. **Why natural.** A relation $g \cdot V = e$ is an explicit formula for g^{-1}, namely $g^{-1} = V$. The commutator is the only sensible target (proving $gh = hg$ is proving $ghg^{-1}h^{-1} = e$), and it contains exactly one g^{-1}, so one substitution clears it. Nothing here is inspired; it is the only substitution the data offers. **Link to the next block.** After substituting, the trailing h^{-1} merges with h^{a_n} and the leading h merges with h^{a_1}, so the outcome is again a word of the shape $gh^{b_1} \cdots gh^{b_n}$.

Block 3. **Cue.** The new exponent list is $a_1 + 1,\ a_2,\ \ldots,\ a_{n-1},\ a_n - 1$: the same list, perturbed only at the two ends. **Why natural.** A perturbation localised at the two ends of a cyclically structured object is almost always a rotation in disguise. The word came from a staircase from $(0, 0)$ to (n, m); adding a step at the front and removing one at the back is precisely what happens when you move the last vertical step round to the front, which is a cyclic rotation of the letter string. **Link to the next block.** So you now look for the exact rotation, and rotations of the block structure correspond to shifts of the index k.

Block 4. **Cue.** a_k is n periodic in k, and $\gcd(m, n) = 1$. **Why natural.** Periodicity means the index lives in $\mathbb{Z}/n\mathbb{Z}$, so the only shifts available are by residues, and coprimality distinguishes exactly one of them: x with $mx \equiv 1$. That is not a guess. The value a_k jumps by one exactly when mk crosses a multiple of n, so the shift that moves a single jump by one place must move mk by exactly 1, which is the shift by m^{-1}. **Link to the next block.** With x named, the shift identity is a computation with floors rather than a search.

Block 5. **Cue.** Bezout gives $mx + ny = 1$, and $-mx = ny - 1$ turns a shift of the index into a shift of the numerator by -1. **Why natural.** Floors are insensitive to adding multiples of n to the numerator, so the ny evaporates and the only residue of the shift is the -1. A numerator lowered by one changes the floor exactly when the numerator was divisible by n, which is the one line identity $\lfloor (N-1)/n \rfloor - \lfloor N/n \rfloor = -[\, n \mid N \,]$. Coprimality converts $n \mid mk$ into $n \mid k$, and the two divisibility indicators are exactly the two end corrections from Block 3. **Link to the next block.** The perturbed word is therefore literally the word with indices shifted by $-x$.

Block 6. **Cue.** The shifted index list, cut at the place where it wraps, is a concatenation PQ whose reverse concatenation QP is the original relator. **Why natural.** A cyclic rotation of a word is by definition the two halves of the word swapped. And the reason the problem closes here is the most elementary group fact of all: $QP = e$ forces $P = Q^{-1}$, hence $PQ = e$. The rotation of a trivial word is trivial. **Link to the next block.** The commutator equals $PQ = e$, which is the conclusion.

Layer 3: WHY

The mathematical truth beneath this route is a statement about words, not about groups: the commutator of g and h, computed modulo the relation, is a cyclic rotation of the relator. Once you know that, the group theory is one line, because a cyclic rotation of a word is a conjugate of it and

conjugates of the identity are the identity. Everything hard in the problem has been moved into the combinatorics of the exponent sequence, and that combinatorics is entirely mechanical once you notice periodicity plus coprimality.

Why is the rotation there at all? Because the word is the Christoffel word of slope m/n, the staircase under the line, and the family of staircases with a fixed slope but a shifted starting phase is a single cyclic orbit. Shifting the phase by one lattice step changes the exponent list in exactly two places, and coprimality guarantees that shifting the phase by one step is the same as shifting the index by m^{-1} modulo n. The problem is engineered so that the commutator performs precisely a one step phase shift. That is why the answer is short, and it is why the topological solution works: the staircase and the straight line segment are the same loop, and a straight loop of coprime slope on a torus generates a copy of $\mathbb{Z}$.

The discernment mechanism to install is this. When the hypothesis is one relation and the conclusion is another relation, do not try to manipulate the group. Write both as words in the free group and ask how the conclusion sits inside the normal closure of the relator, starting with the cheapest possibilities: is it the relator itself, a rotation, a conjugate, a product of two conjugates? The search is finite and short, and the arithmetic in the problem is there to tell you which one to try. Here the arithmetic screams: an exponent sequence that is visibly periodic, a coprimality hypothesis whose only algebraic content is invertibility, and a word whose only asymmetric places are its two ends. Three cues, one shift.

Train it in two ways. First, get fluent at reading an exponent sequence as a lattice path, so that "add one at the front, remove one at the back" becomes visibly a rotation rather than an accident of algebra. Second, drill the reflex $QP = e \Rightarrow PQ = e$ until it is instant, because it is the step that converts "I have the right word" into "I am finished", and candidates lose marks by proving the rotation and then not knowing what to do with it.

Note also what the examiners chose not to give you. They did not say G is finite, so counting arguments are dead. They did not say g or h has finite order, so there is no torsion to exploit. They did give $\gcd(m, n) = 1$, which has no geometric use in the conclusion but is exactly the invertibility you need. When a hypothesis has no visible role in the statement of the conclusion, it is a hypothesis about the method.

Revolving concepts

Floor function. For real t, $\lfloor t \rfloor$ is the greatest integer not exceeding t. The identities that orbit it here: $\lfloor t + N \rfloor = \lfloor t \rfloor + N$ for integer N; $\lfloor (N-1)/n \rfloor - \lfloor N/n \rfloor = -[\, n \mid N \,]$ for integers N and $n \geq 1$; and $\lfloor t + s \rfloor - \lfloor t \rfloor \in \{\lfloor s \rfloor, \lfloor s \rfloor + 1\}$ for $s > 0$, which is why every a_k takes one of two values. See Book 1, field: Number theory, topic: floor functions and lattice point counting.

Beatty and mechanical sequences. For real $\alpha > 0$ the sequence $\lfloor \alpha k \rfloor$ is a Beatty sequence and its increments form a mechanical word. The word is periodic precisely when α is rational, with period the denominator in lowest terms; it is balanced, meaning any two factors of equal length have letter counts differing by at most one. See Book 1, field: Combinatorics, topic: words and balanced sequences.

Christoffel words. The lower Christoffel word of slope m/n with $\gcd(m, n) = 1$ is the letter string of the staircase from $(0, 0)$ to (n, m) that stays weakly below $y = mx/n$ and as close to it as possible. Its conjugacy class is a single cyclic orbit of size $m + n$. The associated one relator group $\langle g, h \mid W \rangle$ is infinite cyclic. This is a genuine knowledge gap for anyone who has not met combinatorics on words, and it is worth installing, because the same words appear in continued fractions, in Sturmian dynamics and in the three distance theorem.

Bezout and modular inverses. $\gcd(m, n) = 1$ if and only if $mx + ny = 1$ has an integer solution, if and only if m is invertible modulo n. See Book 1, field: Number theory, topic: Bezout and the extended Euclidean algorithm.

Free groups, presentations, von Dyck. If F is free on g, h and $W \in F$, then giving elements g, h of

a group G with $W(g,h) = e$ is the same as giving a homomorphism $\langle g, h \mid W\rangle \to G$. Hence any identity provable in the presented group holds in every G. This is why it is legitimate, and often cleaner, to prove such statements once and for all in the universal example.

The recursion, generalised. Write $W_{m,n}$ for the word attached to a coprime pair. Two recursions govern the whole family and both come from the Euclidean algorithm. First, for $m > n$, the increments for (m,n) exceed those for $(m-n,n)$ by one in every position, so

$$W_{m,n}(g,h) = W_{m-n,n}(gh,\, h),$$

which is the substitution $g \mapsto gh$, an automorphism of the free group. Second, reversing the staircase and reflecting in $y = x$ gives

$$W_{m,n}(g,h)^{-1} = W_{n,m}(h^{-1},\, g^{-1}),$$

which exchanges the roles of m and n. Alternating the two steps is exactly the continued fraction expansion of m/n, and the family is indexed by the Stern and Brocot tree. The general lesson for every recursion of this type: when a family is indexed by a coprime pair, the induction that works is the Euclidean one, and each descent step should be realised as a substitution rather than as a computation. The general solution of the underlying first order recursion $s_k = s_{k-1} + a_k$, $s_0 = 0$, is the telescoped $s_k = \lfloor mk/n \rfloor$, and it is the telescoping that gives the letter counts for free. See Book 1, field: Sequences and recurrences, topic: telescoping and Euclidean descent.

Circuit Bank

- **WHEN** one relation $W = e$ in an unknown group and a commutator to prove **DO** try to exhibit the commutator as a cyclic rotation or a conjugate of W **BECAUSE** those are the only consequences of a relation that cost nothing, so if the proof is short it must be one of them
- **WHEN** $\gcd(m,n) = 1$ standing next to $\lfloor mk/n \rfloor$ **DO** write down x with $mx \equiv 1 \pmod{n}$ and compute the effect of shifting the index by x **BECAUSE** the increment sequence has period n, so shifts are the available symmetry, and m^{-1} is the shift that moves one jump by one place
- **WHEN** a word whose exponents change by $+1$ at one end and -1 at the other **DO** suspect a cyclic rotation of the same word and look for the shift that realises it **BECAUSE** moving one letter from the back to the front does exactly this to a staircase word
- **WHEN** $QP = e$ **DO** write $PQ = e$ immediately **BECAUSE** a one sided inverse in a group is a two sided inverse, so rotations of a trivial word are trivial
- **WHEN** a hypothesis with no visible role in the conclusion, such as a coprimality condition **DO** treat it as a statement about the method, and ask which operation it makes invertible **BECAUSE** examiners insert such hypotheses to license exactly one step, and finding that step usually finds the solution

Backtesting

1. First useful cue: the group is arbitrary, so the only permitted moves are conjugation and cyclic rotation of the given relator. That single sentence dictates the whole strategy.
2. Method the expression suggested: the exponents are increments of $\lfloor mk/n \rfloor$, hence periodic with period n, hence acted on by index shifts; coprimality names the shift.
3. Where you would hesitate, and the thought to reinforce: after substituting g^{-1} you are looking at $gh^{a_1+1}gh^{a_2}\cdots gh^{a_n-1}$ and the temptation is to start cancelling. Reinforce: do not cancel, classify. Ask whether the new word belongs to the same family as the old one.

4. The trap the question sets: the Euclidean induction is so obviously suggested by $\gcd(m,n)=1$ that most solvers commit to it, then discover that the case $m<n$ needs the reversal of the staircase, and either wave at it or spend a page making it rigorous. The coprimality hypothesis is pointing at invertibility, not at the algorithm.
5. What makes you recognise it faster next time: the shape "one relation in, one relation out, group unspecified". Every such problem is a word problem in the free group, and every short solution is a rotation or a conjugation.
6. The reusable decision rule: when a periodic sequence built from m has period n and $\gcd(m,n)=1$, the index shift by $m^{-1} \bmod n$ is a symmetry that changes the sequence in exactly two places. Write it down before you do anything else.

Mirror drills

1. Show that $a_k \in \{\lfloor m/n \rfloor, \lfloor m/n \rfloor + 1\}$ for every k, without assuming $\gcd(m,n)=1$. *Answer:* write $t = m(k-1)/n$ and $s = m/n$; then $\lfloor t+s \rfloor - \lfloor t \rfloor = \lfloor \{t\} + s \rfloor$ where $\{t\} \in [0,1)$, and $s \le \{t\} + s < s+1$ pins the floor to two values.
2. In any group, deduce $yx = e$ from $xy = e$. *Answer:* $y = x^{-1}(xy) = x^{-1}$, hence $yx = x^{-1}x = e$. This is the last line of the ideal solution and it must be automatic.
3. Take $m=3$, $n=5$. Compute the a_k, find x, and verify the rotation explicitly. *Answer:* $a = (0,1,0,1,1)$ and $x = 2$ since $3 \cdot 2 = 6 \equiv 1 \pmod 5$. The corrected list $(a_1+1, a_2, a_3, a_4, a_5 - 1) = (1,1,0,1,0)$ equals $(a_4, a_5, a_1, a_2, a_3)$. Here $P = ghgh$ and $Q = g^2hg$, with $QP = g^2hg^2hgh = W$ and PQ the rotated word.

8.5 A5. Every derivative nonnegative forces the function flat

Let $f : \mathbb{R} \to \mathbb{R}$ be an infinitely differentiable function satisfying $f(0) = 0$, $f(1) = 1$, and $f(x) \geq 0$ for all $x \in \mathbb{R}$. Show that there exist a positive integer n and a real number x such that $f^{(n)}(x) < 0$.

Read it out loud

The conclusion is an existence statement with two quantifiers, which is your instruction to negate it and work with the negation. Suppose instead that $f^{(n)}(x) \geq 0$ for every positive integer n and every real x. Adding the given $f \geq 0$, you are assuming $f^{(n)} \geq 0$ for all $n \geq 0$, where $f^{(0)} = f$. Functions like that have a name, absolutely monotonic, and the question becomes: can an absolutely monotonic function on the whole line vanish at 0 and equal 1 at 1? Say it that way and you can already feel the answer. Absolute monotonicity is a rigidity condition; morally these are the functions with nonnegative Taylor coefficients, and a function with nonnegative Taylor coefficients that is flat at a point has nothing left to be.

Where is the strength in the hypotheses? Not in $[0, 1]$. On $[0, 1]$ alone the claim is false: $f(x) = x$ has $f(0) = 0$, $f(1) = 1$ and every derivative nonnegative there. So the proof is forced to use the negative axis, and there is only one way to use it. Since $f' \geq 0$ the function is nondecreasing, so for $x \leq 0$ we have $0 \leq f(x) \leq f(0) = 0$, hence f vanishes identically on the left half line, hence every derivative of every order vanishes at 0. That first move is not a choice; it is the only place the global hypothesis can be spent, and after it the problem reads: an absolutely monotonic function that is infinitely flat at the origin still manages to reach 1 at $x = 1$. Kill that.

The fields embedded here are real analysis (monotonicity, convexity, integration by parts, Taylor with remainder) with the theory of absolutely and completely monotone functions orbiting in the background. What is really being tested is a habit of mind rather than a technique: will you prove a statement about the whole class of such functions instead of about your one function? The class is closed under differentiation, and that closure is the engine of the shortest proof. Prove something about f alone and the engine is unavailable.

Within thirty seconds you should have written: negate; $f \equiv 0$ on $(-\infty, 0]$; $f^{(k)}(0) = 0$ for all k; and the note "the class is closed under d/dx, so any inequality can be fed back through the derivative".

Approaches, ranked worst to best

1. **Hunt for a counterexample.** *Cue:* infinitely differentiable but not analytic. *Circuit:* recall the flat bump function e^{-1/x^2} and try to build one that works. (Worst. It is a reasonable ten seconds of orientation and a terrible ten minutes of work. The flat functions all have derivatives of both signs, but discovering that for one example proves nothing, and the problem asks for a proof about every f.)

2. **Taylor at 0 with Lagrange remainder and a direct bound.** *Cue:* all derivatives controlled in sign. *Circuit:* write $f(x) = f^{(n)}(\xi)x^n/n!$ and try to make the right side small. (Stalls. All you know about $f^{(n)}(\xi)$ is that it is nonnegative, and it may be enormous; nonnegativity gives a lower bound on the remainder, not an upper bound. The information is pointing the wrong way and no amount of effort turns it round.)

3. **Quote Bernstein's theorem.** *Cue:* all derivatives of one sign on a half line. *Circuit:* completely monotone functions are Laplace transforms of nonnegative measures. (Correct and quick if you already own the theorem: apply it to $x \mapsto f(M - x)$, get a measure, note $f(0) = 0$ forces the measure to vanish. But it is a sledgehammer far beyond the syllabus, it cannot be reconstructed in the arena, and a grader will want the statement quoted exactly. Name it as a knowledge gap rather than pretending it is elementary.)

4. **Prove that f is analytic.** *Cue:* a C^∞ function with total control on the signs of remainders. *Circuit:* show the Taylor series converges to f, then use that an analytic function flat at a point is identically zero. (Correct and instructive: nonnegativity of $f - g_n$ bounds the Taylor polynomials above by f, and $f - g_n \le f^{(n+1)}(x)/(n+1)!$ closes it. Three moving parts, and the last step needs the identity theorem, so it costs more ink than it needs to.)
5. **Taylor with integral remainder and a kernel comparison.** *Cue:* repeated integration against a nonnegative $f^{(n)}$. *Circuit:* compare $(1-t)^n$ with $(1-t)^{n-1}$ to gain a factor $1/n$. (Very strong, four lines. With $F(x) = \int_0^x f$, both $f(1)$ and $F(1)$ are pure remainder integrals because all derivatives vanish at 0, and the ratio of the two kernels gives $\int_0^1 f \le f(1)/n$ for every n. It ranks second only because it needs the integral form of the remainder to be at your fingertips.)
6. **Self improving inequality on the class.** *Cue:* a hypothesis about infinitely many derivatives. *Circuit:* name the class, note it is closed under differentiation, and feed the inequality back into itself. (Best. Two line base case, two line inductive step by parts, and the conclusion is a contradiction with the Archimedean property. It uses nothing beyond the fundamental theorem of calculus and integration by parts, and the mechanism generalises to every bootstrap argument you will meet.)

The ideal solution

Suppose, for contradiction, that $f^{(n)}(x) \ge 0$ for every positive integer n and every real x. Since also $f \ge 0$, we have $f^{(n)}(x) \ge 0$ for every $n \ge 0$ and every real x, where $f^{(0)} = f$.

Step 1: the left half line. Let u be any infinitely differentiable function with $u^{(n)} \ge 0$ everywhere for all $n \ge 0$ and $u(0) = 0$. Since $u' \ge 0$, u is nondecreasing, so $u(x) \le u(0) = 0$ for every $x \le 0$; combined with $u \ge 0$ this gives $u(x) = 0$ for all $x \le 0$. Hence $u^{(k)}(x) = 0$ for all $x < 0$ and all $k \ge 0$, and since each $u^{(k)}$ is continuous, letting $x \to 0^-$ gives

$$u^{(k)}(0) = 0 \qquad (k \ge 0).$$

Step 2: the class and its closure. Let

$$S = \{\, u : \mathbb{R} \to \mathbb{R} \text{ infinitely differentiable } :\ u^{(n)} \ge 0 \text{ everywhere for all } n \ge 0,\ u(0) = 0 \,\}.$$

Our f lies in S. If $u \in S$ then every derivative of u' is a derivative of u and so is nonnegative, and $u'(0) = 0$ by Step 1; hence $u' \in S$. Thus S is closed under differentiation.

Step 3: the bootstrap. We claim that for every positive integer k,

$$u(x) \le \frac{x}{k}\, u'(x) \qquad \text{for all } u \in S \text{ and all } x \ge 0.$$

Induct on k. For $k = 1$: since $u'' \ge 0$, the derivative u' is nondecreasing, so for $x \ge 0$,

$$u(x) = u(x) - u(0) = \int_0^x u'(t)\, dt \le \int_0^x u'(x)\, dt = x\, u'(x).$$

Assume the claim for k, for every member of S. Let $u \in S$ and $x \ge 0$. Applying the claim to $u' \in S$ gives $u'(t) \le (t/k)\, u''(t)$ for all $t \ge 0$, so

$$k\, u(x) = \int_0^x k\, u'(t)\, dt \le \int_0^x t\, u''(t)\, dt = \Big[\, t\, u'(t) \,\Big]_0^x - \int_0^x u'(t)\, dt = x\, u'(x) - u(x),$$

using integration by parts and $u(0) = 0$. Rearranging, $(k+1)\, u(x) \le x\, u'(x)$, which is the claim for $k+1$.

Step 4: the contradiction. Apply Step 3 to $u = f$ at $x = 1$:

$$1 = f(1) \le \frac{f'(1)}{k} \qquad \text{for every positive integer } k,$$

that is, $k \le f'(1)$ for every positive integer k. This contradicts the Archimedean property of $\mathbb{R}$, since $f'(1)$ is a fixed real number.

Therefore the assumption fails: there exist a positive integer n and a real x with $f^{(n)}(x) < 0$.

Remark 8.5.1. The proof used $f(1) > 0$ rather than $f(1) = 1$, and it used the hypothesis $f \ge 0$ only on $(-\infty, 0]$. Both observations are worth keeping: they identify exactly which parts of the data carry weight.

Layer 1: WHAT

Block 1. Negate the conclusion: assume all derivatives of all orders are nonnegative everywhere, so $f^{(n)} \ge 0$ for $n \ge 0$.

Block 2. Spend the global hypothesis: f nondecreasing and nonnegative with $f(0) = 0$ forces $f \equiv 0$ on $(-\infty, 0]$, hence $f^{(k)}(0) = 0$ for every k.

Block 3. Name the class S of such functions vanishing at the origin, and check that it is closed under differentiation.

Block 4. Prove the crude inequality $u(x) \le x\,u'(x)$ on $[0, \infty)$ from convexity.

Block 5. Feed the inequality back through the derivative and integrate by parts; each pass improves the constant from $1/k$ to $1/(k+1)$.

Block 6. Evaluate at $x = 1$: a fixed positive number is at most $f'(1)/k$ for every k, which is impossible.

Juxtaposition. The nominal route writes Taylor's theorem at 0, observes $f(x)$ equals its remainder, and then tries to bound the remainder above. It cannot: the hypotheses bound $f^{(n)}$ below, not above. The bifurcation point is Block 3. The zombie keeps proving statements about f; the ideal solver widens the target to a statement about every member of S, and only then is the closure of S under differentiation usable. This is the same move as strengthening an induction hypothesis so that the induction closes, and it is worth noticing that the wider statement is easier than the narrow one.

Layer 2: HOW

Block 1. **Cue.** The conclusion says "there exist n and x", a pure existence claim with no candidate offered. **Why natural.** An existence claim with no construction in sight is a contradiction claim in disguise. Negating it converts two existential quantifiers into two universal ones, and universal hypotheses are the ones you can compute with. After negation you hold an infinite family of inequalities, which is far more usable than a promise that one inequality fails somewhere. **Link to the next block.** The negated hypothesis is a statement about all derivatives, so the next question is what class of functions satisfies it.

Block 2. **Cue.** The hypothesis $f(x) \ge 0$ is stated for all $x \in \mathbb{R}$, although the interesting data $f(0) = 0$ and $f(1) = 1$ live in $[0, 1]$. **Why natural.** Examiners do not widen a hypothesis for decoration. On $[0, 1]$ alone the statement is false, as $f(x) = x$ shows, so the negative axis must be load bearing. There is exactly one thing to do with it: $f' \ge 0$ makes f nondecreasing, and a nondecreasing nonnegative function with $f(0) = 0$ is zero to the left of 0. **Link to the next block.** Vanishing on an interval kills all derivatives there, and continuity carries this to the endpoint, giving infinite order flatness at 0.

Block 3. **Cue.** The hypothesis constrains $f, f', f'', \ldots$ all at once, and the conclusion of Block 2 applies verbatim to f'. **Why natural.** A hypothesis quantified over all derivatives is not a statement about one function; it is membership of a class, and this class is visibly stable under differentiation. That stability is useless while your theorem mentions f only. Restate the goal for all $u \in S$ and the stability becomes a tool: anything proved for the class can be applied to u', which is again in the

class. **Link to the next block.** Now you need a first inequality relating u to u', so that the feedback has something to feed on.

Block 4. **Cue.** $u'' \geq 0$, that is, u is convex, and $u(0) = 0$. **Why natural.** Convexity is the statement that u' is nondecreasing, and the cheapest consequence is that the average of u' over $[0, x]$ is at most its endpoint value. Integrating, $u(x) \leq x\,u'(x)$. This is the crudest possible comparison of a function with its derivative and requires no cleverness; you want the crudest one, because the bootstrap will sharpen it. **Link to the next block.** The inequality has a constant 1 in front of xu', and the whole game is now to improve that constant.

Block 5. **Cue.** The induction hypothesis applies to u', producing $u'(t) \leq (t/k)u''(t)$, and you want to compare u with u'. **Why natural.** To convert a statement about u' and u'' into a statement about u and u', integrate. The integrand $t\,u''(t)$ then asks to be integrated by parts, and the reason that is the right move is that the two pieces it produces are exactly the two quantities in play: the boundary term is $xu'(x)$ and the remaining integral is $u(x)$. The inequality closes on itself, and solving for $u(x)$ gains one unit in the denominator. **Link to the next block.** Iterating, the constant is $1/k$ for every k, so the bound tends to zero while $u(1)$ does not move.

Block 6. **Cue.** A fixed positive number bounded by C/k for every positive integer k. **Why natural.** This is the standard way to contradict a positive value: you never show it is negative, you show it is smaller than every positive number. The Archimedean property finishes it. Note that $f'(1)$ needs no estimate at all; it only needs to be a real number, which it is because f is differentiable. **Link to the next block.** The contradiction closes the proof by contradiction opened in Block 1.

Layer 3: WHY

The truth underneath is a rigidity theorem. A function absolutely monotone on an interval is real analytic there, and on a half line $(-\infty, b)$ it is a Laplace transform of a nonnegative measure; this is Bernstein's theorem, in the reflected form for completely monotone functions. A real analytic function that vanishes to infinite order at a point vanishes on a neighbourhood, and by a connectedness argument on the whole interval. Your f is flat at the origin by Block 2, so if it were absolutely monotone it would be identically zero, contradicting $f(1) = 1$. The ideal solution is nothing but a hand rolled proof of that corollary which never mentions analyticity, and the reason it can avoid analyticity is that it only needs one number, $f(1)$, to be crushed, not the whole function.

The discernment mechanism to install has two halves. First: a hypothesis about infinitely many derivatives is a hypothesis about a class, and the class is almost always stable under the operator that indexes it. Whenever that happens, widen the theorem to the class before proving anything, because otherwise the stability sits unused. This is the analytic twin of strengthening the induction hypothesis, and both are cases of the same principle, that a more general statement can be easier to prove because it has more instances available to feed itself. Second: to contradict a fixed positive quantity, manufacture bounds of the form C/k for all k. Repeated integration is the standard factory for such bounds, since each integration against a monotone nonnegative integrand contributes a factor of a length or of an index, and both the second best and the best approach above are just two ways of running that factory.

What were the examiners nudging toward? Three things, all deliberate. They wrote "for all $x \in \mathbb{R}$" when they only needed to name values at 0 and 1; that is the pointer to the left half line. They wrote "infinitely differentiable" rather than "analytic", telling you that flatness at a point is genuinely possible in general, so no cheap appeal to power series is available and the contradiction must be quantitative. And they set $f(0) = 0$ with $f \geq 0$, which makes the origin a global minimum, so the first derivative there vanishes for free and the flatness cascade starts on its own.

Train the discernment by rehearsing the negation step until it is reflexive, and by collecting bootstrap arguments: Gronwall style estimates, the Cauchy and Schwarz self improvement in Hilbert space arguments, and induction schemes where the hypothesis is applied to a transformed version of the same object. In each, the payoff comes from having stated the result for a class rather than for a specimen.

Revolving concepts

Absolutely monotone functions. A function u on an interval with $u^{(n)} \geq 0$ for all $n \geq 0$ is called absolutely monotonic. Nonnegative constants and e^x are examples; the class is closed under addition, multiplication, composition and differentiation, and u absolutely monotone implies u' absolutely monotone. Reflecting the variable turns absolute monotonicity into complete monotonicity, where $(-1)^n u^{(n)} \geq 0$. See Book 1, field: Real analysis, topic: monotonicity classes.

Bernstein's theorem. If $u : [0,\infty) \to \mathbb{R}$ is continuous, infinitely differentiable on $(0,\infty)$ and $(-1)^n u^{(n)} \geq 0$ for all $n \geq 1$ and all $x > 0$, then there is a nonnegative finite Borel measure μ on $[0,\infty)$ with $u(x) = \int_0^\infty e^{-tx}\, d\mu(t)$ for $x \geq 0$. This is a genuine knowledge gap: it needs measure theory and a compactness argument (Helly selection), and you should file it as a named result rather than pretend to reconstruct it.

Taylor with integral remainder. For u of class C^n,

$$u(x) = \sum_{k=0}^{n-1} \frac{u^{(k)}(0)}{k!} x^k + \frac{1}{(n-1)!} \int_0^x (x-t)^{n-1} u^{(n)}(t)\, dt.$$

When every $u^{(k)}(0)$ vanishes the sum disappears and the function equals its remainder, which is the identity behind the second best approach. See Book 1, field: Real analysis, topic: Taylor's theorem in its Lagrange, Cauchy and integral forms.

Flat functions and Borel's theorem. The function equal to e^{-1/x^2} for $x > 0$ and 0 for $x \leq 0$ is C^∞ with all derivatives vanishing at 0, so C^∞ is strictly weaker than analytic. Borel's theorem says every sequence of reals occurs as the derivative sequence at 0 of some C^∞ function. Together these say that flatness at a point is no obstruction in general, so the sign hypothesis is doing all the work in this problem.

Convexity and the chord bound. $u'' \geq 0$ is equivalent to u' nondecreasing and to u lying below its chords and above its tangents. The consequence used here, $\int_0^x u' \leq x u'(x)$, is the statement that the mean value of a nondecreasing function on $[0,x]$ is at most its right endpoint value.

A nonnegative continuous function with zero integral is zero. Needed if you finish by the integral route: if $u \geq 0$ is continuous on $[0,1]$ and $\int_0^1 u = 0$ then $u \equiv 0$, since positivity at a point gives positivity on a neighbourhood and hence a positive integral.

The recursion, generalised. The bootstrap is a recursion in the constant. Suppose that $u(x) \leq c_k\, x\, u'(x)$ holds for all $u \in S$ and all $x \geq 0$. Applying it to $u' \in S$ and integrating by parts,

$$u(x) = \int_0^x u'(t)\, dt \leq c_k \int_0^x t\, u''(t)\, dt = c_k\big(x u'(x) - u(x)\big),$$

so $u(x) \leq \dfrac{c_k}{1+c_k}\, x\, u'(x)$, that is,

$$c_{k+1} = \frac{c_k}{1+c_k}, \qquad c_1 = 1.$$

Take reciprocals: $1/c_{k+1} = 1/c_k + 1$, an arithmetic progression, so $c_k = 1/k$. That is the general method for every recursion of this Riccati flavour, $c_{k+1} = c_k/(a + b c_k)$: invert to linearise, solve the linear recursion, invert back. The same pattern governs the constants in Gronwall bootstraps and in epsilon improvement arguments, where the point is always that the improved constant tends to a limit which is incompatible with the data. See Book 1, field: Sequences and recurrences, topic: linearising a rational recursion by reciprocals.

Circuit Bank

- **WHEN** a conclusion of the form "there exist n and x with a strict inequality" **DO** negate it at once and work with the universal hypothesis **BECAUSE** existence with no candidate is a contradiction proof, and universally quantified inequalities are the ones you can compute with
- **WHEN** a hypothesis constraining infinitely many derivatives **DO** name the class of functions satisfying it and check whether the class is closed under differentiation **BECAUSE** if it is, every inequality you prove for the class can be applied to u' and fed back, which is unavailable while your theorem mentions only f
- **WHEN** a hypothesis stated on all of $\mathbb{R}$ when the data lives on $[0,1]$ **DO** find the step that uses the wider domain, and expect it to be the first step **BECAUSE** examiners widen a hypothesis only when a counterexample exists on the narrow one, so the widening is a signpost to the load bearing move
- **WHEN** an integrand of the form $t\,u''(t)$ **DO** integrate by parts **BECAUSE** the boundary term returns $xu'(x)$ and the remaining integral returns $u(x)$, so the inequality closes on itself and can be solved for $u(x)$
- **WHEN** a fixed positive quantity you wish to contradict **DO** produce a bound of the form C/k valid for every positive integer k **BECAUSE** the Archimedean property then forces the quantity to be nonpositive, and no estimate of C is required

Backtesting

1. First useful cue: the conclusion is existential with no candidate, so negate it; and $f \geq 0$ is imposed on all of $\mathbb{R}$, so the negative axis must be used.
2. Method the expression suggested: monotonicity from $f' \geq 0$ collapses f on the left half line, giving infinite order flatness at 0; then a comparison of u with u' that improves itself under differentiation.
3. Where you would hesitate, and the thought to reinforce: after proving $u(x) \leq xu'(x)$ you may not see what to do with a constant of 1. Reinforce: the class is closed under d/dx, so apply the inequality to u' and integrate. An inequality about a class is a machine, not a fact.
4. The trap the question sets: Taylor at 0 looks irresistible, but the sign hypotheses bound the remainder from below, so every direct estimate points the wrong way. Candidates burn twenty minutes discovering this.
5. What makes you recognise it faster next time: the shape "all derivatives of one sign". That phrase should immediately produce the words absolutely monotone, the fact that the class is closed under differentiation, and the memory that such functions are analytic.
6. The reusable decision rule: when a hypothesis is closed under an operator, prove the theorem for the class and iterate the operator; the constant obeys a first order recursion, and the recursion is what supplies the factor tending to zero.

Mirror drills

1. Let u be C^∞ on $\mathbb{R}$ with $u \geq 0$, $u' \geq 0$ and $u(0) = 0$. Show $u^{(k)}(0) = 0$ for every $k \geq 0$. *Answer:* monotonicity plus nonnegativity gives $u \equiv 0$ on $(-\infty, 0]$, so $u^{(k)} \equiv 0$ on $(-\infty, 0)$, and each $u^{(k)}$ is continuous at 0.
2. Solve $c_1 = 1$, $c_{k+1} = c_k/(1 + c_k)$. *Answer:* take reciprocals, $1/c_{k+1} = 1/c_k + 1$, so $1/c_k = k$ and $c_k = 1/k$. This is exactly the constant produced by the bootstrap.

3. Under the standing hypotheses (all derivatives nonnegative, $f(0) = 0$), prove the alternative finish $\int_0^1 f \le f(1)/n$ for every $n \ge 1$. *Answer:* put $F(x) = \int_0^x f$, so $F^{(k)}(0) = 0$ for all k and $F^{(n+1)} = f^{(n)}$. Taylor with integral remainder gives $f(1) = \frac{1}{(n-1)!}\int_0^1 (1-t)^{n-1} f^{(n)}(t)\,dt$ and $F(1) = \frac{1}{n!}\int_0^1 (1-t)^n f^{(n)}(t)\,dt$; now $0 \le (1-t)^n \le (1-t)^{n-1}$ on $[0,1]$ and $f^{(n)} \ge 0$.

8.6 A6. Rational squared distances give rational area ratios

Suppose that A, B, C, and D are distinct points, no three of which lie on a line, in the Euclidean plane. Show that if the squares of the lengths of the line segments AB, AC, AD, BC, BD, and CD are rational numbers, then the quotient

$$\frac{\text{area}(\triangle ABC)}{\text{area}(\triangle ABD)}$$

is a rational number.

Read it out loud

Euclidean geometry on the surface, linear algebra over $\mathbb{Q}$ underneath. The data is six numbers and every one of them is a *square* of a length. That word is the loudest cue on the page, because a squared length is a dot product of a vector with itself, and the squared length of a difference is the polarisation of a dot product of two different vectors. So the six rational numbers are saying one thing: fix the vectors $b = \vec{AB}$, $c = \vec{AC}$, $d = \vec{AD}$, and the Gram matrix of the triple (b, c, d), meaning the array of all pairwise dot products, has rational entries. Read the hypothesis that way and half the problem has evaporated, because the hypothesis has been converted from six lengths, which are algebraically awkward, into six inner products, which are the natural coordinates of the situation.

The target is a *ratio* of areas, not an area, and the reason is worth pausing on. The areas themselves need not be rational: an equilateral triangle with all squared side lengths equal to 1 has area $\sqrt{3}/4$. Worse, it is easy to prove that each area has rational square, and that is not enough, because $\sqrt{2}$ and $\sqrt{3}$ both have rational squares while $\sqrt{2}/\sqrt{3}$ is irrational. So the proof needs an object that sees both triangles simultaneously. The product of the two areas is such an object, and since an area is half of the absolute value of a determinant, the product of two determinants is the thing to compute. Determinants multiply, so a product of two determinants is a single determinant, and the single determinant it equals is a mixed Gram matrix whose entries are exactly the rational dot products.

The fields wired in: bilinear forms and Gram determinants; multiplicativity of the determinant; elementary field theory, in that $\mathbb{Q}$ is closed under the four operations and contains the absolute value of each of its elements. What is really being tested is whether you can resist placing coordinates. Coordinates are available and they work, but they introduce square roots that you then have to keep track of by hand, and the reason the statement is true stays hidden.

Within thirty seconds you should have on paper: b, c, d as vectors based at A; the polarisation identity $2\,u \cdot v = |u|^2 + |v|^2 - |u - v|^2$; area $= \frac{1}{2}|\det[\,u\ v\,]|$; and the line $\det[\,b\ c\,] \cdot \det[\,b\ d\,] = \det\left([\,b\ c\,]^{\mathsf{T}}[\,b\ d\,]\right)$.

Approaches, ranked worst to best

1. **Unstructured coordinates.** *Cue:* the phrase "Euclidean plane". *Circuit:* put A at the origin and give every other point two unknown coordinates. (Worst. You acquire six unknowns and six quadratic conditions, the unknowns are irrational, and you have no criterion for what is rational at any stage. The algebra is unbounded and the conclusion never comes into view.)

2. **Adapted coordinates.** *Cue:* the segment AB is common to both triangles. *Circuit:* choose the frame that makes the shared datum trivial, origin at the midpoint of AB, A on the positive horizontal axis. (This is the official first route and it works: with $A = (\sqrt{a}, 0)$ and $B = (-\sqrt{a}, 0)$, adding and subtracting the conditions on AC and BC shows $x_1\sqrt{a}$ and $x_1^2 + y_1^2$ rational, hence x_1^2, x_2^2, x_1x_2, y_1^2, y_2^2 rational, and then the condition on CD forces y_1y_2 rational, so $y_1/y_2 = (y_1y_2)/y_1^2$ is rational. Honest, about a page, and it conceals the mechanism: the products x_1x_2 and y_1y_2 are mixed Gram entries in disguise.)

3. **Law of cosines on the angles at A.** *Cue:* squared lengths plus the formula area $= \frac{1}{2}bc\sin A$. *Circuit:* convert lengths into angles, then use the angle addition formula. (Works. Each $2\,AB\cdot AC\cos\alpha$ is rational, so $\cos^2$ and $\sin^2$ of each angle is rational; expanding $\cos(\alpha+\gamma)$ isolates $\sin\alpha\sin\gamma$ as rational, and the ratio of areas is a quotient of such products. Correct, but you must handle signed angles carefully and the ink cost is roughly double.)
4. **The Cayley and Menger determinant.** *Cue:* all pairwise squared distances are given, which is exactly that matrix. *Circuit:* quote the volume formula and the associated polynomial identity. (Correct and conceptually right, since the Cayley and Menger determinant is the canonical function of the squared distance data. But you need to recall the identity relating $\text{area}(ABC)^2\,\text{area}(ABD)^2$ to AB^2V^2 exactly, and recalling a specific polynomial identity under exam conditions is a poor bet.)
5. **Gram matrices and multiplicativity of the determinant.** *Cue:* squared lengths, and a ratio rather than a value. *Circuit:* polarise into dot products, write areas as determinants, and multiply the two determinants into one. (Best. Four lines, no coordinates, no square roots ever written down, and the generalisation to higher dimensions is immediate. It also makes visible exactly which hypothesis does what, which is what lets you construct the sharpness examples.)

The ideal solution

For $u, v \in \mathbb{R}^2$ write $[\,u\;v\,]$ for the 2×2 matrix with columns u and v, and recall that the triangle with vertices A, $A+u$, $A+v$ has area $\frac{1}{2}\,|\det[\,u\;v\,]|$.

Set $b = B-A$, $c = C-A$, $d = D-A$.

Lemma 8.6.1. *All six of the numbers $b\cdot b$, $c\cdot c$, $d\cdot d$, $b\cdot c$, $b\cdot d$, $c\cdot d$ are rational.*

Proof. The three diagonal entries are AB^2, AC^2, AD^2, rational by hypothesis. For the others use polarisation: for any u, v,

$$2\,u\cdot v = |u|^2 + |v|^2 - |u-v|^2.$$

Taking $(u,v) = (b,c)$ gives $2\,b\cdot c = AB^2 + AC^2 - BC^2$, since $b-c = \vec{CB}$ has squared length BC^2; taking $(u,v) = (b,d)$ gives $2\,b\cdot d = AB^2 + AD^2 - BD^2$; taking $(u,v) = (c,d)$ gives $2\,c\cdot d = AC^2 + AD^2 - CD^2$. Each right hand side is rational. □

Lemma 8.6.2. *For $u, v, w, z \in \mathbb{R}^2$,*

$$\det[\,u\;v\,]\cdot\det[\,w\;z\,] \;=\; \det\begin{pmatrix} u\cdot w & u\cdot z \\ v\cdot w & v\cdot z \end{pmatrix}.$$

Proof. Since $\det M^{\mathsf{T}} = \det M$ and the determinant is multiplicative, $\det[\,u\;\;v\,]\det[\,w\;\;z\,] = \det\big([\,u\;v\,]^{\mathsf{T}}[\,w\;z\,]\big)$, and the (i,j) entry of $[\,u\;v\,]^{\mathsf{T}}[\,w\;z\,]$ is the dot product of the ith column of $[\,u\;v\,]$ with the jth column of $[\,w\;z\,]$. □

Proof of the statement. Because A, B, D are three of the four points and no three of the points are collinear, the vectors b and d are linearly independent, so $\det[\,b\;d\,] \neq 0$. Apply Lemma 2 twice:

$$\det[\,b\;c\,]\cdot\det[\,b\;d\,] = \det\begin{pmatrix} b\cdot b & b\cdot d \\ c\cdot b & c\cdot d \end{pmatrix} = (b\cdot b)(c\cdot d) - (b\cdot d)(b\cdot c),$$

$$\big(\det[\,b\;d\,]\big)^2 = \det\begin{pmatrix} b\cdot b & b\cdot d \\ d\cdot b & d\cdot d \end{pmatrix} = (b\cdot b)(d\cdot d) - (b\cdot d)^2.$$

By Lemma 1 both right hand sides are rational, and the second is nonzero. Hence

$$\frac{\det[\,b\;c\,]}{\det[\,b\;d\,]} \;=\; \frac{\det[\,b\;c\,]\cdot\det[\,b\;d\,]}{\big(\det[\,b\;d\,]\big)^2} \;\in\; \mathbb{Q}.$$

Finally,

$$\frac{\text{area}(\triangle ABC)}{\text{area}(\triangle ABD)} = \frac{\frac{1}{2}\,|\det[\,b\;c\,]|}{\frac{1}{2}\,|\det[\,b\;d\,]|} = \left|\frac{\det[\,b\;c\,]}{\det[\,b\;d\,]}\right|,$$

and the absolute value of a rational number is rational. Note that $\det[\,b\;c\,] \neq 0$ as well, since A, B, C are not collinear, so the quotient is a positive rational. □

Remark 8.6.3. The same three lines prove more. Since $\big(\det[\,b\;c\,]\big)^2 = (b\cdot b)(c\cdot c) - (b\cdot c)^2$ is rational, each area has rational square; combined with the above, both determinants lie in the one dimensional $\mathbb{Q}$ vector space spanned by $\det[\,b\;d\,]$, a number whose square is rational. In other words all four areas $\text{area}(\triangle ABC)$, $\text{area}(\triangle ABD)$, $\text{area}(\triangle ACD)$, $\text{area}(\triangle BCD)$ are rational multiples of one common square root.

Proposition 8.6.4 (the general form)**.** *Let $P_0, \dots, P_k$ and Q be points of $\mathbb{R}^N$ lying in a common k dimensional affine subspace, with $P_0, \dots, P_k$ affinely independent and $P_0, \dots, P_{k-1}, Q$ affinely independent. If every pairwise squared distance among these $k+2$ points is rational, then*

$$\frac{\text{vol}_k(P_0P_1\cdots P_k)}{\text{vol}_k(P_0P_1\cdots P_{k-1}Q)} \in \mathbb{Q}.$$

Proof. Choose an orthonormal basis of the k dimensional direction space of the common affine subspace and use it to write the vectors $v_i = P_i - P_0$ and $w = Q - P_0$ as columns in $\mathbb{R}^k$. Dot products are unchanged by this choice, so by polarisation all pairwise dot products of $v_1, \dots, v_k, w$ are rational. The two matrices $M = [\,v_1\;\cdots\;v_k\,]$ and $M' = [\,v_1\;\cdots\;v_{k-1}\;w\,]$ are $k \times k$, the k volumes are $|\det M|/k!$ and $|\det M'|/k!$, and $\det M \det M' = \det(M^{\mathsf{T}}M')$ is a determinant of rationals, as is $(\det M')^2 = \det((M')^{\mathsf{T}}M') \neq 0$. Divide. □

Layer 1: WHAT

Block 1. Vectorise at the common vertex: $b = B - A$, $c = C - A$, $d = D - A$.

Block 2. Polarise: the six squared lengths become the six dot products, all rational.

Block 3. Write each area as half the modulus of a determinant, so the target is $|\det[\,b\;c\,]/\det[\,b\;d\,]|$.

Block 4. Notice the obstruction: each determinant has rational square, and that is not enough, because $\sqrt{2}$ and $\sqrt{3}$ have rational squares and an irrational ratio.

Block 5. Multiply the two determinants into one: $\det[\,b\;c\,]\det[\,b\;d\,] = \det\big([\,b\;c\,]^{\mathsf{T}}[\,b\;d\,]\big)$, a 2×2 determinant of dot products, hence rational.

Block 6. Divide by $(\det[\,b\;d\,])^2$, which is rational and nonzero because A, B, D are not collinear, then take absolute values.

Juxtaposition. The nominal route places coordinates and grinds; it succeeds, but it spends a page and it hides the mechanism inside the products x_1x_2 and y_1y_2. The sharper split is not at coordinates but at Block 4. Almost every solver reaches "each area has rational square" within two minutes. The zombie then tries to prove each area lies in a fixed quadratic field by chasing the actual value of the square root. The ideal solver refuses to look at either area alone and computes their product, because the product is the only polynomial expression in the given data that involves both triangles at once. That is the bifurcation point: *compute the product, never the individual*.

Layer 2: HOW

Block 1. Cue. Both triangles share the side AB, and every listed datum is a distance from one labelled point to another. **Why natural.** Distances are functions of differences, so the coordinate free encoding of the whole configuration is the set of difference vectors based at one chosen vertex. Choosing A, the vertex shared by both triangles, is forced by the shape of the target: the two triangles differ in exactly one vector, c against d. **Link to the next block.** Once the configuration is three vectors, every piece of the data must be expressible in the inner product of those vectors.

Block 2. **Cue.** The hypothesis is about squares of lengths, not lengths. **Why natural.** A squared length is $u \cdot u$. The remaining data, the squared lengths of BC, BD, CD, are squared lengths of differences, and the identity $|u - v|^2 = |u|^2 + |v|^2 - 2u \cdot v$ is the definition of the inner product rearranged. So the hypothesis is literally the statement that the Gram matrix of (b, c, d) has rational entries; nothing is being deduced here, only translated. This is why the problem says squares: it is handing you a Gram matrix. **Link to the next block.** With the hypothesis in inner product form, the conclusion must also be put in inner product form.

Block 3. **Cue.** Area, in a rationality question. **Why natural.** There are two standard formulas for a triangle's area, one half base times height and one half the modulus of a determinant. The first involves a length and a perpendicular distance, neither of which is polynomial in the given data. The determinant is bilinear in the vectors, hence polynomial. In a rationality question you must express everything as polynomials in the data, so the determinant is the only admissible choice. **Link to the next block.** Now the target is a ratio of determinants and the data is a Gram matrix, so the question becomes how determinants and Gram matrices talk to each other.

Block 4. **Cue.** $\big(\det[\, b\ c\,]\big)^2 = (b \cdot b)(c \cdot c) - (b \cdot c)^2$ is visibly rational, and so is the analogue for d. **Why natural.** This is the first thing everyone computes, and it is a trap worth naming out loud: rational squares do not give a rational ratio. Testing the assertion against $\sqrt{2}$ and $\sqrt{3}$ takes five seconds and saves ten minutes. Recognising that your natural first computation is insufficient is itself a step, and the correct response is to ask what extra quantity would suffice. **Link to the next block.** The missing quantity must involve both triangles, and there is exactly one symmetric candidate, the product of the two determinants.

Block 5. **Cue.** A product of two determinants of the same size. **Why natural.** Determinants are multiplicative, so a product of determinants is never left as a product: it is a single determinant of a product of matrices. Transposing one factor costs nothing since $\det M^{\mathsf{T}} = \det M$, and the transposed product is precisely the matrix of pairwise dot products. So the multiplication converts the only unavailable quantity into an array of quantities you already know to be rational. The cue "product of determinants" firing the move "multiply the matrices" is a circuit worth installing on its own. **Link to the next block.** The product is rational and the square of the denominator is rational, so the ratio is a quotient of rationals.

Block 6. **Cue.** The ratio of areas involves absolute values, and a denominator must be nonzero. **Why natural.** The hypothesis that no three points are collinear appears nowhere else in the argument, so this is where it is spent: it makes b and d independent and $\det[\, b\ d\,] \neq 0$. And $\mathbb{Q}$ is closed under absolute value, so the modulus signs cost nothing. **Link to the next block.** Every hypothesis has now been used exactly once, which is the signal that the proof is complete and tight.

Layer 3: WHY

The first truth is that squared distances and Gram matrices are the same data. Polarisation is invertible over any field in which 2 is invertible, so knowing all pairwise squared distances among $P_0, \ldots, P_k$ is exactly knowing all pairwise inner products of the difference vectors, and rationality transfers both ways. Every problem whose hypothesis is a list of squared distances should therefore be rewritten as a problem about a symmetric matrix before anything else happens. That single reflex converts a geometry problem into a linear algebra problem where rationality is a transparent condition.

The second truth is about square roots. A k volume is the square root of a Gram determinant, so under rational hypotheses the volumes live in quadratic extensions of $\mathbb{Q}$. Knowing α^2 and β^2

are rational tells you nothing about α/β; knowing $\alpha\beta$ is rational tells you everything, because $\alpha/\beta = \alpha\beta/\beta^2$. So in any problem where two irrational quantities are compared, look for the product rather than for the individuals. The mixed Gram determinant exists precisely to produce that product, and Lemma 2 is the statement that the product of two areas is a polynomial in the mutual inner products even though each area separately is not.

The third truth concerns which hypotheses matter, and here it pays to build the counterexamples rather than trust a feeling. First, drop CD^2 from the hypotheses. Take

$$A = (0,0), \quad B = (1,0), \quad C = (1,\sqrt{2}), \quad D = (2,\sqrt{3}).$$

Then $AB^2 = 1$, $AC^2 = 3$, $AD^2 = 7$, $BC^2 = 2$, $BD^2 = 4$ are all rational, no three of the points are collinear, and yet

$$\frac{\text{area}(\triangle ABC)}{\text{area}(\triangle ABD)} = \frac{\sqrt{2}}{\sqrt{3}},$$

which is irrational; correspondingly $CD^2 = 6 - 2\sqrt{6}$ is irrational. So the datum CD^2, which touches neither of the two triangles being compared, is indispensable: it is the only source of the entry $c \cdot d$ in the mixed Gram matrix. Second, drop planarity. In $\mathbb{R}^3$ take

$$A = (0,0,0), \quad B = (1,0,0), \quad C = (1,\sqrt{2},0), \quad D = (2,0,\sqrt{3}),$$

so that $AB^2 = 1$, $AC^2 = 3$, $AD^2 = 7$, $BC^2 = 2$, $BD^2 = 4$, $CD^2 = 6$ are all rational and no three points are collinear, while the two areas are again $\sqrt{2}/2$ and $\sqrt{3}/2$ with irrational ratio. The reason the proof fails here is exactly located: with three dimensional column vectors the matrices $[\,b\ c\,]$ and $[\,b\ d\,]$ are 3×2 and the product of their determinants is not defined, while $\det\left([\,b\ c\,]^{\mathsf{T}}[\,b\ d\,]\right)$ is by the Cauchy and Binet formula a sum of three products rather than a single one. The phrase "in the Euclidean plane" is doing real work, and the proof is honest about where.

The discernment mechanism, stated as a rule: in a rationality question, work only with quantities that are polynomial in the given data, and when a target is not polynomial, pair it with a second non polynomial quantity so that the product is. The examiners nudged toward it by giving squares, by asking for a ratio and not a value, and by including the one squared distance, CD^2, that has no other conceivable use. Three cues, one method. To train it, take any configuration problem with distance data and ask which functions of the data are polynomial: dot products yes, lengths no, cosines no, products of two areas yes, single areas no.

Revolving concepts

Gram matrix and Gram determinant. For vectors $v_1, \ldots, v_k$ in an inner product space, the Gram matrix is G with $G_{ij} = v_i \cdot v_j$. It is symmetric and positive semidefinite; $\det G \geq 0$, with equality exactly when the vectors are dependent; and when the v_i lie in a k dimensional space with coordinate matrix M, $G = M^{\mathsf{T}}M$ and $\det G = (\det M)^2$. The k volume of the parallelepiped they span is $\sqrt{\det G}$, and of the simplex, $\sqrt{\det G}/k!$. See Book 1, field: Linear algebra, topic: inner products and Gram matrices.

Polarisation. $2\,u \cdot v = |u|^2 + |v|^2 - |u - v|^2$, valid in any real inner product space. It says the inner product is recoverable from the norm, so distance data and inner product data determine each other.

Determinant multiplicativity, and Cauchy with Binet. $\det(MN) = \det M \det N$ for square M, N, and $\det M^{\mathsf{T}} = \det M$. For an $N \times k$ matrix M and $k \times N$ matrix M' with $N > k$, the Cauchy and Binet formula gives $\det(M^{\mathsf{T}}M') = \sum_S \det(M_S)\det(M'_S)$, summed over all k element subsets S of rows. This is exactly why the argument is planar: the sum collapses to a single term only when $N = k$. See Book 1, field: Matrices and determinants, topic: multiplicativity and Cauchy with Binet.

Area as a determinant. For $u, v \in \mathbb{R}^2$, the triangle spanned has area $\frac{1}{2}|\det[\,u\ v\,]|$; the signed version is $\frac{1}{2}\det[\,u\ v\,]$ and it is the shoelace formula in disguise. Bilinearity is what makes it usable in arithmetic questions.

The Cayley and Menger determinant. For points $P_0, \dots, P_k$ in $\mathbb{R}^k$ with d_{ij} the distances,

$$\mathrm{vol}_k^2 = \frac{(-1)^{k+1}}{2^k (k!)^2} \det \begin{pmatrix} 0 & d_{01}^2 & \cdots & d_{0k}^2 & 1 \\ d_{10}^2 & 0 & \cdots & d_{1k}^2 & 1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ d_{k0}^2 & d_{k1}^2 & \cdots & 0 & 1 \\ 1 & 1 & \cdots & 1 & 0 \end{pmatrix}.$$

For $k = 2$ this is Heron's formula, in the form $16\,\mathrm{area}^2 = -\det$; for $k = 3$ it is the identity of Piero della Francesca, $288V^2 = \det$. The dimensional recursion is the content: the $(k+2) \times (k+2)$ determinant of a point set vanishes exactly when the points lie in a $(k-1)$ dimensional flat, so the same family of determinants simultaneously computes volumes and detects degeneracy, one dimension at a time. This is the canonical function of squared distance data and it is the reason the hypothesis of the problem was phrased in squares. See Book 1, field: Geometry, topic: distance geometry and the Cayley and Menger determinant.

Rationality of $\mathbb{Q}$ *under the operations.* $\mathbb{Q}$ is a field and is closed under absolute value, so a quotient of nonzero rationals, and the modulus of a rational, are rational. Trivial, but it is the last line of the proof and must be said.

Rational distance sets. It is a genuine fact, not an assumption, that configurations satisfying the hypothesis exist, and one can even force all six distances themselves to be rational: Euler found infinitely many points on a circle with pairwise rational distances, obtainable by pushing rational points of the line to the circle by a linear fractional map and using the Brahmagupta and Pell equation $x^2 - 2y^2 = -1$ to supply infinitely many of them. Whether a dense subset of the plane can have all pairwise distances rational is the Erdos and Ulam problem, still open. Knowing this stops you from wasting time wondering whether the hypothesis is vacuous.

Circuit Bank

- **WHEN** a hypothesis giving you *squared* lengths **DO** polarise at once into dot products and declare the Gram matrix rational **BECAUSE** squared lengths are the diagonal of a Gram matrix and polarisation fills in the off diagonal, so the hypothesis is a matrix statement, not a geometry statement
- **WHEN** an area or a volume inside an arithmetic question **DO** replace it by a determinant, never by a base times height or a product of a length and a sine **BECAUSE** determinants are polynomial in the coordinates and bilinear in the vectors, while lengths and sines are not, and only polynomial quantities behave under rationality
- **WHEN** two irrational quantities whose ratio you must control, each with rational square **DO** compute their product, not their values **BECAUSE** $\alpha/\beta = \alpha\beta/\beta^2$, so the product is the single missing ingredient, and rational squares alone never suffice
- **WHEN** a product of two determinants of equal size **DO** fuse them into one determinant of a matrix product, transposing one factor if that produces inner products **BECAUSE** $\det M \det N = \det(M^{\mathsf{T}} N)$ turns a product of unknowns into an array of knowns
- **WHEN** a hypothesis in the list that seems irrelevant to the conclusion **DO** find the single step that uses it, and if you cannot, build the counterexample that deletes it **BECAUSE** a proof that never uses a given datum is almost always wrong, and constructing the sharpness example is the fastest audit available

Backtesting

1. First useful cue: the data consists of *squares* of lengths, which is a Gram matrix; and a ratio is requested rather than a value, which warns that the individual areas are irrational.

2. Method the expression suggested: vectorise at the shared vertex A, polarise into dot products, write areas as determinants, and fuse the two determinants into one.
3. Where you would hesitate, and the thought to reinforce: after computing that each area has rational square you will want to conclude. Reinforce: $\sqrt{2}$ and $\sqrt{3}$ have rational squares and irrational ratio, so stop and find the object that sees both triangles at once.
4. The trap the question sets: coordinates. They work, they feel safe, and they cost a page while burying the mechanism, so that you learn nothing transferable and have a page of algebra in which to make a slip.
5. What makes you recognise it faster next time: the pairing of "squared distances given" with "ratio of areas asked". That pairing has one answer, mixed Gram determinants, and it is the same answer in every dimension.
6. The reusable decision rule: in a rationality question, restrict yourself to quantities polynomial in the data; if the target is not polynomial, multiply it by a partner until it is, then divide by a square at the end.

Mirror drills

1. Prove $16\,\mathrm{area}(\triangle ABC)^2 = 4\,AB^2AC^2 - (AB^2 + AC^2 - BC^2)^2$, and recognise it as Heron's formula. *Answer:* $4\,\mathrm{area}^2 = \big(\det[\,b\ c\,]\big)^2 = (b\cdot b)(c\cdot c) - (b\cdot c)^2$, then substitute $2\,b\cdot c = AB^2 + AC^2 - BC^2$ and multiply by 4.
2. Find four points in the plane, no three collinear, with $AB^2, AC^2, AD^2, BC^2, BD^2$ all rational but the area ratio irrational, and say which line of the proof breaks. *Answer:* $A = (0,0)$, $B = (1,0)$, $C = (1,\sqrt{2})$, $D = (2,\sqrt{3})$ give ratio $\sqrt{2}/\sqrt{3}$; the entry $c\cdot d$ of the mixed Gram matrix is no longer rational, since $CD^2 = 6 - 2\sqrt{6}$.
3. Let u, v, w be vectors in $\mathbb{R}^2$ with all pairwise dot products rational and $\det[\,u\ v\,] \neq 0$. Show that $\det[\,u\ w\,]/\det[\,u\ v\,]$ and $\det[\,v\ w\,]/\det[\,u\ v\,]$ are rational, and deduce that w is a rational combination of u and v. *Answer:* both ratios are rational by the mixed determinant argument, and Cramer's rule expresses w in the basis (u, v) with exactly those two quotients as coefficients.

8.7 B1. Which Point to Remove for a Balanced Split

Let $\mathcal{P}$ be the set of vectors defined by

$$\mathcal{P} = \left\{ \begin{pmatrix} a \\ b \end{pmatrix} \middle| \, 0 \le a \le 2, \; 0 \le b \le 100, \text{ and } a, b \in \mathbb{Z} \right\}.$$

Find all $\mathbf{v} \in \mathcal{P}$ such that the set $\mathcal{P} \setminus \{\mathbf{v}\}$, obtained by omitting the vector $\mathbf{v}$ from $\mathcal{P}$, can be partitioned into two sets of equal size and equal sum.

Read it out loud

This is a grid of lattice points, three columns wide ($a = 0, 1, 2$) and a hundred and one rows tall ($b = 0$ to 100), and you are asked which single point you can delete so that the remaining 302 points split into two halves of 151 points each, whose vector sums agree. Say the sentence out loud and the costume falls away: this is an equal sum partition question, and the coordinates are irrelevant as coordinates. All that matters is that $\mathcal{P}$ is a finite multiset of pairs of numbers, and you are being asked an if and only if question about which single element, once removed, leaves something splittable.

Every question of the shape "remove one object so the rest splits into two equal halves" starts with the same first move, and it should be in your hand within the first thirty seconds: look at the total. If $\mathcal{P} \setminus \{\mathbf{v}\}$ really does split into S_1 and S_2 of equal sum, then $\Sigma(\mathcal{P}) = \mathbf{v} + \Sigma(S_1) + \Sigma(S_2) = \mathbf{v} + 2\Sigma(S_1)$, so $\mathbf{v}$ is congruent to $\Sigma(\mathcal{P})$ modulo 2 in every coordinate. This costs one computation, summing a rectangular grid, which is just the point count times the mean point, and it turns an apparently open search over 303 candidates into a search over a residue class.

The second field wired into the problem is pure construction: once the necessary condition narrows the candidates, you must exhibit an actual partition for each survivor. The set $\mathcal{P}$ is a rectangle of lattice points, and a rectangle has an obvious symmetry, reflection through its centre. That symmetry is the second cue: it hands you 151 pairs of points with identical vector sum before a single explicit set is written down, and the only genuine work left is patching the one point where the symmetry itself breaks down.

Approaches, ranked worst to best

1. **Search small cases by hand.** *Cue:* the grid is finite, so in principle you could just try. *Circuit:* brute enumeration. This costs unbounded ink for no structural payoff: checking 303 candidates against the possible splits of the remaining 302 points is not a proof, and it never explains why the second coordinate must be even.

2. **Guess a construction for a specific removed point, then generalise the pattern.** *Cue:* small cases of the second coordinate look symmetric. *Circuit:* pattern spotting without an invariant. This can find a working split for particular values eventually, but without the sum invariant first you cannot know you have found every valid $\mathbf{v}$, only that the ones you tried happened to work; the necessary direction of the proof is entirely missing.

3. **Sum invariant for necessity, central involution for sufficiency.** *Cue:* the request is an if and only if, and $\mathcal{P}$ is visibly symmetric about its centre. *Circuit:* sum invariant, then pairing by the involution that fixes the centre. This is the route to take: the invariant argument is three lines and narrows $\mathbf{v}$ to 51 candidates of the form $(1, b)$ with b even, and the symmetry of the grid supplies almost all of the remaining construction for free.

The ideal solution

Write $\Sigma(S)$ for the vector sum of a finite set $S \subset \mathbb{Z}^2$. Since a ranges over $\{0,1,2\}$ (mean 1) and b ranges over $\{0,1,\dots,100\}$ (mean 50), and $|\mathcal{P}| = 3 \cdot 101 = 303$,

$$\Sigma(\mathcal{P}) = 303 \begin{pmatrix} 1 \\ 50 \end{pmatrix} = \begin{pmatrix} 303 \\ 15150 \end{pmatrix}.$$

Necessity. Suppose $\mathcal{P} \setminus \{\mathbf{v}\}$ splits into S_1, S_2 with $|S_1| = |S_2|$ and $\Sigma(S_1) = \Sigma(S_2)$. Then $\Sigma(\mathcal{P}) = \mathbf{v} + 2\Sigma(S_1)$, so

$$\mathbf{v} = \Sigma(\mathcal{P}) - 2\Sigma(S_1) \equiv \begin{pmatrix} 303 \\ 15150 \end{pmatrix} \equiv \begin{pmatrix} 1 \\ 0 \end{pmatrix} \pmod 2.$$

Writing $\mathbf{v} = (a,b)$, the first coordinate must be odd, forcing $a = 1$ since $a \in \{0,1,2\}$; the second coordinate must be even. So $\mathbf{v} = (1,b)$ with $b \in \{0,2,4,\dots,100\}$, exactly 51 candidates.

Sufficiency. Fix such a $\mathbf{v} = (1,b)$. Let $\iota(a,b) = (2-a, 100-b)$, an involution of $\mathcal{P}$ to itself, since $a \in \{0,1,2\} \Rightarrow 2-a \in \{0,1,2\}$ and $b \in \{0,\dots,100\} \Rightarrow 100-b \in \{0,\dots,100\}$; every point p satisfies $p + \iota(p) = (2,100)$. The unique fixed point is $(1,50)$; every other point lies in exactly one pair $\{p, \iota(p)\}$, giving 151 pairs, each summing to $(2,100)$.

Case $b \neq 50$. Then $\mathbf{v}$ is not the fixed point, and its partner is $(1,100-b) \neq \mathbf{v}$. Removing $\mathbf{v}$ breaks this pair, leaving the single leftover point $(1,100-b)$ alongside the untouched fixed point $(1,50)$; these two cannot simply be dropped one to a side, since their sum $(2,150-b)$ is not obviously twice anything useful. Repair this by sacrificing one further pair $\{p, (2,100)-p\}$ of the involution, putting both leftovers and p into S_1, and the partner $(2,100)-p$ alone into S_2, together with 74 further intact pairs assigned wholly to S_1 and 75 further intact pairs assigned wholly to S_2 (any split of the remaining pairs into 74 and 75 will do, since every intact pair contributes the identical sum $(2,100)$ regardless of which side it lands on). This gives $|S_1| = 148+3 = 151$ and $|S_2| = 150+1 = 151$ for any choice of p. Equating the two sums,

$$74(2,100) + (1,100-b) + (1,50) + p = 75(2,100) + \big((2,100) - p\big),$$

gives $2p = 2(2,100) - (2,150-b) = (2,50+b)$, so $p = (1,25+b/2)$, an integer point of $\mathcal{P}$ precisely because b is even, the same evenness the necessity step already forced; it lies in the range $[25,75]$ for the second coordinate since $0 \le b \le 100$. One checks $p \neq \mathbf{v}$ and $p \neq (1,100-b)$ exactly when $b \neq 50$, since either equality reduces to $b = 50$. So $\{p, (2,100)-p\} = \{(1,25+b/2), (1,75-b/2)\}$ is a genuine third pair, distinct from the first, and the construction is complete: S_1 consists of $(1,100-b)$, $(1,50)$, $(1,25+b/2)$, and 74 of the remaining 149 pairs; S_2 consists of $(1,75-b/2)$ and the other 75 of those pairs.

Case $b = 50$. Now $\mathbf{v} = (1,50)$ is the fixed point of ι itself, so removing it leaves exactly the 151 intact pairs, with no leftover singles at all. Splitting 151 pairs into two sides of 151 points each needs an odd number of split pairs, since every intact pair contributes an even amount, 2, to whichever side receives it, while a split pair contributes the odd remainder 1 to each side. A single split pair $\{q, (2,100)-q\}$ would force $q = (1,50)$ for the two sides to balance, by the same algebra as above, but that point has already been removed. So split three pairs instead. The offsets $(1,0)$, $(0,1)$, and $(-1,-1)$ from the centre satisfy $(1,0)+(0,1)+(-1,-1) = (0,0)$, so take the three points $(2,50)$, $(1,51)$, $(0,49)$, whose partners under ι are $(0,50)$, $(1,49)$, $(2,51)$. Put $S_1 = \{(2,50),(1,51),(0,49)\}$ together with 74 of the remaining 148 pairs, and $S_2 = \{(0,50),(1,49),(2,51)\}$ together with the other 74 pairs. Both sides have $3+148 = 151$ points, and since the offsets from $(1,50)$ used on each side cancel, $\Sigma(S_1) = 3(1,50) = \Sigma(S_2)$ once the equal contribution of the 74 shared pairs on each side is added.

In every case a valid partition exists, so

$$\boxed{\mathbf{v} = (1,b) \text{ for } b \in \{0,2,4,\dots,100\}}$$

is exactly the set of vectors with the stated property, 51 vectors in all.

Layer 1: WHAT

Block 1: compute $\Sigma(\mathcal{P})$ from the mean of each coordinate. Block 2: the sum invariant forces $\mathbf{v} \equiv \Sigma(\mathcal{P}) \pmod 2$, giving $\mathbf{v} = (1, b)$ with b even. Block 3: the grid's central involution supplies 151 equal sum pairs and one fixed point. Block 4: for $b \neq 50$, solve a balance equation to find the compensating pair. Block 5: for $b = 50$, use a three term zero sum offset identity instead of a single compensating pair.

The nominal route tries to construct partitions for small values of b by trial and error, hoping to spot which $\mathbf{v}$ work; this never proves the only if direction, and it is easy to miss that $b = 50$ needs separate treatment since small trial values rarely land on the centre. The bifurcation point is the line $\Sigma(\mathcal{P}) = \mathbf{v} + 2\Sigma(S_1)$: writing that equation, instead of attempting a construction first, is what turns an open ended search into a two line necessary condition.

Layer 2: HOW

Block 1. **Cue.** A rectangular grid of lattice points, presented as a Cartesian product of two ranges. **Why natural.** The grid is a Cartesian product of two arithmetic progressions, so its total is the point count times the mean point, the fastest route to $\Sigma(\mathcal{P})$ without adding 303 vectors by hand. **Link to the next block.** Having the exact total ready is what makes the invariant argument in the next block a short computation instead of a guess.

Block 2. **Cue.** The phrase partitioned into two sets of equal size and equal sum. **Why natural.** Any such partition's two halves sum to the same value, so their combined total is twice that value; subtracting the removed vector from the grand total must therefore be even in every coordinate, a congruence condition on the removed vector alone. **Link to the next block.** This collapses 303 candidates to 51, and tells you exactly where to look for a construction, namely the middle column with even second coordinate.

Block 3. **Cue.** The surviving candidates all sit on the line of symmetry $a = 1$, and $\mathcal{P}$ is itself symmetric about the point $(1, 50)$. **Why natural.** A rectangle's only point symmetry is reflection through its centre, and reflecting $\mathcal{P}$ through its centre is a bijection of $\mathcal{P}$ to itself, so pairing each point with its reflection costs nothing extra. **Link to the next block.** This turns the sufficiency direction from inventing 151 points from nothing into redistributing 151 ready made pairs.

Block 4. **Cue.** Removing $\mathbf{v}$ breaks exactly one pair, leaving one orphan point beside the untouched fixed point. **Why natural.** Two leftover points whose sum is not obviously twice anything cannot simply be split one to a side, so the fix has to come from deliberately breaking a second pair and solving for the split point that restores the balance. **Link to the next block.** Setting the two sides' sums equal and solving the resulting linear equation for the unknown split point is what produces the specific value $25 + b/2$, rather than guessing it.

Block 5. **Cue.** Removing the fixed point itself leaves an odd total, 151, to split from an even supply of whole pairs. **Why natural.** Splitting an odd count from whole pairs alone is impossible, and the algebra shows a single broken pair would need its two halves to average to the point already removed; three broken pairs, chosen so their offsets from the centre cancel, is the smallest fix that avoids this. **Link to the next block.** The identity $(1, 0) + (0, 1) + (-1, -1) = (0, 0)$ is the shortest way to find three cancelling offsets, the same style of small linear solve as the previous block, one level more elaborate.

Layer 3: WHY

The problem sits at the interaction of an algebraic invariant, the coordinate sum reduced modulo 2, and a geometric symmetry, the point reflection of a centrally symmetric finite set. Every delete one element so the rest is evenly splittable problem has this two part anatomy: the invariant gives the necessary direction for free, because summing is linear and splitting in half means dividing by two, so a residue condition on the removed element is unavoidable; the construction almost always exploits a symmetry native to the ambient set, because a symmetry of the whole set restricted to a subset is an equal sum pairing at zero cost, leaving only the boundary case, the point that breaks the pattern, to patch by hand. The discernment mechanism: on any equal sum partition problem, write the total and divide by two before anything else, and on any symmetric ambient set, a rectangle, an arithmetic progression, a cyclic group, ask what its automorphisms fix, since an automorphism restricted to a subset is a pairing for free. Train this by asking, on every new equal partition problem, what the symmetry group of the ambient set is and what it fixes, before attempting a bespoke construction. The examiners nudge you toward the centre of the grid by choosing point counts, 3 and 101, that are both odd, guaranteeing a single isolated fixed point rather than a symmetry with none at all.

Revolving concepts

Sum invariants for partition problems: a partition into k equal sum parts forces the total to be divisible by k in every coordinate that is being summed; the same idea underlies the classical question of whether a run of consecutive integers can be split into two equal sum halves. See Book 1, field Combinatorics, topic Invariants.

Point symmetry, an involution ι on a finite set with $\iota \circ \iota = \mathrm{id}$: an involution partitions a finite set into fixed points and two element orbits, and here $\mathcal{P}$ has exactly one fixed point because both dimensions of the grid, 3 and 101, are odd counts. See Book 1, field Discrete structures and graphs, topic Group actions and involutions.

Solving a balance equation for an unknown lattice point: not a named theorem but a standing technique, whenever a construction needs the point whose presence restores an equality, write the equality as an equation in that unknown and solve it rather than guess it. This generalises to any adjust several degrees of freedom to hit a target sum problem, where as many equations as unknowns pin the fix down exactly; here effectively one scalar equation, since the pair's other point is determined once the first is chosen. See Book 1, field Combinatorics, topic Constructive existence proofs.

Circuit Bank

- **WHEN** you are asked which element to delete so the rest splits into two equal size, equal sum parts **DO** write the total sum of the ambient set and use total equals the removed element plus twice one side's sum, before attempting any construction **BECAUSE** splitting into two equal sums is linear, so the deleted element is congruent to the total modulo twice anything, an invariant that costs one computation and rules out almost every candidate
- **WHEN** the ambient set is visibly symmetric, a rectangle, an interval, a cyclic structure **DO** find its involution or automorphism group and pair elements by it before building anything by hand **BECAUSE** a symmetry of the whole set restricts to an equal sum pairing on any symmetric subset, for free
- **WHEN** a natural pairing leaves one or two points unpaired and their sum will not balance **DO** sacrifice one further pair, write the target equality as an equation in the unknown split point, and solve for it **BECAUSE** guessing the compensating point is unreliable, while the balance equation has a unique solution whenever the arithmetic allows one

Backtesting

1. First useful cue: the words equal size and equal sum immediately after partitioned, announcing a sum invariant argument before any construction.
2. Method the expression suggested: sum the whole grid using count times mean, then apply total equals the removed vector plus twice one side's sum.
3. Where you would hesitate, and the thought to reinforce: building S_1, S_2 explicitly for $b \neq 50$, when the two leftover points do not balance; the thought to reinforce is set the target equality as an equation in an unknown point, do not guess it.
4. The trap the question sets: $\mathbf{v} = (1, 50)$ passes the parity necessity test but is the one value where the standard single pair fix is unavailable, since it is the fixed point of the very symmetry the construction relies on.
5. What makes you recognise it faster next time: noticing early that $(1, 50)$ is the centre of the grid, and that the centre of any centrally symmetric ambient set is always the special case in this style of problem, worth checking separately before writing the general construction.
6. The reusable decision rule: for a remove one point, split the rest evenly problem on a set with a point symmetry, always test the centre of that symmetry separately, since it is the unique point with no partner.

Mirror drills

1. Let $T = \{1, 2, \ldots, 2n + 1\}$. Determine which $k \in T$ can be removed so that $T \setminus \{k\}$ splits into two size n subsets of equal sum. *Answer:* write $\Sigma(T) = k + 2\Sigma(S_1)$ and pair j with $2n + 2 - j$ about the centre $n + 1$; k must share the parity forced by $\Sigma(T)$, and the fixed point $k = n + 1$ needs the same three pair repair used above for $b = 50$.
2. A multiset of $2m + 1$ vectors has a centre of symmetry c, an involution $p \mapsto 2c - p$ fixing only c. Show that removing c always leaves a set expressible as m pairs of equal sum, while removing any other point leaves two singletons instead. *Answer:* the involution restricted to the set minus its only fixed point has no fixed points left, so it splits into exactly m two element orbits; removing a non fixed point instead breaks exactly one orbit, leaving two singletons, exactly the situation repaired above.
3. For which even b does the pair $\{(1, 25 + b/2), (1, 75 - b/2)\}$ fail to be distinct from $\{(1, b), (1, 100 - b)\}$ and from the fixed point $(1, 50)$? *Answer:* only $b = 50$, exactly the case already excluded and handled separately, confirming the two cases are exhaustive and non overlapping.

8.8 B2. A Decreasing Coefficient Polynomial Avoids the Unit Disc

Let n be a positive integer, and let $f_n(z) = n + (n-1)z + (n-2)z^2 + \cdots + z^{n-1}$. Prove that f_n has no roots in the closed unit disc $\{z \in \mathbb{C} : |z| \leq 1\}$.

Read it out loud

The coefficients of f_n count down from n to 1 in unit steps as the power counts up from 0 to $n-1$. Say that out loud and it is a much narrower description than some degree $n-1$ polynomial: this is a polynomial whose coefficients form an arithmetic progression. The claim to prove, no roots in the closed unit disc, is a statement about where roots cannot be, which usually means you should bound $|f_n(z)|$ away from zero on that region rather than try to solve for the roots directly.

The first cue, inside thirty seconds: a polynomial whose coefficients form an arithmetic progression is exactly the situation where multiplying by $(z-1)$ (or $(1-z)$) telescopes, because subtracting a copy of the series shifted by one power cancels the constant difference between consecutive coefficients and leaves something far shorter. This is the polynomial cousin of summation by parts, and it is the single most useful move whenever coefficients are linear in the exponent.

The fields embedded here are the algebra of polynomial roots, the triangle inequality together with its equality case, and, as alternative routes, the location theorems for roots of a derivative and for polynomials with monotone positive coefficients. The problem is really testing whether you reach for the transformation that turns an n term sum into a genuinely short expression, rather than trying to bound the raw sum directly, which has no exploitable structure until it is transformed.

Approaches, ranked worst to best

1. **Bound $|f_n(z)|$ directly by the reverse triangle inequality.** *Cue:* it is a sum, and sums admit a lower bound from the reverse triangle inequality. *Circuit:* naive lower bound. The resulting bound, $|f_n(z)| \geq n - \big((n-1) + \cdots + 1\big) = n - \binom{n}{2}$, is negative for $n \geq 3$, so it proves nothing at all; the raw sum carries no structure to exploit until it is transformed.

2. **Quote the Enestrom and Kakeya theorem.** *Cue:* the coefficients are positive and strictly monotone, exactly that theorem's hypothesis. *Circuit:* citing a named location theorem. Correct and fast once the theorem is already installed, but it is a fact beyond A level that must be quoted, not derived on the spot, so it is a fallback rather than the ideal route, and it is flagged below as a genuine knowledge gap.

3. **Gauss and Lucas on the coefficient reversal of f_n.** *Cue:* $g_n(z) = nz^{n-1} + \cdots + 2z + 1$ is visibly the derivative of $1 + z + \cdots + z^n$. *Circuit:* derivative root location plus coefficient reversal. Fully rigorous, but it costs two extra layers: identifying the roots of $1 + z + \cdots + z^n$ as roots of unity, then translating back through the coefficient reversal, for the same conclusion the next approach reaches directly.

4. **Multiply by $(z-1)$ to telescope, then bound with the triangle inequality and interrogate its equality case.** *Cue:* coefficients in arithmetic progression. *Circuit:* telescoping by a shifted subtraction, then the equality case of the triangle inequality. This is the route to take: two elementary facts, no named theorem, no auxiliary polynomial, and it produces the required strict inequality in a handful of lines.

The ideal solution

First, $f_n(1) = n + (n-1) + \cdots + 1 = \dfrac{n(n+1)}{2} > 0$, so $z = 1$ is not a root of f_n.

Multiply f_n by $(z-1)$:

$$(z-1)f_n(z) = (z-1)\sum_{k=0}^{n-1}(n-k)z^k = \sum_{k=0}^{n-1}(n-k)z^{k+1} - \sum_{k=0}^{n-1}(n-k)z^k.$$

Re-index the first sum with $j = k+1$, so j runs from 1 to n, and compare it term by term with the second sum, indexed by $k = j$ running from 0 to $n-1$. For $1 \le j \le n-1$ the coefficient of z^j is $\big(n-(j-1)\big) - (n-j) = 1$; the top term, $j = n$, contributes z^n alone, from the first sum only; the bottom term, $k = 0$, contributes $-n$ alone, from the second sum only. Hence

$$(z-1)f_n(z) = z^n + z^{n-1} + \cdots + z - n.$$

(Check at $n = 2$: $(z-1)(2+z) = z^2 + z - 2$, matching the right side.)

Claim 8.8.1. *f_n has no root with $|z| \le 1$.*

Proof. Suppose $f_n(z_0) = 0$ for some z_0 with $|z_0| \le 1$. Since $f_n(1) \neq 0$, necessarily $z_0 \neq 1$, so the identity above, evaluated at z_0, gives

$$z_0^n + z_0^{n-1} + \cdots + z_0 = n.$$

Take moduli and apply the triangle inequality:

$$n = |z_0^n + \cdots + z_0| \le |z_0|^n + \cdots + |z_0| \le \underbrace{1 + \cdots + 1}_{n} = n,$$

the second inequality using $|z_0| \le 1$. Equality must hold throughout both inequalities. Equality in $|z_0|^k \le 1$ for every $k = 1, \ldots, n$ forces $|z_0| = 1$. Equality in $|w_1 + \cdots + w_n| \le |w_1| + \cdots + |w_n|$ for the nonzero terms $w_k = z_0^k$ forces every w_k to share the argument of $w_1 = z_0$; in particular $\arg(z_0^2) \equiv \arg(z_0) \pmod{2\pi}$, that is $2\arg(z_0) \equiv \arg(z_0)$, so $\arg(z_0) \equiv 0$. Combined with $|z_0| = 1$, this forces $z_0 = 1$, contradicting $z_0 \neq 1$. So no such z_0 exists. □

Hence f_n has no root in the closed unit disc, as required. ■

Layer 1: WHAT

Block 1: multiply by $(z-1)$ to telescope the arithmetic coefficients into $z^n + \cdots + z - n$. Block 2: rule out $z = 1$ directly, needed before dividing the argument by the very factor just introduced. Block 3: bound the telescoped expression's modulus by n using $|z| \le 1$. Block 4: interrogate the equality case of both inequalities and show it forces $z = 1$, already excluded.

The nominal route attempts to bound $f_n(z)$ itself with the reverse triangle inequality; this fails outright, since the resulting lower bound is negative once $n \ge 3$. The bifurcation point is the decision to transform the polynomial, multiplying by $z-1$, before attempting any inequality at all, rather than inequality bashing the given expression as it stands: a coefficient sequence that decreases steadily to a vanishing point is never the shape you bound directly, it is the shape you telescope first.

Layer 2: HOW

Block 1. **Cue.** Coefficients $n, n-1, \ldots, 1$, differing by the constant 1 at every step. **Why natural.** Multiplying by $(z-1)$ and shifting one copy of the series against the other cancels every interior coefficient difference, because that difference is the same constant throughout; a non arithmetic coefficient sequence would not collapse this cleanly. **Link to the next block.** Having a two term expression, $z^n + \cdots + z - n$, in hand is what makes an inequality worth attempting at all; the original n term sum was never tight enough to bound usefully.

Block 2. **Cue.** You are about to divide an equation by the factor $(z-1)$ you just introduced. **Why natural.** Multiplying by $(z-1)$ can manufacture a spurious root at $z=1$ that was never a root of f_n; the only safe way to use the telescoped identity is to check $z=1$ separately first, using $f_n(1)$ directly, an easy positive number. **Link to the next block.** Once $z=1$ is excluded, the implication that $f_n(z_0)=0$ forces $z_0^n+\cdots+z_0=n$ is valid for exactly the values of z_0 still to be ruled out.

Block 3. **Cue.** The hypothesis $|z_0|\le 1$, sitting unused until now. **Why natural.** The triangle inequality converts a sum you cannot evaluate into a bound you can compute, and $|z_0|\le 1$ is exactly what caps each term $|z_0|^k$ at 1, giving the clean bound n. **Link to the next block.** The bound matches the required value n exactly, and an exact match is the signal to look at the equality case rather than declare the matter closed.

Block 4. **Cue.** The chain $n\le n$ is not itself a contradiction, it is an equality, and an equality inside an inequality always encodes extra information. **Why natural.** The triangle inequality is tight only when every term points the same way in the complex plane, and $|z_0|^k=1$ for every k only when $|z_0|=1$; combining both equality conditions pins z_0 down to the single point $z_0=1$. **Link to the next block.** That single point was already excluded, so the chain closes into a contradiction, which is the entire content of the proof.

Layer 3: WHY

The telescoping route is the ideal one because it is the unique single algebraic move that turns a coefficient sequence with constant difference into a sequence with only two nonzero terms; any other single move, differentiating, substituting z by $1/z$, splitting into real and imaginary parts, leaves an n term expression, and the triangle inequality can bound n terms but never bound them tightly, exactly as the naive first approach showed. The discernment mechanism: look at consecutive coefficients, and if their difference is constant, the polynomial is a piece of the derivative of a geometric series, so multiplying by $(z-1)$ or differentiating a geometric sum are the two doors into that structure. Once the expression is down to two terms, the triangle inequality is worth trying, and whenever the bound obtained matches the target value exactly, chase the equality case rather than stop, since an exact match is either the whole proof, if equality turns out to be impossible, or a genuine extra root you have missed. The examiners nudge you here by letting the coefficients run all the way down to 1 rather than stopping early, which is exactly what makes $f_n(1)>0$ a clean, separate fact worth checking before anything else; a coefficient pattern that stopped short would not offer this free exclusion of $z=1$.

Revolving concepts

Polynomials with coefficients in arithmetic progression, and telescoping under multiplication by $z-1$: the general statement, for any sequence c_k, is that $(z-1)\sum c_k z^k$ equals $\sum(c_{k-1}-c_k)z^k$ plus boundary terms, the polynomial analogue of summation by parts. Whenever the difference $c_{k-1}-c_k$ is itself simple, constant, geometric, or otherwise short, the same multiply and telescope move applies, not only to arithmetic sequences. See Book 1, field Polynomials, topic Coefficient transforms.

The triangle inequality for complex numbers and its equality case, all terms sharing a common argument: this orbits every problem where a sum of complex numbers must be bounded and the bound is claimed to be tight or not, and the two sided use here, bounding then interrogating equality, is the standard pattern for such problems. See Book 1, field Inequalities, topic Triangle inequality in the complex plane.

Roots of unity as the roots of $1+z+\cdots+z^n$, the Gauss and Lucas theorem locating the roots of a derivative inside the convex hull of the roots of the original polynomial, coefficient reversal and

reciprocal roots, and the Enestrom and Kakeya theorem for polynomials with monotone positive coefficients. These four facts power the two alternative routes above; the Gauss and Lucas theorem and the Enestrom and Kakeya theorem both sit beyond A level and are flagged as knowledge gaps, the first because it is a genuinely useful, frequently quotable location theorem for the roots of a derivative, the second because it answers exactly this shape of question, monotone positive coefficients, in one line once memorised, and deserves to be installed even though the ideal solution above does not need it. See Book 1, field Complex numbers, topic Root location theorems.

Circuit Bank

- **WHEN** a polynomial's coefficients form an arithmetic progression, or any sequence with a short, simple difference pattern **DO** multiply by $(z-1)$ (or $(1-z)$) and telescope before attempting any inequality **BECAUSE** the shifted subtraction cancels the interior structure of the coefficients and leaves a short expression that an inequality can actually bite on
- **WHEN** you are about to multiply an equation by a factor such as $(z-1)$ to simplify it **DO** check separately, first, whether the value that factor vanishes at is already excluded **BECAUSE** multiplying by a vanishing factor can manufacture roots that were never there, and the simplified identity is only valid once that value is ruled out on its own
- **WHEN** a bound obtained from the triangle inequality exactly matches the target value **DO** do not stop; write down the precise equality condition of every inequality used and solve for what it forces **BECAUSE** an exact match is only a proof once equality is shown impossible, and the only way to know that is to chase the equality case to its end

Backtesting

1. First useful cue: coefficients running down from n to 1 in unit steps, the signature of an arithmetic progression.
2. Method the expression suggested: multiply by $(z-1)$ to telescope, then bound the result with the triangle inequality.
3. Where you would hesitate, and the thought to reinforce: on reaching $|z_0^n + \cdots + z_0| \leq n$ and seeing it match the required value n exactly, when it is tempting to call the bound sufficient; the thought to reinforce is that an exact match means chase the equality case, it is not yet a contradiction.
4. The trap the question sets: multiplying by $(z-1)$ is only valid once $z=1$ has been excluded separately, and it is easy to forget this and treat the telescoped identity as fully equivalent to the original one.
5. What makes you recognise it faster next time: seeing that the reverse triangle inequality applied to the raw, untransformed sum gives a useless negative bound is itself the signal that the polynomial needs transforming before any inequality is worth writing down.
6. The reusable decision rule: whenever consecutive coefficients differ by a constant, or any short and simple pattern, multiply by $(z-1)$ and telescope before reaching for any inequality on the original expression.

Mirror drills

1. Show that $g(z) = 1 + 2z + 3z^2 + \cdots + nz^{n-1}$ has no root with $|z| \leq \frac{n-1}{n+1}$. *Hint:* g is the derivative of $1 + z + \cdots + z^n$; either telescope directly by a shifted subtraction, or reduce to the Gauss and Lucas theorem on the roots of unity, as in the third approach above.

2. For which z with $|z| = 1$ does $1 + z + z^2 + \cdots + z^{n-1}$ attain $\left|1 + z + \cdots + z^{n-1}\right| = n$ exactly? *Answer:* only $z = 1$, by the identical equality case argument used above.

3. Derive the general identity $(z-1)\sum_{k=0}^{n-1} c_k z^k = c_{n-1}z^n - c_0 + \sum_{k=1}^{n-1}(c_{k-1} - c_k)z^k$ for an arbitrary coefficient sequence $c_0, \ldots, c_{n-1}$, and check it reduces to the identity used above when $c_k = n - k$. *Hint:* shift the index in one of the two sums exactly as in the ideal solution, and match coefficients power by power.

8.9 B3. The Power of Two Tower Divisibility Chain

Find all positive integers $n < 10^{100}$ for which, simultaneously, n divides 2^n, $n-1$ divides $2^n - 1$, and $n-2$ divides $2^n - 2$.

Read it out loud

Write the three conditions uniformly as $(n-j) \mid (2^n - j)$ for $j = 0, 1, 2$. Since $(2^n - n) - (2^n - j) = j - n = -(n-j)$, each condition is equivalent to $(n-j) \mid (2^n - n)$. So the question, read out loud, is not three unrelated facts: it is asking for the n such that three consecutive integers, n, $n-1$, $n-2$, all divide the single fixed quantity $2^n - n$. The field wired into this is elementary number theory, factorisation for the condition depending on n alone, and the multiplicative order of 2 for the two that involve arithmetic modulo a shifted value of n.

The first cue, inside thirty seconds: of the three conditions, $n \mid 2^n$ is far more restrictive than the other two, because it is a statement about n's prime factorisation alone, with no exponential arithmetic required. A number dividing a pure power of 2 can itself have no odd prime factor, so $n \mid 2^n$ forces n to be a power of 2 outright. That is the condition to resolve first: always resolve the condition that pins down the unknown's shape before touching a condition that needs arithmetic modulo the unknown.

Once $n = 2^m$, the remaining two conditions both take the shape $2^k - 1$ divides $2^\ell - 1$, governed by exactly one lemma, that the order of 2 modulo $2^k - 1$ is exactly k. That lemma is applied twice, once to pin down m, once more, after cancelling a common factor of 2 from the third condition, to pin down the exponent ℓ with $m = 2^\ell$. The bound $n < 10^{100}$ looks like it wants heavy machinery, but it is arithmetic dressing: one logarithm converts it into a bound on m small enough to finish by listing candidates.

Approaches, ranked worst to best

1. **Search directly for n by trial division.** *Cue:* the bound 10^{100} looks like it invites a computer search. *Circuit:* brute enumeration. Hopeless on its own terms: 10^{100} candidates cannot be tested one at a time, and even a sieve gains nothing without first knowing n must be a power of 2.
2. **Analyse $(n-1) \mid (2^n - 1)$ or $(n-2) \mid (2^n - 2)$ before pinning down the shape of n.** *Cue:* these two look more special, since they subtract 1 and 2. *Circuit:* multiplicative order modulo an arbitrary integer. Technically possible but far harder: the order of 2 modulo a general integer has no clean formula, whereas the order modulo a Mersenne number $2^k - 1$ is exactly k; tackling these first throws away the simplification the first condition supplies for free.
3. **Resolve $n \mid 2^n$ first, then apply the Mersenne order lemma twice, peeling one layer of the power tower at each remaining condition, then convert the bound by a logarithm.** *Cue:* the three conditions are visibly nested in restrictiveness. *Circuit:* power of two factorisation, then the recurring order lemma, then a logarithm estimate. This is the route to take: each step uses the cheapest tool available for that layer, and the recursive self similarity of the three conditions is used explicitly rather than fought.

The ideal solution

Lemma 8.9.1. *A positive integer k divides 2^k if and only if k is a power of 2, including $k = 1 = 2^0$.*

Proof. If $k = 2^r$ then $2^r \mid 2^{2^r}$, since $r \leq 2^r$ for every $r \geq 0$. Conversely, if k has an odd prime factor p, then $p \mid k \mid 2^k$ would force $p \mid 2^k$, impossible, since 2^k has no prime factor other than 2. □

Lemma 8.9.2. *For positive integers k, ℓ, $2^k - 1$ divides $2^\ell - 1$ if and only if k divides ℓ.*

Proof. Write $\ell = qk + r$ with $0 \le r < k$. Modulo $2^k - 1$, $2^k \equiv 1$, so $2^\ell = (2^k)^q 2^r \equiv 2^r$, giving $2^\ell - 1 \equiv 2^r - 1 \pmod{2^k - 1}$. If $k \mid \ell$ then $r = 0$ and $2^\ell - 1 \equiv 0$. Conversely if $2^k - 1 \mid 2^\ell - 1$ then $2^k - 1 \mid 2^r - 1$; but $0 \le 2^r - 1 < 2^k - 1$ since $0 \le r < k$, forcing $2^r - 1 = 0$, that is $r = 0$, that is $k \mid \ell$. □

Now solve the problem. By the power of two lemma applied with $k = n$, the first condition forces $n = 2^m$ for some integer $m \ge 0$. Since $2^{10} > 10^3$, $2^{340} = (2^{10})^{34} > (10^3)^{34} = 10^{102} > 10^{100}$, so $n < 10^{100}$ forces $m < 340$, that is $m \le 339$. The case $m = 0$, that is $n = 1$, already fails the second condition, since $n - 1 = 0$ and 0 divides only 0, while $2^n - 1 = 1 \ne 0$; so in fact $1 \le m \le 339$.

For the second condition, $n - 1 = 2^m - 1$, so the Mersenne order lemma (with $k = m$, $\ell = n = 2^m$) shows $(n - 1) \mid (2^n - 1)$ holds if and only if $m \mid 2^m$, which by the power of two lemma (with $k = m$) holds if and only if m is itself a power of 2. Write $m = 2^p$ for some integer $p \ge 0$.

For the third condition, $n - 2 = 2^m - 2 = 2(2^{m-1} - 1)$ and $2^n - 2 = 2(2^{n-1} - 1)$, both even since $m \ge 1$ makes $n = 2^m$ even; cancelling the common factor 2, $(n - 2) \mid (2^n - 2)$ is equivalent to $(2^{m-1} - 1) \mid (2^{n-1} - 1)$. If $m = 1$ (that is $p = 0$, $n = 2$), this must be checked directly rather than through the lemma, since $m - 1 = 0$: here $n - 2 = 0$ and $2^n - 2 = 2 \ne 0$, so the condition fails, and $p = 0$ is excluded. For $m \ge 2$ (that is $p \ge 1$), the Mersenne order lemma (with $k = m - 1$, $\ell = n - 1 = 2^m - 1$) shows the condition holds if and only if $(m - 1) \mid (2^m - 1)$, that is, since $m = 2^p$, if and only if $(2^p - 1) \mid (2^{2^p} - 1)$. Applying the Mersenne order lemma once more (with $k = p$, $\ell = 2^p$) this holds if and only if $p \mid 2^p$, which by the power of two lemma (with $k = p$) holds if and only if p is itself a power of 2. Write $p = 2^t$ for some integer $t \ge 0$.

So all three conditions together force

$$n = 2^m = 2^{2^p}, \qquad p = 2^t \ge 1.$$

The bound $m \le 339$ becomes $2^p \le 339$; since $2^8 = 256 \le 339 < 512 = 2^9$, this restricts p to the powers of 2 not exceeding 8, namely $p \in \{1, 2, 4, 8\}$ (the next power, $p = 16$, gives $2^{16} = 65536 > 339$). The four resulting values of n are

$$n = 2^{2^1} = 4, \qquad n = 2^{2^2} = 16, \qquad n = 2^{2^4} = 65536, \qquad n = 2^{2^8} = 2^{256}.$$

Each satisfies $n < 10^{100}$: the largest, $2^{256} = (2^{10})^{25} \cdot 2^6 < (10^3)^{25} \cdot 100 = 10^{77} < 10^{100}$. Direct verification confirms each value satisfies all three original conditions: for instance $n = 4$ gives $4 \mid 16$, $3 \mid 15$, $2 \mid 14$; $n = 16$ gives $16 \mid 65536$, $15 \mid 65535$, $14 \mid 65534$; the same chain of reasoning run forward verifies $n = 65536$ and $n = 2^{256}$.

$$\boxed{n \in \{4,\ 16,\ 65536,\ 2^{256}\}}$$

Layer 1: WHAT

Block 1: state and prove that an integer divides a power of 2 exactly when it is itself a power of 2. Block 2: state and prove that $2^k - 1$ divides $2^\ell - 1$ exactly when k divides ℓ. Block 3: apply Block 1 to the first condition, fixing $n = 2^m$, and bound m by a logarithm. Block 4: apply Block 2 then Block 1 to the second condition, fixing $m = 2^p$. Block 5: cancel a factor of 2 in the third condition, then apply Block 2 twice and Block 1, fixing $p = 2^t$. Block 6: intersect with the bound on m to get the finite list $p \in \{1, 2, 4, 8\}$ and the four values of n.

The nominal route analyses $(n - 1) \mid (2^n - 1)$ or $(n - 2) \mid (2^n - 2)$ directly, without first fixing $n = 2^m$, which forces reasoning about the order of 2 modulo an arbitrary integer, a genuinely hard, unstructured problem with no clean formula. The bifurcation point is the decision to resolve $n \mid 2^n$, the condition depending on n alone with no shift, before either of the other two, because it is the tightest, cheapest constraint available.

Layer 2: HOW

Block 1. **Cue.** The condition $n \mid 2^n$. **Why natural.** A number dividing a pure power of 2 can itself carry no other prime factor, so n's entire structure collapses to a single exponent, $n = 2^m$; this is the cheapest possible reduction, using only unique factorisation, before any exponential arithmetic is needed. **Link to the next block.** Having fixed $n = 2^m$, the other two conditions can now be rephrased purely in terms of m.

Block 2. **Cue.** $(n-1) \mid (2^n - 1)$, and $n - 1 = 2^m - 1$ is now a Mersenne type number. **Why natural.** Powers of 2 modulo $2^m - 1$ repeat with period exactly m, since $2^m \equiv 1$ and no smaller positive power can be congruent to 1, so $2^n \equiv 1 \pmod{2^m - 1}$ exactly when $m \mid n$. **Link to the next block.** Since $n = 2^m$, the condition $m \mid n$ becomes $m \mid 2^m$, the same test as Block 1 applied one level down, to m instead of n.

Block 3. **Cue.** $m \mid 2^m$. **Why natural.** This is the Block 1 lemma again, verbatim, with m in the role of n; recognising the repetition saves re-deriving the factorisation argument from scratch. **Link to the next block.** Writing $m = 2^p$ produces a new unknown, p, and the third condition must now be re-expressed in terms of it.

Block 4. **Cue.** $(n-2) \mid (2^n - 2)$, and both sides are visibly even once $m \geq 1$. **Why natural.** Dividing a divisibility statement by a common factor is always legitimate and always worth doing when both sides share one; here it strips the condition to $(2^{m-1} - 1) \mid (2^{n-1} - 1)$, the same Mersenne shape as Block 2, shifted down by one throughout. **Link to the next block.** Applying the Block 2 lemma twice in a row, once relating $m - 1$ and $n - 1$, once more, using $n - 1 = 2^m - 1$ and $m = 2^p$, to relate p and 2^p, forces p itself to be a power of 2, the same test a third time.

Block 5. **Cue.** The bound $n < 10^{100}$, sitting untouched since the problem statement. **Why natural.** An astronomical bound on n is only useful once translated, by one logarithm estimate, into a bound on the much smaller exponent m, since the finite list of candidates actually needed is a list of values of p, not of n. **Link to the next block.** Intersecting $p \in \{1, 2, 4, 8, 16, \ldots\}$ with $2^p \lesssim 339$ leaves exactly four values, and the problem is finished by exhibiting n for each.

Layer 3: WHY

This problem resolves into a clean finite list rather than an intractable search because all three conditions, stripped of their surface differences, are instances of the identical arithmetic fact applied at nested depths: an integer divides a power of 2 exactly when it is itself a power of 2, and a Mersenne number $2^k - 1$ divides another exactly when the exponents are in a divisibility relation, itself the multiplicative cousin of the first fact, since the greatest common divisor of $2^k - 1$ and $2^\ell - 1$ is $2^{\gcd(k,\ell)} - 1$, mirroring $\gcd(k, \ell)$ directly. Once the three given conditions are rewritten in the uniform shifted form $(n - j) \mid (2^n - n)$ for $j = 0, 1, 2$, the recursive self similarity stops being a coincidence: the problem is manufactured by nesting the same test three times, at n, then at the exponent m that test produces, then at the exponent p that test produces again. The discernment mechanism: whenever a family of divisibility conditions involves the same base raised to the same exponent at several shifted moduli, look for the order of that base modulo each modulus, and whenever an integer condition looks like a power of two test, suspect the same test recurring one level down once the previous answer is substituted in. Train this by writing each condition in a uniform shifted form before analysing any of them separately, since the uniform form is what reveals whether conditions are genuinely nested or only superficially similar. The examiners nudge you toward peeling one layer at a time by choosing three conditions of visibly increasing algebraic

distance from n, namely n itself, then $n-1$, then $n-2$, which all but announces that each should be resolved before the next is attempted.

Revolving concepts

Power of two divisibility: an integer k divides 2^k exactly when k is a power of 2. This one fact recurs three times in this problem, at n, then at m, then at p; the recursion to flag is that whenever the same structural test must be satisfied by a value and then by the exponent that value produces, the answer comes out as a power tower, with one level of the tower per recurrence of the test. See Book 1, field Number theory, topic Divisors of powers.

The multiplicative order of 2 modulo a Mersenne number: 2^k-1 divides $2^\ell-1$ exactly when k divides ℓ. The orbiting generalisation is $\gcd(2^k-1, 2^\ell-1) = 2^{\gcd(k,\ell)}-1$, proved by the same division algorithm argument, run to produce a greatest common divisor rather than a yes or no answer; replacing 2 by any integer $a > 1$ gives the same equivalence, a^k-1 divides $a^\ell-1$ exactly when k divides ℓ, and this generalises further to the order of an element in any group where that order is finite. See Book 1, field Number theory, topic Multiplicative order.

Converting an astronomical bound into a bound on an exponent by taking a logarithm: not a theorem so much as a standing habit, whenever a bound is given as $n < 10^K$ and n is known to equal b^m for a fixed base b, take $\log_b$ of both sides to get a bound on m small enough to enumerate; the exact value of K almost never matters beyond fixing how many candidates survive. See Book 1, field Real analysis, topic Logarithmic estimates.

Circuit Bank

- **WHEN** several structural conditions are imposed simultaneously on the same unknown **DO** resolve the condition that pins the unknown's shape down most tightly, using the least machinery, before touching any condition that needs arithmetic modulo the unknown itself **BECAUSE** substituting the resolved shape into the remaining conditions turns them from statements about an arbitrary integer into statements about one new parameter, almost always far easier
- **WHEN** a divisibility statement has the shape 2^k-1 divides $2^\ell-1$, or the same with any fixed base in place of 2 **DO** invoke, or re-derive on the spot, the order lemma: this holds exactly when k divides ℓ **BECAUSE** the powers of the base modulo 2^k-1 cycle with period exactly k, by the division algorithm applied to the exponent
- **WHEN** a divisibility statement's two sides share an obvious common factor **DO** cancel it immediately, before analysing the original, unreduced statement **BECAUSE** ca divides cb exactly when a divides b, for nonzero c, and the reduced statement is always at least as easy, usually revealing the same structure one level down
- **WHEN** a bound is given as an enormous power of ten applied to a quantity already known to be a power of a fixed base **DO** take the logarithm in that base to convert the bound into a small, enumerable bound on the exponent **BECAUSE** the exact numerical bound almost never matters beyond determining how many candidates in a short, structured list survive

Backtesting

1. First useful cue: n divides 2^n is a condition on n alone with no shift, visibly more restrictive than the other two.
2. Method the expression suggested: unique factorisation resolves the first condition instantly; the multiplicative order of 2 modulo a Mersenne number resolves the other two once the first has fixed n's shape.

3. Where you would hesitate, and the thought to reinforce: at the third condition, when $n - 2 = 2^m - 2$ does not immediately look like a Mersenne number; the thought to reinforce is factor out the common power of 2 first, then compare shapes.
4. The trap the question sets: the small cases $n = 1$ and $n = 2$ pass the power of two style test for their own condition but fail the next condition, because zero divides only zero; both must be checked and excluded explicitly, not waved away.
5. What makes you recognise it faster next time: noticing that all three conditions can be rewritten uniformly as $(n - j) \mid (2^n - j)$ for $j = 0, 1, 2$, which announces that the same test must recur at nested depths before a single line of casework is written.
6. The reusable decision rule: when a problem imposes the same kind of condition at several shifted arguments, rewrite all of them in one uniform shape first, resolve the most restrictive instance, and substitute forward; expect the answer to come out as a tower, one layer per condition.

Mirror drills

1. Prove the Mersenne order lemma again from scratch, then use it to find all n with $n \mid 2^n$ and $(n-1) \mid (2^n - 1)$ only, dropping the third condition. *Answer:* exactly $n = 2^{2^p}$ for every integer $p \geq 0$, an infinite family, since only the third condition in the original problem cuts this down further.
2. If $n < 10^{50}$ and n is known to equal 3^m for some integer $m \geq 0$, what is the largest possible value of m? *Hint:* take $\log_3$ of 10^{50}; since $\log_3 10$ is a little above 2.09, $m \leq 104$, the same logarithm to exponent translation used to bound m above.
3. Find all positive integers n with $n \mid 2^n$ and $n \mid 3^n$ simultaneously. *Hint:* the first condition forces n to be a power of 2, the second forces n to be a power of 3; only $n = 1$ can be both, the same resolve the most restrictive shape first circuit applied to two conditions that turn out jointly very restrictive.

8.10 B4. Chebyshev disguise turns Fibonacci angles periodic

Given a real number a, we define a sequence by $x_0 = 1$, $x_1 = x_2 = a$, and

$$x_{n+1} = 2x_n x_{n-1} - x_{n-2} \qquad (n \geq 2).$$

Prove that if $x_n = 0$ for some n, then the sequence is periodic.

Read it out loud

Say the recursion aloud before you write anything: two, times a term, times the term before it, minus the term three places back. Not a sum of multiples of earlier terms, so it is not a linear recursion and no characteristic equation is coming. The only nonlinearity in sight is a single product of two consecutive terms, and it arrives with a factor of 2 bolted to the front. There is exactly one place in your installed mathematics where a doubled product of two things equals a sum of two other things of the same species: the product to sum identity

$$2\cos A \cos B = \cos(A + B) + \cos(A - B).$$

Rearranged, that reads $\cos(A + B) = 2\cos A \cos B - \cos(A - B)$, which is the given recursion with $A + B$, A, B, $A - B$ in the four slots. The factor of 2 is not decoration. It is the examiner's fingerprint.

Now check the initial conditions against that reading. If $x_n = \cos\theta_n$, the recursion demands $\theta_{n+1} = \theta_n + \theta_{n-1}$ and $\theta_{n-2} = \theta_n - \theta_{n-1}$, which is the same statement twice. So $\theta_n = F_n b$ for a Fibonacci-type sequence F_n and a single free angle b. Test it: $x_0 = 1$ wants $\theta_0 = 0$, so $F_0 = 0$; $x_1 = x_2 = a$ wants $\theta_1 = \theta_2$, so $F_1 = F_2$. That is precisely the standard Fibonacci sequence $0, 1, 1, 2, 3, 5, \ldots$. The problem has handed you $x_0 = 1$ and the doubled equality $x_1 = x_2$ for no reason other than to make the Fibonacci indexing come out exactly right. Two given values, two confirmations.

The fields wired into this problem: sequences and recurrences (a nonlinear recursion linearised by a change of coordinates), trigonometry and Chebyshev polynomials (the linearising map), and number theory (the Pisano period, that is, the Fibonacci sequence read modulo an integer). The word *periodic* in the conclusion is the number theory cue: real sequences are almost never periodic by accident, so periodicity has to come from a finite state space somewhere, and the only finite object in this problem will be $\mathbb{Z}/d\mathbb{Z}$.

Within thirty seconds you should have written on your paper: $2\cos A \cos B = \cos(A + B) + \cos(A - B)$; the guess $x_n = \cos(F_n b)$ with $a = \cos b$; and the caveat that $a = \cos b$ needs $|a| \leq 1$, so the case $|a| > 1$ must be disposed of separately. That caveat is the only real bookkeeping in the whole problem, and candidates who forget it lose points for an incomplete case analysis rather than for missing the idea.

Approaches, ranked worst to best

1. **Polynomial brute force.** *Cue:* every x_n is a polynomial in a, so compute them. *Circuit:* expand and hunt for a pattern. This is the zombie route. You get $x_3 = 2a^2 - 1$, $x_4 = 4a^3 - 3a$, $x_5 = 16a^5 - 20a^3 + 5a$, and then x_6 has degree 8, x_7 degree 13: the degrees are the Fibonacci numbers, so the ink cost grows exponentially and you never reach a general statement. It ranks last as a plan, but notice that it is not worthless as a diagnostic: those three polynomials are T_2, T_3 and T_5, the Chebyshev polynomials, and recognising them hands you the same substitution the fast route gets for free. Brute force here is the slow road to the correct cue.

2. **Pigeonhole on the state vector.** *Cue:* "periodic" means finitely many states. *Circuit:* state plus finiteness gives eventual periodicity. The recursion is a map on triples, $(x_{n-2}, x_{n-1}, x_n) \mapsto (x_{n-1}, x_n, 2x_n x_{n-1} - x_{n-2})$, so if the triples lived in a finite set you would be done. They do

not; they live in $\mathbb{R}^3$. This route stalls immediately because you have no boundedness and, even with boundedness, no finiteness. It ranks above brute force only because the instinct is correct: periodicity will come from a finite set, and your job is to find which one.

3. **Conserved cubic.** *Cue:* a three term nonlinear recursion of Markov type usually hides an invariant. *Circuit:* look for a symmetric function of consecutive terms that the map preserves. Set $I_n = x_{n-1}^2 + x_n^2 + x_{n+1}^2 - 2x_{n-1}x_nx_{n+1}$. Substituting $x_{n+2} = 2x_{n+1}x_n - x_{n-1}$ gives

$$x_{n+2}^2 - 2x_nx_{n+1}x_{n+2} = x_{n+2}\left(x_{n+2} - 2x_nx_{n+1}\right) = -x_{n-1}x_{n+2} = x_{n-1}^2 - 2x_{n-1}x_nx_{n+1},$$

so $I_{n+1} = I_n$, and $I_1 = 1 + a^2 + a^2 - 2a^2 = 1$. Thus $x_{n-1}^2 + x_n^2 + x_{n+1}^2 - 2x_{n-1}x_nx_{n+1} = 1$ always. This is correct, it is the Cayley cubic surface, and it is the algebraic shadow of the identity $\cos^2 A + \cos^2 B + \cos^2(A+B) - 2\cos A\cos B\cos(A+B) = 1$. It ranks third because it proves the right kind of rigidity but not the conclusion: an invariant confines the orbit to a surface, and a surface is not a finite set. Keep it in your pocket as confirmation that the cosine parametrisation is the truth and not a guess.

4. **Trace coordinates in SL_2.** *Cue:* the conserved cubic is a Fricke trace identity. *Circuit:* model x_n as half the trace of a word in two matrices. Every relation above is the shadow of $\mathrm{tr}(MN) + \mathrm{tr}(MN^{-1}) = \mathrm{tr}(M)\mathrm{tr}(N)$ for $M, N \in SL_2$. Correct, general, and far heavier machinery than a competition answer needs. Ranked fourth: right church, wrong pew.

5. **Cosine parametrisation with Fibonacci angles.** *Cue:* $2\times$ product minus a third term, plus $x_0 = 1$ and $x_1 = x_2$. *Circuit:* product to sum linearises the recursion, turning multiplication of values into addition of angles. Best. It converts the nonlinear recursion into $\theta_{n+1} = \theta_n + \theta_{n-1}$, converts the hypothesis $x_N = 0$ into the statement that b is a rational multiple of 2π, and converts periodicity of (x_n) into periodicity of $(F_n \bmod d)$, which is free because a bijection of a finite set has purely periodic orbits. Every hypothesis in the problem gets used exactly once.

The ideal solution

Step 1: the case $|a| > 1$ never has a zero term. We claim $1 = |x_0| \le |x_1| \le |x_2| \le \cdots$. The first two inequalities hold because $|x_1| = |x_2| = |a| > 1 = |x_0|$. Suppose $n \ge 2$ and $1 = |x_0| \le |x_1| \le \cdots \le |x_n|$. Then, using $|x_{n-2}| \le |x_n|$ and $|x_{n-1}| \ge 1$,

$$|x_{n+1}| = |2x_nx_{n-1} - x_{n-2}| \ge 2|x_n||x_{n-1}| - |x_{n-2}| \ge 2|x_n||x_{n-1}| - |x_n| = |x_n|\left(2|x_{n-1}| - 1\right) \ge |x_n|.$$

By induction $|x_n| \ge 1$ for every n, so $x_n \ne 0$ for every n and the hypothesis of the problem is never satisfied. The implication we must prove is therefore vacuously true when $|a| > 1$, and we may assume $|a| \le 1$ from here on.

Step 2: parametrise. Since $|a| \le 1$ there is $b \in [0, \pi]$ with $a = \cos b$. Let $F_0 = 0$, $F_1 = F_2 = 1$ and $F_{n+1} = F_n + F_{n-1}$ for $n \ge 1$ be the Fibonacci numbers.

Claim 8.10.1. *$x_n = \cos(F_n b)$ for every $n \ge 0$.*

Proof. For $n = 0$, $\cos(F_0 b) = \cos 0 = 1 = x_0$; for $n = 1, 2$, $\cos(F_n b) = \cos b = a = x_n$. Let $m \ge 2$ and suppose the formula holds for all indices up to m. Since $m \ge 2$ we have $F_m - F_{m-1} = F_{m-2}$ and $F_m + F_{m-1} = F_{m+1}$, so

$$2x_mx_{m-1} = 2\cos(F_m b)\cos(F_{m-1}b) = \cos((F_m - F_{m-1})b) + \cos((F_m + F_{m-1})b) = \cos(F_{m-2}b) + \cos(F_{m+1}b).$$

Hence $x_{m+1} = 2x_mx_{m-1} - x_{m-2} = \cos(F_{m+1}b)$, which closes the induction. □

Step 3: turn the hypothesis into rationality. Suppose $x_N = 0$ for some N. Since $x_0 = 1$ we have $N \ge 1$, so $F_N \ge 1$. From $\cos(F_N b) = 0$ we get $F_N b = \frac{(2j+1)\pi}{2}$ for some integer $j \ge 0$, hence

$$b = \frac{c}{d}(2\pi), \qquad c = 2j + 1 \in \mathbb{Z}, \quad d = 4F_N \in \mathbb{Z}_{>0}.$$

This is the only place the hypothesis is used, and it is used to buy one thing: an integer denominator d.

Step 4: the terms depend only on F_n modulo d. If $F_m \equiv F_n \pmod d$ then $\frac{cF_m}{d} - \frac{cF_n}{d} = c \cdot \frac{F_m - F_n}{d}$ is an integer, so

$$x_m = \cos\left(\frac{cF_m}{d}\,2\pi\right) = \cos\left(\frac{cF_n}{d}\,2\pi\right) = x_n.$$

Step 5: Fibonacci modulo d is purely periodic. Work in the finite set $(\mathbb{Z}/d\mathbb{Z})^2$ and let $v_n = (F_n, F_{n+1}) \bmod d$. Then $v_{n+1} = T(v_n)$ where $T(u,v) = (v, u+v)$. The map T is a bijection of $(\mathbb{Z}/d\mathbb{Z})^2$, with inverse $T^{-1}(u,v) = (v-u, u)$. Since $(\mathbb{Z}/d\mathbb{Z})^2$ is finite, the sequence $v_0, v_1, v_2, \ldots$ must repeat a value, say $v_{n_1} = v_{n_2}$ with $n_1 < n_2$; applying T^{-1} exactly n_1 times gives $v_0 = v_\ell$ where $\ell = n_2 - n_1 \geq 1$. Applying T repeatedly then gives $v_{n+\ell} = v_n$ for all $n \geq 0$, so in particular

$$F_{n+\ell} \equiv F_n \pmod d \qquad (n \geq 0).$$

Step 6: conclude. By Steps 4 and 5, $x_{n+\ell} = x_n$ for every $n \geq 0$ with $\ell \geq 1$ fixed. The sequence is periodic. ■

The invertibility of T in Step 5 is not a formality. Without it you would only get *eventual* periodicity, and the problem asks for periodicity of the sequence, meaning from the start. State it.

Layer 1: WHAT

Block 1. Split on $|a| > 1$ against $|a| \leq 1$, because the parametrisation you are about to use only exists in the second case.

Block 2. Kill $|a| > 1$ by proving $|x_n|$ is nondecreasing and starts at 1: the hypothesis $x_n = 0$ is then never met.

Block 3. Write $a = \cos b$ and guess $x_n = \cos(F_n b)$, forced by the shape $2 \cdot (\text{product}) - (\text{third term})$.

Block 4. Verify the guess by induction, using $F_m - F_{m-1} = F_{m-2}$ to make the product to sum identity land on the right indices.

Block 5. Convert $x_N = 0$ into $b/(2\pi) = c/d$ with $d = 4F_N$ an integer.

Block 6. Observe x_n depends only on F_n mod d.

Block 7. Fibonacci modulo d is purely periodic because the shift is a bijection of a finite set. Done.

The nominal route runs Block 1, then expands $x_3, x_4, x_5, \ldots$ as polynomials in a, then tries to solve $x_N(a) = 0$ for a and check periodicity for each root. It dies at the moment the degrees start growing like F_n: a general argument is impossible from there because you would need to control the roots of infinitely many polynomials of unbounded degree.

The bifurcation point is one line: the line where you look at $2x_n x_{n-1} - x_{n-2}$ and decide whether it is an instruction to multiply out or a name for an identity you already own. Everything downstream is determined by that decision.

Layer 2: HOW

Block 1. Cue. The substitution you are reaching for is $a = \cos b$, and cosine has range $[-1, 1]$. **Why natural.** You never start a case split for its own sake. You start it because the tool you have chosen has a domain, and you must check the input lies in it. Naming the excluded case first, before you use the tool, is what stops a grader writing "incomplete" on an otherwise perfect script. **Link to the next block.** Having isolated $|a| > 1$, you now need to kill it, and killing it means showing no term can be zero.

Block 2. **Cue.** You must prove a term is never 0, that is, bounded away from 0. The natural quantity is $|x_n|$, and the recursion has $2x_nx_{n-1}$ in it, a product of two things you are already trying to bound below. **Why natural.** If all previous $|x_k| \geq 1$, then the product term $2|x_n||x_{n-1}|$ is at least twice the size of either factor, and the triangle inequality lets the subtracted $|x_{n-2}|$ eat at most one copy of $|x_n|$. So growth beats subtraction, with the exact bookkeeping $2|x_n||x_{n-1}| - |x_n| = |x_n|(2|x_{n-1}| - 1) \geq |x_n|$. Notice that the induction hypothesis has to be the whole chain $|x_0| \leq \cdots \leq |x_n|$, not just the last inequality, because you use $|x_{n-2}| \leq |x_n|$ and $|x_{n-1}| \geq 1$ at the same time. Strong induction, not ordinary induction. **Link to the next block.** With that case closed, $|a| \leq 1$ is now a standing assumption and the substitution is legal.

Block 3. **Cue.** Doubled product of two consecutive terms, minus one earlier term. Plus $x_0 = 1$ and $x_1 = x_2$. **Why natural.** This is the definition-level reason and it is worth saying exactly: cos is the real part of the character $\theta \mapsto e^{i\theta}$, so multiplying two cosines mixes the sum and the difference of the angles, and $2\cos A \cos B = \cos(A+B) + \cos(A-B)$ is nothing more than $\frac{1}{2}(e^{iA} + e^{-iA})(e^{iB} + e^{-iB})$ expanded. A recursion that says "value at index $n+1$ equals twice the product of the values at n and $n-1$, minus the value at $n-2$" is therefore an instruction that the indices behave additively: $\theta_{n+1} = \theta_n + \theta_{n-1}$ and $\theta_{n-2} = \theta_n - \theta_{n-1}$. Those two conditions are the same condition, which is why the guess is consistent rather than overdetermined. The initial data then pins the angle sequence to $F_n b$. **Link to the next block.** Once the guess exists, verification is a routine induction, and you should do it explicitly because the index arithmetic is where errors hide.

Block 4. **Cue.** An explicit formula proposed for all n, with a three term recursion behind it. **Why natural.** The induction needs $F_m - F_{m-1} = F_{m-2}$, which requires $m \geq 2$, exactly the range in which the recursion is given. That alignment is a sanity check on your indexing: if the given recursion had started at $n \geq 3$ you would have a different Fibonacci offset. Check the three base cases $n = 0, 1, 2$ separately; do not wave at them. **Link to the next block.** Now the sequence is completely understood as a function of one real number b, and the hypothesis can be translated.

Block 5. **Cue.** $x_N = 0$ means $\cos(\text{something}) = 0$. **Why natural.** Zeros of cosine are the half-odd multiples of π. That is a rigidity statement: an equation of the form $\cos\theta = 0$ does not merely constrain θ, it forces θ/π to be a half-odd-integer, hence forces θ to be a rational multiple of π. Rational means there is a denominator, and a denominator is an integer, and an integer is the finiteness you needed in Block 7. This is the whole reason the problem gives you a vanishing term rather than, say, a bounded sequence. **Link to the next block.** The denominator $d = 4F_N$ is the finite modulus. Now push everything into $\mathbb{Z}/d\mathbb{Z}$.

Block 6. **Cue.** $x_n = \cos(2\pi c F_n/d)$ and cosine has period 2π. **Why natural.** The value only depends on cF_n/d modulo 1, hence only on F_n modulo d. In other words, the map "integer $\mapsto$ term" factors through $\mathbb{Z}/d\mathbb{Z}$. Factoring an infinite index set through a finite quotient is the standard way to manufacture periodicity, and you should recognise it as the same manoeuvre as proving that decimal expansions of rationals recur. **Link to the next block.** All that is left is periodicity of Fibonacci modulo d.

Block 7. **Cue.** A sequence in a finite set defined by a fixed one-step rule on pairs. **Why natural.** Two facts, both needed. Finiteness gives a repeat by pigeonhole. Invertibility of the shift $T(u,v) = (v, u+v)$, whose inverse is $(u,v) \mapsto (v-u, u)$, upgrades that repeat from "eventually periodic" to "purely periodic", which is what lets you claim $x_{n+\ell} = x_n$ for all $n \geq 0$ including $n = 0$. The determinant of T as a matrix over $\mathbb{Z}/d\mathbb{Z}$ is -1, a unit, which is the structural reason invertibility holds for every modulus. **Link to the next block.** Combine with Block 6 and the proof closes.

Layer 3: WHY

The mathematical truth under this problem is that the recursion is an addition law written in a bad chart. On the circle, angles add. The function cos is a two-to-one chart from the circle to the interval $[-1, 1]$, and when you push the addition law through that chart it stops looking like addition and starts looking like $2xy - z$. Every "nonlinear recursion with a mysterious pattern" of this family is the same phenomenon: $x_{n+1} = 2x_n x_{n-1} - x_{n-2}$, $x_{n+1} = x_n^2 - 2$, $x_{n+1} = 2x_n x_{n-1} - x_{n-1}$ and their cousins are all the group law $\theta \mapsto \theta + \varphi$ in disguise. So the correct instinct on seeing such a recursion is not "find a pattern" but "find the chart".

Which chart? The functions f satisfying $f(A+B) + f(A-B) = 2f(A)f(B)$ for all real A, B are exactly $f = \cos(\lambda\cdot)$, $f = \cosh(\lambda\cdot)$, the constant $f \equiv 1$ and the degenerate $f \equiv 0$. That is d'Alembert's functional equation, and it explains both halves of the case analysis at once: for $|a| \le 1$ the chart is cos, and for $|a| > 1$ it is cosh, since $2\cosh A \cosh B = \cosh(A+B) + \cosh(A-B)$ is the same identity. Indeed, with $a = \varepsilon \cosh t$ and $\varepsilon = \pm 1$ you can check that $x_n = \varepsilon^{F_n} \cosh(F_n t)$, so $|x_n| = \cosh(F_n t) \ge 1$. That is the conceptual version of Step 1, and it is worth knowing even though the three line induction is faster to write in the arena. The lesson is that the case split is not an annoyance bolted onto the problem; it is the two branches of the same parametrisation, separated by whether the group is compact.

The discernment mechanism, stated so you can train it: count the nonlinearity and count the degrees of freedom. Here the nonlinearity is a single bilinear term with a coefficient of exactly 2, and there is exactly one free parameter a in the whole problem. One parameter means the answer is a one-parameter family, and one-parameter families that satisfy quadratic addition rules are trigonometric or hyperbolic. If the coefficient had been 3 instead of 2, no identity would fit and you would be in a different problem entirely. So read coefficients as fingerprints, not as noise.

The examiners nudged in three places, and you should be able to name all three after the fact. First, $x_0 = 1$, which is $\cos 0$ and pins $F_0 = 0$. Second, $x_1 = x_2$, which pins $F_1 = F_2$ and is otherwise a strange thing to impose. Third, the conclusion says *periodic*, not *bounded* or *eventually periodic*, and a real sequence is periodic only when it factors through something finite, which tells you in advance that an integer modulus is coming. Three nudges, one method.

Training protocol. Take ten nonlinear recursions with small integer coefficients, compute four terms of each as polynomials in the seed, and try to identify the polynomial family: Chebyshev T_n, Chebyshev U_n, Fibonacci-like, or none. You are not training the algebra, you are training the recognition that $2a^2 - 1$ and $4a^3 - 3a$ are not random cubics but T_2 and T_3. Recognition of small named polynomials is a memory task, and memory tasks are the cheap part of the Putnam.

Revolving concepts

Chebyshev polynomials. T_n is defined by $T_n(\cos\theta) = \cos(n\theta)$, equivalently by $T_0 = 1$, $T_1 = x$, $T_{n+1} = 2xT_n - T_{n-1}$. The orbiting identities are $T_m T_n = \frac{1}{2}(T_{m+n} + T_{|m-n|})$ and $T_m \circ T_n = T_{mn}$. In this problem $x_n = T_{F_n}(a)$, which is why the computed terms were T_2, T_3, T_5. Learn the first five by sight: 1, x, $2x^2 - 1$, $4x^3 - 3x$, $8x^4 - 8x^2 + 1$, $16x^5 - 20x^3 + 5x$. See Book 1, field: Polynomials, topic: Chebyshev families.

d'Alembert's functional equation. The equation $f(A+B) + f(A-B) = 2f(A)f(B)$ with f continuous and not identically zero has solutions $\cos(\lambda x)$, $\cosh(\lambda x)$ and 1. This is the theorem that turns "I guessed cosine" into "cosine is forced". See Book 1, field: Functional equations, topic: cosine functional equation.

The generalised recursion, and how to solve any member of the family. Suppose a sequence satisfies $x_{n+1} = 2x_n x_{n-1} - x_{n-2}$ with $x_0 = 1$. Then $x_n = f(\theta_n)$ where f solves d'Alembert and θ satisfies the *linear* recursion $\theta_{n+1} = \theta_n + \theta_{n-1}$. That is the general principle and it is worth stating in full generality, because it applies to every recursion in this family:

(1) Detect the shape $x_{n+1} = 2x_n x_{n-k} - x_{n-2k}$ or similar, in which the four index positions satisfy

$p + q = r + s$ in the sense that the outer pair sums to the inner pair.

(2) Substitute $x_n = \cos\theta_n$; the recursion becomes the linear recursion on angles obtained by reading the index arithmetic.

(3) Solve the linear recursion by its characteristic equation. For $\theta_{n+1} = \theta_n + \theta_{n-1}$ the solution is $\theta_n = F_n\theta_1 + F_{n-1}\theta_0$, and with $\theta_0 = 0$ this is $\theta_n = F_n b$. For $\theta_{n+1} = 2\theta_n$ (which is the recursion $x_{n+1} = 2x_n^2 - 1$) it is $\theta_n = 2^n b$.

(4) A term vanishes if and only if the corresponding angle is an odd multiple of $\pi/2$, which forces b to be a rational multiple of 2π with a computable denominator d.

(5) Periodicity of (x_n) then reduces to periodicity of the integer sequence (θ_n/b) modulo d. For a linear recursion with integer coefficients and invertible transfer matrix over $\mathbb{Z}/d\mathbb{Z}$, that periodicity is automatic and pure. For Fibonacci the transfer matrix is $\begin{pmatrix} 0 & 1 \\ 1 & 1 \end{pmatrix}$ with determinant -1, a unit modulo every d.

Pisano period. The Fibonacci sequence modulo m is purely periodic; its period $\pi(m)$ is the Pisano period. Facts worth carrying: $\pi(2) = 3$, $\pi(4) = 6$, $\pi(8) = 12$, $\pi(10) = 60$; $\pi(\mathrm{lcm}(m, n)) = \mathrm{lcm}(\pi(m), \pi(n))$ for coprime m, n; and $\pi(m) \leq 6m$ always. You never need the exact value on a Putnam, only that it exists and is finite. See Book 1, field: Number theory, topic: linear recurrences modulo m.

The conserved cubic. $x_{n-1}^2 + x_n^2 + x_{n+1}^2 - 2x_{n-1}x_n x_{n+1} = 1$ for this sequence. Recognise the surface $x^2 + y^2 + z^2 - 2xyz = 1$: it is the Cayley cubic, the character variety of the free group on two generators mapped to SL_2, and the identity $\cos^2 A + \cos^2 B + \cos^2(A + B) - 2\cos A \cos B \cos(A + B) = 1$ is its trigonometric parametrisation. The Markov equation $x^2 + y^2 + z^2 = 3xyz$ is its close relative and the "Vieta jumping" technique lives on the same idea: a quadratic in one variable with the other two fixed has two roots, and the map swapping them is an involution on solutions. See Book 1, field: Number theory, topic: Vieta jumping and Markov triples.

Pure versus eventual periodicity. A sequence produced by iterating a map T on a finite set is always eventually periodic; it is purely periodic if and only if the starting state lies in a cycle, which is guaranteed when T is a bijection. This distinction is worth points on the Putnam more often than it looks. See Book 1, field: Discrete structures and graphs, topic: functional graphs, rho shape and cycles.

Circuit Bank

- **WHEN** a recursion containing $2\,x_n x_{n-1}$ minus an earlier term **DO** substitute $x_n = \cos\theta_n$ and read off the linear recursion on the angles **BECAUSE** $2\cos A \cos B = \cos(A + B) + \cos(A - B)$ is the only identity in which a doubled product of two values equals a sum of two values of the same function, so the coefficient 2 identifies the chart
- **WHEN** the substitution you want has a restricted domain, such as cos needing $|a| \leq 1$ **DO** open with the excluded case and dispose of it before using the substitution **BECAUSE** a solution that silently assumes its parametrisation exists is incomplete, and the excluded case is usually the easy one: here $|a| > 1$ gives a growth induction in three lines
- **WHEN** a hypothesis of the form "some term equals zero" applied to a trigonometric parametrisation **DO** convert it into "the angle is a rational multiple of π" and extract the integer denominator **BECAUSE** zeros of cos are rigid, at odd multiples of $\pi/2$, so vanishing is a rationality statement, and rationality is where the finite modulus comes from
- **WHEN** you need periodicity of a real sequence **DO** find a finite set the sequence factors through, then use pigeonhole on a bijection **BECAUSE** real sequences are never periodic by accident; the periodicity must be inherited from a finite quotient such as $\mathbb{Z}/d\mathbb{Z}$
- **WHEN** pigeonhole gives you a repeated state **DO** check whether the step map is invertible

before claiming pure periodicity **BECAUSE** a noninvertible map gives a rho shaped orbit with a tail, so you would only earn eventual periodicity, and the question asked for periodicity from the start

Backtesting

1. First useful cue: the coefficient 2 in front of the product $x_n x_{n-1}$, reinforced by $x_0 = 1$ and $x_1 = x_2$. Those three data points say cosine, Fibonacci indexing, and nothing else.
2. Method the expression suggested: linearise by parametrisation. Replace values by angles so that a quadratic recursion becomes an additive one, then do all the work on the angles.
3. Where you would hesitate, and the thought to reinforce: at the end, after proving $x_n = \cos(F_n b)$, it is tempting to say "and clearly the sequence repeats". It does not clearly repeat; b is a general real number and $(F_n b \bmod 2\pi)$ is equidistributed for almost every b. The hypothesis $x_N = 0$ is what forces b into the rational multiples of 2π. Reinforce the thought: *use every hypothesis, and if a hypothesis has not been used yet, that is where the remaining difficulty lives*.
4. The trap the question sets: two traps. Forgetting that $a = \cos b$ requires $|a| \leq 1$, which costs a case; and claiming eventual periodicity when the question asked for periodicity, which costs the invertibility sentence in Step 5.
5. What makes you recognise it faster next time: memorise the first six Chebyshev polynomials by sight. If $x_3 = 2a^2 - 1$ appears on your page and you do not immediately think $\cos 2b$, you have already lost ten minutes.
6. The reusable decision rule: when a nonlinear recursion has a doubled product and small integer coefficients, look for a one-parameter group whose addition law it encodes, and do the whole problem on the parameter, not on the values.

Mirror drills

1. Let $y_0 = a$ and $y_{n+1} = y_n^2 - 2$. Find a closed form and describe when the sequence is periodic. *Answer:* write $a = t + t^{-1}$ (allowing $|t| = 1$ exactly when $|a| \leq 2$); then $y_n = t^{2^n} + t^{-2^n}$. With $t = e^{i\theta}$, $y_n = 2\cos(2^n\theta)$, periodic exactly when $\theta/2\pi$ is rational with odd denominator times a power of two handled by the doubling map.
2. Prove that $x_{n-1}^2 + x_n^2 + x_{n+1}^2 - 2x_{n-1}x_n x_{n+1}$ is constant for the sequence of this problem, without using the cosine formula. *Answer:* substitute $x_{n+2} = 2x_{n+1}x_n - x_{n-1}$ into the expression at index $n+1$ and factor $x_{n+2}(x_{n+2} - 2x_n x_{n+1}) = -x_{n-1}x_{n+2}$; the expression collapses to its value at index n.
3. Show that the Fibonacci sequence modulo m is purely periodic for every $m \geq 1$, and compute the period for $m = 8$. *Answer:* the pair map $(u, v) \mapsto (v, u + v)$ has matrix of determinant -1, hence is invertible modulo m, so orbits are pure cycles. Modulo 8 the sequence runs $0, 1, 1, 2, 3, 5, 0, 5, 5, 2, 7, 1$ and then repeats, so the period is 12.

8.11 B5. Symmetrised Jacobian forces global injectivity

Let $f = (f_1, f_2)$ be a function from $\mathbb{R}^2$ to $\mathbb{R}^2$ with continuous partial derivatives $\frac{\partial f_i}{\partial x_j}$ that are positive everywhere. Suppose that

$$\frac{\partial f_1}{\partial x_1}\frac{\partial f_2}{\partial x_2} - \frac{1}{4}\left(\frac{\partial f_1}{\partial x_2} + \frac{\partial f_2}{\partial x_1}\right)^2 > 0$$

everywhere. Prove that f is one-to-one.

Read it out loud

A map from the plane to the plane, continuously differentiable, all four partial derivatives positive, and one extra inequality. Prove it is injective. Injectivity is a global statement about two arbitrary points, so before anything else write down what you are actually going to prove: for all $p \neq q$ in $\mathbb{R}^2$, $f(p) \neq f(q)$. Write the two points down as symbols now, because the entire solution is going to be built on the segment between them.

Now read the extra inequality out loud, slowly. It says: (entry 11) times (entry 22), minus the square of the *average* of entries 12 and 21. Abbreviate $f_{ij} = \partial f_i/\partial x_j$ and put the Jacobian on the page:

$$Df = \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix}, \qquad S := \tfrac{1}{2}\left(Df + (Df)^{\mathsf{T}}\right) = \begin{pmatrix} f_{11} & \frac{f_{12}+f_{21}}{2} \\ \frac{f_{12}+f_{21}}{2} & f_{22} \end{pmatrix}.$$

The hypothesis is exactly $\det S > 0$. It is not $\det Df > 0$, which would be $f_{11}f_{22} - f_{12}f_{21}$, and the difference between those two expressions is the whole problem. The examiners went out of their way to write $\frac{1}{4}(f_{12} + f_{21})^2$ rather than $f_{12}f_{21}$, which is a strictly larger quantity by $\frac{1}{4}(f_{12} - f_{21})^2 \geq 0$, so the hypothesis you are given is strictly stronger than positivity of the Jacobian determinant. When an examiner hands you a hypothesis that is deliberately stronger than the obvious one, the extra strength is load bearing, and your job is to find the argument that uses it.

Together with $f_{11} > 0$, the condition $\det S > 0$ says that S is positive definite (Sylvester's criterion: leading principal minors positive). Positive definiteness of the symmetric part of Df is a named condition: it says f is a *strictly monotone* map, the multivariable analogue of "has positive derivative". In one variable, positive derivative gives strictly increasing gives injective, and the middle step is the mean value theorem. So the shape of the solution is fixed before you have done any work: manufacture a function of one real variable whose derivative is $v^{\mathsf{T}}Sv$, and apply one-variable calculus to it.

Fields wired in: multivariable differential calculus (chain rule along a path, and the fact that continuous partials imply total differentiability), quadratic forms and definiteness from linear algebra, and the theory of monotone operators, which is the correct general home for this statement. What is really being tested is whether you know that local information does not become global information for free. Within thirty seconds you should have written Df, its symmetric part S, the remark that the hypothesis is $\det S > 0$ with $f_{11} > 0$ hence S positive definite, and the two arbitrary points p and q.

Approaches, ranked worst to best

1. **Coordinates and elimination.** *Cue:* $f(p) = f(q)$ is two equations, so try to solve them. *Circuit:* set up simultaneous equations and eliminate. There are no formulas for f_1 and f_2; only their derivatives are constrained. Nothing can be eliminated. Ranked last because it does not even start.

2. **Inverse function theorem.** *Cue:* positive partials, positive determinant, smooth map. *Circuit:* Jacobian nonsingular gives a local diffeomorphism. This one is dangerous, because it is true

and it feels like a solution. Since $\frac{1}{4}(f_{12}+f_{21})^2 \geq f_{12}f_{21}$ by the arithmetic mean and geometric mean inequality applied to the two positive numbers f_{12} and f_{21}, the hypothesis gives

$$\det Df = f_{11}f_{22} - f_{12}f_{21} \geq f_{11}f_{22} - \tfrac{1}{4}(f_{12}+f_{21})^2 > 0,$$

so f is a local diffeomorphism at every point. That establishes *local* injectivity and stops. It does not imply global injectivity, and the standard counterexample is one you must have loaded: $F(x,y) = (e^x \cos y, e^x \sin y)$, the complex exponential, has $\det DF = e^{2x} > 0$ everywhere and satisfies $F(x, y+2\pi) = F(x,y)$. A script that says "the Jacobian is nonzero, hence f is injective" scores zero. Ranked second precisely so that you can see the trap and step over it.

3. **Monotonicity in each variable separately.** *Cue:* all four partials are positive, so each f_i is increasing in each argument. *Circuit:* use the one-variable mean value theorem coordinate by coordinate. Suppose $f(p) = f(q)$ with $v = q - p \neq 0$. If $v_1 \geq 0$ and $v_2 \geq 0$, then moving from p to q increases x_1 weakly, increases x_2 weakly, and at least one strictly, so $f_1(q) > f_1(p)$: contradiction. The same argument with signs reversed handles $v_1 \leq 0$, $v_2 \leq 0$. So the only surviving case is v_1 and v_2 strictly of opposite sign. This is a genuine reduction and it earns partial credit, and more importantly it tells you what the strange hypothesis is for: it exists to handle the mixed directions, where the increase in one coordinate fights the decrease in the other. Ranked third because it is correct but incomplete, and the completion turns out to subsume it.

4. **Convexity via a potential.** *Cue:* positive definite symmetric matrix of second derivatives is a Hessian condition. *Circuit:* if $f = \nabla\varphi$ with φ strictly convex then f is injective. If Df happened to be symmetric, then f would be the gradient of some φ with positive definite Hessian, so φ would be strictly convex and $\nabla\varphi$ injective. Here Df need not be symmetric, so this does not apply directly; but it is the right picture, and the proof of "gradient of a strictly convex function is injective" is exactly the proof we want, with the potential deleted. Ranked fourth: correct instinct, needs an unavailable hypothesis.

5. **Strict monotonicity along the joining segment.** *Cue:* a global conclusion about two arbitrary points, plus a hypothesis that is a quadratic form. *Circuit:* pair the difference of images with the difference of points and differentiate along the straight path. Best. Define $g(t) = v \cdot f(p + tv)$ with $v = q - p$. Then $g'(t) = v^{\mathsf{T}} Df\, v = v^{\mathsf{T}} S v > 0$, so $g(1) > g(0)$, so $v \cdot f(q) \neq v \cdot f(p)$, so $f(q) \neq f(p)$. Every hypothesis is used, nothing is wasted, and the argument is five lines. It also generalises verbatim to $\mathbb{R}^n$.

The ideal solution

Throughout write $f_{ij} = \partial f_i / \partial x_j$, and for a point $p \in \mathbb{R}^2$ and a vector $v = (v_1, v_2)$ set

$$Q_p(v) = f_{11}(p)\, v_1^2 + \big(f_{12}(p) + f_{21}(p)\big)\, v_1 v_2 + f_{22}(p)\, v_2^2.$$

Lemma 8.11.1. *For every $p \in \mathbb{R}^2$ and every $v \neq 0$ we have $Q_p(v) > 0$.*

Proof. Fix p and abbreviate $A = f_{11}(p)$, $B = f_{12}(p) + f_{21}(p)$, $C = f_{22}(p)$. By hypothesis $A > 0$ and $\Delta := AC - \frac{1}{4}B^2 > 0$. Multiplying by the positive number A and completing the square,

$$A\, Q_p(v) = A^2 v_1^2 + ABv_1v_2 + ACv_2^2 = \big(Av_1 + \tfrac{1}{2}Bv_2\big)^2 + \big(AC - \tfrac{1}{4}B^2\big)\, v_2^2 = \big(Av_1 + \tfrac{1}{2}Bv_2\big)^2 + \Delta\, v_2^2.$$

Both terms on the right are nonnegative, so $A\, Q_p(v) \geq 0$. If $A\, Q_p(v) = 0$ then $\Delta v_2^2 = 0$, and since $\Delta > 0$ this forces $v_2 = 0$; the first term then reads $A^2 v_1^2 = 0$, and since $A > 0$ this forces $v_1 = 0$. Hence $A\, Q_p(v) = 0$ only for $v = 0$, so $A\, Q_p(v) > 0$ for $v \neq 0$, and dividing by $A > 0$ gives $Q_p(v) > 0$. □

Now let $p, q \in \mathbb{R}^2$ with $p \neq q$, put $v = q - p \neq 0$, and let

$$\gamma(t) = p + tv \quad (t \in [0,1]), \qquad g(t) = v \cdot f(\gamma(t)) = v_1 f_1(\gamma(t)) + v_2 f_2(\gamma(t)).$$

Each f_i has partial derivatives that exist and are continuous on all of $\mathbb{R}^2$, hence each f_i is (totally) differentiable, and the chain rule applies along the path γ, whose derivative is the constant vector v. Therefore g is differentiable on $[0,1]$ with

$$g'(t) = v_1 \left(\nabla f_1(\gamma(t)) \cdot v\right) + v_2 \left(\nabla f_2(\gamma(t)) \cdot v\right) = v^{\mathsf{T}} Df(\gamma(t))\, v = Q_{\gamma(t)}(v),$$

where the middle equality is just the expansion $v^{\mathsf{T}} Df\, v = f_{11}v_1^2 + f_{12}v_1v_2 + f_{21}v_2v_1 + f_{22}v_2^2$. By the Lemma applied at the point $\gamma(t)$ with the fixed nonzero vector v, we get $g'(t) > 0$ for every $t \in [0,1]$.

Hence g is strictly increasing on $[0,1]$: for $0 \leq s < t \leq 1$ the mean value theorem gives $g(t) - g(s) = g'(\xi)(t-s) > 0$. In particular $g(1) > g(0)$, that is,

$$v \cdot f(q) > v \cdot f(p), \qquad \text{equivalently} \qquad (q-p) \cdot (f(q) - f(p)) > 0.$$

If $f(q)$ were equal to $f(p)$ the left side would be 0, a contradiction. Therefore $f(p) \neq f(q)$ whenever $p \neq q$, which is to say f is one-to-one. ■

Two remarks that a full-marks script may state in one line each. First, only the symmetric part of Df ever appears, since $v^{\mathsf{T}} Df\, v = v^{\mathsf{T}} \left(\frac{1}{2}(Df + (Df)^{\mathsf{T}})\right) v$; this is why the hypothesis was written with $\frac{1}{2}(f_{12} + f_{21})$ rather than with f_{12} and f_{21} separately. Second, the positivity of f_{12} and f_{21} was never used, and the positivity of f_{11} and f_{22} is nearly free: the hypothesis forces $f_{11}f_{22} > \frac{1}{4}(f_{12} + f_{21})^2 \geq 0$, so f_{11} and f_{22} are nonzero with a common sign, which by continuity and connectedness of $\mathbb{R}^2$ is constant; if that sign were negative you would apply the argument to $-f$, which is injective if and only if f is.

Layer 1: WHAT

Block 1. Name the target: take arbitrary $p \neq q$ and aim to prove $f(p) \neq f(q)$.

Block 2. Read the hypothesis as positive definiteness of the symmetrised Jacobian $S = \frac{1}{2}(Df + (Df)^{\mathsf{T}})$.

Block 3. Prove the pointwise Lemma: $v^{\mathsf{T}} S v > 0$ for $v \neq 0$, by multiplying by f_{11} and completing the square.

Block 4. Build the scalar function $g(t) = v \cdot f(p + tv)$ on the segment from p to q, with $v = q - p$.

Block 5. Differentiate: $g'(t) = v^{\mathsf{T}} Df(\gamma(t)) v = Q_{\gamma(t)}(v) > 0$.

Block 6. Conclude $g(1) > g(0)$, that is $(q-p) \cdot (f(q) - f(p)) > 0$, so the images differ.

The nominal route stops at Block 2 with the words "so the Jacobian determinant is positive and f is locally invertible", then attempts to patch local invertibility into global invertibility by a connectedness or covering space argument that either fails or quietly assumes what it is proving. The bifurcation point is the sentence after you compute $\det Df > 0$. At that line you either reach for the inverse function theorem, which is a local statement and cannot produce a global conclusion, or you reach for the segment from p to q, which is a global object and can. The physical tell that the second choice is the right one: the inverse function theorem does not use the difference between $\frac{1}{4}(f_{12} + f_{21})^2$ and $f_{12}f_{21}$, and an unused hypothesis means a wrong route.

Layer 2: HOW

Block 1. Cue. The word "one-to-one". **Why natural.** Injectivity is by definition a statement about a *pair* of points. Writing p and q on the page immediately, before touching the hypothesis, turns an abstract property into two concrete objects you can join, and the only canonical thing you can do with two points of $\mathbb{R}^2$ is join them by a segment. This is not a guess; it is the only construction

available. **Link to the next block.** Now you need a hypothesis that says something about the segment, so look at what the hypothesis controls: derivatives, that is, infinitesimal behaviour along directions.

Block 2. Cue. The expression $\frac{1}{4}(f_{12}+f_{21})^2$ with an average inside a square. **Why natural.** A quadratic form $Av_1^2+Bv_1v_2+Cv_2^2$ is positive definite exactly when $A>0$ and $4AC-B^2>0$, and $4AC-B^2>0$ is $AC-\frac{1}{4}B^2>0$. The problem has literally printed the definiteness criterion for the form with $A=f_{11}$, $B=f_{12}+f_{21}$, $C=f_{22}$. The reason the cross term appears as a sum rather than as a product is that a quadratic form only sees the symmetric part of a matrix, since $v^{\mathsf{T}}Mv$ is a scalar and equals its own transpose $v^{\mathsf{T}}M^{\mathsf{T}}v$. So the moment you see $f_{12}+f_{21}$ you should be looking for the expression $v^{\mathsf{T}}Df\,v$ somewhere in your future. **Link to the next block.** The hypothesis is about $v^{\mathsf{T}}Df\,v$, so the solution must produce that expression; the only way a Jacobian gets squeezed between two copies of the same vector is a directional derivative in the direction v, dotted again with v.

Block 3. Cue. A quadratic in two variables that must be shown positive. **Why natural.** Completing the square is the only method that proves positivity of a quadratic form with an explicit certificate rather than an appeal. Multiply by A rather than dividing by it, because that keeps everything polynomial and makes the equality case transparent: the identity $A\,Q=(Av_1+\frac{1}{2}Bv_2)^2+\Delta v_2^2$ can be checked by eye. The equality analysis is where the strictness comes from, so do not skip it: you need $Q>0$, not $Q\geq 0$. **Link to the next block.** With pointwise positivity in hand, you can afford to let the point move along a path, since the bound holds at every point.

Block 4. Cue. You have two points, a hypothesis about derivatives, and a need for a global conclusion. **Why natural.** The mean value theorem does not exist for vector valued functions: there is no c with $f(q)-f(p)=Df(c)(q-p)$ in general, because the two components would need the same c. The repair is universal and worth installing permanently: pair with a fixed vector to make a scalar function of one real variable, then use one-variable calculus. Which fixed vector? Pair with $v=q-p$ itself, because that is the only vector in the problem and because it makes g' equal to the quadratic form $v^{\mathsf{T}}Df\,v$ that the hypothesis controls. Pairing with any other vector w would give $w^{\mathsf{T}}Df\,v$, a bilinear expression with no sign information. **Link to the next block.** Now differentiate, carefully.

Block 5. Cue. A composition of a C^1 map with a straight line. **Why natural.** Continuity of the partial derivatives is stated in the problem for a reason: existence of partial derivatives alone does not imply differentiability, and the chain rule needs differentiability, not just partials. So the phrase "with continuous partial derivatives" in the statement is not padding, it is the licence for this step, and a full-marks script says so in half a sentence. **Link to the next block.** Positivity of g' on $[0,1]$ gives strict monotonicity of g by the one-variable mean value theorem.

Block 6. Cue. $g(1)>g(0)$. **Why natural.** Unwind the definition: $g(1)-g(0)=v\cdot(f(q)-f(p))=(q-p)\cdot(f(q)-f(p))$. Equality of the images would make this zero. So injectivity drops out, and notice you have proved something stronger than asked, namely strict monotonicity of f as an operator. Stating the stronger conclusion costs nothing and shows the grader you know what you proved. **Link to the next block.** Done.

Layer 3: WHY

The mathematical truth is a single sentence: *a strictly monotone map is injective*, where strictly monotone means $\langle f(q)-f(p),\,q-p\rangle>0$ for all $p\neq q$. Injectivity follows because the inner

product of a nonzero conclusion with anything cannot vanish; if $f(q) = f(p)$ then the inner product is 0, contradiction. All the analysis in this problem exists to convert a pointwise hypothesis about Df into that global inequality, and the conversion is the fundamental theorem of calculus along the segment:

$$(q-p)\cdot(f(q)-f(p)) = \int_0^1 (q-p)^{\mathsf{T}}\, Df(p+t(q-p))\,(q-p)\,dt.$$

Read that identity as the correct multivariable replacement for "$f(b)-f(a) = f'(c)(b-a)$". The mean value theorem in its equality form dies in several variables; its integral form survives, and everything you want to do with derivatives across a gap must go through the integral form or through a scalar reduction like the one used above.

Why is positivity of the symmetric part the right condition, rather than positivity of $\det Df$? Because $\det Df > 0$ is a statement about the map's local behaviour up to orientation, and orientation preserving local diffeomorphisms can wrap the plane around itself infinitely often, as the complex exponential does. Positive definiteness of S is a statement about *every direction at once*: it says that moving in any direction v increases the component of f in that same direction v. That is a rigidity you cannot wrap. Concretely, the complex exponential fails the test: for $F(x,y) = (e^x\cos y, e^x\sin y)$ the symmetric part of DF at $(0,\pi)$ has $F_{11} = e^x\cos y = -1 < 0$, so S is not positive definite there. Every counterexample to the naive route is killed by the correct hypothesis, which is exactly what you expect when the hypothesis is the right one.

The discernment mechanism, stated as a rule you can run in the arena: match the scope of the hypothesis to the scope of the conclusion. A pointwise hypothesis plus a pointwise conclusion is a computation. A pointwise hypothesis plus a global conclusion always requires an integration or a connectedness argument, and in practice that means constructing a path and studying a function of one real variable on it. Ask yourself: what is the one-variable function here? If you cannot name it, you are not yet solving the problem. In this problem it is $g(t) = v\cdot f(p+tv)$, and once you have written it the rest is bookkeeping.

There is a second discernment rule, cheaper and just as reliable: *account for every hypothesis*. The problem gave you a strictly stronger condition than $\det Df > 0$. If your route only uses $\det Df > 0$, your route cannot be the intended one, and on a Putnam the intended one is nearly always the only complete one. This is why you should check, at the moment you commit to a method, whether any given condition is going unused. Here the unused-hypothesis alarm fires exactly at the bifurcation point.

How to train it. Take five theorems of the form "local property implies global property" and, for each, find where the extra hypothesis enters: a continuous locally constant function on a connected space is constant (connectedness), a local diffeomorphism that is proper is a covering map (properness, this is Hadamard's global inverse theorem), a function with zero derivative on an interval is constant (the interval), a locally injective map need not be injective (no extra hypothesis, no theorem). Building that table once makes the trap in this problem visible on sight rather than after ten minutes.

Cultural note worth carrying: the question of when $\det Df$ being nonvanishing forces global injectivity is deep. Over the complex numbers with polynomial f, the statement that $\det Df$ a nonzero constant implies f is invertible is the Jacobian conjecture, open since 1939. So when you feel that "the Jacobian is positive, surely it is injective" ought to be provable, you are feeling the pull of an open problem. Trust the counterexample instead.

Revolving concepts

Quadratic forms and definiteness. For a symmetric S, the form $Q(v) = v^{\mathsf{T}}Sv$ is positive definite when $Q(v) > 0$ for all $v \neq 0$. Sylvester's criterion: positive definite exactly when every leading principal minor is positive, which in the two by two case is $S_{11} > 0$ and $\det S > 0$. Equivalently for

$Av_1^2 + Bv_1v_2 + Cv_2^2$: $A > 0$ and $B^2 - 4AC < 0$, the familiar discriminant condition. See Book 1, field: Linear algebra, topic: quadratic forms and Sylvester's criterion.

Symmetric and skew parts. Any square M splits as $\frac{1}{2}(M + M^\mathsf{T}) + \frac{1}{2}(M - M^\mathsf{T})$, and $v^\mathsf{T} M v = v^\mathsf{T} \left(\frac{M+M^\mathsf{T}}{2}\right) v$ because the skew part contributes $v^\mathsf{T} K v = -v^\mathsf{T} K v = 0$. This is the single fact that explains the shape of the problem's hypothesis. See Book 1, field: Matrices and determinants, topic: symmetric and skew symmetric decomposition.

Continuous partials imply differentiability. If all partial derivatives of $F : \mathbb{R}^n \to \mathbb{R}$ exist in a neighbourhood of p and are continuous at p, then F is differentiable at p. The converse of the naive statement is false: $F(x, y) = xy/(x^2 + y^2)$ extended by 0 has both partials at the origin but is not even continuous there. The problem's phrase "with continuous partial derivatives" is what licenses the chain rule. See Book 1, field: Real analysis, topic: total differentiability versus partial derivatives.

Mean value theorem, vector valued failure and its repairs. For $F : \mathbb{R} \to \mathbb{R}^m$ with $m \geq 2$ there is in general no c with $F(b) - F(a) = F'(c)(b - a)$; take $F(t) = (\cos t, \sin t)$ on $[0, 2\pi]$. The repairs are (i) the integral form $F(b) - F(a) = \int_a^b F'$, (ii) the inequality $\|F(b) - F(a)\| \leq (b - a) \sup \|F'\|$, and (iii) pairing with a fixed vector to return to the scalar case, which is what this solution does. See Book 1, field: Real analysis, topic: mean value theorems in several variables.

Monotone operators. A map $f : \mathbb{R}^n \to \mathbb{R}^n$ is monotone if $\langle f(q) - f(p), q - p \rangle \geq 0$ for all p, q, and strictly monotone if the inequality is strict for $p \neq q$. Strictly monotone implies injective. If f is C^1, then f is monotone if and only if the symmetric part of Df is positive semidefinite everywhere, and positive definiteness everywhere gives strict monotonicity by the integral identity above. The gradient of a convex function is monotone, and this is the special case where Df is symmetric. See Book 1, field: Real analysis, topic: monotone operators and convexity.

Global inverse theorems. The inverse function theorem is local. Hadamard's global inverse theorem says a local diffeomorphism $\mathbb{R}^n \to \mathbb{R}^n$ that is proper (preimages of compact sets are compact) is a global diffeomorphism. The condition in this problem gives properness as well, since $\langle f(q) - f(p), q - p \rangle > 0$ with a definiteness bound forces $\|f(q)\| \to \infty$, so in fact f here is a bijection onto $\mathbb{R}^2$; the problem only asks for the injective half. See Book 1, field: Real analysis, topic: local versus global inverse theorems.

Generalisation to n dimensions. The identical proof shows: if $f : \mathbb{R}^n \to \mathbb{R}^n$ is C^1 and $Df(x) + Df(x)^\mathsf{T}$ is positive definite for every x, then f is injective. There is no recursion or induction anywhere in this problem; the generalisation is in the dimension, not in an index.

Circuit Bank

- **WHEN** a hypothesis containing $\partial f_1/\partial x_2 + \partial f_2/\partial x_1$ as a symmetric sum **DO** write the symmetrised Jacobian $\frac{1}{2}(Df + (Df)^\mathsf{T})$ and test it for positive definiteness **BECAUSE** a quadratic form $v^\mathsf{T} M v$ sees only the symmetric part of M, so a hypothesis that symmetrises the off diagonal entries is announcing that a quadratic form in a single direction v is coming
- **WHEN** you must prove a smooth map on all of $\mathbb{R}^n$ is globally injective **DO** take arbitrary $p \neq q$, set $v = q - p$, and study $g(t) = v \cdot f(p + tv)$ on $[0, 1]$ **BECAUSE** injectivity is a statement about pairs, the segment is the only canonical object joining a pair, and pairing with v converts the vector problem into a one-variable problem where the mean value theorem is legal
- **WHEN** the inverse function theorem seems to finish an injectivity problem **DO** stop, and test your argument against $(x, y) \mapsto (e^x \cos y, e^x \sin y)$ **BECAUSE** that map has positive Jacobian determinant everywhere and is periodic, so any argument deriving global injectivity from a nonvanishing Jacobian alone is wrong
- **WHEN** a quadratic $Av_1^2 + Bv_1v_2 + Cv_2^2$ must be shown strictly positive with $A > 0$ **DO**

multiply by A and write $AQ = (Av_1 + \frac{1}{2}Bv_2)^2 + (AC - \frac{1}{4}B^2)v_2^2$ **BECAUSE** multiplying rather than dividing keeps the identity polynomial and makes the equality case visible, which is what turns a nonnegativity statement into a strict positivity statement

- **WHEN** you have chosen a method and one stated hypothesis is still unused **DO** treat that as an alarm and change method **BECAUSE** Putnam hypotheses are minimal by construction, so an unused hypothesis nearly always means your route proves something weaker than what was asked

Backtesting

1. First useful cue: the printed expression $f_{11}f_{22} - \frac{1}{4}(f_{12} + f_{21})^2$, which is the determinant of the symmetrised Jacobian and the discriminant condition for a positive definite binary quadratic form.
2. Method the expression suggested: produce $v^{\mathsf{T}} Df\, v$ somewhere, which means a directional derivative in direction v paired again with v, which means the function $g(t) = v \cdot f(p + tv)$ on the segment joining the two points.
3. Where you would hesitate, and the thought to reinforce: after computing $\det Df > 0$ it feels like the inverse function theorem should close the problem. Reinforce: *local hypotheses give local conclusions; to go global you must integrate along a path or use compactness*. Keep the complex exponential counterexample in memory as the permanent refutation.
4. The trap the question sets: the positivity of all four partials is a decoy that pushes you toward coordinatewise monotonicity, which handles only the directions in the first and third quadrants; and the Jacobian determinant is a decoy that pushes you toward the inverse function theorem. Neither uses the actual hypothesis.
5. What makes you recognise it faster next time: memorise the phrase "positive definite symmetric part equals strictly monotone equals injective". This problem is that theorem in dimension two with the definition spelled out.
6. The reusable decision rule: to prove a global property of a differentiable map, build a one-variable function along a segment whose derivative is exactly the quantity the hypothesis controls.

Mirror drills

1. Prove that if $f : \mathbb{R}^n \to \mathbb{R}^n$ is C^1 and $Df(x) + Df(x)^{\mathsf{T}}$ is positive definite at every x, then f is injective. *Answer:* identical proof; $g(t) = v \cdot f(p + tv)$ has $g'(t) = v^{\mathsf{T}} Df\, v = \frac{1}{2}v^{\mathsf{T}}(Df + Df^{\mathsf{T}})v > 0$ for $v \neq 0$.
2. Exhibit a C^∞ map $\mathbb{R}^2 \to \mathbb{R}^2$ with $\det Df > 0$ everywhere that is not injective, and say which hypothesis of B5 it violates. *Answer:* $(x, y) \mapsto (e^x \cos y, e^x \sin y)$, with $\det Df = e^{2x} > 0$ and period 2π in y; it violates positivity of $f_{11} = e^x \cos y$, hence the definiteness of the symmetric part.
3. Let $\varphi : \mathbb{R}^n \to \mathbb{R}$ be C^2 with positive definite Hessian everywhere. Prove $\nabla\varphi$ is injective, and deduce that a strictly convex function has at most one stationary point. *Answer:* apply drill 1 to $f = \nabla\varphi$, whose Jacobian is the (symmetric) Hessian; injectivity of $\nabla\varphi$ means $\nabla\varphi(x) = 0$ has at most one solution.

8.12 B6. Exponential weights bound restricted composition counts

Let S be the set of sequences of length 2018 whose terms are in the set $\{1,2,3,4,5,6,10\}$ and sum to 3860. Prove that the cardinality of S is at most

$$2^{3860} \cdot \left(\frac{2018}{2048}\right)^{2018}.$$

Read it out loud

Count the sequences of a fixed length, from a fixed finite alphabet of positive integers, with a fixed total. Then bound that count. Say the bound aloud in its factored form: something to the power of the *sum*, times something to the power of the *length*. That is

$$2^{3860}\left(\frac{2018}{2048}\right)^{2018} = (2^{\,\text{sum}}) \cdot \left(\lambda^{\,\text{length}}\right), \qquad \lambda = \frac{2018}{2048}.$$

A bound whose two exponents are the two parameters of the problem is not a bound somebody derived by estimating binomial coefficients. It is a bound that came out of a product, because products are the only things that turn into powers, and it has one factor per term of the sequence plus one factor per unit of sum.

Now attack the constant. The bound contains $2048 = 2^{11}$, which appears nowhere in the problem, so it was manufactured. Reduce: $\frac{2018}{2048} = \frac{1009}{1024} = \frac{1009}{2^{10}}$. And $2^{10} = 1024$ is exactly what you get when you put the largest alphabet element, 10, into an exponent of 2. So compute the one quantity that the alphabet and the number 2 can produce together:

$$\sum_{j\in\{1,2,3,4,5,6,10\}} 2^{-j} = \left(\frac{1}{2}+\frac{1}{4}+\frac{1}{8}+\frac{1}{16}+\frac{1}{32}+\frac{1}{64}\right)+\frac{1}{1024} = \frac{63}{64}+\frac{1}{1024} = \frac{1008+1}{1024} = \frac{1009}{1024}.$$

That is λ, on the nose. The problem is now over, and you have not yet written a single line of proof. Everything after this is bookkeeping: you must find the argument in which the number $\sum_j 2^{-j}$ gets raised to the power 2018 and compared against 2^{-3860} times a count.

The fields wired in: combinatorics (compositions with restricted parts), generating functions (the product $\left(x+x^2+\cdots+x^6+x^{10}\right)^{2018}$, whose coefficient of x^{3860} is the answer), and probability in the guise of the exponential moment method, also called the Chernoff bound. What is really being tested is whether you know that a coefficient of a power series with nonnegative coefficients is bounded by the value of the series at any positive argument, divided by the corresponding power. That single sentence is the problem.

Within thirty seconds, write down: $\lambda = 2018/2048 = 1009/1024 = \sum_{j\in A} 2^{-j}$, and the plan "weight each sequence by $2^{-(\text{its sum})}$ and sum over all sequences".

Approaches, ranked worst to best

1. **Exact count.** *Cue:* it is a finite count, so compute it. *Circuit:* stars and bars, or expand the generating function. You cannot expand $\left(x+\cdots+x^6+x^{10}\right)^{2018}$ by hand, and even the crude relaxation to all positive integer parts gives, by stars and bars, $\binom{3859}{2017} \approx 10^{1158}$. The target is about 10^{1149}. So the relaxation overshoots by a factor of about 10^9 and the route is dead, not merely painful. Ranked last. It does earn one thing: computing the trivial bound tells you the margin is thin, and thin margins mean the intended argument is tight and structural rather than wasteful.

2. **Multinomial and Stirling.** *Cue:* classify by how many of each letter you use. *Circuit:* maximise a multinomial coefficient subject to two linear constraints. Writing k_j for the number of terms

equal to j, the count is $\sum \binom{2018}{k_1,\ldots,k_7}$ over all (k_j) with $\sum k_j = 2018$ and $\sum jk_j = 3860$. You would then bound each multinomial by its entropy exponential, maximise by Lagrange multipliers, and count the number of terms in the sum. This works in principle, and it is the honest large deviations calculation. It ranks second to last because Stirling error terms, the number of summands and the constrained maximisation together cost pages, and every one of those pages is a chance to lose a point.

3. **Double induction on length and sum.** *Cue:* a recursion on the last term is available. *Circuit:* strengthen the statement to all (k, n) and induct. Define $a(k, n)$ as the number of sequences of length k from the alphabet summing to n. Splitting on the last term gives $a(k, n) = \sum_{j \in A} a(k-1, n-j)$, and the strengthened claim $a(k, n) \leq 2^n \lambda^k$ then propagates:

$$\sum_{j \in A} 2^{n-j} \lambda^{k-1} = 2^n \lambda^{k-1} \sum_{j \in A} 2^{-j} = 2^n \lambda^{k-1} \cdot \lambda = 2^n \lambda^k.$$

This is a complete and correct solution, and it is the one printed in the official file. It ranks third because you have to invent the strengthened statement (the original statement at a single (k, n) is not inductable), organise a two variable induction, and check the degenerate cases where $n \leq k$ or n is out of range. All of that work disappears in the next two routes, and the reason it disappears is that the induction is nothing but the distributive law applied one factor at a time.

4. **Probabilistic Chernoff.** *Cue:* a count of length 2018 sequences with a fixed sum is a probability. *Circuit:* apply Markov's inequality to $e^{\lambda X}$. Let $X_1, \ldots, X_{2018}$ be independent and uniform on the alphabet. Then $|S| = 7^{2018} \Pr\left[\sum X_i = 3860\right]$, and $\Pr[\sum X_i = 3860] \leq 2^{3860} \mathbb{E}\left[2^{-\sum X_i}\right] = 2^{3860} \left(\frac{1}{7} \sum_j 2^{-j}\right)^{2018}$ by Markov applied to the positive random variable $2^{-\sum X_i}$. Correct, and identical in content to the best route, but it dresses a finite sum as an expectation and pays a 7^{2018} that immediately cancels. Ranked fourth: same idea, extra costume.

5. **Weight and compare.** *Cue:* the bound factors as 2^{sum} times λ^{length}, and $\lambda = \sum_j 2^{-j}$. *Circuit:* attach a positive weight to every object, sum over a superset, and keep only the terms you care about. Best. Give the sequence s the weight $2^{-\sigma(s)}$ where $\sigma(s)$ is its sum. Sum over *all* sequences of length 2018 from the alphabet: by the distributive law the total is $\left(\sum_j 2^{-j}\right)^{2018} = \lambda^{2018}$. Every sequence in S contributes exactly 2^{-3860}, and all weights are positive, so $|S|\, 2^{-3860} \leq \lambda^{2018}$. Six lines, no induction, no error terms, no probability. It also exposes the free parameter: the number 2 was never sacred, and replacing it by $1/x$ gives a whole family of bounds, of which the problem's is the tightest.

The ideal solution

Write $A = \{1, 2, 3, 4, 5, 6, 10\}$, let $T = A^{2018}$ be the set of *all* sequences of length 2018 with terms in A, and for $s = (s_1, \ldots, s_{2018}) \in T$ write $\sigma(s) = s_1 + \cdots + s_{2018}$. Thus $S = \{s \in T : \sigma(s) = 3860\}$.

Attach to each $s \in T$ the weight

$$w(s) = 2^{-\sigma(s)} = \prod_{i=1}^{2018} 2^{-s_i} > 0.$$

Step 1: the total weight of T. Expanding a product of sums by the distributive law, the 7^{2018} terms of

$$\prod_{i=1}^{2018} \left(\sum_{j \in A} 2^{-j} \right)$$

are exactly the products $\prod_{i=1}^{2018} 2^{-s_i}$, one for each choice of $(s_1, \ldots, s_{2018}) \in T$. Hence

$$\sum_{s \in T} w(s) = \left(\sum_{j \in A} 2^{-j} \right)^{2018}.$$

Step 2: evaluate the constant.

$$\sum_{j \in A} 2^{-j} = \frac{1}{2} + \frac{1}{4} + \frac{1}{8} + \frac{1}{16} + \frac{1}{32} + \frac{1}{64} + \frac{1}{1024} = \left(1 - \frac{1}{64}\right) + \frac{1}{1024} = \frac{1008}{1024} + \frac{1}{1024} = \frac{1009}{1024} = \frac{2018}{2048}.$$

Step 3: keep only the sequences you want. Every weight is positive and $S \subseteq T$, so

$$\sum_{s \in S} w(s) \;\le\; \sum_{s \in T} w(s) = \left(\frac{2018}{2048} \right)^{2018}.$$

Every $s \in S$ has $\sigma(s) = 3860$ and therefore $w(s) = 2^{-3860}$ exactly, so the left side equals $|S| \cdot 2^{-3860}$. Hence

$$|S| \cdot 2^{-3860} \le \left(\frac{2018}{2048} \right)^{2018},$$

and multiplying by $2^{3860} > 0$ gives

$$\boxed{|S| \;\le\; 2^{3860} \left(\frac{2018}{2048} \right)^{2018}.}$$

■

(The inequality in Step 3 is in fact strict, since $T \setminus S$ is nonempty: the constant sequence of ones lies in T and has sum $2018 \neq 3860$. The problem only asks for "at most", so nothing more is needed.)

Remark 8.12.1. The number 2 is a free parameter in disguise. For any real $x > 0$, set $P(x) = x + x^2 + x^3 + x^4 + x^5 + x^6 + x^{10}$. The same three steps with the weight $x^{\sigma(s)}$ give

$$|S|\, x^{3860} \le \sum_{s \in T} x^{\sigma(s)} = P(x)^{2018}, \qquad \text{so} \qquad |S| \le \frac{P(x)^{2018}}{x^{3860}} \quad \text{for every } x > 0.$$

The problem hands you the case $x = \frac{1}{2}$. Which x is best? Minimising log of the right side, that is $2018 \log P(x) - 3860 \log x$, gives the stationarity condition

$$\frac{xP'(x)}{P(x)} = \frac{3860}{2018}.$$

At $x = \frac{1}{2}$ we have $xP'(x) = \sum_{j \in A} j\, 2^{-j} = \frac{1930}{1024}$ and $P(x) = \frac{1009}{1024}$, so the left side is $\frac{1930}{1009}$, and the right side is $\frac{3860}{2018} = \frac{1930}{1009}$. They agree exactly. So $x = \frac{1}{2}$ is precisely the optimal choice, and the number 3860, which plays no other role in the problem, was chosen to make it so.

Layer 1: WHAT

Block 1. Read the bound as a product: $(2^{\text{sum}}) \cdot \left(\lambda^{\text{length}}\right)$, one factor per unit of sum and one per term.

Block 2. Reduce 2018/2048 to 1009/1024 and recognise it as $\sum_{j \in A} 2^{-j}$.

Block 3. Define the weight $w(s) = 2^{-\sigma(s)}$ on the set T of all sequences from the alphabet.

Block 4. Compute $\sum_{s \in T} w(s) = \left(\sum_j 2^{-j} \right)^{2018}$ by the distributive law.

Block 5. Throw away every sequence not in S; the weights are positive so the sum only drops.

Block 6. On S the weight is the constant 2^{-3860}, so the surviving sum is $|S|2^{-3860}$. Rearrange.

The nominal route tries to count S, or to bound each multinomial coefficient in a classification of S, and drowns. The bifurcation point is the instant you look at 2018/2048 and decide whether it is a ratio to be estimated or a value to be recognised. Recognition costs one line of arithmetic and finishes the problem; estimation costs the rest of the three hours. A related and equally sharp way to put it: the nominal route counts the objects, the ideal route *weighs* them, and weighing is easier because weight is multiplicative across coordinates while counting with a constraint is not.

Layer 2: HOW

Block 1. **Cue.** A bound of the form $\alpha^n \beta^k$ where n is the sum and k is the length. **Why natural.** Exponentials in a parameter come from products with one factor per unit of that parameter. A length 2018 sequence has 2018 coordinates, so β^{2018} says "one factor per coordinate", which is the signature of independence across coordinates: the sum over all sequences of a product of per-coordinate quantities factors into a product of sums. And 2^{3860}, with 3860 the constrained total, says "one factor per unit of sum", which is the signature of a weight of the form $x^{\sigma(s)}$. Both exponents therefore point at the same object before you have done any work. **Link to the next block.** If the answer is a product of per-coordinate sums, you need to know what one coordinate contributes, so evaluate β and try to see it as a sum over the alphabet.

Block 2. **Cue.** The constant $2048 = 2^{11}$, which is not in the problem, sitting under 2018. **Why natural.** Manufactured constants are messages. Cancel the factor 2 to get 1009/1024, and note $1024 = 2^{10}$ with 10 the largest alphabet element; that pairing is what tells you to try $\sum_j 2^{-j}$ rather than, say, $\sum_j 2^j$ or a geometric sum with a different ratio. The arithmetic then confirms it exactly, and exact confirmation of a guessed identity is proof that you have found the intended structure, not a coincidence. **Link to the next block.** You now know the per-coordinate contribution is $\sum_j 2^{-j}$, so the weight on a whole sequence must be $\prod_i 2^{-s_i} = 2^{-\sigma(s)}$.

Block 3. **Cue.** You want to bound $|S|$, and you know a quantity attached to sequences that factorises. **Why natural.** Counting a constrained set is hard because the constraint couples the coordinates. Weighting is easy because a weight of the form $x^{\sigma(s)}$ decouples: it is a product of one factor per coordinate. So replace the hard question "how many s have $\sigma(s) = 3860$" by the easy question "what is the total weight of everything", and let the weight itself detect the constraint. This is the same manoeuvre as introducing a generating function, stripped of the formal power series language. **Link to the next block.** Now compute the easy total.

Block 4. **Cue.** A sum over a Cartesian product of a product of per-coordinate terms. **Why natural.** Distributivity, nothing more: expanding $\prod_{i=1}^{k} \left(\sum_{j \in A} c_j \right)$ produces exactly one term for each choice function $i \mapsto s_i$. Say this explicitly in the write up; it is the only step in the proof with any content, and a grader wants to see that you know why the product of sums equals the sum over sequences. **Link to the next block.** You now have the total over the superset T, and you want a statement about the subset S.

Block 5. **Cue.** A subset of a set on which a positive function is summed. **Why natural.** If $w > 0$ then $\sum_S w \leq \sum_T w$ for $S \subseteq T$. This is the entire inequality of the problem, and it is the discrete form of Markov's inequality: one nonnegative term of a sum is at most the whole sum. The positivity of the weight is what makes discarding terms safe, so check it once and say so. **Link to the next block.** On S the weight is constant, so the left side collapses to a multiple of $|S|$.

Block 6. Cue. Constant weight on the constrained set. **Why natural.** The constraint $\sigma(s) = 3860$ was chosen by the problem, and it is precisely what makes w constant on S: the weight was designed to be a function of the sum alone. That is the moment the design pays off. Rearranging is trivial. **Link to the next block.** Finished, and the free parameter question remains for the remark.

Layer 3: WHY

The truth beneath this problem is one inequality, and it has three names depending on who is speaking. A combinatorialist says: if a power series has nonnegative coefficients, then for $x > 0$ each coefficient satisfies $a_n x^n \le \sum_m a_m x^m$, so $a_n \le x^{-n} F(x)$. A probabilist says: for a nonnegative random variable Y and $t > 0$, $\Pr[Y \ge t] \le \mathbb{E}[Y]/t$, which is Markov's inequality, and applying it to $Y = x^{\sum X_i}$ produces the Chernoff bound. A student who has not met either says: one term of a sum of positive numbers is at most the sum. All three are the same sentence, and the third is enough to write a full marks solution. Install the sentence, not the vocabulary.

The second truth is that the bound has a free parameter, and free parameters must be optimised. The family $|S| \le x^{-n} P(x)^k$ holds for every $x > 0$; the best member is at $xP'(x)/P(x) = n/k$, which says: choose the weight so that the *weighted average value of a single term* equals the average value n/k that the constraint demands. That is exponential tilting. You are reweighting the uniform measure on the alphabet until its mean lands exactly on the target mean, and at that tilt the bound is as tight as this method can be. The general statement is Cramér's theorem in large deviations: $\frac{1}{k} \log \Pr\left[\frac{1}{k}\sum X_i \approx c\right] \to -I(c)$, where the rate function I is the Legendre transform of the logarithm of the moment generating function. The tilt condition above is exactly the equation defining the Legendre transform's stationary point. Knowing this tells you in advance that the method is not merely valid but essentially sharp, which is confirmed by the numbers: the true value of $|S|$ is about 8.05×10^{1146}, and the bound is about 10^{1149}, a factor of roughly 150 across a quantity with more than a thousand digits.

The examiners nudged in exactly two places, and both are numbers. First, 2048, which is not natural to the problem and exists only so that 2018/2048 reduces to $\sum_j 2^{-j}$. Second, 3860, which is otherwise irrelevant and exists only so that $x = 1/2$ is the exact optimum of the family of bounds. When two of the constants in a problem have no combinatorial meaning, they are keys, and the discernment habit to train is: *before estimating anything, try to recognise every constant in the target expression as an exact value of some natural quantity built from the data*. Reverse engineering the answer is not cheating; on a bounding problem it is the primary method, because the shape of the bound tells you which inequality was used to produce it.

How to train the recognition. Take twenty published bounds of the form C^n or $\alpha^n \beta^k$ and, for each, work out what quantity the base is: an alphabet sum, a moment generating function at a specific point, an entropy exponential $2^{H(p)n}$, or a spectral radius. After twenty you will see 1009/1024 and think "that is $\sum 2^{-j}$ over the alphabet" as fast as you see 0.5 and think one half.

Knowledge gaps, stated honestly

This problem has an unusual profile and you should be clear about it. The theory needed to *write* the solution is essentially nil: the distributive law, a geometric sum, and the observation that a subsum of positive terms is at most the whole sum. Nothing there is beyond A-level. The theory needed to *find* the solution is entirely beyond A-level, and it is the following, in order of priority.

(1) Ordinary generating functions as bookkeeping. The single fact that $\left(\sum_{j \in A} x^j\right)^k = \sum_n a(k,n) x^n$, where $a(k,n)$ counts length k sequences from A summing to n. Without this, the weight $2^{-\sigma(s)}$ looks arbitrary; with it, the weight is just evaluation at $x = 1/2$.

(2) Markov's inequality and the exponential moment (Chernoff) method: bound a tail or a coefficient by evaluating a positive series at a point of your choosing. This is the single most

reusable estimation technique in competition combinatorics and probability, and it is not taught anywhere in the A-level or Further Maths syllabus.

(3) Exponential tilting and Legendre duality, that is, the fact that the free parameter should be chosen so that the tilted mean equals the target mean, and that this choice makes the bound sharp to first order in the exponent. This is Cramér's theorem territory.

(4) Optional but worth knowing exists: Cauchy's coefficient formula and the saddle point method, which convert the same inequality into an asymptotic equality and would produce the constant 8.05×10^{1146} rather than an upper bound.

If you have not installed (1) and (2), you will not find this problem in the arena no matter how long you sit with it, and no amount of cleverness substitutes. That is what a knowledge gap means, and this is one.

Revolving concepts

Compositions with restricted parts. A composition of n into k parts from a set A is exactly an element of A^k summing to n; the count is $a(k,n) = [x^n]\left(\sum_{j\in A} x^j\right)^k$. Unrestricted positive parts give $\binom{n-1}{k-1}$ by stars and bars, which here is $\binom{3859}{2017}$ and is the crude comparison bound. See Book 1, field: Combinatorics, topic: compositions and stars and bars.

Generating functions and the recursion, generalised. The recursion behind the official solution is $a(k,n) = \sum_{j\in A} a(k-1, n-j)$, obtained by deleting the last term. This is a linear recursion in k with a shift in n, and its general solution method is exactly the generating function: multiply by x^n, sum over n, and get $G_k(x) = P(x)G_{k-1}(x)$ with $G_0 = 1$, hence $G_k(x) = P(x)^k$. Every recursion of the form "the count at stage k is a fixed nonnegative combination of shifted counts at stage $k-1$" solves this way, and the accompanying bound is the general fact

$$a(k,n) \le x^{-n}P(x)^k \quad \text{for every } x > 0, \qquad \text{optimised at } \frac{xP'(x)}{P(x)} = \frac{n}{k}.$$

Carry that display; it is the reusable object, and this problem is its instance with $P(x) = x + \cdots + x^6 + x^{10}$, $k = 2018$, $n = 3860$, $x = 1/2$. See Book 1, field: Generating functions, topic: products, coefficient extraction and coefficient bounds.

Markov's inequality and Chernoff bounds. For $Y \ge 0$ and $t > 0$, $\Pr[Y \ge t] \le \mathbb{E}[Y]/t$. Applied to $Y = e^{\theta\sum X_i}$ with independent X_i, it gives $\Pr\left[\sum X_i \ge a\right] \le e^{-\theta a}\left(\mathbb{E}e^{\theta X}\right)^k$ for every $\theta > 0$, and independence is what turns the expectation of the product into the product of expectations. See Book 1, field: Probability, topic: concentration inequalities.

Exponential tilting and the Legendre transform. Given a distribution p on A, the tilted distribution $p_\theta(j) \propto p(j)e^{\theta j}$ has mean $\frac{d}{d\theta}\log \mathbb{E}[e^{\theta X}]$. Choosing θ so that this mean equals the target n/k minimises the Chernoff bound, and the minimum value is $e^{-kI(n/k)}$ with I the Legendre transform of $\log \mathbb{E}[e^{\theta X}]$. See Book 1, field: Probability, topic: large deviations and Cramer's theorem.

Weighted counting as a general habit. Whenever a constraint couples coordinates, look for a weight that is multiplicative across coordinates and that is constant on the constrained set. Then the total weight is computable by factorisation, and the constrained count falls out by comparison. The same habit proves $\binom{2n}{n} \le 4^n$, bounds the number of self avoiding walks, and underlies the probabilistic method. See Book 1, field: Combinatorics, topic: weighted counting and the first moment method.

Circuit Bank

- **WHEN** a bound of the shape $\alpha^{\text{sum}}\beta^{\text{length}}$ **DO** look for a per-coordinate weight, because only products give exponents in a parameter **BECAUSE** the sum over a Cartesian product of a multiplicative weight factorises into a product of identical per-coordinate sums, which is where β^k comes from

- **WHEN** an unexplained constant in the target, such as 2048 or 1009/1024 **DO** try to realise it exactly as a natural quantity built from the data, here $\sum_{j\in A} 2^{-j}$ **BECAUSE** competition bounds are reverse engineered from the intended proof, so the constants in the answer identify the inequality that produced them
- **WHEN** you must bound the number of solutions of a linear constraint over a finite alphabet **DO** weight each object by x^{sum}, sum over everything, and discard the terms you do not need **BECAUSE** all weights are positive, so a subsum is at most the whole sum; this is Markov's inequality in its discrete form and needs no probability language
- **WHEN** a valid bound still contains a free positive parameter **DO** optimise it, and check whether the problem's numbers make the optimum a round value **BECAUSE** the family $a(k,n) \le x^{-n}P(x)^k$ is tightest at $xP'(x)/P(x) = n/k$, and a problem that hands you $x = 1/2$ has usually chosen n to sit exactly at that stationary point
- **WHEN** a counting problem where a constraint couples the coordinates **DO** stop counting and start weighing **BECAUSE** counting with a coupled constraint does not factorise, but a multiplicative weight does, and the constraint can be recovered afterwards because the weight is constant on the constrained set

Backtesting

1. First useful cue: the factored shape of the bound, 2^{3860} against the sum and $(2018/2048)^{2018}$ against the length, followed immediately by the reduction $2018/2048 = 1009/1024 = \sum_{j\in A} 2^{-j}$.
2. Method the expression suggested: weighted counting. Attach $2^{-\sigma(s)}$ to every sequence, factorise the total over all sequences, then discard everything outside S.
3. Where you would hesitate, and the thought to reinforce: the instinct to estimate $|S|$ directly, or to set up the multinomial classification. Reinforce the thought: *a bound with exponents equal to the problem's parameters was produced by a product, so find the product*. Also reinforce that discarding positive terms is a legal and powerful move, not a lossy accident.
4. The trap the question sets: the number 3860 invites you to think it matters combinatorially. It does not; it appears in the proof only as the exponent of the constant weight on S. Its real role is to make $x = 1/2$ the optimal tilt, which you can only see after the fact.
5. What makes you recognise it faster next time: build the reflex that any target of the form $x^{-n}P(x)^k$ is a coefficient bound. Whenever a bound has two exponents and the problem has two parameters, write the generating function down within the first minute.
6. The reusable decision rule: to bound the number of objects satisfying an additive constraint, weight every object by x^{sum} for a positive x of your choosing, sum over all objects by factorisation, and compare; then optimise x if the problem has not chosen it for you.

Mirror drills

1. Prove $\binom{2n}{n} \le 4^n$ in one line by weighted counting. *Answer:* $\binom{2n}{n}$ is one term of $\sum_k \binom{2n}{k} = (1+1)^{2n} = 4^n$, and all terms are positive. Same circuit with $x = 1$.
2. Let $b(k,n)$ be the number of sequences of length k with terms in $\{1,2\}$ summing to n. Prove $b(k,n) \le x^{-n}(x+x^2)^k$ for all $x > 0$, and find the optimal x when $n = 3k/2$. *Answer:* the same weighting argument; the tilt condition $\frac{x(1+2x)}{1+x} = \frac{n}{k} = \frac{3}{2}$ gives $2x^2 + x = \frac{3}{2} + \frac{3}{2}x$, that is $4x^2 - x - 3 = 0$, so $x = 1$, which is correct because $n/k = 3/2$ is the mean of the uniform distribution on $\{1,2\}$.

3. State and prove the Chernoff bound: for independent X_i and $\theta > 0$, $\Pr\left[\sum_{i=1}^{k} X_i \geq a\right] \leq e^{-\theta a}\left(\mathbb{E}e^{\theta X}\right)^k$. *Answer:* the event is $e^{\theta \sum X_i} \geq e^{\theta a}$; apply Markov to the positive random variable $e^{\theta \sum X_i}$ and factor the expectation using independence.

Chapter 9

Putnam 2017

9.1 A1. the smallest set generated by squares and a shift

Let S be the smallest set of positive integers such that

(a) 2 is in S,

(b) n is in S whenever n^2 is in S, and

(c) $(n+5)^2$ is in S whenever n is in S.

Which positive integers are not in S?

(The set S is smallest in the sense that S is contained in any other set satisfying (a), (b) and (c).)

Read it out loud

The phrase that should catch your eye is "the smallest set such that". You are not given S as a formula or a recurrence: you are given a seed, 2, and two rules, and S is the least set that swallows the seed and stays closed under those rules. That word, smallest, is the whole command term: S is the intersection of every closed set containing the seed, which turns "which integers are missing" into two separate, finite tasks rather than one open search.

Look at what the rules do. Rule (c) shifts n up by 5 then squares, producing much larger numbers. Rule (b) takes a known perfect square and extracts its root, producing smaller numbers. One rule climbs, one descends, and both are built to meet on perfect squares, so expect to compose them.

The constant 5 is not decoration. Because (c) always adds a multiple of the same constant before squaring, and squaring is exactly the operation for which divisibility by a prime transfers cleanly ($p \mid m \iff p \mid m^2$), the residue of the seed mod 5 is the first thing to write down. The other loose end is the fixed point of squaring, $1^2 = 1$, which deserves a separate glance whenever a rule and its near inverse coexist.

Approaches, ranked worst to best

1. **raw forward generation.** *Cue:* the request to name missing integers invites listing. *Circuit:* repeated application of (b) and (c) to whatever is currently known. (this only ever proves a number IS in S by exhibiting a chain; it can never certify a number is NOT in S, since minimality is a global property, not a property of what you have generated so far.)

2. **induction on "$n \notin S$" without first naming a closed candidate.** *Cue:* the request for a clean description. *Circuit:* strong induction on the size of n. (workable in principle, but only after reinventing the residue invariant inside the inductive step, so it hides the same insight the pincer states openly.)

3. **guess a closed superset, then generate to match it.** *Cue:* the word "smallest" attached to closure conditions. *Circuit:* the two inequality pincer for smallest closed sets: guess T, prove T closed ($S \subseteq T$), then chase the seed through the rules until every part of T is reached ($T \subseteq S$). (correct and short: each half is mechanical once the residue invariant names T.)

The ideal solution

Claim. The positive integers not in S are exactly 1 and the multiples of 5.

Let T be the set of positive integers other than 1 and the multiples of 5.

Half 1: $S \subseteq T$. It suffices that T satisfies (a), (b), (c), since S is contained in every set that does. Condition (a): $2 \in T$. Use throughout that 5 is prime, so $5 \mid m \iff 5 \mid m^2$ for any positive integer m, and that $m^2 = 1 \iff m = 1$.

(b) If $n^2 \in T$ then $n^2 \neq 1$ and $5 \nmid n^2$, so $n \neq 1$ and $5 \nmid n$, giving $n \in T$.

(c) If $n \in T$ then $n \neq 1, 5 \nmid n$, so $n \geq 2$, hence $n+5 \geq 7$ and $(n+5)^2 \geq 49 \neq 1$. Also $5 \mid (n+5) \iff 5 \mid n$, false, so $5 \nmid (n+5)^2$. Hence $(n+5)^2 \in T$.

So T is closed under (a), (b), (c), giving $S \subseteq T$.

Half 2: $T \subseteq S$. First derive a free shift: if $n \in S$ then $(n+5)^2 \in S$ by (c), and applying (b) to the integer $n+5$ gives $n+5 \in S$. So

$$n \in S \implies n+5 \in S, \quad \text{hence} \quad n \in S \implies n+5k \in S \text{ for every } k \geq 0. \tag{d}$$

Chase the seed using (c), (d), (b) in turn:

$$2 \xrightarrow{(c)} 49 \xrightarrow{(c)} 54^2 \xrightarrow{(d),+220} 56^2 \xrightarrow{(b)} 56 \xrightarrow{(d),+65} 121 \xrightarrow{(b)} 11 \xrightarrow{(d),+5} 16 \xrightarrow{(b)} 4 \xrightarrow{(d),+5} 9 \xrightarrow{(b)} 3,$$

and, continuing from 16,

$$16 \xrightarrow{(d),+20} 36 \xrightarrow{(b)} 6.$$

(Each (d) jump adds a multiple of 5: $220 = 5 \cdot 44$, $65 = 5 \cdot 13$, $20 = 5 \cdot 4$.) This exhibits $2, 3, 4, 6 \in S$, one representative of each nonzero residue class mod 5 ($2, 3, 4, 1$ respectively), and each is the least member of its class inside T (the class of 1 starts at 1, excluded, then continues at 6). By (d), every larger member of each class also lies in S, so $T \subseteq S$.

Combining both halves, $S = T$:

the positive integers not in S are 1 and every positive multiple of 5.

Layer 1: WHAT

Block 1: read "smallest set satisfying closure conditions" as the two inequality pincer. Block 2: extract the invariant, divisibility by the prime 5, and flag the fixed point $n = 1$. Block 3: build T from the invariant, verify closure, giving $S \subseteq T$. Block 4: compose (c) then (b) into the free shift rule (d). Block 5: chase the seed to land on $3, 4, 6$ (with 2), one per nonzero residue class. Block 6: read off that these are minimal class representatives, invoke (d) for the rest, combine into $S = T$.

The nominal route never leaves Block 1: it enumerates from 2, notices the pattern, and stops at a guess. The bifurcation point is the moment right after that guess, where the ideal route asks how to rule out every other exclusion and is forced into Blocks 2 and 3, while the zombie route never asks and so never builds T.

Layer 2: HOW

Block 1. **Cue.** S is defined by "smallest" plus closure rules, not a formula. **Why natural.** A least closed set is the intersection of all closed supersets, which is only useful once you exhibit one such superset and one generated subset. **Link to the next block.** Choosing the pincer raises the question of which invariant makes a superset easy to verify.

Block 2. **Cue.** Rule (c) shifts by the fixed constant 5 before squaring. **Why natural.** A shift by a multiple of 5 preserves residue mod 5, and 5 prime gives $5 \mid m \iff 5 \mid m^2$, so both rules preserve "not a multiple of 5" in the same direction. **Link to the next block.** The fixed point $n = 1$ is the only remaining loose end, so it joins the invariant to define T.

Block 3. **Cue.** T (everything except 1 and multiples of 5) is now concrete. **Why natural.** Verifying (a), (b), (c) for T is one line each, because the invariant was chosen to be preserved by both rules. **Link to the next block.** Closure only gives $S \subseteq T$; the harder, generative half remains.

Block 4. **Cue.** (b) and (c) are near inverses, one squares, one extracts a root. **Why natural.** Composing them collapses to plain addition: $(n+5)^2$ is already the perfect square (b) needs, so no further squaring is required. **Link to the next block.** Membership now propagates by $+5$ for free, shrinking the task to one representative per residue class.

Block 5. **Cue.** Four residue classes each need one representative, reachable by computation alone. **Why natural.** Every arrow is forced: (c) is used to manufacture a new square, (d) to close the gap between two squares (or two representatives) by a multiple of 5, (b) the instant a square appears. **Link to the next block.** The numbers landed on, $2, 3, 4, 6$, are exactly the least members of their classes in T.

Block 6. **Cue.** The four representatives cover residues $1, 2, 3, 4$ and residue 0 is excluded entirely, matching T. **Why natural.** Rule (d) turns one membership fact into an entire upward class, so four facts expand into all of T. **Link to the next block.** Both inclusions now hold, giving $S = T$ and the boxed answer.

Layer 3: WHY

The route is ideal because the problem is a closure operator on the positive integers, and every closure operator generated by a seed and finite rules sits between the same two bounds: any closed superset bounds it above, any explicit derivation bounds it below. Forward search only ever touches the lower bound, since minimality is a property of every closed superset at once, not of one growing set. The discernment mechanism is syntactic: see "the smallest set such that" attached to closure rules and commit to a two sided inclusion before anything else.

The invariant, divisibility by 5, is forced by matching the shift constant in (c) to a property both rules preserve, and primality is exactly what makes $p \mid m \iff p \mid m^2$ hold with no extra hypothesis. Had the shift been composite, say 4, that same style of invariant would fail ($4 \mid m^2$ does not force $4 \mid m$, take $m = 2$), and you would need the prime factor 2 instead. Train this by factorising the shift constant and testing its prime factors as candidate invariants, never the shift constant itself.

Revolving concepts

Smallest set satisfying closure conditions. The intersection of every set containing the generators and closed under the rules, equivalently everything reachable from the generators by finitely many rule applications. The double inclusion technique here generalises directly to generated subgroups, generated ideals, and the smallest sigma algebra containing a collection of sets. See Book 1, field: Logic and set theory, topic: closure systems.

Prime divisibility of powers. For a prime p, $p \mid m \iff p \mid m^k$ for any $k \geq 1$, from unique factorisation. This is why primality of the shift constant, not merely its value, is load bearing. See Book 1, field: Number theory, topic: prime factorisation.

Fixed points of squaring. Among integers, $n^2 = n$ only at $n = 0, 1$. Whenever a rule and its near inverse coexist, the fixed points of the rule are the first place sporadic exceptions hide. See Book 1, field: Number theory, topic: fixed points.

Additive closure and residue classes. A subset of the positive integers closed under adding a fixed c is, within each residue class mod c it meets, an up set: everything from some threshold onward. See Book 1, field: Sequences and recurrences, topic: numerical semigroups. Generalised recursion: whenever closure gives $n \in S \Rightarrow n + c \in S$ and representatives $r_1, \ldots$ are realised for the nonzero residues mod c actually present, S equals the union of $\{n \equiv r_i \pmod{c} : n \geq r_i\}$; the infinite description collapses to a finite list of realised classes and their least elements, the pattern to hunt for in any shift plus closure problem.

Circuit Bank

- **WHEN** a set is defined as smallest satisfying closure conditions **DO** split into two inclusions, a guessed closed superset from the rules' arithmetic, then explicit generation from the seed **BECAUSE** minimality gives the upper inclusion for free once any closed candidate is verified, while the lower inclusion needs an explicit derivation closure alone cannot supply
- **WHEN** a closure rule shifts by a fixed constant k before a power or root **DO** test divisibility by each prime factor of k, not by k itself **BECAUSE** $p \mid m \iff p \mid m^j$ holds cleanly only for primes, so composite constants need their prime factors as the invariant
- **WHEN** two closure rules resemble an operation and its near inverse **DO** compose them directly and simplify **BECAUSE** such compositions often collapse a complicated rule into a much simpler net operation that does the remaining work

Backtesting

1. First useful cue: "the smallest set of positive integers such that", attached to closure rules.
2. Method the expression suggested: the shift constant 5 is prime, so test divisibility by 5 immediately.
3. Where you would hesitate: after guessing the exclusions from a few generated examples, unsure how to rule out every other one; reinforce that minimality only bites once a closed candidate is named and verified.
4. The trap the question sets: generating more and more elements of S by hand, which feels like proof but never touches the direction that certifies exclusions.
5. What makes you recognise it faster next time: treating "smallest set satisfying closure conditions" as a fixed phrase that always calls for the two inclusion pincer.
6. The reusable decision rule: match the shift constant's prime factors to an invariant, verify a candidate superset closed, then chase the seed until one representative per needed residue class is exhibited.

Mirror drills

1. Repeat the construction with seed 2 and the same rules but shift 3. *Answer:* the upper bound argument transfers unchanged since 3 is prime and $2 \neq 1, 3 \nmid 2$, so 1 and the multiples of 3 are excluded provided the seed's orbit reaches both nonzero residues mod 3.
2. Let S be the smallest set of integers containing 0, closed under $n \mapsto n + 2$ and $n \mapsto -n$. *Answer:* S is the even integers: they are closed under both rules and contain 0, and every even integer is reached from 0 by repeated $+2$ or its negation.
3. Reseed the original problem at 3 instead of 2, shift still 5. *Answer:* the upper bound half is unchanged, since it only used that the seed is not 1 and not a multiple of 5; only the explicit chase needs rechecking to confirm the new orbit still reaches every nonzero residue class mod 5.

9.2 A2. a quotient recursion that is secretly linear

Let $Q_0(x) = 1$, $Q_1(x) = x$, and

$$Q_n(x) = \frac{(Q_{n-1}(x))^2 - 1}{Q_{n-2}(x)}$$

for all $n \geq 2$. Show that, whenever n is a positive integer, $Q_n(x)$ is equal to a polynomial with integer coefficients.

Read it out loud

The definition hands you a division, and a division you are told will always come out even. That is the first thing to say out loud: Q_n is only manifestly a rational function, and the entire content of the problem is that the denominator Q_{n-2} always divides the numerator $Q_{n-1}^2 - 1$ exactly. Whenever a Putnam recursion is stated as a quotient and asks you to prove the result is a polynomial, the examiners are testing whether you reach for polynomial long division (slow, case heavy) or for the algebraic reason the division is exact.

Look at the shape of the numerator: "square of the previous term, minus one". That specific shape, $u_{n-1}^2 -$ (constant) divided by u_{n-2}, is not arbitrary. It is exactly the identity satisfied by two consecutive terms of an ordinary linear three term recurrence, the same shape as Fibonacci's $F_{n-1}F_{n+1} - F_n^2 = (-1)^n$. That resemblance is the guiding light: within the first thirty seconds you should be asking "what linear recurrence, sharing these seeds, would make this quotient identity true by construction," not "how do I divide these polynomials."

The seeds themselves are a further cue: $Q_0 = 1, Q_1 = x$ are exactly the seeds of the simplest possible linear recurrence $u_n = xu_{n-1} - u_{n-2}$, the recurrence behind Chebyshev-like polynomials. That recurrence is manifestly integer coefficient by direct substitution, so if it can be shown to obey Q's own recursion, the whole problem collapses to checking one identity about a sequence with no denominators at all.

Approaches, ranked worst to best

1. **direct polynomial division term by term.** *Cue:* the recursion is stated as an explicit quotient. *Circuit:* compute $Q_2, Q_3, Q_4, \ldots$ by hand, verifying each division is exact. (this only confirms finitely many cases, the algebra grows quickly since each numerator involves the square of an ever larger polynomial, and it gives no reason the pattern continues forever.)
2. **induct on $\deg Q_n = n$ and manipulate the given recursion directly.** *Cue:* avoiding an invented auxiliary sequence. *Circuit:* divide two consecutive instances of the defining relation to expose a telescoping constant ratio, then unwind it into Q's own three term recurrence. (correct and used by some solvers, but it needs the degree induction as a prerequisite before any division of rational expressions is licensed, and discovering that the ratio $(Q_n + Q_{n-2})/Q_{n-1}$ is constant is a harder thing to notice cold than the route below.)
3. **invent a linear auxiliary sequence with the same seeds and identify it with Q.** *Cue:* the quotient's shape matches a Cassini type identity for a linear recurrence. *Circuit:* define $P_n = xP_{n-1} - P_{n-2}$, prove $P_{n-1}^2 - P_nP_{n-2} = 1$ by a companion matrix determinant, then bootstrap $Q_n = P_n$. (best: every hard step lands on P, which has no denominators to worry about, and the determinant identity is a one line induction.)

The ideal solution

Define a second sequence of polynomials by

$$P_0(x) = 1, \quad P_1(x) = x, \quad P_n(x) = xP_{n-1}(x) - P_{n-2}(x) \ (n \geq 2).$$

Since each P_n ($n \geq 2$) is x times an integer coefficient polynomial minus another, P_n has integer coefficients for every $n \geq 0$, by an immediate induction. It suffices to show $Q_n = P_n$ for all $n \geq 0$.

Step 1: a Cassini identity for P. Extend the recurrence backward by setting $P_{-1} = 0$, consistent with $P_1 = xP_0 - P_{-1}$. For $n \geq 1$ define

$$M_n = \begin{pmatrix} P_{n-1} & P_n \\ P_{n-2} & P_{n-1} \end{pmatrix}, \qquad T = \begin{pmatrix} x & -1 \\ 1 & 0 \end{pmatrix}.$$

A direct check of entries, using $P_n = xP_{n-1} - P_{n-2}$ in the top right and top left, gives $TM_n = M_{n+1}$ for every $n \geq 1$. Since $\det T = x \cdot 0 - (-1) \cdot 1 = 1$, induction on n gives $\det M_n = \det M_1$ for all $n \geq 1$. Directly, $\det M_1 = P_0^2 - P_1 P_{-1} = 1 - 0 = 1$. So

$$(P_{n-1})^2 - P_n P_{n-2} = \det M_n = 1 \qquad \text{for every } n \geq 1. \tag{$*$}$$

Step 2: $P_n \neq 0$. By induction, $\deg P_0 = 0, \deg P_1 = 1$, and for $n \geq 2$, $\deg(xP_{n-1}) = n$ strictly exceeds $\deg P_{n-2} = n - 2$, so no cancellation occurs and $\deg P_n = n$. In particular P_n is never the zero polynomial.

Step 3: bootstrap $Q_n = P_n$. Induct on n. The base cases $Q_0 = 1 = P_0$ and $Q_1 = x = P_1$ hold by definition. Suppose $n \geq 2$ and $Q_{n-1} = P_{n-1}, Q_{n-2} = P_{n-2}$. By ($*$) with this value of n, $(P_{n-1})^2 - 1 = P_n P_{n-2}$ identically as polynomials, and $P_{n-2} \neq 0$ by Step 2, so dividing gives

$$Q_n = \frac{(Q_{n-1})^2 - 1}{Q_{n-2}} = \frac{(P_{n-1})^2 - 1}{P_{n-2}} = P_n,$$

an exact polynomial quotient. This completes the induction: $Q_n = P_n$ for every $n \geq 0$, and since P_n has integer coefficients, so does Q_n.

$Q_n(x) = P_n(x)$, defined by $P_0 = 1, P_1 = x, P_n = xP_{n-1} - P_{n-2}$, which manifestly has integer coefficients.

Layer 1: WHAT

Block 1: recognise the quotient's shape, "square of previous term minus a constant, over two back", as a Cassini type identity belonging to some linear recurrence. Block 2: guess the linear recurrence $P_n = xP_{n-1} - P_{n-2}$ sharing Q's seeds; note it is integer coefficient by construction. Block 3: reduce the whole problem to one identity, $P_{n-1}^2 - P_n P_{n-2} = 1$. Block 4: prove that identity via a companion matrix of determinant 1. Block 5: separately confirm $P_n \neq 0$ via a degree induction, so the division in Step 3 of the solution is licensed. Block 6: bootstrap $Q_n = P_n$ by strong induction, using Steps 1 and 2 at each stage.

The nominal route stalls inside Block 1: it computes $Q_2, Q_3, \ldots$ directly, confirms integrality case by case, and never converts the pattern into a proof for general n. The bifurcation point is the moment you ask "what recurrence produces exactly this quotient as an identity" instead of "let me compute the next term and check."

Layer 2: HOW

Block 1. Cue. The recursion divides a quadratic expression in the previous term by the term two back. **Why natural.** This is precisely the algebraic shape of $u_{n-1}^2 - u_n u_{n-2} =$ constant, the identity any linear two term recurrence with unit determinant companion matrix satisfies. **Link to the next block.** If such a linear recurrence exists and shares Q's seeds, it is a candidate for Q itself.

Block 2. Cue. $Q_0 = 1, Q_1 = x$ are the seeds of the simplest possible linear recurrence in x. **Why natural.** Defining $P_n = xP_{n-1} - P_{n-2}$ with these seeds gives a sequence with integer coefficients by direct induction, no division involved anywhere. **Link to the next block.** Whether $P = Q$ now rests entirely on whether P obeys Q's own defining relation.

Block 3. **Cue.** P obeys Q's relation exactly when $P_{n-1}^2 - 1 = P_n P_{n-2}$ identically. **Why natural.** Uniqueness of a recurrence's solution given initial data means matching both the seeds and the recursion forces the sequences to coincide. **Link to the next block.** Proving this single polynomial identity is now the entire remaining task.

Block 4. **Cue.** The target identity is exactly a 2×2 determinant of consecutive terms. **Why natural.** Packaging P_{n-1}, P_n, P_{n-2} into a matrix M_n and noting $TM_n = M_{n+1}$ for the companion matrix T turns the identity into $\det M_n = (\det T)^{n-1} \det M_1$, and $\det T = 1$ collapses this to a constant. **Link to the next block.** Before dividing by P_{n-2} in the bootstrap, its nonvanishing needs separate justification.

Block 5. **Cue.** The division $Q_n = (Q_{n-1}^2 - 1)/Q_{n-2}$ is only meaningful if $Q_{n-2} \neq 0$. **Why natural.** A quick degree induction shows $\deg P_n = n$, since xP_{n-1} strictly dominates P_{n-2} in degree, so P_n is never zero. **Link to the next block.** With both the identity and the nonvanishing in hand, the induction identifying Q_n with P_n can run at every step.

Block 6. **Cue.** Two sequences share initial values and the same defining recursion. **Why natural.** Substituting the inductive hypothesis into Q's definition and using $(*)$ turns the division into the exact quotient P_n, with no remainder. **Link to the next block.** The induction closes, and integer coefficients transfer from P to Q immediately.

Layer 3: WHY

The route is ideal because a quotient recursion of the shape $u_n = (u_{n-1}^2 - c)/u_{n-2}$ is never really about division at all: it is the Cassini invariant of some linear recurrence in disguise, and linear recurrences with integer or polynomial coefficients are always polynomial by construction. The discernment mechanism is to see past the division to the quadratic identity it encodes, and to recall that such identities are the fingerprints of companion matrices with determinant 1 (giving a genuine constant) or determinant -1 (giving an alternating sign, as in Fibonacci). Once that fingerprint is recognised, the problem stops being about Q at all and becomes a short, self contained fact about P, which is easier on every axis: no denominators, no need to verify nonvanishing beyond a degree count, and a proof by induction that never touches an inequality or a case split.

Training this discernment means building the reflex: whenever a recursion's right hand side is "(square of one term, minus a constant) divided by another term two steps back", immediately write down the guess $u_n = au_{n-1} - u_{n-2}$ (or the appropriate two term linear form matching the seeds) and check its Cassini identity, before attempting any direct manipulation of the given quotient.

Revolving concepts

Cassini type identities via companion matrices. For $u_n = au_{n-1} + bu_{n-2}$, the matrices $M_n = \begin{pmatrix} u_{n-1} & u_n \\ u_{n-2} & u_{n-1} \end{pmatrix}$ satisfy $M_{n+1} = TM_n$ for the companion matrix $T = \begin{pmatrix} a & b \\ 1 & 0 \end{pmatrix}$, so $\det M_n = (\det T)^{n-1} \det M_1$. See Book 1, field: Linear algebra, topic: companion matrices. Generalised recursion: whenever $\det T = 1$, $u_{n-1}^2 - u_n u_{n-2}$ is the same constant for every n; whenever $\det T = -1$, it alternates sign, recovering Fibonacci's $F_{n-1}F_{n+1} - F_n^2 = (-1)^n$. Any Putnam sequence defined by a quotient of this shape should immediately be tested against this template, with a, b read off from matching the quotient's coefficients.

Uniqueness of solutions to a second order recurrence. Two sequences obeying the same recursion and agreeing on the first two terms are identical for all later terms, by direct induction. See Book 1, field: Sequences and recurrences, topic: recurrence uniqueness.

Degree growth under a dominant linear term. If one term of a linear recurrence has strictly larger degree than the other at every step, no cancellation of leading terms is possible, so the degree grows by exactly the dominant term's degree each step. See Book 1, field: Polynomials, topic: degree induction.

The Laurent phenomenon. Beyond A level: recursions of the shape $u_n u_{n-2} = u_{n-1}^2 - 1$ (and more general exchange relations) are toy instances of the Laurent phenomenon from cluster algebra theory, where iterating such a birational recursion always produces Laurent polynomials, here genuine polynomials, in the initial data. Not needed for the elementary proof, but it is the reason this style of problem recurs: the integrality is not a coincidence of this one recursion, it is a shadow of a much larger structure. See Book 1, field: Abstract algebra, topic: cluster algebras.

Circuit Bank

- **WHEN** a recursion divides (square of the previous term minus a constant) by the term two back **DO** guess the linear recurrence sharing the same seeds and test its Cassini identity via a companion matrix **BECAUSE** such quotients are exactly the fingerprint of a linear recurrence with unit or minus one determinant companion matrix, so the division is guaranteed exact once that identity is confirmed
- **WHEN** two sequences share initial values and a recursive rule **DO** prove them equal by strong induction rather than manipulating either sequence's closed form **BECAUSE** uniqueness of solutions to a well posed recursion makes this the shortest possible route to equality
- **WHEN** a division needs to be justified as exact, or a term needs to be shown nonzero **DO** run a quick degree induction rather than inspecting coefficients **BECAUSE** a strictly dominant term at each step rules out cancellation, giving the degree, and hence nonvanishing, for free

Backtesting

1. First useful cue: the recursion divides a quadratic expression in the previous term by the term two back.
2. Method the expression suggested: match the quotient's shape to the Cassini identity of a guessed linear recurrence sharing the same seeds.
3. Where you would hesitate: unsure why the division is guaranteed exact before proving anything; reinforce that the Cassini identity, once proved for the linear sequence, supplies exact divisibility automatically.
4. The trap the question sets: computing Q_2, Q_3, Q_4 directly and trying to induct on the quotient itself, fighting an increasingly complicated rational expression.
5. What makes you recognise it faster next time: the signature "square minus a constant, divided by two back" as the fixed fingerprint of a Cassini identity, the same fingerprint behind Fibonacci's identity.
6. The reusable decision rule: replace a quotient defined recursion with a guessed linear recursion on the same seeds, prove its Cassini identity via a unit determinant companion matrix, then identify the two sequences by induction.

Mirror drills

1. Prove Fibonacci's identity $F_{n-1}F_{n+1} - F_n^2 = (-1)^n$. *Answer:* same companion matrix argument with $T = \begin{pmatrix} 1 & 1 \\ 1 & 0 \end{pmatrix}$, $\det T = -1$, so $\det M_n = (-1)^{n-1} \det M_1$ alternates sign.

2. Let $R_0 = 2, R_1 = x, R_n = xR_{n-1} - R_{n-2}$. Find the constant value of $R_{n-1}^2 - R_n R_{n-2}$. *Answer:* the same companion matrix has $\det T = 1$, and $\det M_1 = R_0^2 - R_1 R_{-1} = 4 - x^2$ (using $R_{-1} = x$ from the backward extension), so the constant is $4 - x^2$ for every n.

3. A numerical sequence has $u_0 = 1, u_1 = 2$, $u_n = (u_{n-1}^2 - 1)/u_{n-2}$. Identify the hidden linear recurrence and closed form. *Answer:* matching $u_2 = (4-1)/1 = 3$ to $au_1 - u_0$ gives $a = 2$, so $u_n = 2u_{n-1} - u_{n-2}$, the arithmetic progression recursion, giving $u_n = n + 1$.

9.3 A3. turning an unused zero integral into a sign

Let a and b be real numbers with $a < b$, and let f and g be continuous functions from $[a, b]$ to $(0, \infty)$ such that $\int_a^b f(x)\,dx = \int_a^b g(x)\,dx$ but $f \neq g$. For every positive integer n, define

$$I_n = \int_a^b \frac{(f(x))^{n+1}}{(g(x))^n}\,dx.$$

Show that $I_1, I_2, I_3, \ldots$ is an increasing sequence with $\lim_{n\to\infty} I_n = \infty$.

Read it out loud

Say the hypotheses back to yourself in order. You are told two integrals are equal, and separately that the two functions producing them are not the same function. The first statement, $\int f = \int g$, is algebraically identical to $\int_a^b (f - g)\,dx = 0$, and that is a fact you have been handed for free and have not yet used anywhere. The second statement, $f \neq g$, is currently just a label; it only becomes useful once continuity converts "differs somewhere" into "differs, with one sign, on a whole subinterval". Both hypotheses are inert until you build an expression that actually contains $(f - g)$.

Now look at what you are asked to prove: a sequence indexed by an exponent is increasing and unbounded. Whenever a sequence like $I_n = \int \phi^{n+1}/\psi^n$ is indexed by a power, the natural object to build is not I_n itself but the difference $I_{n+1} - I_n$, because the numerators of two consecutive terms are just one extra factor of f or g apart, and a difference of two things that are "the same shape, one exponent apart" is exactly where the factorisation $a^k - b^k = (a - b)(a^{k-1} + \cdots + b^{k-1})$ lives.

Put the two observations together and the plan writes itself within the first thirty seconds: form $I_{n+1} - I_n$, pull out a common power of f/g, insert the zero you were handed, and watch a difference of equal powers of f and g appear, ready to be factored down to $(f - g)^2$ times something manifestly positive.

Approaches, ranked worst to best

1. **attempt a closed form or a direct size estimate for I_n.** *Cue:* the problem hands you an explicit integral formula. *Circuit:* substitution or crude bounding of f^{n+1}/g^n. (there is no closed form for arbitrary continuous f, g, and this approach never engages with the two hypotheses at all, so it stalls immediately with nothing to say about monotonicity or growth.)

2. **Cauchy-Schwarz on I_n, I_{n+1}, I_{n+2} to get log-convexity.** *Cue:* a sequence indexed by an exponent, resembling a moment sequence. *Circuit:* $I_n I_{n+2} \geq I_{n+1}^2$, then a ratio test using an auxiliary term $I_0 = \int f$. (this works and is a standard powerful circuit, but it needs the extra term I_0, a separate small argument that $I_1 > I_0$, and a ratio telescoping step; more machinery than is needed here.)

3. **insert the zero integral into the difference $I_{n+1} - I_n$ and factor.** *Cue:* $\int (f - g) = 0$ is an unused hypothesis, and the sequence's exponent structure invites a difference of equal powers. *Circuit:* the identity $a^k - b^k = (a - b)(a^{k-1} + \cdots + b^{k-1})$ applied to $f^{n+1} - g^{n+1}$. (best: handles both the increasing claim and, after one more short continuity argument, the divergence claim, with no auxiliary sequence and no inequality beyond positivity of an integral.)

The ideal solution

Step 1: the sequence is increasing. Fix $n \geq 1$. Since $g > 0$ throughout $[a, b]$,

$$I_{n+1} - I_n = \int_a^b \left(\frac{f^{n+2}}{g^{n+1}} - \frac{f^{n+1}}{g^n} \right) dx = \int_a^b \frac{f^{n+1}}{g^{n+1}} (f - g)\,dx.$$

Because $\int_a^b (f-g)\,dx = 0$, subtracting this (zero) quantity changes nothing:

$$I_{n+1} - I_n = \int_a^b \left(\frac{f^{n+1}}{g^{n+1}} - 1\right)(f-g)\,dx = \int_a^b \frac{f^{n+1}-g^{n+1}}{g^{n+1}}(f-g)\,dx.$$

Using $f^{n+1} - g^{n+1} = (f-g)\sum_{k=0}^{n} f^k g^{n-k}$,

$$I_{n+1} - I_n = \int_a^b \frac{(f-g)^2 \sum_{k=0}^{n} f^k g^{n-k}}{g^{n+1}}\,dx.$$

The integrand is continuous (as f, g are continuous and $g > 0$) and nonnegative everywhere, since it is a square times a sum of positive terms divided by a positive quantity. It is not identically zero: as $f \neq g$ on $[a,b]$, some x_0 has $f(x_0) \neq g(x_0)$, so the integrand is strictly positive at x_0, hence (by continuity) on a subinterval of positive length around x_0. A continuous, nonnegative function on $[a,b]$ that is positive on a subinterval of positive length has a strictly positive integral. Hence $I_{n+1} > I_n$ for every $n \geq 1$, so $I_1 < I_2 < I_3 < \cdots$.

Step 2: the sequence diverges. First, it cannot be that $f \leq g$ throughout $[a,b]$: if it were, $h = g - f$ would be continuous, nonnegative, with $\int_a^b h\,dx = \int g - \int f = 0$; by the same positivity fact used above, this forces $h \equiv 0$, contradicting $f \neq g$. So there is some $t \in [a,b]$ with $f(t) > g(t)$. Fix c with $1 < c < f(t)/g(t)$; since f/g is continuous, there is a nondegenerate closed subinterval $J \subseteq [a,b]$ containing t on which $f(x) \geq c\,g(x)$ for all $x \in J$, hence also $f(x)^n \geq c^n g(x)^n$ on J for every positive integer n. Then, discarding the rest of $[a,b]$ (where the integrand is still positive, so only makes the bound weaker),

$$I_n = \int_a^b \frac{f^{n+1}}{g^n}\,dx \geq \int_J \frac{f^{n+1}}{g^n}\,dx = \int_J f\left(\frac{f}{g}\right)^n dx \geq c^n \int_J f\,dx.$$

Since f is continuous and strictly positive on the nondegenerate interval J, $d := \int_J f\,dx$ is a fixed positive number, independent of n. As $c > 1$, $c^n \to \infty$, so

$$I_n \geq c^n d \longrightarrow \infty \qquad \text{as } n \to \infty.$$

This proves both required claims:

$$\boxed{I_1 < I_2 < I_3 < \cdots \quad \text{and} \quad \lim_{n\to\infty} I_n = \infty.}$$

Layer 1: WHAT

Block 1: notice $\int (f-g)\,dx = 0$ is an unused hypothesis waiting to be inserted into a constructed difference. Block 2: form $I_{n+1} - I_n$ and pull out the common factor f^{n+1}/g^{n+1}, reducing to $\int (f^{n+1}/g^{n+1})(f-g)\,dx$. Block 3: insert the zero to rewrite the bracket as $f^{n+1}/g^{n+1} - 1$, then factor $f^{n+1} - g^{n+1}$ to expose $(f-g)^2$ times a manifestly positive sum. Block 4: invoke "continuous, nonnegative, not identically zero implies strictly positive integral" to conclude $I_{n+1} > I_n$. Block 5: reuse the same positivity fact in a proof by contradiction to rule out $f \leq g$ everywhere, producing a point, then (by continuity) an interval J, where $f \geq cg$ for some $c > 1$. Block 6: bound I_n below by discarding everything outside J, giving $I_n \geq c^n d \to \infty$.

The nominal route never leaves Block 1: it tries to bound or estimate I_n directly, by size comparison or a blind inequality, and never builds an expression containing $(f-g)$, so the zero-integral hypothesis sits unused throughout. The bifurcation point is the decision to compute a difference of consecutive terms specifically because the hypothesis is stated as an equality of two integrals, rather than working with a single I_n in isolation.

Layer 2: HOW

Block 1. **Cue.** The hypothesis is phrased as two integrals being equal. **Why natural.** This is algebraically the same as saying the integral of the difference is zero, a fact that can be inserted into any later expression without changing its value. **Link to the next block.** Holding this fact in reserve, the next task is to build an expression that actually contains $(f - g)$.

Block 2. **Cue.** The sequence is indexed by an exponent, and consecutive terms I_n, I_{n+1} share the shape $f^\bullet/g^\bullet$ one power apart. **Why natural.** Subtracting consecutive terms and factoring out the shared power turns the comparison into a single integral of $(f^{n+1}/g^{n+1})(f - g)$, rather than two separate unrelated integrals. **Link to the next block.** This integral has no obvious sign yet; the reserved zero from Block 1 is what will fix that.

Block 3. **Cue.** The bracket f^{n+1}/g^{n+1} has no fixed sign relative to 1 at a glance. **Why natural.** Subtracting the reserved zero converts it to $f^{n+1}/g^{n+1} - 1$, and the numerator $f^{n+1} - g^{n+1}$ always factors as $(f - g)$ times a sum of positive terms, exposing $(f - g)^2$ once multiplied against the outer $(f - g)$. **Link to the next block.** A square is now sitting in the integrand; only its non-identical-vanishing needs to be argued.

Block 4. **Cue.** The integrand is now visibly a square times positive quantities, divided by a positive quantity. **Why natural.** $f \neq g$ guarantees the square is positive somewhere, and continuity upgrades "somewhere" to "on a subinterval", which is exactly what a nonzero integral requires. **Link to the next block.** Monotonicity is settled; divergence still needs a genuinely new idea, since a merely increasing sequence need not be unbounded.

Block 5. **Cue.** Divergence needs more than "somewhere positive"; it needs growth that compounds with n. **Why natural.** Ruling out $f \leq g$ everywhere (by the same positivity fact, in a proof by contradiction) produces a single point where $f > g$, and continuity of f/g upgrades that point to an entire interval J with a fixed ratio bound $c > 1$. **Link to the next block.** An interval with a uniform ratio bound is precisely what lets the exponent n do compounding work.

Block 6. **Cue.** On J, every factor of f/g in the integrand is at least $c > 1$. **Why natural.** Discarding the rest of the domain only weakens a lower bound (the integrand is nonnegative there), so I_n is bounded below by c^n times the fixed constant $\int_J f$. **Link to the next block.** Since $c > 1$, this lower bound tends to infinity, finishing the proof.

Layer 3: WHY

The route is ideal because a hypothesis stated as "two integrals are equal" is never really about equality, it is a licence to insert a specific zero into any expression you build afterward, and the only expressions worth building are ones that already contain $(f - g)$, since that is the only place the licence can be spent. Once you commit to constructing a difference of consecutive sequence terms (forced by the fact that you are asked about monotonicity in the first place), the exponents guarantee a difference of equal powers of f and g appears, and the factorisation of $a^k - b^k$ is the unique elementary tool that turns such a difference into a signed quantity. The discernment mechanism is to treat $\int(f - g) = 0$ as inert until an $(f - g)$ literally appears in your working, and to actively engineer its appearance by subtracting consecutive terms rather than analysing a single term.

The divergence half reuses the very same real analysis fact, positivity of the integral of a nonzero nonnegative continuous function, but in the contrapositive direction: instead of using it to sign an integral, it is used to rule out a global inequality ($f \leq g$ everywhere), which is what manufactures the single strict inequality point that continuity then inflates into an interval. Training this discernment

means recognising that this one lemma about continuous nonnegative functions is doing all the real work in the problem twice over, and that whenever a "differs somewhere" hypothesis needs to become "differs substantially, on a set of positive length", continuity of the relevant ratio or difference is the only tool that makes the upgrade.

Revolving concepts

Positivity of the integral of a nonnegative continuous function. If h is continuous and nonnegative on $[a, b]$ with $h(t_0) > 0$ for some t_0, then $h > \frac{1}{2}h(t_0)$ on a subinterval of positive length around t_0, so $\int_a^b h\,dx > 0$; equivalently, a nonnegative continuous function with zero integral is identically zero. See Book 1, field: Real analysis, topic: positivity of integrals.

Difference of like powers. For real a, b and integer $k \geq 1$, $a^k - b^k = (a-b)(a^{k-1} + a^{k-2}b + \cdots + b^{k-1})$. See Book 1, field: Polynomials, topic: factorisation identities. Generalised recursion: whenever a difference of two equal, high powers of two quantities appears or can be engineered to appear, this identity converts it into (the base difference) times (a manifestly signed sum when the quantities share a sign), the standard route to manufacturing a squared, hence signed, quantity from a bare difference.

Continuity upgrades a point to an interval. If ϕ is continuous at t and $\phi(t) > c$, then $\phi(x) > c$ throughout some subinterval around t (one sided if t is an endpoint), since nearby values cannot jump away from $\phi(t)$. See Book 1, field: Real analysis, topic: local behaviour of continuous functions.

Cauchy-Schwarz and log-convexity of parametrised integrals. For positive integrands, Cauchy-Schwarz applied to $\sqrt{f^{n+1}/g^n} \cdot \sqrt{f^{n+3}/g^{n+2}} = f^{n+2}/g^{n+1}$ gives $I_n I_{n+2} \geq I_{n+1}^2$, so $\log I_n$ is convex in n and the ratios I_{n+1}/I_n are nondecreasing; this is the mechanism behind the alternate solution and a standard route to monotonicity or growth results for any sequence of integrals indexed by an exponent. See Book 1, field: Inequalities, topic: log-convexity via Cauchy-Schwarz.

Circuit Bank

- **WHEN** a hypothesis states two integrals (or sums) are equal **DO** immediately rewrite it as "the integral (or sum) of the difference is zero" and hold it in reserve **BECAUSE** that reserved zero can be inserted into any later expression built from the same difference without changing its value, which is otherwise the only use such a hypothesis has

- **WHEN** a sequence of integrals is indexed by an exponent and monotonicity is asked for **DO** compute the difference of consecutive terms and factor out the shared power **BECAUSE** consecutive terms of an exponent indexed sequence differ by exactly one extra factor, which is precisely the setting where a difference of equal powers, and hence the factorisation $a^k - b^k = (a-b)(\cdots)$, becomes available

- **WHEN** a difference of two equal powers of positive quantities appears **DO** factor it as $(a-b)$ times a sum of positive cross terms **BECAUSE** this is the one identity that turns a bare power difference into a manifestly signed quantity, and combined with an outer $(a-b)$ it produces a square, the cheapest possible guaranteed sign

- **WHEN** a growth or divergence claim needs more than a single point of strict inequality **DO** upgrade the point to a subinterval via continuity, then bound the integral below using only that subinterval **BECAUSE** a single point contributes nothing to an integral, but continuity guarantees the strict inequality persists on a set of positive length, which is enough to build a genuine geometric lower bound

Backtesting

1. First useful cue: the hypothesis states two integrals are equal, an unused zero waiting to be inserted somewhere.
2. Method the expression suggested: form the difference of consecutive terms of the exponent indexed sequence and factor out the shared power of f/g.
3. Where you would hesitate: after reaching $\int (f^{n+1}/g^{n+1})(f-g)\,dx$, unsure how to pin down a sign; reinforce that inserting the reserved zero and factoring $f^{n+1} - g^{n+1}$ always exposes a square.
4. The trap the question sets: trying to bound I_n directly by size comparisons or a blind inequality, which leaves the zero integral hypothesis completely unused and the monotonicity claim unreachable.
5. What makes you recognise it faster next time: seeing "two integrals equal, functions not identical" as a fixed signature that wants a difference built and a zero inserted, with the continuity upgrade reused whenever a single point of strict inequality needs to become an interval.
6. The reusable decision rule: whenever $\int (f-g) = 0$ is given, build the next expression as a difference containing $(f-g)$, insert the zero, factor, and reach for continuity the moment a single point of strict inequality needs to grow into an interval.

Mirror drills

1. Let h be continuous and nonnegative on $[a,b]$ with $h(t_0) > 0$ for some $t_0 \in [a,b]$. Show $\int_a^b h\,dx > 0$, and deduce that a nonnegative continuous function with zero integral is identically zero. *Answer:* by continuity $h > \frac{1}{2}h(t_0)$ on a subinterval of positive length around t_0, so $\int h$ exceeds $\frac{1}{2}h(t_0)$ times that length; the second claim is the contrapositive.
2. With positive reals $a_1, \ldots, a_m$ and $b_1, \ldots, b_m$ satisfying $\sum a_i = \sum b_i$ but $(a_i) \neq (b_i)$, show $\sum_i a_i^{n+1}/b_i^n$ is strictly increasing in n. *Answer:* the identical computation with sums replacing integrals: the difference of consecutive terms telescopes to $\sum_i (a_i - b_i)^2 (\text{positive terms})/b_i^{n+1}$, positive since not all $a_i = b_i$.
3. Under the hypotheses of the main problem, show directly that $\int_a^b (f-g)^2\,dx > 0$. *Answer:* $f \neq g$ gives a point of strict inequality, continuity upgrades it to a subinterval, and a nonnegative continuous function positive on a subinterval has positive integral.

9.4 A4. A sliding window splits the class

A class with $2N$ students took a quiz, on which the possible scores were $0, 1, \ldots, 10$. Each of these scores occurred at least once, and the average score was exactly 7.4. Show that the class can be divided into two groups of N students in such a way that the average score for each group was exactly 7.4.

Read it out loud

Say the structure before you say the numbers. You are handed a multiset of $2N$ integers drawn from the block $\{0, 1, \ldots, 10\}$, told that every value in that block is actually used, and told that the total is pinned at $7.4 \cdot 2N$. You are asked to cut that multiset into two halves of equal size and equal sum. So this is an existence statement about an exact value, and existence statements about exact values are almost never proved by building the object. They are proved by moving a parameter and catching the target as it goes past.

The fields wired into it: a scrap of elementary number theory (the average is not an integer, so it is quietly doing arithmetic work on N), order combinatorics (sorting is the only structure a bare multiset has), and the discrete intermediate value theorem, which is what is really being tested.

Within thirty seconds three lines should be on your paper. First, the total is $7.4 \cdot 2N = \frac{74N}{5}$, and it is a sum of integers, so 5 divides N and the per group target $7.4N = \frac{37N}{5}$ is itself an integer. Second, it is enough to find N of the scores summing to $7.4N$, because the complementary N then sum to $14.8N - 7.4N = 7.4N$ for free; that halves the problem and costs nothing. Third, sort.

Now look at the hypothesis that reads like scene setting: each of the scores 0 to 10 occurred at least once. That is not scene setting, it is the engine. Take thirteen students scoring 0 and thirty seven scoring 10. That is 50 students with total 370, average exactly 7.4, and no valid split exists at all, because every group sum is then a multiple of 10 while the target is 185. The fact you must use is that the set of achieved values is a block of consecutive integers. Any proof of yours that never uses it is wrong, and you should hunt for the error rather than celebrate.

Approaches, ranked worst to best

1. **Classify by the eleven multiplicities.** *Cue:* the scores only run from 0 to 10, so "there are finitely many shapes". *Circuit:* case analysis on the score distribution. You are staring at an eleven dimensional lattice of multiplicity vectors cut by one linear equation, with N unbounded. There is no finite check to run. Induction on N dies for the same reason: deleting two students can destroy "every score occurs" and it moves the target, so the hypothesis does not survive the induction step.

2. **Generating function coefficient extraction.** *Cue:* "choose a subset with a prescribed sum". *Circuit:* read off $[y^N x^{7.4N}] \prod_j (1 + yx^{a_j})$. To finish you must prove that one coefficient is positive, and there is no positivity mechanism available here: no unimodality you can invoke, no symmetry that forces a nonzero term. The generating function restates the problem in heavier notation and hands you nothing.

3. **Reserve one student of each score, then fine tune.** *Cue:* "each of these scores occurred at least once" reads as "you have been given a full set of small change". *Circuit:* park the eleven compulsory students, run a window on the rest, repair the remainder with a small subset of the parked ones. This works, and it is essentially the second official route. The cost is an extra lemma (every integer from 15 to 40 is a sum of five distinct elements of $\{1, \ldots, 10\}$) plus a page of index bookkeeping. More ink, same conclusion.

4. **Exchange argument.** *Cue:* "two groups, equal sums". *Circuit:* descend on the integer potential $|\Sigma A - \Sigma B|$. Start from any split into halves A and B. The potential is a nonnegative integer

with the same parity as the total, hence even. If it is positive, then $\Sigma A > \Sigma B$ forces some $a \in A$ to exceed some $b \in B$; take such a pair with $a - b$ minimal, and the consecutive block hypothesis forces $a - b = 1$, since otherwise the value $a - 1$ occurs somewhere in the class and produces a strictly smaller pair. Swap that pair and the potential drops by exactly 2. Descent in $\mathbb{N}$ terminates, at 0. This is a genuinely good solution and it generalises well; it sits below the winner only because the minimal pair step is a second argument bolted on.

5. **Sliding window plus discrete intermediate value theorem.** *Cue:* exact target, and the only structure available is order. *Circuit:* parametrise a family of candidate groups that changes by one student at a time, and catch the target on the way past. Sort, slide a window of N consecutive scores along the sorted list, and observe that the two extreme windows average to the target. Monotonicity does the searching. One inequality and one deletion finish it.

The ideal solution

Write $a_1 \le a_2 \le \cdots \le a_{2N}$ for the scores in nondecreasing order, so every a_j lies in $\{0, 1, \ldots, 10\}$ and every element of $\{0, 1, \ldots, 10\}$ occurs among them. Put $\Sigma = \sum_{j=1}^{2N} a_j = 7.4 \cdot 2N = \frac{74N}{5}$.

Step 0 (integrality, and halving the goal). Σ is a sum of integers, so 5 divides $74N$ and hence 5 divides N. Therefore

$$T := 7.4N = \frac{37N}{5} \in \mathbb{Z}, \qquad \Sigma = 2T.$$

It is enough to produce N of the scores with sum exactly T: the remaining N scores then sum to $\Sigma - T = T$, and each group has average $T/N = 7.4$.

Step 1 (the window family). For $0 \le i \le N$ let

$$s_i = \sum_{j=i+1}^{i+N} a_j,$$

the sum of the N consecutive terms starting at position $i + 1$. Then

$$s_{i+1} - s_i = a_{i+N+1} - a_{i+1} \ge 0$$

by sortedness, so $s_0 \le s_1 \le \cdots \le s_N$. Also $s_0 + s_N = \Sigma = 2T$, since the two extreme windows are the bottom half and the top half. Hence $s_0 \le T \le s_N$.

Step 2 (the ends are distinct). The scores are not all equal, because 0 and 10 both occur. Suppose $s_0 = s_N$. Then every s_i is equal, so $a_{j+N} = a_j$ for $1 \le j \le N$. In particular $a_1 = a_{N+1} \ge a_N \ge a_1$, so $a_1 = a_2 = \cdots = a_{N+1}$, and then $a_{N+j} = a_j = a_1$ for $1 \le j \le N$ makes all $2N$ scores equal. Contradiction, so $s_0 < s_N$.

Step 3 (discrete intermediate value theorem). Let i be the largest index in $\{0, 1, \ldots, N\}$ with $s_i \le T$; it exists because $s_0 \le T$. If $i = N$ then $s_N \le T \le s_N$ gives $s_N = T$, whence $s_0 = 2T - s_N = T = s_N$, contradicting Step 2. So $i \le N - 1$, and by maximality

$$s_i \le T < s_{i+1} = s_i + a_{i+N+1} - a_{i+1}.$$

Step 4 (the value that must be deleted). Consider the $N + 1$ consecutive terms $a_{i+1}, \ldots, a_{i+N+1}$, whose total is $s_i + a_{i+N+1}$. If we delete one term of value v we are left with N terms of total T precisely when

$$v = s_i + a_{i+N+1} - T,$$

which is an integer by Step 0. The two inequalities of Step 3 now pin it down: $s_i \le T$ gives $v \le a_{i+N+1}$, and $T < s_i + a_{i+N+1} - a_{i+1}$ gives $v > a_{i+1}$. Thus

$$a_{i+1} < v \le a_{i+N+1}.$$

Step 5 (the block lemma).

Lemma 9.4.1. *For any $1 \le p \le q \le 2N$, every integer w with $a_p \le w \le a_q$ occurs among $a_p, a_{p+1}, \dots, a_q$.*

Proof. Since $0 \le a_p \le w \le a_q \le 10$, the integer w lies in $\{0, 1, \dots, 10\}$ and therefore occurs in the class, say $w = a_t$. If $t < p$ then $a_t \le a_p \le w = a_t$, so $a_p = w$ and w occurs at position p. If $t > q$ then symmetrically $a_q = w$. Otherwise $p \le t \le q$ already. In every case w occurs in the run. □

This is the only place the hypothesis "each score occurred at least once" is used, and it is exactly the place where it is indispensable.

Step 6 (finish). Apply the lemma with $p = i + 1$ and $q = i + N + 1$: by Step 4 the integer v satisfies $a_{i+1} \le v \le a_{i+N+1}$, so $v = a_t$ for some t with $i + 1 \le t \le i + N + 1$. Let the first group consist of the N students at positions $\{i + 1, \dots, i + N + 1\} \setminus \{t\}$; their total is

$$(s_i + a_{i+N+1}) - v = T.$$

The remaining N students have total $\Sigma - T = T$. Both groups therefore have average $\boxed{7.4}$, as required.

Layer 1: WHAT

Block 1. Clear the decimal: $\Sigma = \frac{74N}{5}$ is an integer, so $5 \mid N$ and $T = 7.4N$ is an integer.

Block 2. Halve the goal: find one group of N with sum T; the complement is then automatic.

Block 3. Sort, because the only structure a multiset carries is order, and the special hypothesis is a statement about the sorted list having no gaps.

Block 4. Build the window sums $s_0 \le \cdots \le s_N$. Their two ends add to $2T$, so the target is trapped between them.

Block 5. Take the last window at or below T. The step to the next window is a single swap of size $a_{i+N+1} - a_{i+1}$, so the shortfall $T - s_i$ is smaller than that swap.

Block 6. Widen the window to $N + 1$ terms and ask which single value has to be deleted. It is forced: $v = s_i + a_{i+N+1} - T$, and the two inequalities put v strictly inside the range of the run.

Block 7. The block lemma says the run realises every integer between its endpoints, so v is actually there. Delete it and finish.

Juxtaposition. The nominal route fixes the target and then tries to build a group that hits it, greedily or by case analysis on the multiplicities, or attempts induction on N. Every version of that spends its effort on the wrong object. The bifurcation point is Block 4: the zombie chooses *a subset*, the ideal route chooses *a one parameter family of subsets* and lets monotonicity do the choosing. Everything after Block 4 is bookkeeping.

Layer 2: HOW

Block 1. **Cue.** An average stated as 7.4, a non-integer, attached to an integer valued quantity. **Why natural.** A non-integer average is never decoration. Multiply out at once: $\Sigma = \frac{74N}{5}$ must be an integer, so $5 \mid N$. That single line is what later makes v an integer, and without it Step 4 produces a fractional value and the whole proof collapses. **Link to the next block.** Knowing T is an integer, you can state the goal purely in integers.

Block 2. **Cue.** "Divided into two groups of N ... average for each group". **Why natural.** Two groups of equal size partitioning a fixed total is one condition, not two: fix one half and the other is determined. Always strip a partition condition down to a single subset condition before you start work. **Link to the next block.** You now need one subset of size N with a prescribed sum, which is a search over a structured family.

Block 3. **Cue.** A multiset of integers with a hypothesis about *which values occur*. **Why natural.** Sorting is free and it converts "which values occur" into "the sorted list has no gaps", which is a usable local statement about consecutive terms. Sorting also makes any window sum monotone in its starting index, which is the property the whole proof runs on. **Link to the next block.** With the list sorted, a family of subsets indexed by one integer becomes available.

Block 4. **Cue.** You need one subset of size N hitting an exact value, and you have a total order. **Why natural.** Take the windows of N consecutive terms. Consecutive windows differ by exactly one swap, so the sums move monotonically and in controlled steps; and the first and last windows are the bottom half and the top half, whose sums add to $\Sigma = 2T$. Hence $s_0 \le T \le s_N$ with no computation at all. This is the discrete intermediate value theorem being set up. **Link to the next block.** The target lies between the two ends of a monotone integer sequence, so some step straddles it.

Block 5. **Cue.** A monotone sequence crossing a target. **Why natural.** Choose the last index at or below the target. That is the only choice that gives you a two sided bound: $s_i \le T$ from the definition, $T < s_{i+1}$ from maximality. Also check the degenerate end, $i = N$, since a proof that silently assumes $i < N$ is incomplete; the symmetry $s_0 + s_N = 2T$ kills that case in one line. **Link to the next block.** You now have $0 \le T - s_i < a_{i+N+1} - a_{i+1}$, which is a bound of exactly one swap.

Block 6. **Cue.** A shortfall smaller than one swap. **Why natural.** A window of N terms cannot be adjusted; a window of $N + 1$ terms with one term deleted can. Widening by one is the cheapest possible extra degree of freedom, and the deleted value is then forced rather than chosen. The two inequalities translate directly into $a_{i+1} < v \le a_{i+N+1}$, which is the statement that the required value lies inside the range of the run. **Link to the next block.** It remains to know that the run contains every integer in its own range.

Block 7. **Cue.** "Each of these scores occurred at least once", not yet used. **Why natural.** The achieved values form the block $\{0, 1, \ldots, 10\}$, so a value missing from a run would have to be missing from the whole class, and it is not. Sandwiching the position of an occurrence against the ends of the run finishes it. This is where the decorative looking hypothesis pays for itself. **Link to the next block.** Delete one occurrence of v and the two groups are built.

Layer 3: WHY

Three mathematical truths sit underneath this route.

The first is the discrete intermediate value theorem: if $s_0, \ldots, s_m$ are integers with $s_0 \le T \le s_m$ and $|s_{i+1} - s_i| \le d$ for all i, then some i has $s_i \le T < s_{i+1}$ and hence $0 \le T - s_i < d$. You cannot hope for equality in general, only for "within one step". That is the correct expectation to carry into the arena, and it tells you in advance that you will need one further degree of freedom to close the last gap.

The second is where that degree of freedom comes from. Enlarging the window from N to $N + 1$ terms turns "the window sum is close to T" into "the window sum minus a deletable term equals T", and the value to be deleted is then determined, not chosen. What remains is a pure availability question, and availability is exactly what a block of consecutive values guarantees. The setters supplied that block and gave it to you as a sentence about a quiz. The counterexample with only the scores 0 and 10 shows that no proof can avoid it.

The third is the cheapest of the three and the most often forgotten: an equal split of a fixed total is a single condition. Half the candidates in the arena will carry two symbols ΣA and ΣB through the whole proof.

Discernment mechanism. Ask three questions in order. Is the required object exact rather than approximate? Then do not construct it, parametrise a family. Does the family move in bounded steps? Then a discrete intermediate value theorem gives you "within one step". Is there a hypothesis in the question that supplies precisely one step of slack? If yes, that hypothesis is what the last line of your proof will consume.

How to train it. After every problem you solve, list the hypotheses and mark the line of your proof that uses each one. An unused hypothesis is either a signal that your proof is wrong or a signal that you have proved a stronger theorem, and you must know which. Second habit: whenever you sort, immediately ask what the gaps in the sorted list can be, because "monotone with gaps at most 1" is a far stronger object than "monotone".

Revolving concepts

Discrete intermediate value theorem. If integers $s_0, \dots, s_m$ satisfy $s_0 \le T \le s_m$ and $s_{i+1} - s_i \le d$ for all i, then there is an i with $0 \le T - s_i < d$; if moreover every step is 0 or 1, then some $s_i = T$ exactly. This last case is worth memorising: unit steps give exact landing, and much of combinatorics is arranging for unit steps. See Book 1, field: Combinatorics, topic: discrete intermediate value theorem.

Sliding windows and the general swap recursion. The family s_i obeys the first order recursion $s_{i+1} = s_i + a_{i+N+1} - a_{i+1}$, whose increments are single element swaps. Generalised form: if $X_0, X_1, \dots, X_m$ are finite sets with $|X_{i+1} \triangle X_i| = 2$ (one element out, one element in) and f is a sum over the set, then $f(X_{i+1}) - f(X_i)$ is a difference of two element values, so the sequence $f(X_i)$ has steps bounded by the spread of the values. Any such family plus a sign change at the ends yields a near solution, and the same recursion is the engine behind necklace splitting arguments and behind the standard proof that a cyclic sequence of partial sums must cross its own mean. See Book 1, field: Combinatorics, topic: exchange and window families.

Blocks of consecutive values. If the value set of a finite multiset of integers is a block of consecutive integers, then every run of its sorted list realises every integer between the endpoints of that run. Generalisation: if the value set is contained in an arithmetic progression of common difference d and contains every term of that progression between its extremes, the same statement holds for values congruent to a_p modulo d. This is the "no gaps" property, and it is the standard partner of any window argument that needs to land exactly.

Descent on an integer potential. The exchange approach uses that $|\Sigma A - \Sigma B|$ is a nonnegative integer decreasing by exactly 2 at each swap; a strictly decreasing sequence in $\mathbb{N}$ terminates. General recursion: any local search with a potential Φ : states $\to \mathbb{N}$ that strictly decreases at every legal move terminates, and the terminal state satisfies the local optimality condition you want. Parity is what tells you the terminal value is 0 rather than 1: here $\Sigma A - \Sigma B \equiv \Sigma A + \Sigma B \pmod 2$ and the total is even. See Book 1, field: Combinatorics, topic: monovariants and descent.

Integrality from a rational average. An average of p/q in lowest terms over n integers forces $q \mid n$. Here $7.4 = 37/5$ over $2N$ students gives $5 \mid 2N$, hence $5 \mid N$. See Book 1, field: Number theory, topic: divisibility from means.

Circuit Bank

- **WHEN** an average or mean given as a non-integer **DO** clear the denominator immediately and read off the divisibility it forces **BECAUSE** the exact value of the target is what makes an integer argument possible, and the divisibility is usually load bearing later
- **WHEN** a partition of a fixed total into two parts with a prescribed property **DO** throw away one part and impose the condition on a single subset **BECAUSE** the complement is determined, so two symbols where one will do is pure cost

- **WHEN** an exact target and a total order on the data **DO** build a one parameter family of candidate sets that changes by one element at a time, and apply the discrete intermediate value theorem **BECAUSE** a monotone integer sequence with bounded steps crossing a target must come within one step of it
- **WHEN** a shortfall smaller than one step **DO** enlarge the candidate set by one element and delete the forced value instead **BECAUSE** widening by one converts a near miss into an exact hit, and the value to delete becomes determined rather than chosen
- **WHEN** a hypothesis that looks like scene setting **DO** find the single line of your proof that fails without it, or construct the counterexample **BECAUSE** setters do not plant free hypotheses, and an unused hypothesis usually means a broken proof

Backtesting

1. First useful cue: the non-integer average 7.4 attached to integer scores. It forces $5 \mid N$ and an integer target, and it is also the reason no student can score the average.
2. Method the expression suggested: sort and slide a window of N consecutive scores, because an exact target plus a total order is the standard trigger for a discrete intermediate value argument.
3. Where you would hesitate, and the thought to reinforce: at the moment the window sum lands strictly below the target, you feel stuck because a window of N terms cannot be nudged. The thought to install is "take $N+1$ terms and delete one", which is the cheapest available extra degree of freedom.
4. The trap the question sets: treating "each of these scores occurred at least once" as flavour. Without it the statement is false, as the class of thirteen zeros and thirty seven tens shows.
5. What makes you recognise it faster next time: the pairing of an exactness demand with a hypothesis that the achieved values form a consecutive block. That pairing is the signature of a window argument that must land exactly.
6. The reusable decision rule: for an exact subset sum, parametrise rather than construct, get within one step by monotonicity, and spend the "no gaps" hypothesis on the final step.

Mirror drills

1. Let $2N$ integers lie in $\{0, 1, \ldots, 10\}$, let every one of those eleven values occur, and let the total be even. Show the multiset splits into two halves of equal sum. *Answer:* identical proof with $T = \Sigma/2$; nothing used 7.4 beyond the integrality of T, which is now a hypothesis.
2. Let $a_1 \leq \cdots \leq a_{2N}$ be integers in $\{0, 1, \ldots, d\}$. Show that some N of them have sum within $d/2$ of half the total. *Answer:* window sums have steps at most d and satisfy $s_0 + s_N = \Sigma$, so the last s_i at or below $\Sigma/2$ is within d of it; the better constant follows by comparing s_i and s_{i+1} and taking the closer one.
3. Find a multiset of four integers with even total that cannot be split into two pairs of equal sum, and identify which hypothesis of the problem it violates. *Answer:* $\{1, 1, 1, 3\}$ has total 6 but pair sums only 2 and 4; its value set $\{1, 3\}$ is not a block of consecutive integers.

9.5 A5. Three card players and a rotating lead

Each of the integers from 1 *to* n *is written on a separate card, and then the cards are combined into a deck and shuffled. Three players,* A*,* B*, and* C*, take turns in the order* $A, B, C, A, \ldots$ *choosing one card at random from the deck. (Each card in the deck is equally likely to be chosen.) After a card is chosen, that card and all higher-numbered cards are removed from the deck, and the remaining cards are reshuffled before the next turn. Play continues until one of the three players wins the game by drawing the card numbered* 1*.*

Show that for each of the three players, there are arbitrarily large values of n *for which that player has the highest probability among the three players of winning the game.*

Read it out loud

Read the process, not the players. A deck of n cards; someone draws card k uniformly; the deck becomes $\{1, \ldots, k-1\}$. So the state of the game is a single integer that strictly decreases, the transition is "replace n by a uniform element of $\{0, 1, \ldots, n-1\}$", and the game ends when the state hits 0. The three players are not really three objects: they are one object, namely the number of draws taken so far, read modulo 3. That is the sentence to say out loud, and it is the whole problem.

The fields wired in: probability (first step analysis on a self similar process), generating functions (the number of draws is the only statistic that matters, so mark it with a variable), complex numbers (a residue modulo 3 is extracted by a roots of unity filter), and a scrap of real analysis at the end (a divergent series with vanishing terms).

Notice what the conclusion is really asking. It is not asking who wins. It is asking for infinitely many sign changes in a three way comparison, that is, for a quantity that keeps rotating. The moment you read "arbitrarily large values of n for which that player has the highest probability", you should expect an angle that grows without bound in steps that shrink to zero, because that is the only cheap mechanism that produces infinitely many rotations.

Within thirty seconds you should have written the first step recursion with the roles rotated,

$$a_n = \frac{1}{n}\sum_{k=1}^{n} c_{k-1}, \quad b_n = \frac{1}{n}\sum_{k=1}^{n} a_{k-1}, \quad c_n = \frac{1}{n}\sum_{k=1}^{n} b_{k-1},$$

with $a_0 = b_0 = 0$ and $c_0 = 1$, where a_n, b_n, c_n are the winning probabilities of the player to move, the next player and the third player. The convention $c_0 = 1$ says that when the deck empties the player who has just drawn (the third player in the rotated labelling) has won. Do not then start computing terms. Look at what you have: three coupled sequences with a cyclic symmetry, each an average of the previous ones.

Approaches, ranked worst to best

1. **Compute and observe.** *Cue:* "show there are arbitrarily large n". *Circuit:* tabulate a_n, b_n, c_n and watch the lead change. This is not a proof, and here it is actively misleading. The lead changes at $n = 3$, then 14, then 145, then 1608: each handover needs n to be multiplied by roughly 11, so a table up to $n = 100$ shows two handovers and then a long flat stretch that looks like convergence. A candidate who trusts the table concludes the wrong thing.

2. **Direct estimates on the recursion.** *Cue:* three coupled sequences. *Circuit:* bound $a_n - b_n$ by hand and hunt for sign changes. You are trying to control a rotation with inequalities in one real variable at a time. Every bound you prove is a projection of a two dimensional spiral onto a line, so the estimates fight each other and nothing closes.

3. **Order type state machine.** *Cue:* "at most one of the three differences vanishes". *Circuit:* define the six possible orderings of (a_n, b_n, c_n) and prove the transitions cycle. This is the third official

route, and it works with no complex numbers at all. It costs several preliminary lemmas (that none of the three partial sums can diverge or converge to a nonzero limit) and roughly two pages. Fully elementary, fully rigorous, expensive.

4. **Normalise the difference vector and rotate it.** *Cue:* $x_n + y_n + z_n = 0$ for the three differences. *Circuit:* project onto a circle and show the angle increments are a null sequence with divergent sum. This is the first official route and it is the right picture. It reaches the answer by inventing the circle by hand rather than by letting algebra hand it to you.

5. **Generating function in the number of draws, evaluated at a cube root of unity.** *Cue:* the winner is determined by the number of draws modulo 3. *Circuit:* mark the number of draws with X, telescope the averaging recursion, then apply a roots of unity filter. The recursion for the generating function collapses to $P_n = \frac{X+n-1}{n} P_{n-1}$, so $P_n(X) = X(X+1)\cdots(X+n-1)/n!$ in one line, and substituting a primitive cube root of unity turns the whole triple into a single complex number whose argument is a visible sum of arctangents. The rotation is then not a picture to be invented but a product of complex numbers to be read off.

The ideal solution

Let M_n be the number of draws in the game played with a deck of n cards, and let

$$P_n(X) = \mathbb{E}\left[X^{M_n}\right] = \sum_{m \geq 0} \Pr[M_n = m]\, X^m,$$

with $P_0(X) = 1$. The game always terminates, since the deck strictly shrinks at every draw. Player A takes draws $1, 4, 7, \ldots$, player B takes draws $2, 5, 8, \ldots$ and player C takes draws $3, 6, 9, \ldots$, and the winner is the player who takes the final draw, so with

$$a_n = \Pr[M_n \equiv 1], \quad b_n = \Pr[M_n \equiv 2], \quad c_n = \Pr[M_n \equiv 0] \pmod 3$$

the three winning probabilities are a_n, b_n, c_n and $a_n + b_n + c_n = 1$.

Step 1 (the generating function). With n cards, the first draw is card k with probability $1/n$, after which the deck is exactly $\{1, \ldots, k-1\}$ and the remaining play is an independent copy of the game with $k-1$ cards. Hence $M_n = 1 + M_{k-1}$ in distribution, conditionally on the first card being k, and

$$P_n(X) = \frac{X}{n} \sum_{k=1}^{n} P_{k-1}(X) \qquad (n \geq 1).$$

Multiply by n and subtract the same identity for $n-1$ multiplied by $n-1$:

$$nP_n(X) - (n-1)P_{n-1}(X) = XP_{n-1}(X),$$

so $P_n(X) = \frac{X+n-1}{n} P_{n-1}(X)$ for $n \geq 2$; and $P_1(X) = XP_0(X) = X$ agrees with the same formula at $n = 1$. Therefore

$$P_n(X) = \prod_{j=1}^{n} \frac{X+j-1}{j} = \frac{X(X+1)\cdots(X+n-1)}{n!}.$$

As a check, $P_n(1) = 1$.

Step 2 (roots of unity filter). Let $\omega = e^{2\pi i/3}$, so $\omega^3 = 1$ and $1 + \omega + \omega^2 = 0$. Put

$$z_n = P_n(\omega) = \mathbb{E}\left[\omega^{M_n}\right] = c_n + a_n\omega + b_n\omega^2.$$

Using $\operatorname{Im}\omega = \frac{\sqrt{3}}{2}$ and $\operatorname{Im}\omega^2 = -\frac{\sqrt{3}}{2}$, and $\omega^3 = 1$,

$$\operatorname{Im}(z_n) = \tfrac{\sqrt{3}}{2}(a_n - b_n), \quad \operatorname{Im}(\omega z_n) = \tfrac{\sqrt{3}}{2}(c_n - a_n), \quad \operatorname{Im}(\omega^2 z_n) = \tfrac{\sqrt{3}}{2}(b_n - c_n),$$

the last two because $\omega z_n = b_n + c_n\omega + a_n\omega^2$ and $\omega^2 z_n = a_n + b_n\omega + c_n\omega^2$. Every factor $\omega + j$ with j a nonnegative integer is nonzero, since ω is not real, so $z_n \neq 0$ and $\theta_n := \arg z_n$ is defined. Reading the three displayed identities as statements about which half planes z_n, ωz_n, $\omega^2 z_n$ lie in:

$$a_n > b_n \iff \theta_n \in (0, \pi), \quad c_n > a_n \iff \theta_n \in (-\tfrac{2\pi}{3}, \tfrac{\pi}{3}), \quad b_n > c_n \iff \theta_n \in (\tfrac{2\pi}{3}, \tfrac{5\pi}{3}),$$

all modulo 2π. Intersecting the pairs of conditions,

$$A \text{ strictly leads} \iff \theta_n \in \left(\tfrac{\pi}{3}, \pi\right), \quad B \text{ strictly leads} \iff \theta_n \in \left(\pi, \tfrac{5\pi}{3}\right), \quad C \text{ strictly leads} \iff \theta_n \in \left(-\tfrac{\pi}{3}, \tfrac{\pi}{3}\right),$$

three open arcs of length $\frac{2\pi}{3}$ each, again modulo 2π.

Step 3 (the argument grows without bound in vanishing steps). From $z_1 = \omega$ and $z_{n+1} = \frac{\omega+n}{n+1} z_n$, and since $n + 1 > 0$ and $\omega + n = \left(n - \frac{1}{2}\right) + \frac{\sqrt{3}}{2} i$ has positive real part for $n \geq 1$,

$$\theta_n = \frac{2\pi}{3} + \sum_{j=1}^{n-1} \delta_j, \qquad \delta_j = \arg(\omega + j) = \arctan \frac{\sqrt{3}}{2j - 1} \in \left(0, \tfrac{\pi}{2}\right],$$

is a valid choice of $\arg z_n$, because the argument of a product is the sum of the arguments modulo 2π. The increments satisfy $\delta_j \to 0$; and for $j \geq 2$ we have $0 < \frac{\sqrt{3}}{2j-1} \leq 1$, so by concavity of arctan on $[0, 1]$,

$$\delta_j \geq \frac{\pi}{4} \cdot \frac{\sqrt{3}}{2j - 1} \geq \frac{\sqrt{3}\,\pi}{8j},$$

and $\sum_j \delta_j$ diverges by comparison with the harmonic series. Thus (θ_n) is strictly increasing, tends to $+\infty$, and has increments tending to 0.

Step 4 (the landing lemma).

Lemma 9.5.1. *Let (θ_n) be real, increasing, with $\theta_n \to +\infty$ and $\theta_{n+1} - \theta_n \to 0$. Then for every $\alpha \in \mathbb{R}$ and every $\varepsilon > 0$ there are infinitely many n with $\theta_n \in (\alpha + 2\pi k, \alpha + 2\pi k + \varepsilon)$ for some integer k; that is, $\theta_n \bmod 2\pi$ lies in the arc $(\alpha, \alpha + \varepsilon)$ for infinitely many n.*

Proof. Choose M with $\theta_{n+1} - \theta_n < \varepsilon$ for all $n \geq M$. Let k be any integer with $\alpha + 2\pi k > \theta_M$. Since $\theta_n \to +\infty$ there is a least index $n(k)$ with $\theta_{n(k)} > \alpha + 2\pi k$, and $n(k) > M$ because $\theta_M \leq \alpha + 2\pi k$. Minimality gives $\theta_{n(k)-1} \leq \alpha + 2\pi k$, and $n(k) - 1 \geq M$ gives $\theta_{n(k)} < \theta_{n(k)-1} + \varepsilon$. Hence $\alpha + 2\pi k < \theta_{n(k)} < \alpha + 2\pi k + \varepsilon$. Letting $k \to \infty$ makes $\theta_{n(k)} \to \infty$, so the indices $n(k)$ are unbounded and there are infinitely many of them. □

Step 5 (conclusion). Apply the lemma to each of the three arcs of Step 2, each of length $\varepsilon = \frac{2\pi}{3} > 0$. For each player there are infinitely many n with θ_n in that player's arc, and hence infinitely many n at which that player's winning probability is strictly the largest of the three. In particular such n are arbitrarily large. □

Two checks worth doing in the arena. At $n = 3$ the formula gives $P_3(X) = \frac{X(X+1)(X+2)}{6} = \frac{X^3+3X^2+2X}{6}$, so $(a_3, b_3, c_3) = (\frac{1}{3}, \frac{1}{2}, \frac{1}{6})$, matching the recursion directly, and $\theta_3 = \frac{2\pi}{3} + \frac{\pi}{3} + \frac{\pi}{6} = \frac{7\pi}{6}$, which lies in B's arc, as it should. At $n = 2$, $\theta_2 = \pi$ sits exactly on a boundary ray, and indeed $a_2 = b_2 = \frac{1}{2}$ is a tie; boundary rays are ties, which is why the arcs were taken open.

Layer 1: WHAT

Block 1. Identify the state: a deck of size n, transition to a uniform smaller size, absorbing at 0. First step analysis with the roles rotated.

Block 2. Refuse to carry three sequences. The winner is a function of one integer, the number of draws, read modulo 3.

Block 3. Mark that integer with X and form $P_n(X) = \mathbb{E}[X^{M_n}]$. The averaging recursion telescopes after multiplying by n, giving $P_n = \frac{X+n-1}{n} P_{n-1}$ and hence the rising factorial $X(X + 1) \cdots (X + n - 1)/n!$.

Block 4. Extract residues modulo 3 by evaluating at ω. The triple (a_n, b_n, c_n) becomes the single complex number $z_n = c_n + a_n\omega + b_n\omega^2$.

Block 5. Translate "who leads" into "which 120 degree sector contains z_n".

Block 6. The argument of a product is the sum of arguments, so θ_n is a sum of arctangents: increasing, increments to zero, sum divergent.

Block 7. A monotone divergent sequence with vanishing steps enters every arc infinitely often. Done.

Juxtaposition. The nominal route sets up the same recursion, computes twenty terms, sees A then B then C, and then tries to prove sign changes for $a_n - b_n$ by direct estimates on the coupled recursion. It dies because the object it is estimating is the projection of a rotation. The bifurcation point is Block 2: the decision to stop tracking three probabilities and track instead the one integer statistic that determines all three. Everything downstream is forced by that decision, including the appearance of ω, which is simply the standard tool for reading an integer modulo 3.

Layer 2: HOW

Block 1. **Cue.** "that card and all higher-numbered cards are removed" and "the remaining cards are reshuffled". **Why natural.** Reshuffling is the setters telling you the process is memoryless: the future depends on the deck size alone. So write the one step recursion at once, and note the roles rotate, since the player who has just drawn becomes the third player of the new game. **Link to the next block.** You now have three coupled averaging recursions with cyclic symmetry.

Block 2. **Cue.** Three players in a fixed cyclic order, and a winner determined by whose turn the last draw was. **Why natural.** The player index is the draw count modulo 3. Carrying three sequences is carrying the same information three times. Ask instead for the distribution of the draw count M_n, from which all three probabilities follow. **Link to the next block.** One scalar random variable replaces the triple.

Block 3. **Cue.** A recursion of the shape $u_n = \frac{1}{n}\sum_{k<n} u_k$, and a random variable whose distribution you want. **Why natural.** Two moves, both standard. First, an averaging recursion is cleared by multiplying by n and subtracting the previous relation; that always turns a Cesàro type recursion into a first order one. Second, a distribution supported on the nonnegative integers is carried by its generating function, and $M_n = 1 + M_{k-1}$ becomes multiplication by X. The product $\prod_j \frac{X+j-1}{j}$ then also tells you that M_n is a sum of independent Bernoulli variables with parameters $1/j$, which is worth noticing. **Link to the next block.** You have an exact closed form for the whole distribution, valid for every n.

Block 4. **Cue.** You want probabilities of residues modulo 3, and you hold a polynomial. **Why natural.** That is precisely the roots of unity filter. Evaluating at ω compresses the three residue classes into one complex number, and it is the only evaluation that keeps all the information you need while discarding everything you do not. **Link to the next block.** The state of the game at size n is now a single point of the complex plane.

Block 5. **Cue.** A triple of nonnegative reals summing to 1, mapped linearly to $\mathbb{C}$. **Why natural.** The map sends the centroid to 0 and the three extreme cases to $\omega, \omega^2, 1$; comparisons between the coordinates therefore become half plane conditions on z_n, ωz_n, $\omega^2 z_n$, and intersecting pairs of half planes gives sectors. There is nothing to invent: multiplication by ω is exactly the cyclic relabelling of the players. **Link to the next block.** "Each player leads infinitely often" has become "the argument visits each of three arcs infinitely often".

Block 6. Cue. A product formula and a question about arguments. **Why natural.** Arguments add over products, and the factors $\omega + j$ have positive real part, so their arguments are honest arctangents in $(0, \pi/2]$ with no branch anxiety. The size $\arctan \frac{\sqrt{3}}{2j-1} \approx \frac{\sqrt{3}}{2j}$ tells you the whole story: the steps shrink like $1/j$, so they tend to zero, and their sum diverges like $\frac{\sqrt{3}}{2} \log n$. **Link to the next block.** A monotone divergent angle with vanishing steps cannot skip an arc.

Block 7. Cue. "arbitrarily large values of n". **Why natural.** That phrasing asks for infinitely many indices, not for a limit, so a soft argument suffices: a sequence which climbs past every multiple of 2π in steps eventually smaller than the arc length must land inside every translate of the arc. Write the lemma out separately; it is short and it is the only analysis in the solution.

Layer 3: WHY

The first truth is that a cyclic three way comparison is a statement about the characters of $\mathbb{Z}/3$. A probability vector (a, b, c) with $a + b + c = 1$ carries exactly two degrees of freedom, and the natural coordinates on that two dimensional space are the values of the two nontrivial characters, that is, $c + a\omega + b\omega^2$ and its conjugate. In those coordinates the cyclic relabelling of the players becomes multiplication by ω, a rotation by 120 degrees. Any argument phrased in terms of $a_n - b_n$ alone is fighting this structure; any argument phrased in terms of z_n inherits it for free.

The second truth is that an averaging recursion becomes a first order recursion after one multiplication and one subtraction, and a first order recursion in the complex plane is a spiral: $z_{n+1} = \frac{\omega+n}{n+1} z_n$ multiplies the modulus by $\frac{\sqrt{n^2-n+1}}{n+1} < 1$ and adds $\arctan \frac{\sqrt{3}}{2n-1}$ to the argument. The lead rotates because the multiplier $\omega + n$ has positive argument for every n; the rotation never stops because $\sum \frac{1}{n}$ diverges; and it slows down because the argument increments shrink like $1/n$. All three features of the answer are visible in one line of the recursion.

The third truth is a soft one that is worth naming because it appears constantly: divergence with vanishing steps forces recurrence. You do not need equidistribution, you do not need any measure theory, you need only that a sequence which increases past every bound in steps smaller than the target arc cannot jump over the arc. Reach for it whenever a problem asks for infinitely many n with a sign or an order reversed.

Discernment mechanism. When a process has k symmetric roles and the outcome depends on a step count modulo k, pair the state with the characters of $\mathbb{Z}/k$: that is, evaluate the generating function at k-th roots of unity. When the answer must be "infinitely often", look for an unbounded angle with vanishing increments. When a recursion averages all earlier terms, clear it by n and subtract.

What the setters nudged, and how. They gave three players rather than two, because with two players the same generating function gives $P_n(-1) = 0$ for every $n \geq 2$ and the game is exactly fair, so nothing to prove. They reshuffled the deck, to tell you the process is memoryless. They asked for "arbitrarily large n" rather than "all large n", which is a direct instruction not to look for a limit. And the numbers punish the impatient: the handovers occur at $n = 3, 14, 145, 1608, \ldots$, each roughly $e^{4\pi/(3\sqrt{3})} \approx 11.2$ times the last, so a table of the first hundred values suggests a false conclusion.

Revolving concepts

First step analysis. For an absorbing Markov chain, condition on the first transition and write $u_n = \sum_k p_{nk} u_k$. Here the extra content is the role rotation, encoded by the convention $a_0 = b_0 = 0$, $c_0 = 1$. See Book 1, field: Probability, topic: first step analysis.

Cesàro recursions. Any recursion of the form $nu_n = \sum_{k=0}^{n-1} f(u_k)$ is converted to a first order recursion by writing the same identity at $n-1$ and subtracting, giving $nu_n - (n-1)u_{n-1} = f(u_{n-1})$. This is the single most useful manipulation for recursions defined by averages, and it applies

verbatim to sequences, to generating functions and to matrices. See Book 1, field: Sequences and recurrences, topic: averaging recursions.

Roots of unity filter. For $f(X) = \sum_m \alpha_m X^m$ and $\zeta = e^{2\pi i/k}$,

$$\sum_{m \equiv r \ (k)} \alpha_m = \frac{1}{k} \sum_{j=0}^{k-1} \zeta^{-jr} f(\zeta^j).$$

Here $k = 3$ and only $j = 1$ carries information beyond $f(1) = 1$. See Book 1, field: Complex numbers, topic: roots of unity filter.

Rising factorials, Stirling numbers and random permutations. $P_n(X) = \frac{X(X+1)\cdots(X+n-1)}{n!} = \prod_{j=1}^{n} \left(\frac{j-1}{j} + \frac{X}{j}\right)$. The second form says M_n is a sum of independent Bernoulli variables with parameters $1, \frac{1}{2}, \ldots, \frac{1}{n}$; the first says the coefficients are unsigned Stirling numbers of the first kind divided by $n!$. So the number of draws in this game has the distribution of the number of cycles of a uniformly random permutation of $\{1, \ldots, n\}$, with mean the harmonic number $H_n \approx \log n$ and variance also about $\log n$. That is the honest explanation of the slow rotation: the draw count is concentrated near $\log n$, so its residue modulo 3 drifts as $\log n$ advances. See Book 1, field: Combinatorics, topic: Stirling numbers of the first kind.

Generalised recursion. The complex recursion $z_{n+1} = \frac{\omega+n}{n+1} z_n$ with $z_1 = \omega$ has the exact solution $z_n = \frac{\Gamma(n+\omega)}{\Gamma(\omega)\,\Gamma(n+1)}$, and the standard asymptotic $\Gamma(n+\alpha)/\Gamma(n+\beta) \sim n^{\alpha-\beta}$ gives $z_n \sim n^{\omega-1}/\Gamma(\omega)$. Since $\omega - 1 = -\frac{3}{2} + \frac{\sqrt{3}}{2} i$, the modulus decays like $n^{-3/2}$ and the argument grows like $\frac{\sqrt{3}}{2} \log n$: a logarithmic spiral. The same pattern, $\prod_{j=1}^{n} \frac{j+\alpha-1}{j} = \binom{n+\alpha-1}{n}$, is the general solution of any recursion $u_n = \frac{n+\alpha-1}{n} u_{n-1}$ and is worth installing once and for all. This is a knowledge gap for you: Gamma function asymptotics are not in the Further Mathematics syllabus, and the finite product plus arctangent sum used in the ideal solution is exactly the elementary substitute.

Divergence with vanishing steps. If $\theta_n \uparrow \infty$ and $\theta_{n+1} - \theta_n \to 0$, then $\{\theta_n \bmod 2\pi\}$ is dense in the circle and meets every arc infinitely often. This is much weaker than equidistribution, which for $\theta_n = \frac{\sqrt{3}}{2} \log n$ actually fails in the usual Weyl sense. Know the difference and claim only what you need. See Book 1, field: Real analysis, topic: recurrence of slowly varying sequences.

Circuit Bank

- **WHEN** a process with k symmetric roles taken in cyclic order, the winner decided by the number of steps modulo k **DO** track the number of steps, form its generating function, and evaluate at k-th roots of unity **BECAUSE** a residue class is extracted by characters, and the cyclic relabelling of the roles becomes multiplication by a root of unity
- **WHEN** a recursion in which the new term is the average of all previous terms **DO** multiply by n, write the same relation at $n-1$, and subtract **BECAUSE** averaging recursions telescope into first order recursions, and first order recursions have product solutions
- **WHEN** a triple of nonnegative reals summing to 1 whose ordering you must compare **DO** map it to $c + a\omega + b\omega^2$ in the complex plane **BECAUSE** the map turns the three comparisons into three half planes and the three "who leads" regions into 120 degree sectors
- **WHEN** a demand for "arbitrarily large n" with some order reversed **DO** find an angle that increases without bound with increments tending to zero **BECAUSE** such a sequence enters every arc infinitely often, which is exactly infinitely many reversals and needs no equidistribution
- **WHEN** a table of small cases that seems to settle the behaviour **DO** estimate the scale on which the phenomenon operates before trusting it **BECAUSE** here the handovers are spaced geometrically with ratio about 11, so any finite table is silent about what follows

Backtesting

1. First useful cue: the deck is reshuffled and only the top card matters, so the state is the deck size alone and the winner is the draw count modulo 3.
2. Method the expression suggested: the averaging recursion $u_n = \frac{1}{n}\sum_{k<n} u_k$ demanded the multiply and subtract move, and the residue modulo 3 demanded a cube root of unity.
3. Where you would hesitate, and the thought to reinforce: after obtaining $z_n = P_n(\omega)$ you may not see what to do with a complex number. The thought to install is that a complex number is a modulus and an argument, that the modulus is irrelevant to a comparison, and that the argument is a sum because z_n is a product.
4. The trap the question sets: computing a table and believing it. The first two handovers are at $n = 3$ and $n = 14$, the third is at $n = 145$ and the fourth at $n = 1608$.
5. What makes you recognise it faster next time: any cyclic competition between three symmetric agents is a $\mathbb{Z}/3$ problem, and $\mathbb{Z}/3$ problems are done with ω, not with three sequences.
6. The reusable decision rule: compress a symmetric multi agent state into one complex number via a character, convert the recursion into complex multiplication, and read the conclusion off the argument.

Mirror drills

1. Two players A and B play the same card game with n cards. Who is favoured? *Answer:* now A wins if and only if the draw count is odd, so $b_n - a_n = \mathbb{E}[(-1)^{M_n}] = P_n(-1) = \frac{(-1)\cdot 0\cdot 1\cdots(n-2)}{n!} = 0$ for every $n \geq 2$, so the game is exactly fair for all $n \geq 2$; only $n = 1$ favours A.
2. Show that the number of draws in the n card game has the same distribution as the number of cycles of a uniformly random permutation of $\{1, \ldots, n\}$. *Answer:* both have probability generating function $X(X+1)\cdots(X+n-1)/n!$; equivalently both are sums of independent Bernoulli variables with parameters $1/j$ for $j = 1, \ldots, n$.
3. Let $\theta_n = \sum_{j=1}^{n} \arcsin\frac{1}{j}$. Show that $e^{i\theta_n}$ enters every arc of the unit circle infinitely often. *Answer:* the increments tend to 0 and $\arcsin\frac{1}{j} \geq \frac{1}{j}$ so the sum diverges; apply the landing lemma.

9.6 A6. Icosahedron edge colourings via linear algebra

The 30 edges of a regular icosahedron are distinguished by labelling them $1, 2, \ldots, 30$. *How many different ways are there to paint each edge red, white, or blue such that each of the 20 triangular faces of the icosahedron has two edges of the same colour and a third edge of a different colour? (The top matter on each exam paper included the logo of the Mathematical Association of America, which is itself an icosahedron.)*

Read it out loud

Say what the object is. There are 3^{30} colourings, roughly 2×10^{14}, and 20 constraints, one per face, each involving only three of the thirty edges. So this is a constraint satisfaction count with a fixed sparse incidence structure: a solid, a set of 30 variables in a 3 element set, and 20 local conditions. Nothing about it is asking you to be clever about icosahedra. It is asking you what the local condition really says.

Now say the local condition out loud. "Two edges of the same colour and a third of a different colour" is the negation of "all three the same or all three different". Those two forbidden patterns are the two orbits of the colour permutation group on triples that are somehow uniform, and any time you meet "all equal or all distinct" with exactly three colours, you should try to make the colours into the field $\mathbb{F}_3$ and look at the sum. Over $\mathbb{F}_3$, all equal gives $3u = 0$, and all distinct gives $0 + 1 + 2 = 0$, while two equal and one different, say u, u, w with $w \neq u$, gives $2u + w = w - u \neq 0$. The condition is exactly "the sum of the three edge labels is nonzero".

That is the entire content of the problem, and the rest is linear algebra. The fields wired in: finite fields and vector spaces over them, rank and nullity, and a small amount of polyhedral combinatorics ($V = 12$, $E = 30$, $F = 20$, every vertex has degree 5 and its five neighbours form a pentagon, every edge lies in exactly two faces).

Within thirty seconds you should have written the equivalence above, the map

$$T : \mathbb{F}_3^E \to \mathbb{F}_3^F, \qquad T(x)_f = \sum_{e \in f} x_e,$$

and the sentence "count the admissible targets, then multiply by the size of a fibre". And you should have noticed the trap: the edges are labelled 1 to 30, so nothing is to be counted up to symmetry. The icosahedron in the logo is decoration.

Approaches, ranked worst to best

1. **Burnside and the rotation group.** *Cue:* the word icosahedron, and a note about a logo. *Circuit:* count orbits of colourings under the symmetry group. This answers a different question. The edges carry labels 1 to 30, so a colouring is a function on a labelled set and two colourings that differ by a rotation are different. Any candidate who divides by 60 has thrown away the sentence that was put there to stop them.

2. **Layer by layer transfer matrix.** *Cue:* the icosahedron is a pentagonal antiprism with two pyramids on it. *Circuit:* propagate a state vector over 3^5 states through five layers. This is a genuine algorithm, and by hand it is a page of 243 dimensional bookkeeping to compute what one rank computation gives you. You would in effect be doing linear algebra without admitting it.

3. **Inclusion and exclusion over the bad faces.** *Cue:* "each of the 20 faces" is a conjunction of conditions. *Circuit:* sum over subsets S of faces with sign, counting colourings whose faces in S are all bad. Each term requires the number of solutions of a subsystem, which depends on the rank of that subsystem, and the ranks are not uniform across the 2^{20} subsets. The sum does not collapse.

4. **Peel the two polar caps.** *Cue:* an icosahedron splits as five faces, a pentagon, ten faces, a pentagon, five faces. *Circuit:* colour the equatorial band freely, then complete each ring. Colour the ten equatorial edges freely, note that each of the five faces containing exactly one edge of a polar pentagon then determines that edge in 2 ways, and finally show each polar fan of five spokes admits exactly 2^5 completions. It reaches $3^{10} \cdot 2^{20}$ correctly, but the last step is the same odd cycle fact that the winning route uses, so this is the winning route with extra scaffolding.
5. **Recode into $\mathbb{F}_3$ and count fibres.** *Cue:* three colours, and a condition invariant under permuting the colours. *Circuit:* make the condition linear, prove the linear map is onto, and count as (number of admissible images) times (size of the kernel). The condition becomes $T(x)_f \neq 0$ for every f; there are 2^{20} admissible images; T is onto, so every fibre has 3^{30-20} elements. Two lines of counting once surjectivity is proved, and surjectivity is one local gadget.

The ideal solution

Let E be the set of the 30 edges and F the set of the 20 faces. Identify the three colours with the three elements of $\mathbb{F}_3$, the field with three elements, in any fixed way, so that a colouring is a vector $x \in \mathbb{F}_3^E$. Define

$$T : \mathbb{F}_3^E \longrightarrow \mathbb{F}_3^F, \qquad T(x)_f = \sum_{e \subset f} x_e,$$

the sum being over the three edges of the face f. Each coordinate of $T(x)$ is a sum of coordinates of x, so T is $\mathbb{F}_3$-linear.

Lemma 9.6.1. *A face f has two edges of one colour and a third edge of a different colour if and only if $T(x)_f \neq 0$.*

Proof. The three edge colours of f fall into exactly one of three mutually exclusive cases. If all three are equal to u, then $T(x)_f = 3u = 0$. If all three are distinct, they are $0, 1, 2$ in some order and $T(x)_f = 0 + 1 + 2 = 0$. Otherwise exactly two are equal, say the colours are u, u, w with $w \neq u$, and then $T(x)_f = 2u + w = w - u \neq 0$, because $2 = -1$ in $\mathbb{F}_3$ and $w \neq u$. So the required pattern holds exactly when the sum is nonzero. □

Write $\mathbb{F}_3^\times = \{1, 2\}$ for the two nonzero elements. The colourings we must count are therefore

$$\mathcal{N} = \left| \left\{ x \in \mathbb{F}_3^E : T(x) \in (\mathbb{F}_3^\times)^F \right\} \right| = \left| T^{-1}\left((\mathbb{F}_3^\times)^F \right) \right|.$$

Lemma 9.6.2. *T is surjective.*

Proof. It suffices to show that every standard basis vector of $\mathbb{F}_3^F$ lies in the image, since those span.

Fix a vertex v. In the icosahedron v has degree 5 and its five neighbours form a pentagon, so the five edges at v can be listed as $e_1, \ldots, e_5$ in cyclic order in such a way that the five faces containing v are exactly $f_1, \ldots, f_5$, where f_i has edges e_i, e_{i+1} and one edge of that pentagon (indices modulo 5). Every edge at v lies in exactly two faces, and both of them contain v; hence if x is supported on $\{e_1, \ldots, e_5\}$ then $T(x)$ is supported on $\{f_1, \ldots, f_5\}$ and

$$T(x)_{f_i} = x_{e_i} + x_{e_{i+1}} \qquad (i = 1, \ldots, 5).$$

Take $x_{e_1} = 2$, $x_{e_2} = 1$, $x_{e_3} = 2$, $x_{e_4} = 1$, $x_{e_5} = 2$ and $x_e = 0$ for every other edge. Then

$$T(x)_{f_1} = 2 + 1 = 0, \quad T(x)_{f_2} = 1 + 2 = 0, \quad T(x)_{f_3} = 2 + 1 = 0, \quad T(x)_{f_4} = 1 + 2 = 0,$$

while $T(x)_{f_5} = x_{e_5} + x_{e_1} = 2 + 2 = 1$. So $T(x)$ is the standard basis vector at the face f_5.

Finally, let f be any face and v any vertex of f. The two edges of f at v are consecutive in the cyclic order at v, so the labelling above may be started so that they are e_5 and e_1, that is, so that $f = f_5$. Hence every standard basis vector of $\mathbb{F}_3^F$ is in the image, and T is surjective. □

The alternating pattern $2, 1, 2, 1, 2$ is not a lucky guess: on a cycle of odd length the assignment that alternates cannot close up, and the defect is forced to appear at exactly one place. Formally, the map $\mathbb{F}_3^5 \to \mathbb{F}_3^5$ sending $(x_1, \ldots, x_5)$ to $(x_1 + x_2, \ldots, x_5 + x_1)$ has matrix $I + S$ with S the cyclic shift, whose determinant is $\prod_{\zeta^5=1}(1+\zeta) = 1 - (-1)^5 = 2 \neq 0$ in $\mathbb{F}_3$. So the five spokes at a vertex can realise *any* prescribed values on the five faces at that vertex, and in particular a single 1.

Now finish. Since T is a surjective linear map from a 30 dimensional space to a 20 dimensional space, rank and nullity give $\dim \ker T = 30 - 20 = 10$, so $|\ker T| = 3^{10}$. Every fibre $T^{-1}(w)$ with $w \in \mathbb{F}_3^F$ is a coset of $\ker T$ and therefore has exactly 3^{10} elements. The set of admissible targets has $|(\mathbb{F}_3^\times)^F| = 2^{20}$ elements. Hence

$$\mathcal{N} = 2^{20} \cdot 3^{10} = \boxed{61917364224}.$$

Layer 1: WHAT

Block 1. Read the face condition as the negation of "all equal or all distinct".

Block 2. Three colours plus a condition invariant under permuting them: recode the colours as $\mathbb{F}_3$ and test the sum. Both forbidden patterns have sum 0, the allowed pattern has sum nonzero.

Block 3. The condition is now $T(x)_f \neq 0$ for all f, with T linear. The problem has become: how many vectors does a linear map send into a prescribed product set.

Block 4. Fibres of a linear map are cosets of the kernel, so the count is (number of admissible targets in the image) times $|\ker T|$. All that is missing is the rank.

Block 5. Prove surjectivity by hitting one standard basis vector with a local gadget: the five spokes at a vertex, alternating $2, 1, 2, 1, 2$ around an odd cycle.

Block 6. Rank and nullity, then arithmetic: $2^{20} \cdot 3^{10}$.

Juxtaposition. The nominal route treats this as a hard enumeration: peel the solid into layers, define a state on each pentagon, and push a transfer matrix through five stages, or worse, start worrying about the symmetry group because the word icosahedron and a logo were mentioned. The bifurcation point is Block 2. Before it you have a combinatorial condition on three coloured objects; after it you have a system of 20 linear inequations over a field, and the only remaining question in the whole problem is the rank of a 20×30 matrix. Notice what the ideal route never does: it never enumerates a single face pattern, never uses a symmetry group, and never draws the solid beyond one vertex.

Layer 2: HOW

Block 1. **Cue.** "two edges of the same colour and a third edge of a different colour". **Why natural.** A condition on a triple that does not name any particular colour is invariant under all 6 permutations of the colours, so it is a union of orbits of unordered patterns. There are only three such orbits: all equal, all distinct, exactly two equal. Say them out loud and the condition becomes a complement, which is nearly always easier to make algebraic than the original. **Link to the next block.** You now want one algebraic function of a triple that vanishes on the two forbidden orbits.

Block 2. **Cue.** Exactly three colours, and a symmetric condition. **Why natural.** The cheapest colour invariant functions are symmetric functions, and the cheapest of those is the sum. The sum only becomes a colour invariant statement once the colours live in a ring, so put them in $\mathbb{F}_3$. Then $3 = 0$ kills the monochromatic case, and $0 + 1 + 2 = 0$ kills the rainbow case; that these two coincide is special to $\mathbb{F}_3$, so check it rather than assume it. The remaining case gives $w - u \neq 0$. **Link to the next block.** The 20 face conditions are now 20 linear inequations.

Block 3. **Cue.** Twenty conditions, each a linear form in the edge variables, each forbidding one value. **Why natural.** Do not solve the system. Name the map T that packages all 20 forms at once,

so that the condition reads "$T(x)$ has no zero coordinate". Packaging is what lets you count instead of solve. **Link to the next block.** The count is now a statement about a preimage.

Block 4. **Cue.** A preimage of a set under a linear map between finite spaces. **Why natural.** Fibres of a linear map are cosets of the kernel, hence all of the same size; so $|T^{-1}(S)| = |S \cap \operatorname{im} T| \cdot |\ker T|$. Every counting question of this shape reduces to two numbers: how many admissible targets are hit, and how large the kernel is. Both are easy once the rank is known, which is why the rank is now the only open question. **Link to the next block.** You need to decide whether T is onto.

Block 5. **Cue.** You need one standard basis vector in the image, and the incidence structure is local and highly symmetric. **Why natural.** Surjectivity means hitting every basis vector, and by the symmetry of the solid it is enough to hit one, provided the construction can be positioned at any face. So look for the smallest gadget whose image is supported on few faces: the five edges at a single vertex, because every one of them lies in two faces that both contain that vertex. On those five faces the map becomes "sum of cyclically adjacent pairs" on a 5 cycle, and the alternating vector $2, 1, 2, 1, 2$ cancels four of the five sums. It cannot cancel all five, and that is the point: 5 is odd. **Link to the next block.** Surjectivity gives the rank, and the rank gives the kernel.

Block 6. **Cue.** A surjective linear map $\mathbb{F}_3^{30} \to \mathbb{F}_3^{20}$. **Why natural.** Rank and nullity give a kernel of dimension 10; there are 2^{20} vectors in $\mathbb{F}_3^{20}$ with all coordinates nonzero; multiply. State the answer as $2^{20}3^{10}$ and give the decimal only as a courtesy.

Layer 3: WHY

The first truth is a coincidence of $\mathbb{F}_3$ that you must verify and then remember: for three elements u, v, w of $\mathbb{F}_3$, the sum $u + v + w$ vanishes precisely when they are all equal or all distinct. Over $\mathbb{F}_5$ the analogous statement is false, and over $\mathbb{F}_2$ there is nothing to say. This is exactly the fact behind the card game SET, where three cards form a set when each attribute is all equal or all distinct across the three cards, that is, when the attribute vectors sum to zero in $\mathbb{F}_3^4$. The problem is a SET condition on the faces of a solid, negated.

The second truth is the counting principle. If T is linear over a finite field and S is any subset of the target, then $|T^{-1}(S)| = |S \cap \operatorname{im} T| \cdot |\ker T|$, because nonempty fibres are cosets of the kernel. This turns a hard looking count over 3^{30} objects into two easy counts. The whole difficulty of such problems concentrates into deciding $\operatorname{im} T$, and here $\operatorname{im} T$ is everything.

The third truth is why the rank is full, and it is a parity statement. The natural obstruction to surjectivity of a face and edge incidence map is that each edge lies in exactly two faces, so

$$\sum_{f \in F} T(x)_f = \sum_{e \in E} 2x_e.$$

In characteristic 2 the right hand side vanishes identically, the image lies in a hyperplane, and the rank drops to 19; the two colour version of this problem would be governed by that obstruction. In characteristic 3 the factor 2 is invertible and no such relation exists. A second potential obstruction is a two colouring of the faces: if the faces could be split into two classes so that every edge meets one face of each class, then the sum over one class would always equal the sum over the other and the rank would again drop. That is precisely what happens for the octahedron, whose dual graph is the cube and hence bipartite. The dual of the icosahedron is the dodecahedron, which contains 5 cycles and is not bipartite, so that obstruction is absent too. Both obstructions are about even structures, and both are defeated by the same feature: the icosahedron has odd vertex degree and odd dual girth.

Discernment mechanism. When a colouring condition is local and invariant under permuting the colours, ask whether the colour set can be made into a group or a field so that the condition becomes

linear. When a count is "how many x with $T(x)$ avoiding something", never solve for x; count the targets and multiply by the kernel. When you need a linear map defined by an incidence structure to be surjective, look for the smallest local gadget whose image is supported on a single basis vector, and always check the characteristic first, because incidence maps carry parity relations for free.

How to train it. Take every counting problem you meet with a fixed small colour set and ask what the colours would have to be for the condition to be a linear equation. Then take every incidence matrix you meet and compute its rank over $\mathbb{F}_2$ and over $\mathbb{F}_3$ separately; the difference between the two is where the interesting structure lives.

Revolving concepts

The field $\mathbb{F}_3$ and the SET condition. $\mathbb{F}_3 = \mathbb{Z}/3\mathbb{Z}$ with $2 = -1$. For $u, v, w \in \mathbb{F}_3$: $u + v + w = 0$ if and only if $u = v = w$ or $\{u, v, w\} = \{0, 1, 2\}$. Generalisation: in $\mathbb{F}_3^d$, three vectors sum to zero if and only if in every coordinate they are all equal or all distinct; this is the rule of the game SET with $d = 4$, and the maximum size of a SET free collection is the cap set problem. See Book 1, field: Abstract algebra, topic: small finite fields.

Fibres of a linear map over a finite field. For $T : \mathbb{F}_q^m \to \mathbb{F}_q^n$ linear and S a subset of $\mathbb{F}_q^n$, $|T^{-1}(S)| = |S \cap \operatorname{im} T| \cdot q^{m - \operatorname{rank} T}$. This single formula covers a very large family of Putnam counting problems, including counting solutions of linear systems over $\mathbb{F}_p$ and counting tuples with a prescribed weighted sum modulo a prime. See Book 1, field: Linear algebra, topic: counting over finite fields.

Rank and nullity, and the two standard obstructions. $\operatorname{rank} T + \dim \ker T = \dim$ of the source. For a face and edge incidence map on a triangulated surface, the two obstructions to full rank are the characteristic 2 relation $\sum_f T(x)_f = 2 \sum_e x_e$ and a bipartition of the faces in which every edge meets one face of each part. For the icosahedron over $\mathbb{F}_3$ neither occurs and the rank is 20; over $\mathbb{F}_2$ the rank is 19; for the octahedron over $\mathbb{F}_3$ the dual graph is bipartite and the rank is 7 rather than 8.

The odd cycle recursion, generalised. Around a cycle of length n the system $x_i + x_{i+1} = t_i$ for $i = 1, \ldots, n$ (indices modulo n) is solved by the recursion $x_{i+1} = t_i - x_i$, which after going once around gives

$$x_1 = (t_n - t_{n-1} + \cdots \pm t_1) + (-1)^n x_1.$$

If n is odd this determines x_1 uniquely in any field of characteristic other than 2, so the system has exactly one solution for every right hand side, which is the same statement as $\det(I + S) = 1 - (-1)^n \neq 0$. If n is even the map is singular and the right hand sides must satisfy the alternating sum condition. This is the general form of the local gadget, it is the reason the pentagonal vertex figure of the icosahedron is exactly what the proof needs, and it is worth installing on its own because "odd cycles are invertible, even cycles have an alternating obstruction" recurs constantly in combinatorics. See Book 1, field: Linear algebra, topic: circulant matrices.

Icosahedron combinatorics. $V = 12$, $E = 30$, $F = 20$, $V - E + F = 2$; every vertex has degree 5 with the five neighbours forming a pentagon; every edge lies in exactly two faces; the dual is the dodecahedron. General identity for a triangulation: $3F = 2E$, so the count in the ideal solution reads $2^F 3^{E-F} = 2^F 3^{F/2} = (2\sqrt{3})^F$ for any triangulation of a surface on which the analogous T is surjective. See Book 1, field: Discrete structures and graphs, topic: polyhedral combinatorics.

Circuit Bank

- **WHEN** a local condition on three objects coloured from a set of three colours, invariant under permuting the colours **DO** identify the colours with $\mathbb{F}_3$ and compute the sum **BECAUSE** all equal and all distinct both give sum zero over $\mathbb{F}_3$, so the condition becomes a single linear inequation

- **WHEN** a count of the form "how many x with $T(x)$ in a prescribed set", T linear over a finite field **DO** count the admissible targets inside the image and multiply by the size of the kernel **BECAUSE** nonempty fibres are cosets of the kernel and so all have the same size
- **WHEN** a need for surjectivity of a map defined by an incidence structure **DO** construct one local gadget whose image is a single standard basis vector, then move it by relabelling **BECAUSE** surjectivity is exactly the statement that every basis vector is hit, and incidence maps are local so gadgets are small
- **WHEN** any face and edge incidence map **DO** check the characteristic before doing anything else **BECAUSE** each edge lies in two faces, so in characteristic 2 the face sums satisfy a universal relation and the map cannot be onto
- **WHEN** objects in a counting problem explicitly labelled, as in "the edges are distinguished by labelling them 1 to 30" **DO** count functions and never quotient by a symmetry group **BECAUSE** labelled counting and orbit counting are different questions, and the labelling sentence is there to tell you which one is being asked

Backtesting

1. First useful cue: "two edges of the same colour and a third edge of a different colour" with exactly three colours available. That is a colour symmetric condition, hence a candidate for an $\mathbb{F}_3$ recoding.
2. Method the expression suggested: recode the colours as $\mathbb{F}_3$, convert the 20 face conditions into $T(x)_f \neq 0$, and count by fibres.
3. Where you would hesitate, and the thought to reinforce: at surjectivity. The instinct is to build a Hamiltonian path in the dual graph and prescribe values along it, which works but is long. The thought to install is that surjectivity only needs one basis vector plus the freedom to reposition it, so look for the smallest local gadget first.
4. The trap the question sets: the word icosahedron and the note about the logo invite a symmetry argument, but the edges are labelled and no quotient is wanted. A second trap is to forget that the count of admissible target vectors is 2^{20} and not $2 \cdot 20$ or $3^{20} - 2^{20}$.
5. What makes you recognise it faster next time: three colours plus "all equal or all distinct" is always $\mathbb{F}_3$; and a sparse local constraint over a small field is always a rank computation.
6. The reusable decision rule: make the condition linear, then count the image and the kernel rather than the solutions.

Mirror drills

1. Show that the analogous map for two colours cannot be surjective on any polyhedron. *Answer:* every edge lies in exactly two faces, so $\sum_f T(x)_f = 2\sum_e x_e = 0$ over $\mathbb{F}_2$ and the image lies in a hyperplane; for the icosahedron the rank over $\mathbb{F}_2$ is 19 rather than 20.
2. The same colouring question for the regular octahedron (12 edges, 8 faces). Why does the counting shortcut break, and what is the answer? *Answer:* the dual graph is the cube, which is bipartite, so the sum of the four face constraints in one part always equals the sum in the other and the rank is 7, not 8; the kernel has size 3^5 and exactly 86 of the 2^8 all nonzero targets lie in the image, giving $86 \cdot 3^5 = 20898$, which is not $2^8 3^4$.
3. Prove that for odd n the system $x_i + x_{i+1} = t_i$, indices modulo n, has exactly one solution over $\mathbb{F}_3$ for every right hand side, and describe what happens for even n. *Answer:* substitute $x_{i+1} = t_i - x_i$ around the cycle to get $x_1 = A + (-1)^n x_1$; for odd n this determines x_1 since 2

is invertible, and for even n it is solvable only when the alternating sum A vanishes, and then there is a one parameter family.

9.7 B1. why a dilation always finds the meeting point

Let L_1 and L_2 be distinct lines in the plane. Prove that L_1 and L_2 intersect if and only if, for every real number $\lambda \neq 0$ and every point P not on L_1 or L_2, there exist points A_1 on L_1 and A_2 on L_2 such that $\overrightarrow{PA_2} = \lambda \overrightarrow{PA_1}$.

Read it out loud

Two distinct lines in the plane, a point off both of them, a nonzero real number, and a vector equation linking a point on each line back to that outside point. Say it in plain terms: for every centre P you could name and every nonzero scaling factor λ you could name, you can find a point on L_1 and a point on L_2 so that the second is obtained from the first by scaling about P by λ. That is not a sentence about lines meeting somewhere out in the plane, it is a sentence about a transformation, because output minus centre equals λ times input minus centre is not a clue pointing at a transformation, it is the transformation, written out in vectors instead of in the word dilation.

The fields wired into this problem are plane geometry, parallelism and intersection, and affine transformations, dilations, also called homotheties. The universal quantifiers for every λ and for every P are doing real work: they tell you the claim must survive an arbitrary choice of centre and an arbitrary nonzero ratio, which is exactly the freedom a dilation has. Within the first thirty seconds you should have written down that $\overrightarrow{PA_2} = \lambda \overrightarrow{PA_1}$ means A_2 is the image of A_1 under the dilation centred at P with ratio λ. Everything else follows from what dilations do to lines.

Approaches, ranked worst to best

1. **Coordinate bash.** *Cue:* explicit line equations and an explicit point $P = (p, q)$. *Circuit:* solve two simultaneous linear equations in the parameter of A_1. (why it ranks here: you must case split on vertical against non-vertical lines to avoid dividing by an undefined slope, the algebra for the two directions of the iff gets written out twice with different letters, and nothing about the argument explains why the statement is true, it only checks that it is).

2. **Normal-vector linear system.** *Cue:* a line is the solution set of $\vec{v} \cdot \vec{x} = c$. *Circuit:* write both membership conditions as one linear system in the unknown displacement and read off solvability from linear independence of the two normal vectors. (why it ranks here: it is coordinate-free and correct, and it makes the link to not parallel explicit through linear independence, but it still asks you to set up and solve a genuine linear system rather than simply naming a transformation whose properties you already know).

3. **Dilation, also called a homothety, centred at P.** *Cue:* the shape of the vector equation itself, output minus centre equals a scalar times input minus centre. *Circuit:* name the transformation, then use the one property of dilations that matters here, a dilation sends a line not through its centre to a parallel line. (why it ranks here: the forward direction becomes two sentences, the backward direction becomes one sentence by choosing the single easiest value $\lambda = 1$, and both directions lean on the same one fact about dilations instead of two separate pieces of machinery).

The ideal solution

Recall the basic dichotomy for two distinct lines in the plane: L_1 and L_2 either intersect in exactly one point, or are parallel, in which case, being distinct, they are disjoint. You will prove the two directions of the biconditional separately, and both will turn on a single fact about the map $h(X) = P + \lambda(X - P)$, the dilation centred at P with ratio $\lambda \neq 0$: it sends any line L with $P \notin L$ to a

line $h(L)$ parallel to L, and it is invertible, with inverse $h^{-1}(Y) = P + \lambda^{-1}(Y - P)$, itself a dilation centred at P.

Forward direction. Suppose L_1 and L_2 intersect. Fix any $\lambda \neq 0$ and any point P not on L_1 or L_2. Let h be the dilation centred at P with ratio λ, and consider $h(L_1)$. Since $P \notin L_1$, the line $h(L_1)$ is parallel to L_1. Because L_1 and L_2 intersect, they are not parallel, so L_1 and L_2 have different direction vectors; consequently $h(L_1)$, sharing its direction with L_1, also has a different direction from L_2, so $h(L_1)$ is not parallel to L_2 either, and the two lines $h(L_1)$ and L_2 meet in exactly one point. Call that point A_2, and set $A_1 = h^{-1}(A_2)$. Then $A_1 \in L_1$, because $A_2 \in h(L_1)$, and by the defining equation of h,

$$A_2 = P + \lambda(A_1 - P), \qquad \text{that is} \qquad \overrightarrow{PA_2} = \lambda\overrightarrow{PA_1},$$

exactly as required. Since P and λ were arbitrary subject only to the stated restrictions, this proves the forward direction in full.

Backward direction, by contrapositive. Suppose L_1 and L_2 do not intersect, that is, since they are distinct, suppose they are parallel and hence disjoint, $L_1 \cap L_2 = \varnothing$. Pick any point P not on L_1 or L_2 (such a point exists, since two lines cannot cover the whole plane) and take $\lambda = 1$. The dilation centred at P with ratio 1 is the identity map, so the required equation $\overrightarrow{PA_2} = 1 \cdot \overrightarrow{PA_1}$ reads $A_2 - P = A_1 - P$, that is $A_2 = A_1$. For points $A_1 \in L_1$ and $A_2 \in L_2$ to satisfy this you would need a common point of L_1 and L_2, but $L_1 \cap L_2 = \varnothing$. So for this choice of P and λ, no such A_1, A_2 exist, and the right-hand side of the biconditional fails. This proves the contrapositive of the backward direction, hence the backward direction itself.

Both directions are now established, so L_1 and L_2 intersect if and only if the stated property holds. ■

Layer 1: WHAT

Block 1: read the vector equation $\overrightarrow{PA_2} = \lambda\overrightarrow{PA_1}$ as the definition of a dilation centred at P with ratio λ, not as two coordinate equations to be solved. Block 2: recall that a dilation sends a line missing its centre to a parallel line, and that it is invertible for $\lambda \neq 0$. Block 3, forward direction: apply the dilation to L_1, use that not parallel is preserved under taking a parallel image, to force $h(L_1)$ to meet L_2, and pull the intersection point back through h^{-1}. Block 4, backward direction: specialise to the single value $\lambda = 1$, where the dilation collapses to the identity and the required equation collapses to $A_1 = A_2$, which is impossible on disjoint lines.

The nominal route sets up explicit coordinates for L_1, L_2 and P, writes A_1 with a parameter, imposes $A_2 = P + \lambda(A_1 - P) \in L_2$, and solves for the parameter, splitting into cases according to whether either line is vertical, then repeating a mirror computation for the converse. The bifurcation point is the very first line of work: whether you translate $\overrightarrow{PA_2} = \lambda\overrightarrow{PA_1}$ into coordinates or into the single word dilation. Naming the transformation is what lets both directions of the proof finish in a few lines with no case split at all.

Layer 2: HOW

Block 1. Cue. The expression $\overrightarrow{PA_2} = \lambda\overrightarrow{PA_1}$ has the exact shape displacement of output from a fixed point equals a scalar times displacement of input from the same fixed point. **Why natural.** That shape is not similar to the definition of a dilation, it is the definition of a dilation; there is no gap between the given equation and the transformation, so naming it costs nothing and buys the whole toolkit of known dilation properties. **Link to the next block.** Once A_2 is understood as $h(A_1)$, the question of whether A_1 exists becomes the question of whether $h(L_1)$ meets L_2, a question about two lines, which is what you already know how to answer.

Block 2. **Cue.** L_1 and L_2 intersect, that is, are not parallel, that is, have different direction vectors. **Why natural.** A dilation scales every displacement vector by the same factor λ, so it scales a line's direction vector by λ and leaves the direction unchanged up to that scalar; a line and its dilated image are always parallel to each other. **Link to the next block.** Since $h(L_1)$ keeps L_1's direction, and L_1's direction differs from L_2's, $h(L_1)$ and L_2 have different directions too, so they meet in exactly one point, which is exactly the pair of points the problem asks you to produce.

Block 3. **Cue.** The problem demands A_1 on L_1 itself, not merely a point whose image lies on L_2. **Why natural.** $\lambda \neq 0$ guarantees h is invertible, with inverse the dilation of ratio $1/\lambda$ at the same centre P, so pulling the found intersection point back through h^{-1} lands on L_1 automatically. **Link to the next block.** This closes the forward direction for an arbitrary P, λ, so the natural next question is whether the converse needs the same generality, or whether one well-chosen instance suffices.

Block 4. **Cue.** The property to be falsified is a for every λ, for every P statement. **Why natural.** Negating a universal quantifier only needs one witness, so instead of repeating a perpendicular-component argument for a general λ, you should ask which single value of λ makes the equation collapse to something trivially false; $\lambda = 1$ turns the dilation into the identity, and the identity's fixed condition $A_1 = A_2$ is visibly impossible on disjoint lines. **Link to the next block.** There is no further block: the single witness P, with $\lambda = 1$, already breaks the property, completing the contrapositive and the proof.

Layer 3: WHY

The truth underneath this problem is that intersect and parallel are not two independent facts about a pair of distinct lines, they are the two branches of a single dichotomy governed entirely by direction vectors, and every affine transformation that scales vectors by a nonzero constant, which is what a dilation is by definition, respects that dichotomy exactly, because it multiplies every direction vector by the same nonzero number and so can never turn a different directions pair into a same direction pair or the reverse. That is why the forward direction needs no case work: parallelism is dilation invariant, full stop.

The discernment mechanism to train is this: whenever a problem statement is built from a vector equation of the shape output minus a fixed point equals λ times input minus the same fixed point, with the fixed point and scalar both allowed to range freely, treat the fixed point as a candidate centre and the scalar as a candidate ratio of a named transformation before doing anything else. The examiners chose to write the condition in vectors rather than say the word dilation outright precisely so that recognising the transformation is the content of the problem; once you see it, the rest is bookkeeping. The second habit to train is quantifier judo: a for all claim you must prove needs the general argument, but a for all claim you must refute needs only its worst instance, and hunting for the value of the free parameter that trivialises the defining equation, here $\lambda = 1$, turning a dilation into the identity, is almost always less ink than adapting the general construction into a nonexistence argument.

Revolving concepts

Dilation, or homothety: the map $h_{P,\lambda}(X) = P + \lambda(X - P)$ for a centre P and ratio $\lambda \neq 0$. It fixes P, scales every displacement from P by λ, sends lines not through P to parallel lines and lines through P to themselves, composes with another dilation of the same centre by multiplying ratios, $h_{P,\lambda} \circ h_{P,\mu} = h_{P,\lambda\mu}$, and is therefore invertible with inverse $h_{P,1/\lambda}$. See Book 1, field: Geometry, topic: affine transformations.

Parallel against intersecting lines in the plane: two distinct lines either share a direction vector up to scaling, parallel, hence disjoint, or have linearly independent direction vectors, intersecting in

exactly one point; this is the coordinate-geometry fact underlying every step of both directions here. See Book 1, field: Geometry, topic: lines and direction vectors.

Lines as level sets of a linear functional: $L = \{\vec{x} : \vec{v} \cdot \vec{x} = c\}$ for a nonzero normal $\vec{v}$; two such lines are parallel exactly when their normals are scalar multiples of each other, and the alternative reinterpretation route solves the problem's system by checking that the two normals are linearly independent, which is the same dichotomy stated in the language of linear algebra rather than transformations. See Book 1, field: Linear algebra, topic: rank and linear independence in two dimensions.

There is no recursive or self-referential structure in this problem: nothing here is defined in terms of earlier terms of itself, so there is no recursion to generalise for this particular question.

Circuit Bank

- **WHEN** a vector equation has the shape output minus a named fixed point equals a nonzero scalar times input minus the same fixed point **DO** name it immediately as a dilation centred at that point with that ratio, and reach for the known dilation-to-line properties rather than coordinates **BECAUSE** the equation is not evidence of a dilation, it is the definition of one, so recognising it is free, and skipping straight to coordinates throws away that toolkit
- **WHEN** you must disprove a claim that begins for every $\lambda \neq 0$ and every point P **DO** search first for the single easiest value of the free parameter that makes the defining condition collapse to something visibly impossible, here $\lambda = 1$ collapsing a dilation to the identity, rather than adapting the general-case construction **BECAUSE** one witness refutes a universal statement, and the trivialising value of a parameter usually costs one line instead of several
- **WHEN** a dilation's image of a line needs to be compared with another line for parallelism **DO** compare directions, not equations, since a dilation multiplies every direction vector by the same nonzero scalar and so preserves the parallel or not-parallel status of any pair of lines it is applied to **BECAUSE** direction-vector comparison is coordinate-free and immune to the vertical-line edge case that plagues slope-based coordinate arguments

Backtesting

1. First useful cue: the vector equation $\overrightarrow{PA_2} = \lambda \overrightarrow{PA_1}$ with P and λ both universally quantified.
2. Method the expression suggested: rename A_2 as the image of A_1 under the dilation centred at P with ratio λ, and work with what dilations do to lines.
3. Where you would hesitate, and the thought to reinforce: wondering whether the backward direction also needs to handle a general λ; reinforce that negating for every λ only needs one value, and $\lambda = 1$ is always worth trying first because it turns the transformation into the identity.
4. The trap the question sets: it invites a coordinate case split on vertical against non-vertical lines, which never actually resolves cleanly and hides the one-line argument behind unnecessary algebra.
5. What makes you recognise it faster next time: seeing the syntactic pattern output minus a named point equals scalar times input minus the same point, and reading it as a transformation's defining equation on sight, before any other processing of the problem.
6. The reusable decision rule: when a fixed point and a free nonzero scalar act uniformly on the whole plane inside a vector identity, default to transformation language, dilation here, translation or rotation for other shapes of equation, over raw coordinates.

Mirror drills

1. Two distinct circles have centres $O_1 \neq O_2$ and radii $r_1 \neq r_2$. Show there is a point P and a ratio λ so that the dilation centred at P with ratio λ sends the first circle onto the second. *Answer:* take $\lambda = r_2/r_1$ and solve $\overrightarrow{PO_2} = \lambda \overrightarrow{PO_1}$ for P on the line through O_1 and O_2; the same output-minus-centre-equals-scalar-times-input pattern pins down the external centre of similitude.

2. For which nonzero values of λ does the dilation $X \mapsto P + \lambda(X - P)$ send every line through P to itself as a set? *Answer:* every $\lambda \neq 0$, since points collinear with the centre stay collinear with the centre after scaling; this is the companion fact to a line missing the centre goes to a parallel line.

3. With $L_1 : x = 0$ and $L_2 : y = 1$, $P = (2, 2)$, and $\lambda = 3$, find $A_1 \in L_1$ and $A_2 \in L_2$ with $\overrightarrow{PA_2} = 3\overrightarrow{PA_1}$. *Answer:* write $A_1 = (0, t)$, so $A_2 = P + 3(A_1 - P) = (-4, 3t - 4)$; forcing the y-coordinate to be 1 gives $t = 5/3$, so $A_1 = (0, 5/3)$ and $A_2 = (-4, 1)$.

9.8 B2. sums of consecutive integers and divisor counting

Suppose that a positive integer N can be expressed as the sum of k consecutive positive integers

$$N = a + (a+1) + (a+2) + \cdots + (a+k-1)$$

for $k = 2017$ but for no other values of $k > 1$. Considering all positive integers N with this property, what is the smallest positive integer a that occurs in any of these expressions?

Read it out loud

Read this: a positive integer N has a representation as a sum of 2017 consecutive positive integers, and that representation is the only one, for any run length $k > 1$, that N possesses. The question does not want a single example, it wants the smallest starting value a across every N with this exact uniqueness property. The phrase for $k = 2017$ but for no other values of $k > 1$ is the whole problem: it is an exact-count condition, exactly one valid run length, dressed up as a sum. The field wired in here is elementary number theory, specifically the arithmetic of factorising $2N$ and counting divisors; underneath the summation this is a divisor-counting problem, not a series problem.

Within the first thirty seconds you should have converted $N = a + (a+1) + \cdots + (a+k-1)$ into the algebraic identity $2N = k(2a+k-1)$, because every subsequent idea, the parity of the two factors, the bijection with divisors, the uniqueness count, lives inside that one rewritten equation. The specific numeral 2017 is not decoration either; a run length that is a large prime is exactly the kind of number that can only occupy one very restricted role in a factor pair, and noticing that is most of the route to the answer.

Approaches, ranked worst to best

1. **Fix $k = 2017$ and search.** *Cue:* the problem names a specific run length. *Circuit:* substitute $k = 2017$ into $N = a + \cdots + (a + 2016)$, then test small values of a by hand to see when no other k seems to work. (why it ranks here: checking no other $k > 1$ works for even one candidate N means checking infinitely many other values of k individually, which is not a finite computation and never becomes a proof, only a list of unverified examples).
2. **Parity of the two factors of $2N$.** *Cue:* $2N = k(2a+k-1)$ splits into a pair of factors of opposite parity. *Circuit:* for the fixed value $k = 2017$, deduce that the complementary factor $2a + 2016$ must have no odd divisor greater than one, that is, must be a pure power of two. (why it ranks here: this is close to the fastest possible route once stated, but stated on its own it does not explain why no odd divisor greater than one is the correct translation of no other k works, it asserts the punchline rather than deriving it from a count).
3. **Bijection between valid k and divisors of the odd part of N.** *Cue:* exactly one k works is an exact-count statement. *Circuit:* split $2N = 2^s \cdot u$ with u odd, show every opposite-parity factor pair of $2N$ corresponds to exactly one divisor of u, and translate the hypothesis into a statement about how many divisors u has. (why it ranks here: it derives the previous approach's punchline instead of asserting it, turns an infinite case check into a single divisor count, and hands you both the necessary condition and the reason it is also sufficient in one pass).

The ideal solution

You will show the answer is $a = 16$.

Step 1: rewrite the sum as a factorisation of $2N$. If $N = a + (a+1) + \cdots + (a+k-1)$ with $a \geq 1$ and $k \geq 2$, then $N = ka + \binom{k}{2}$, so

$$2N = k(2a+k-1).$$

Write $m = 2a + k - 1$. Since $a \geq 1$, $m \geq k + 1 > k$; since $m - k = 2a - 1$ is odd, k and m have opposite parity. Conversely, given any factorisation $2N = km$ with $1 \leq k < m$ and k, m of opposite parity, setting $a = (m - k + 1)/2$ recovers an integer $a \geq 1$, the numerator is even because $m - k$ is odd, and positive because $m > k$, and this a satisfies $2N = k(2a + k - 1)$. So the valid run lengths $k \geq 2$ for a fixed N correspond exactly, one for one, to the ways of splitting $2N$ into two factors of opposite parity with the smaller factor called k; the case $k = 1$ corresponds to the trivial split $m = 2N$, $k = 1$, and is excluded by hypothesis.

Step 2: count those splittings by counting divisors. Write $2N = 2^s u$ with u odd and $s \geq 1$, the exponent of 2 in $2N$ is at least 1 since $2N$ is even. In any opposite-parity factorisation $2N = km$, exactly one of k, m is odd and the other even; the odd one, being coprime to 2, must be a divisor d of u, and the even one is then forced to be $2^s(u/d)$, since it must absorb every factor of 2 in $2N$. So the opposite-parity factorisations of $2N$ correspond exactly to the divisors d of u, via $\{k, m\} = \{d, 2^s u/d\}$ and $k = \min(d, 2^s u/d)$.

For $d = 1$: $k = \min(1, 2^s u) = 1$, the excluded trivial case. For any divisor $d \geq 2$ of u: since $d \leq u$, the complementary factor satisfies $2^s u/d \geq 2^s \geq 2$, so $k = \min(d, 2^s u/d) \geq 2$ is a genuinely valid run length. Different divisors $d \geq 2$ give different values of k: if k is odd then $k = d$, the odd member of the pair, which determines d; if k is even then $k = 2^s u/d$, which likewise determines $d = 2^s u/k$. So the map from d to k is a bijection from the set of divisors of u that are at least 2 onto the set of valid run lengths $k > 1$. This set has size $\tau(u) - 1$, where $\tau(u)$ is the number of positive divisors of u.

Step 3: impose the hypothesis. You are told the valid run lengths $k > 1$ form the single-element set $\{2017\}$, so $\tau(u) - 1 = 1$, that is, $\tau(u) = 2$: the odd part u of $2N$ is prime. The run length $k = 2017$ arises from some divisor d of u with $\min(d, 2^s u/d) = 2017$; since 2017 is odd, it can only be the odd member of the pair, the other member $2^s u/d$ is even because $s \geq 1$, so $d = 2017$, forcing $2017 \mid u$. Combined with u prime, this gives $u = 2017$ exactly. You should check 2017 is itself prime: since $2017 < 45^2 = 2025$, it suffices to test the primes up to 43, namely $2, 3, 5, 7, 11, 13, 17, 19, 23, 29, 31, 37, 41, 43$; direct division shows none of them divides 2017.

Step 4: solve for a and minimise. With $u = 2017$ and $d = 2017$ giving $k = 2017$, the complementary factor is $m = 2^s u/d = 2^s$, and $k < m$ requires $2017 < 2^s$; the least such s is $s = 11$, since $2^{10} = 1024 < 2017 < 2048 = 2^{11}$. Then

$$a = \frac{m - k + 1}{2} = \frac{2^s - 2016}{2} = 2^{s-1} - 1008,$$

which is strictly increasing in s, so it is minimised at the least admissible $s = 11$:

$$a = 2^{10} - 1008 = 1024 - 1008 = \boxed{16}.$$

At $s = 11$: $2N = 2^{11} \cdot 2017 = 2048 \cdot 2017$, $k = 2017$, $m = 2048$, and indeed $N = 2017 \cdot 1024$ with $a = 16$ gives exactly one representation as a sum of more than one consecutive positive integer, the length-2017 one.

Layer 1: WHAT

Block 1: convert the consecutive sum into $2N = k(2a + k - 1)$ and read off that valid k's correspond to opposite-parity factor pairs of $2N$. Block 2: split $2N = 2^s u$ and see that an opposite-parity factor pair is exactly a choice of divisor d of the odd part u. Block 3: turn exactly one $k > 1$ works, namely 2017, into the divisor count $\tau(u) - 1 = 1$, hence u prime, hence, since $2017 \mid u$, $u = 2017$. Block 4: solve $2^s > 2017$ for the least s, giving $a = 2^{s-1} - 1008$ minimised at $s = 11$, so $a = 16$.

The nominal route fixes $k = 2017$ from the start and tries to rule out every competitor run length $k = 2, 3, 4, \ldots$ one at a time by direct divisibility checks on N, an approach that can produce examples of N that work for small cases but can never certify that no other k works, because there are infinitely many candidates to rule out. The bifurcation point sits right after writing

$2N = k(2a + k - 1)$: do you keep k fixed and hunt through values of a, or do you fix the arithmetic of $2N$, its power-of-two part and its odd part, and count, all at once, every k that could possibly work. The reroute is to stop asking whether this particular k fails and start asking how many k succeed in total, which is a single divisor count rather than an unbounded sequence of checks.

Layer 2: HOW

Block 1. **Cue.** The problem phrases a sum of consecutive positive integers with an explicit run length k and starting value a. **Why natural.** Expanding the arithmetic series and doubling to clear the fraction is the standard first move for any sum of consecutive integers statement, and it immediately exposes $2a + k - 1$, a quantity whose parity is forced, opposite to k's, purely because $2a - 1$ is odd. **Link to the next block.** Once $2N = k(2a + k - 1)$ is written down, the natural next question is which pairs of positive integers can appear as k and $2a + k - 1$, which is a question about factor pairs of the single fixed number $2N$.

Block 2. **Cue.** The two factors k and $m = 2a + k - 1$ are forced to have opposite parity, and $2N$ is a specific even number. **Why natural.** Any even number splits uniquely into a power of two and an odd part, and an opposite-parity factor pair must put the entire power-of-two part on one side, leaving only the odd part free to be divided between the two factors; that freedom is exactly a choice of divisor of the odd part. **Link to the next block.** This turns how many valid k are there into how many divisors does the odd part have, a completely different and finite kind of question from the one you started with.

Block 3. **Cue.** The hypothesis says the count of valid $k > 1$ is exactly one. **Why natural.** A set of divisors of u that are at least 2 has exactly one element precisely when u has exactly two divisors in total, which is the definition of u being prime; and the given $k = 2017$, being odd, can only have arisen from the divisor $d = 2017$ of u, pinning u down completely once you know it is prime. **Link to the next block.** With $u = 2017$ fixed, every remaining unknown, a, and the free exponent s, is now determined by a single inequality and a single formula, so the problem has become pure optimisation over one integer parameter.

Block 4. **Cue.** a is expressed as an increasing function of the single remaining free parameter s, subject to the constraint $2^s > 2017$. **Why natural.** A strictly increasing function of an integer parameter is minimised at the smallest admissible value of that parameter, so the whole optimisation collapses to solving one inequality, $2^s > 2017$, by comparing powers of two you already know. **Link to the next block.** There is no further block: substituting the least admissible $s = 11$ gives the numeric answer directly.

Layer 3: WHY

The fact underneath every step is a named classical result, sometimes called the polite number theorem: the number of ways to write a positive integer N as a sum of two or more consecutive positive integers equals the number of divisors of the odd part of N, minus one; in particular N has no such representation at all exactly when N is a power of two. This problem is that theorem wearing a disguise, asking not how many representations in general but for which N is that count equal to exactly one, achieved at the specific length 2017. Once you know the theorem exists, you know in advance that the answer must reduce to a statement about the odd part of N having very few divisors, and you go looking for the bijection that proves it rather than inventing case-by-case divisibility arguments.

The discernment mechanism to train is this: whenever a problem asks about the number of ways to express N as a sum, a product, or a partition of some constrained shape, and then attaches

an exact count or a uniqueness condition to that number of ways, expect the true content to be a divisor-counting identity for some auxiliary integer built from N, here the odd part u of $2N$. The examiners nudge you toward this by making the target run length a specific prime, 2017; that choice is not decorative, a prime run length is exactly the kind of number that can only occupy the role of the one and only nontrivial divisor of the auxiliary odd part, which is what forces u itself to be exactly that prime once you know u has just two divisors. Recognising a specific numeral in a competition problem as chosen because it is prime is itself a trainable reflex: check primality of any distinguished constant before doing anything else with it.

Revolving concepts

Sum of consecutive positive integers as a factor pair: $N = a + \cdots + (a + k - 1)$ for $a \geq 1$, $k \geq 2$, if and only if $2N = km$ for some factorisation into $1 \leq k < m$ of opposite parity, via $m = 2a + k - 1$ and $a = (m - k + 1)/2$. See Book 1, field: Number theory, topic: consecutive integer sums.

The polite number theorem: the count of representations of N as a sum of two or more consecutive positive integers equals $\tau(u) - 1$, where u is the odd part of N, equivalently of $2N$, and τ counts positive divisors; N is representable in no such way exactly when N is a power of two. See Book 1, field: Number theory, topic: polite numbers and odd divisors.

Splitting an integer into its power-of-two part and odd part, $n = 2^s u$ with u odd: every factorisation of n into factors of prescribed opposite parity is obtained by choosing which divisor of u plays the odd role, with the even role forced; this all the twos go to one side idea recurs whenever a problem mixes parity constraints with factorisation. See Book 1, field: Number theory, topic: two-adic valuation.

The divisor-counting function τ and the fact that $\tau(n) = 2$ exactly when n is prime: the bridge between exactly one nontrivial representation and the auxiliary number is prime. See Book 1, field: Number theory, topic: the divisor function.

There is no recursive sequence inside this problem itself, the run $a, a + 1, \ldots, a + k - 1$ is consecutive integers, not a recursively defined sequence, so there is no recursion to generalise here beyond noting that the same split into power-of-two part and odd part, then count divisors of the odd part method reappears verbatim in any problem about the number of ways to write N as a sum, or a difference of two squares, or a product of two factors of prescribed parity.

Circuit Bank

- **WHEN** a problem talks about writing N as a sum of k consecutive positive integers, for one or several values of k **DO** rewrite immediately as $2N = k(2a + k - 1)$ and read the two factors as an opposite-parity pair **BECAUSE** this single identity converts every question about consecutive-integer sums into a question about factorisations of the fixed number $2N$
- **WHEN** a problem attaches an exact count, or a uniqueness statement, to the number of ways some representation of N exists **DO** look for a bijection between the valid representations and the divisors of an auxiliary integer built from N, very often its odd part, then translate the count into a statement about that auxiliary integer's divisor count **BECAUSE** exactly one, or exactly some small number, is almost always secretly a statement that this integer has exactly two, or a few, divisors, and the divisor function is the only standard tool that produces exact small counts
- **WHEN** an even number needs to be split into two factors of specified, opposite parity **DO** factor out the full power of two first, $n = 2^s u$ with u odd, then note the even factor must swallow the entire 2^s, leaving only divisors of u free to distribute **BECAUSE** this avoids ever having to case split on how many twos end up on which side; there is only ever one place for them to go

- **WHEN** a distinguished numeral in a competition problem, here 2017, needs to be identified as prime or composite before you can use it **DO** trial-divide by every prime up to its square root and state the check explicitly, rather than assuming primality because the number looks like it should be prime **BECAUSE** a full-marks proof cannot rest on an unverified numerical claim, and the size of Putnam numerals is always small enough for a short, explicit trial division

Backtesting

1. First useful cue: sum of k consecutive positive integers, for $k = 2017$ but no other $k > 1$.
2. Method the expression suggested: rewrite as $2N = k(2a + k - 1)$, an opposite-parity factor pair of the fixed number $2N$.
3. Where you would hesitate, and the thought to reinforce: not seeing how to convert no other k works into something finite; reinforce that the polite number theorem turns representation counts into divisor counts of the odd part, so hunt for that bijection rather than testing individual k's.
4. The trap the question sets: it invites you to try to rule out $k = 2, 3, 4, \ldots$ one at a time, which is an unbounded search that never becomes a completed proof.
5. What makes you recognise it faster next time: seeing sum of consecutive positive integers and a count or uniqueness condition together, and firing the divisor-of-the-odd-part reflex before writing anything else.
6. The reusable decision rule: whenever a competition problem fixes a fussy-looking specific numeral, a large prime such as 2017, inside a divisibility or factorisation context, check whether that numeral's primality is the entire point, and verify it explicitly by trial division.

Mirror drills

1. How many ways can 45 be written as a sum of two or more consecutive positive integers? *Answer:* 45 is odd so all six of its divisors, $1, 3, 5, 9, 15, 45$, are odd, giving $6 - 1 = 5$ representations.
2. Show that a positive integer N has no representation as a sum of two or more consecutive positive integers if and only if N is a power of two. *Answer:* the odd part of N has only the divisor 1 exactly when N is a power of two, so the representation count $\tau(u) - 1$ is zero exactly then.
3. Find the smallest N that can be written as a sum of consecutive positive integers, length greater than 1, in exactly three different ways. *Answer:* you need the odd part of N to have exactly four divisors; the smallest odd number with four divisors is $15 = 3 \times 5$, and $N = 15$ itself works: $15 = 7 + 8 = 4 + 5 + 6 = 1 + 2 + 3 + 4 + 5$.

9.9 B3. power series digits and a denominator clash

Suppose that $f(x) = \sum_{i=0}^{\infty} c_i x^i$ is a power series for which each coefficient c_i is 0 or 1. Show that if $f(2/3) = 3/2$, then $f(1/2)$ must be irrational.

Read it out loud

Read the hypothesis literally: every coefficient of this power series is a 0 or a 1, nothing else. That single sentence tells you $f(1/2) = \sum c_i 2^{-i}$ is not merely some real number built from a series, it is the binary expansion of a real number, digit for digit, with c_i sitting in the 2^{-i} place. The problem is not really about power series at all beneath the surface, it is about binary expansions and the arithmetic fact that governs which of them represent rational numbers. The two evaluation points, $2/3$ and $1/2$, are chosen so that one of them, $1/2$, reads off the digits directly, and the other, $2/3$, is where the actual numerical information, the given value $3/2$, lives; your job is to move information from one point to the other.

Because the conclusion is must be irrational, the shape of the whole solution is fixed before you compute anything: this will be a proof by contradiction, since there is no way to directly construct or bound an irrational number from a series about which you know almost nothing except its value at a different point. The fields wired in are real analysis, base expansions and rationality, and generating functions, turning a periodic coefficient sequence into a genuine rational function of x. Within the first thirty seconds you should have written down that if $f(1/2)$ is rational then its binary digit sequence (c_i) is eventually periodic, because that is the only way rational ever manifests itself as a concrete, usable structural fact about a digit sequence.

Approaches, ranked worst to best

1. **Guess a specific sequence.** *Cue:* the hypothesis $f(2/3) = 3/2$ looks like it should pin down (c_i). *Circuit:* try to construct, digit by digit, one sequence of 0s and 1s achieving $f(2/3) = 3/2$, then compute $f(1/2)$ for that sequence. (why it ranks here: the sequence achieving $f(2/3) = 3/2$ is not unique, so even a successful computation for one example proves nothing about every sequence satisfying the hypothesis; this cannot become a proof of a universal claim).

2. **Direct estimation of $f(1/2)$.** *Cue:* $f(1/2)$ is a sum of a series with bounded terms. *Circuit:* bound $f(1/2)$ between two nearby rationals using the constraint $f(2/3) = 3/2$, hoping to squeeze it away from every rational number. (why it ranks here: bounding a number between two rationals that are close together never rules out that number itself being rational, so this line of attack has no natural way to finish and stalls before reaching a contradiction).

3. **Contrapositive via eventual periodicity of the binary expansion.** *Cue:* coefficients restricted to 0 or 1 is exactly a binary digit sequence, and must be irrational calls for contradiction. *Circuit:* assume $f(1/2)$ is rational, invoke that a real number is rational if and only if its base-2 expansion is eventually periodic, rewrite $f(x)$ as a genuine rational function of x using that periodicity, then evaluate at $x = 2/3$ and compare the parity of denominators against the given value $3/2$. (why it ranks here: it replaces an infinite series with a finite algebraic expression, and the contradiction it produces is a clean, checkable fact about odd against even denominators rather than an open-ended estimate).

The ideal solution

Suppose, for contradiction, that $f(1/2)$ is rational.

Step 1: extract periodicity. The number $f(1/2) = \sum_{i=0}^{\infty} c_i 2^{-i}$, with each $c_i \in \{0, 1\}$, is precisely the number whose base-2 expansion has digit c_i in position i. A real number is rational if and

only if its base-2 expansion is eventually periodic: the remainders that occur in the long-division algorithm for a fraction with denominator q lie in the finite set $\{0, 1, \ldots, q-1\}$, so some remainder must recur, and from that point on the digits repeat exactly as they did after the first occurrence of that remainder. Since $f(1/2)$ is assumed rational, there exist integers $n \geq 0$ and $m \geq 1$ such that

$$c_i = c_{i+m} \qquad \text{for all } i \geq n.$$

Step 2: convert periodicity into a rational function. Split off the pre-period terms and sum the repeating block as a geometric series in x^m:

$$f(x) = \sum_{i=0}^{n-1} c_i x^i + x^n \sum_{j=0}^{\infty} x^{jm} \sum_{i=0}^{m-1} c_{n+i} x^i = P(x) + \frac{x^n Q(x)}{1 - x^m},$$

where $P(x) = \sum_{i=0}^{n-1} c_i x^i$ and $Q(x) = \sum_{i=0}^{m-1} c_{n+i} x^i$ are honest polynomials with $0/1$ coefficients, and the identity is valid for $|x| < 1$, in particular at $x = 2/3$.

Step 3: evaluate at $x = 2/3$ and track denominators. The value $P(2/3) = \sum_{i=0}^{n-1} c_i (2/3)^i$, once written over a common denominator, equals an integer divided by 3^{n-1}; in lowest terms its denominator divides 3^{n-1}, hence is odd. Similarly $Q(2/3) = \sum_{i=0}^{m-1} c_{n+i} (2/3)^i$ equals an integer divided by 3^{m-1}. The second summand of $f(2/3)$ is

$$\frac{(2/3)^n Q(2/3)}{1 - (2/3)^m} = \frac{2^n 3^m}{3^{n+m-1}(3^m - 2^m)} \sum_{i=0}^{m-1} c_{n+i} 2^i 3^{m-1-i},$$

whose denominator, before any cancellation, is $3^{n+m-1}(3^m - 2^m)$. Since 3^m is odd and 2^m is even for every $m \geq 1$, the quantity $3^m - 2^m$ is odd; and 3^{n+m-1} is a power of an odd number. So this denominator is a product of two odd numbers, hence odd, and remains odd after any cancellation, since cancelling only removes factors that were already present in both numerator and denominator, never introduces the prime 2 into the denominator.

So $f(2/3) = P(2/3) + (\text{second summand})$ is a sum of two rational numbers each with odd denominator in lowest terms. Writing both over the common denominator $3^{n+m-1}(3^m - 2^m)$, itself odd, shows the sum, before reduction, again has odd denominator; reduction cannot turn an odd denominator into an even one. So $f(2/3)$, written in lowest terms, has odd denominator.

Step 4: contradiction. You are given $f(2/3) = 3/2$, and $3/2$ is already in lowest terms with denominator 2, which is even. A single rational number cannot simultaneously have odd denominator, from Step 3, and even denominator, from the hypothesis, in its own unique lowest-terms representation. This contradiction shows the assumption of Step 1 was false: $f(1/2)$ is irrational. ■

Layer 1: WHAT

Block 1: read $c_i \in \{0, 1\}$ as this is a binary expansion, so $f(1/2)$ is a specific real number's base-2 digit string. Block 2: since the conclusion is must be irrational, commit to a proof by contradiction and assume $f(1/2)$ is rational. Block 3: invoke the rational if and only if eventually periodic lemma to get $c_i = c_{i+m}$ for $i \geq n$. Block 4: turn that periodicity into a closed-form rational function $f(x) = P(x) + x^n Q(x)/(1 - x^m)$. Block 5: evaluate the closed form at $x = 2/3$ and show every term, hence the whole sum, has odd denominator. Block 6: contradict the given even denominator of $3/2$.

The nominal route tries to pin down an explicit sequence (c_i) satisfying $f(2/3) = 3/2$, perhaps by a greedy digit-by-digit construction reminiscent of a non-integer base expansion, and then attempts to read off properties of $f(1/2)$ directly from that one sequence. The bifurcation point comes immediately after writing suppose $f(1/2)$ is rational: the zombie move is to try to solve for the digits themselves, which spirals into an underdetermined, non-terminating construction, while

the reroute is to extract only the structural consequence of rationality, eventual periodicity, and carry that structure to the other evaluation point, where a concrete number is already known, rather than trying to pin down the sequence completely.

Layer 2: HOW

Block 1. **Cue.** Every coefficient of the power series is restricted to $\{0,1\}$, nothing else. **Why natural.** A series with only $0/1$ coefficients evaluated at $1/2$ is not an arbitrary real number built from a sum, it is the literal definition of a base-2 expansion with those digits, so any question about $f(1/2)$ is secretly a question about binary expansions. **Link to the next block.** Since the target conclusion is a rationality statement about a digit string, the natural tool is the standard fact linking rationality of a number to periodicity of its digit string.

Block 2. **Cue.** The problem demands you show a quantity must be irrational. **Why natural.** Irrationality is a negative, unbounded statement, a number is not expressible as any ratio of integers, and negative unbounded statements are proved by assuming the opposite and deriving a contradiction with something you can compute, since there is no positive construction of not a fraction. **Link to the next block.** Assuming $f(1/2)$ is rational immediately hands you, via the periodicity lemma, a finite amount of combinatorial data, the numbers n and m and the finitely many digits between them, to work with instead of an infinite unknown sequence.

Block 3. **Cue.** The periodicity you now have, $c_i = c_{i+m}$ for $i \geq n$, describes the coefficient sequence, not yet a usable number. **Why natural.** An eventually periodic sequence's generating function is always exactly a polynomial, the pre-period part, plus a geometric-series tail, because summing a periodic block infinitely many times is precisely what the identity for a geometric series in x^m computes. **Link to the next block.** This closed form $f(x) = P(x) + x^n Q(x)/(1 - x^m)$ can now be evaluated at any point inside the radius of convergence, in particular at the point $x = 2/3$ where the problem has already told you the exact value.

Block 4. **Cue.** The two rational evaluation points, $2/3$, denominator a power of 3, and the target value $3/2$, denominator 2, involve different primes in their denominators. **Why natural.** Every term in the closed form, evaluated at $2/3$, produces a denominator built only from powers of 3 and from $3^m - 2^m$, and $3^m - 2^m$ is odd because 3^m is odd and 2^m is even; no term can ever contribute a factor of 2 to a denominator, since $2/3$ itself has no 2 in its denominator to begin with. **Link to the next block.** Once every term is confirmed odd-denominatored, the sum is odd-denominatored, and this collides head-on with the given value $3/2$, whose denominator is unavoidably even; that collision is the entire content of the contradiction.

Layer 3: WHY

The mathematical truth beneath this problem is that rationality is not a property you can see directly in an infinite sum, it is a property of the structure of the coefficient sequence, periodicity, and structure is exactly the kind of thing that survives being transported to a different evaluation point, whereas a raw numerical value does not. Assuming $f(1/2)$ is rational buys you periodicity; periodicity buys you a rational-function formula for f as a function of x, valid everywhere inside the disc of convergence; and a rational-function formula is the one kind of object whose behaviour at $x = 2/3$ you can fully control and compare against a known target. The deeper reason the contradiction lands is arithmetic: $2/3$ only ever introduces the prime 3 into denominators, its own denominator, and the denominators $3^m - 2^m$ that arise from summing its powers, so nothing forces a factor of 2 into any denominator along the way, yet the hypothesis $f(2/3) = 3/2$ insists that a

factor of 2 must be there. The clash is unavoidable once periodicity is assumed, which is exactly why the contradiction is airtight rather than merely suggestive.

The discernment mechanism to train: whenever a problem gives you a series' value at two different rational points and asks you to deduce the rational or irrational nature of one from the given value at the other, expect the mechanism to be a denominator mismatch, most often a mismatch of which primes divide the denominator, obtained by assuming the structural consequence of rationality at the unknown point, periodicity, or more generally that the object is a rational function, and pushing that structure through to the known point. Whenever coefficients are restricted to a small finite set like $\{0, 1\}$, read that restriction immediately as these are digits in some base, because that reading is what makes the rationality lemma available at all; without it there is no route from rational value to periodic sequence.

Revolving concepts

Rationality and eventually periodic base-b expansions: a real number is rational exactly when its expansion in any fixed integer base is eventually periodic; the proof is a pigeonhole argument on the finitely many possible remainders in long division by the denominator. See Book 1, field: Real analysis, topic: base expansions of rational numbers.

Generating function of an eventually periodic sequence: if $c_i = c_{i+m}$ for all $i \geq n$, then the generating function of (c_i) equals a polynomial plus x^n times a polynomial over $1 - x^m$, built from the pre-period and one period block; more generally, any sequence satisfying a linear recurrence from some point on has a generating function equal to a polynomial divided by the reversed characteristic polynomial of that recurrence, obtained the same way every time, multiply the recurrence by x^i, sum over all valid i, and solve the resulting equation for the generating function. Eventual periodicity is exactly the special case of the recurrence $c_{i+m} = c_i$, whose characteristic polynomial $x^m - 1$ has every m-th root of unity as a root, giving the denominator $1 - x^m$. See Book 1, field: Generating functions, topic: rational generating functions from linear recurrences.

Denominator bookkeeping by prime: a sum or product of rational numbers each having odd denominator, more generally each having denominator coprime to a fixed prime p, again has denominator coprime to p, because a common denominator built from odd numbers is odd, and reducing a fraction never introduces a new prime factor into the denominator that was not already dividing the original denominator. See Book 1, field: Number theory, topic: valuations and denominators.

The identity $3^m - 2^m$ is always odd for $m \geq 1$: a special case of a power of an odd number minus a power of an even number is odd, used here as the one nonobvious denominator factor that must be checked by hand. See Book 1, field: Number theory, topic: parity arguments.

Power series with $0/1$ coefficients as indicator generating functions: $f(x) = \sum_{i \in A} x^i$ for $A = \{i : c_i = 1\}$, a subset of the nonnegative integers, so $f(1/2)$ is literally the real number whose binary expansion marks the elements of A. See Book 1, field: Functional equations, topic: generating functions of sets.

Circuit Bank

- **WHEN** a power series has coefficients restricted to $\{0, 1\}$ **DO** read $f(1/b)$ for any integer base $b \geq 2$ as the literal base-b digit expansion of a real number with those digits, not as an abstract sum **BECAUSE** the restriction to two symbols is exactly the definition of a digit string in that base, and every fact about rationality of digit strings becomes available the instant you make this identification
- **WHEN** you must prove a stated quantity must be irrational **DO** argue by contradiction: assume it is rational, extract the strongest available structural consequence, periodicity of a digit string, algebraicity of a series, and push that structure to wherever a concrete numerical

value is already known **BECAUSE** irrationality has no positive constructive proof in general, so the standard route is always to assume rationality, extract structure, and collide that structure with a fact you can compute

- **WHEN** a coefficient sequence is shown to be eventually periodic **DO** split the generating function into a finite pre-period polynomial plus a period-block polynomial times $x^n/(1-x^m)$ **BECAUSE** this is the unique way to sum an eventually periodic series in closed form, converting an infinite object into two polynomials and one geometric factor
- **WHEN** two rational evaluations of the same object need to be compared for a contradiction **DO** track which prime appears in each side's denominator rather than compute an exact common denominator **BECAUSE** the contradiction usually needs only one bit of information, whether a fixed prime divides the denominator or not, and full computation is unnecessary ink

Backtesting

1. First useful cue: coefficients restricted to $\{0,1\}$, combined with a given exact value at one point and a rationality question at another.
2. Method the expression suggested: contrapositive, plus the rational if and only if eventually periodic base expansion lemma.
3. Where you would hesitate, and the thought to reinforce: wondering whether rational implies eventually periodic needs its own proof inside the write-up; reinforce that this is a standard, nameable lemma about long division remainders, and it should be stated as such rather than waved through as obvious.
4. The trap the question sets: it tempts you to try to construct or guess the actual sequence (c_i) satisfying $f(2/3) = 3/2$, which is neither unique nor necessary for the proof.
5. What makes you recognise it faster next time: seeing each coefficient is 0 or 1 and immediately relabelling the series as a base-2 digit string rather than as an unstructured sequence.
6. The reusable decision rule: whenever an irrationality conclusion must follow from a given rational value at a different point, hunt for a prime that can appear in one side's denominator but not the other's, after first extracting the periodicity or rational-function structure that rationality would force.

Mirror drills

1. Prove the converse direction used in Step 2: if (c_i) is eventually periodic with period m from index n, then $\sum c_i 2^{-i}$ is rational. *Answer:* group into blocks of length m exactly as in Step 2; at $x = 1/2$ the closed form is manifestly a ratio of integers, since $P(1/2)$ and $Q(1/2)$ are rational and $1 - 2^{-m}$ is a nonzero rational.
2. If $g(x) = \sum d_i x^i$ has $d_i \in \{0,1\}$ and $g(1/2) = g(1/3)$, and this common value is rational, what primes can appear in its denominator? *Answer:* at $x = 1/2$ the closed form only ever produces denominators that are powers of 2 or factors of $2^m - 1$; at $x = 1/3$ it only ever produces powers of 3 or factors of $3^m - 1$; a common rational value forces the denominator to divide both kinds simultaneously, a strong restriction worth chasing with the same bookkeeping used above.
3. True or false: if f has coefficients in $\{0,1\}$ and $f(1/2)$ is rational, then f is a rational function of x. *Answer:* true, this is exactly the periodic-to-rational-function lemma of Step 2; the converse, that a rational function with a power series expansion having $0/1$ coefficients must

have eventually periodic coefficients, follows from the theory of linear recurrences via partial fractions.

9.10 B4. Evaluating a period four logarithmic series

Evaluate the sum

$$\sum_{k=0}^{\infty}\left(3\cdot\frac{\ln(4k+2)}{4k+2}-\frac{\ln(4k+3)}{4k+3}-\frac{\ln(4k+4)}{4k+4}-\frac{\ln(4k+5)}{4k+5}\right)$$
$$=3\cdot\frac{\ln 2}{2}-\frac{\ln 3}{3}-\frac{\ln 4}{4}-\frac{\ln 5}{5}+3\cdot\frac{\ln 6}{6}-\frac{\ln 7}{7}-\frac{\ln 8}{8}-\frac{\ln 9}{9}+3\cdot\frac{\ln 10}{10}-\cdots.$$

(As usual, $\ln x$ *denotes the natural logarithm of* x*.)*

Read it out loud

There is one sequence here and one weight. The sequence is $b_n=\dfrac{\ln n}{n}$, evaluated at every positive integer in turn. The weight is a pattern that repeats with period 4: weight 3 when $n\equiv 2 \pmod 4$, weight -1 on the other three residues. Say that sentence to yourself before you write anything else, because the second display in the question is the honest form of the object and the first display is only a bracketing of it. In symbols, set

$$\epsilon_n=\begin{cases}3 & n\equiv 2 \pmod 4\\ -1 & \text{otherwise,}\end{cases}$$

and the series is $\sum_n \epsilon_n b_n$, started at $n=2$. It costs nothing to start at $n=1$ instead, because $b_1=\ln 1=0$, and starting at $n=1$ is what makes the weight a genuinely periodic sequence rather than a periodic sequence with a hole in it. That is the first thing you should have on paper inside thirty seconds.

The second thing on paper is $3-1-1-1=0$. The weight has mean zero over one period. That single fact is doing all the work: it is why the series converges at all (the terms b_n only decay like $\ln n/n$, so no weight of nonzero mean could give a convergent series), and it is why the bracketed version converges absolutely while the unbracketed version does not. The fields wired into this problem are therefore real analysis first (conditional versus absolute convergence, and the discipline of never rearranging a conditionally convergent series) and analytic number theory second (a Dirichlet series with periodic coefficients, evaluated by differentiating at $s=1$).

The third thing on paper is a question, not a calculation: where can a closed form possibly come from? The only structural identity available to ln is $\ln(mn)=\ln m+\ln n$. Since the modulus is 4 and the special residue is 2, the only multiplicative relation in sight is doubling. Compute what doubling does to b:

$$b_{2m}=\frac{\ln(2m)}{2m}=\frac{\ln 2}{2m}+\frac{1}{2}\cdot\frac{\ln m}{m}=\frac{\ln 2}{2m}+\frac{1}{2}b_m.$$

That is an exact self-similarity: the doubling map sends b to half of itself plus a harmonic term carrying a factor $\ln 2$. So $\ln 2$ can enter the answer in exactly one way, through that harmonic term, and the harmonic term will have to be summed against an alternating sign to produce anything finite. Alternating harmonic sums to $\ln 2$. So before you have proved anything you already expect $(\ln 2)^2$, and a two second numerical check on the first few brackets (0.4805 to four places against $(\ln 2)^2=0.48045\ldots$) confirms it. The rest of the work is to get there without committing a rearrangement crime.

What is really being tested is whether you can hold a conditionally convergent object still. The graders on this problem were not looking for the value; the value is guessable. They were looking for a candidate who manipulates only finite partial sums and takes limits exactly once, at the end.

Approaches, ranked worst to best

1. **Numerical reconnaissance and recognition.** *Cue:* "Evaluate the sum" with no free parameter, so the answer is one number. *Circuit:* compute, recognise, then go and prove. Summing a few thousand brackets gives 0.48045, and $\sqrt{0.48045} = 0.6931 = \ln 2$. This is worth zero marks as a proof and is nevertheless compulsory: you should never spend twenty minutes proving a value you have not first checked. It ranks last only as a route to a written solution.

2. **Expand every logarithm and hope.** *Cue:* $\ln(4k+2) = 2\ln 2 + \ln(2k+1)$, $\ln(4k+4) = 2\ln 2 + \ln(k+1)$, and so on. *Circuit:* split every term into its $\ln 2$ part and its residual part. This is the zombie route. You end up with a heap of pieces, each of which is individually divergent ($\sum 1/n$, $\sum \ln n/n$), and you must track the cancellation between divergent objects by hand across a page of algebra. It can be made to work with asymptotic expansions and the Euler to Mascheroni constant, but the ink cost is enormous and every regrouping is a place to lose the marks.

3. **Telescoping the block.** *Cue:* the block has four terms whose weights sum to zero, and any mean zero block is a combination of consecutive differences. *Circuit:* write $a_k = b_k - b_{k+1}$ and express the block in the a basis. You find $3b_{4k+2} - b_{4k+3} - b_{4k+4} - b_{4k+5} = 3a_{4k+2} + 2a_{4k+3} + a_{4k+4}$, note that $\sum_{k\geq 1} a_k$ telescopes to 0 and converges absolutely, and then grind out the remaining $\ln 2$ by pairing $a_{4k+2} + a_{4k+3}$. This is a complete and correct route and it is the one the official solution takes first. It ranks third because it needs two separate ideas (the difference basis, and then a second regrouping to release the $\ln 2$) where the winning route needs one.

4. **Dirichlet series and differentiation at $s = 1$.** *Cue:* a periodic weight of mean zero times $\ln n/n$ is exactly $-\frac{d}{ds}$ of a Dirichlet series at $s = 1$. *Circuit:* build $L(s) = \sum \epsilon_n n^{-s}$, factor it, differentiate. This is the best explanation of the answer and you must know it (see Layer 3). As a written solution it ranks fourth because of the analysis rent: you owe the grader convergence of $L(s)$ for $s > 0$, the identity $L(s) = -\zeta(s)(1 - 2^{1-s})^2$ proved for $s > 1$ and then extended to $s = 1$, and a justification of term by term differentiation. That is a page of standard but unavoidable analysis for a value you can get in five lines.

5. **Decompose the weight, not the terms.** *Cue:* the weight is a period 4 integer sequence, and the doubling self-similarity $b_{2m} = \frac{\ln 2}{2m} + \frac{1}{2}b_m$ is the only multiplicative structure available. *Circuit:* write ϵ as an alternating pattern plus twice an alternating pattern pulled back along $n \mapsto 2n$. This is the ideal route. You never touch an infinite sum until the last line, the unknown value of $\sum(-1)^{n-1}b_n$ cancels against itself and never has to be evaluated, and the $(\ln 2)^2$ appears as a product of two separate copies of $\ln 2$, one from $\ln(2m) = \ln 2 + \ln m$ and one from the alternating harmonic series.

The ideal solution

Write $b_n = \dfrac{\ln n}{n}$ for $n \geq 1$, so that $b_1 = 0$, and let

$$\epsilon_n = \begin{cases} 3, & n \equiv 2 \pmod 4, \\ -1, & \text{otherwise.} \end{cases}$$

The Kth partial sum of the given series is

$$T_K = \sum_{k=0}^{K-1} \left(3b_{4k+2} - b_{4k+3} - b_{4k+4} - b_{4k+5}\right) = \sum_{n=2}^{4K+1} \epsilon_n b_n = \sum_{n=1}^{4K+1} \epsilon_n b_n,$$

the last equality because $b_1 = 0$.

Claim 9.10.1. *For every $n \geq 1$,*

$$\epsilon_n \;=\; -(-1)^{n-1} \;+\; \begin{cases} 2\,(-1)^{m-1}, & n = 2m, \\ 0, & n \text{ odd.} \end{cases}$$

Proof. Three cases. If n is odd the right side is -1, and $n \not\equiv 2 \pmod 4$, so $\epsilon_n = -1$. If $n = 2m$ with m odd, then $n \equiv 2 \pmod 4$ and the right side is $-(-1)^{2m-1} + 2(-1)^{m-1} = 1 + 2 = 3 = \epsilon_n$. If $n = 2m$ with m even, then $n \equiv 0 \pmod 4$ and the right side is $1 - 2 = -1 = \epsilon_n$. □

The even indices n with $1 \leq n \leq 4K+1$ are exactly $n = 2m$ with $1 \leq m \leq 2K$, so the Claim gives the finite identity

$$T_K \;=\; -\sum_{n=1}^{4K+1} (-1)^{n-1} b_n \;+\; 2\sum_{m=1}^{2K} (-1)^{m-1} b_{2m}.$$

Now use $b_{2m} = \dfrac{\ln 2}{2m} + \dfrac{1}{2} b_m$, which is just $\ln(2m) = \ln 2 + \ln m$ divided by $2m$:

$$2\sum_{m=1}^{2K} (-1)^{m-1} b_{2m} \;=\; (\ln 2) \sum_{m=1}^{2K} \frac{(-1)^{m-1}}{m} \;+\; \sum_{m=1}^{2K} (-1)^{m-1} b_m.$$

Abbreviate $A_N = \sum_{n=1}^{N} (-1)^{n-1} b_n$ and $H_N = \sum_{m=1}^{N} \frac{(-1)^{m-1}}{m}$. Combining the last two displays,

$$T_K \;=\; -A_{4K+1} \;+\; A_{2K} \;+\; (\ln 2)\, H_{2K}, \qquad K \geq 1. \tag{$*$}$$

Every step so far is an identity between finite sums; nothing has been rearranged.

It remains to take limits in $(*)$. The function $x \mapsto \frac{\ln x}{x}$ has derivative $\frac{1 - \ln x}{x^2} < 0$ for $x > e$, so the sequence $(b_n)_{n \geq 3}$ is strictly decreasing, and $b_n \to 0$. By the alternating series test applied from $n = 3$ onwards (the terms $n = 1, 2$ form a finite head and do not affect convergence), A_N converges to a finite limit A. Also $H_N \to \ln 2$, the alternating harmonic series. Therefore

$$\lim_{K \to \infty} T_K \;=\; -A + A + (\ln 2)(\ln 2) \;=\; \boxed{(\ln 2)^2}.$$

Two points a grader will want addressed explicitly, both of which are cheap.

First, $\lim_K T_K$ is exactly what the question asks for, since the bracketing displayed in the problem *is* the definition of the sum. In fact the bracketed series converges absolutely, so the value is bracket independent in the strongest sense: writing the block as

$$3b_{4k+2} - b_{4k+3} - b_{4k+4} - b_{4k+5} = (b_{4k+2} - b_{4k+3}) + (b_{4k+2} - b_{4k+4}) + (b_{4k+2} - b_{4k+5}),$$

each difference is at most $3 \max_{x \geq 4k+2} \left| \frac{d}{dx} \frac{\ln x}{x} \right| = O\left(\frac{\ln k}{k^2} \right)$ by the mean value theorem, and $\sum_k \frac{\ln k}{k^2} < \infty$.

Second, the value A of the alternating series $\sum_{n \geq 1} (-1)^{n-1} \frac{\ln n}{n}$ was never needed. It is a genuine constant (it equals $\gamma \ln 2 - \frac{(\ln 2)^2}{2}$, which you are not expected to know and did not use); all that mattered was that it exists, so that $-A_{4K+1}$ and A_{2K} have the same limit and cancel.

Layer 1: WHAT

Block 1. Name the object: one sequence $b_n = \ln n / n$, one period 4 weight ϵ_n of mean zero, sum $= \sum \epsilon_n b_n$ with $b_1 = 0$ so that the sum may start at $n = 1$.

Block 2. Ask where $\ln 2$ can come from. Only from $\ln(2m) = \ln 2 + \ln m$. So record the doubling self-similarity $b_{2m} = \frac{\ln 2}{2m} + \frac{1}{2} b_m$ and decide to build the solution around the map $n \mapsto 2n$.

Block 3. Decompose the *weight* in terms of that map: $\epsilon_n = -(-1)^{n-1}$ plus $2(-1)^{m-1}$ on the even indices $n = 2m$. Verify by three residue cases.

Block 4. Substitute. The $\frac{1}{2}b_m$ half of the self-similarity regenerates the same alternating sum you subtracted; the $\frac{\ln 2}{2m}$ half becomes $\ln 2$ times an alternating harmonic sum.

Block 5. Take limits of the three finite sequences in $(*)$, each of which you have proved converges. The unknown constant cancels, the answer is $(\ln 2)^2$.

Juxtaposition. The nominal route also starts by splitting logarithms, but it splits them *inside the terms*: $\ln(4k+2) = 2\ln 2 + \ln(2k+1)$, $\ln(4k+3)$ stays, $\ln(4k+4) = 2\ln 2 + \ln(k+1)$, $\ln(4k+5)$ stays. You then have four families of terms and a pile of $\ln 2$ multiples, none of which sums separately, and you attempt to reassemble $\sum_n \ln n/n$ out of the pieces while the harmonic parts diverge underneath you. The ideal route splits the same logarithm but only after deciding which indices to apply it to, and it decides that by decomposing the weight.

Bifurcation point. The exact line where the routes split is the line after you have written ϵ_n down. The zombie asks "what are the terms?". You ask "what is the weight, as a sequence, in terms of the maps under which b is self-similar?". Everything downstream is forced.

Layer 2: HOW

Block 1. Cue. Two displays are given, and the second one has no brackets. That is the setter telling you the object is a weighted sum over all n, and that the bracketing is a convergence courtesy rather than part of the definition. **Why natural.** Whenever a series is presented in blocks of equal length, the honest object is (weight sequence) times (term sequence), and the block length is the period of the weight. Writing ϵ_n separates the two so that you can work on each independently. Recording $b_1 = 0$ is not cosmetic: it is what lets the weight be a clean periodic sequence starting at $n = 1$, which is what every later step needs. **Link to the next block.** With the two factors separated, the next question is forced: which of the two carries the structure that could produce a closed form?

Block 2. Cue. The answer must be a closed form, and the only closed form constants available from logarithms of integers are ln of small integers. The modulus is 4, so the small integer in play is 2. **Why natural.** ln turns multiplication into addition. Therefore the only maps on indices that interact usefully with $b_n = \ln n/n$ are the multiplications $n \mapsto dn$, under which $b_{dn} = \frac{\ln d}{dn} + \frac{1}{d}b_n$. This is not a trick; it is the defining property of ln applied to the only place it can be applied. Note the shape of the result: an exact affine self-similarity, with a harmonic remainder that carries the constant $\ln d$. **Link to the next block.** The self-similarity is stated in terms of $n \mapsto 2n$, so if you want to use it you must arrange for the weight to be expressed in terms of $n \mapsto 2n$ too.

Block 3. Cue. ϵ is -1 almost everywhere and $+3$ on one residue class, so $\epsilon_n = -1 + 4 \cdot [\, n \equiv 2 \pmod 4 \,]$; and the condition $n \equiv 2 \pmod 4$ says exactly "$n = 2m$ with m odd". **Why natural.** That reading of the residue class is the whole move. "$n \equiv 2 \pmod 4$" looks like a statement about n; it is really a statement about $m = n/2$, namely that m is odd. Once the special class is described in terms of m, the alternating sign $(-1)^{m-1}$ appears by itself, because "m odd" is the support of $\frac{1+(-1)^{m-1}}{2}$. Matching constants then gives $\epsilon_n = -(-1)^{n-1} + 2(-1)^{m-1}[n = 2m]$, which you confirm in three lines. **Link to the next block.** Now both factors are expressed through the same map, so substitution will make the two alternating sums collide.

Block 4. Cue. You have $-\sum(-1)^{n-1}b_n$ and $+2\sum(-1)^{m-1}b_{2m}$ side by side, and you know what b_{2m} is. **Why natural.** Substituting the self-similarity is the only available move, and it is a good one because the coefficient 2 in front is exactly what turns $2 \cdot \frac{1}{2}b_m$ into b_m, restoring the alternating sum you subtracted. That is not luck: the 2 came from ϵ, and ϵ was built to make the two pieces match, because you chose the decomposition with the doubling identity already in view. **Link to the next block.** What survives is $2 \cdot \frac{\ln 2}{2m}$ summed with alternating signs, which is $\ln 2$ times the alternating harmonic series.

Block 5. Cue. Three finite sums, three separate limits. **Why natural.** Because you kept everything finite, taking limits is legal the instant you prove each sequence converges, and the proofs are one line each (alternating series test for A_N; the standard alternating harmonic value for H_N). Had you rearranged the infinite series instead, you would have needed absolute convergence of a series that does not converge absolutely, and the Riemann rearrangement theorem says the result could have been anything. **Link to the next block.** Nothing follows; you are done. Then check the value numerically, which you had already done before starting.

Layer 3: WHY

The mathematical truth underneath is that periodic weights against n^{-s} are Dirichlet series, and multiplying a Dirichlet series by a factor of the form $1 - c\,d^{-s}$ is exactly the coefficient level operation "subtract c times the value at index n/d". That is the operation you performed in Block 3, in the opposite direction.

Concretely, set $L(s) = \sum_{n\geq 1} \epsilon_n n^{-s}$. For $s > 1$ everything converges absolutely and you may split by residue class:

$$L(s) = -\zeta(s) + 4\sum_{j\geq 0}(4j+2)^{-s} = -\zeta(s) + 4\cdot 2^{-s}(1-2^{-s})\zeta(s) = \zeta(s)\big(-1 + 2^{2-s} - 2^{2-2s}\big),$$

using $\sum_{j\geq 0}(2j+1)^{-s} = (1-2^{-s})\zeta(s)$. Now put $u = 2^{1-s}$: the bracket is $-1 + 2u - u^2 = -(1-u)^2$. Hence

$$L(s) = -\zeta(s)\big(1-2^{1-s}\big)^2 = -\eta(s)\big(1-2^{1-s}\big), \qquad \eta(s) := \sum_{n\geq 1}\frac{(-1)^{n-1}}{n^s} = (1-2^{1-s})\zeta(s).$$

Since $\frac{d}{ds}n^{-s} = -(\ln n)n^{-s}$, the sum you want is $S = -L'(1)$. The factor $1 - 2^{1-s}$ vanishes at $s = 1$ with derivative $\ln 2$, and $\eta(1) = \ln 2$, so

$$L'(1) = -\eta'(1)\cdot 0 - \eta(1)\ln 2 = -(\ln 2)^2, \qquad S = (\ln 2)^2.$$

Read the same computation through ζ instead: ζ has a simple pole at $s = 1$ with residue 1, and $(1-2^{1-s})^2$ has a double zero there with leading coefficient $(\ln 2)^2$; pole times double zero is a simple zero, and $-L'(1)$ picks out (residue) times (leading coefficient) $= (\ln 2)^2$. The answer is a perfect square because the weight is, up to sign, a Dirichlet convolution square.

That is the discernment mechanism, and it generalises cleanly. Given $\sum_n \epsilon_n \frac{\ln n}{n}$ with ϵ periodic of period q and mean zero, form $L(s) = \sum \epsilon_n n^{-s}$ and ask whether $L(s)/\zeta(s)$ factors into terms $1 - c\,d^{-s}$ with $d \mid q$. If it factors with a *double* zero at $s = 1$ you will get an elementary answer with no Euler to Mascheroni constant, and the elementary proof is the coefficient level shadow of the factorisation; if it has only a *simple* zero, γ will appear (for the period 3 weight $(-1,-1,2)$ you get $L(s) = \zeta(s)(3^{1-s}-1)$, a single zero, and the sum is $\gamma\ln 3 - \frac{(\ln 3)^2}{2}$). So the factorisation tells you in advance how hard the answer will be.

To train the discernment: take periodic patterns of mean zero and practise writing them as integer combinations of the constant weight **1**, the alternating weight $(-1)^{n-1}$, and their pullbacks along $n \mapsto n/d$ for the divisors d of the period. Ten minutes of that makes Block 3 automatic. The examiners nudged toward it in the only way available on a printed paper: they wrote the second, unbracketed display so that you would see the pattern $3, -1, -1, -1$ repeating, and they chose the weights so that the pattern is -1 plus a spike, which is the easiest possible pattern to decompose.

Revolving concepts

The sequence $\ln n/n$**.** Definition level facts: $\frac{d}{dx}\frac{\ln x}{x} = \frac{1-\ln x}{x^2}$, so it increases on $(0,e)$ and decreases on (e,∞) with maximum $1/e$ at $x = e$; among integers the maximum is at $n = 3$. Hence $(b_n)_{n\geq 3}$

decreases to 0, which is precisely the hypothesis of the alternating series test. See Book 1, field: Real analysis, topic: monotonicity and the alternating series test.

Absolute versus conditional convergence. $\sum|\epsilon_n b_n|$ diverges (it dominates $\sum\frac{\ln n}{n}$ up to a constant), so the series is conditionally convergent at best; the bracketed series is absolutely convergent. The Riemann rearrangement theorem says a conditionally convergent real series can be rearranged to any real value or to $\pm\infty$, so a solution that rearranges without justification is not merely informal, it is wrong. See Book 1, field: Real analysis, topic: rearrangement.

Dirichlet's test and the convergence of $\sum\epsilon_n b_n$. If ϵ is periodic with mean zero then its partial sums are bounded, and $b_n \downarrow 0$ eventually, so $\sum\epsilon_n b_n$ converges. This is the general statement of which the alternating series test is the period 2 case. See Book 1, field: Real analysis, topic: Dirichlet's test.

Alternating harmonic series. $\sum_{m\geq 1}\frac{(-1)^{m-1}}{m} = \ln 2$. From the Maclaurin series of $\ln(1+x)$ on $(-1,1)$ plus Abel's theorem at $x=1$, or directly from $H_{2N} = \sum_{n=N+1}^{2N}\frac{1}{n} \to \int_0^1\frac{dt}{1+t}$. See Book 1, field: Real analysis, topic: alternating harmonic series.

Dirichlet series with periodic coefficients. $L(s) = \sum\epsilon_n n^{-s}$ converges for $s > 0$ when ϵ has mean zero and for $s > 1$ always; $\zeta(s) = \sum n^{-s}$ has a simple pole at $s = 1$ with residue 1; $\eta(s) = (1-2^{1-s})\zeta(s)$; $\sum_{n\text{ odd}} n^{-s} = (1-2^{-s})\zeta(s)$. The move $\frac{d}{ds}n^{-s} = -(\ln n)n^{-s}$ is why every "$\ln n$ over n" sum is a disguised derivative. See Book 1, field: Number theory, topic: Dirichlet series.

The self-similarity, generalised. There is no recursion in the sequence sense here, but there is an exact self-similarity under index scaling, and it is the engine, so generalise it fully. For any $d \geq 2$ and any real s,

$$\frac{1}{(dn)^s} = d^{-s}\cdot\frac{1}{n^s}, \qquad \frac{\ln(dn)}{(dn)^s} = d^{-s}\left(\frac{\ln n}{n^s} + \frac{\ln d}{n^s}\right),$$

and more generally $\frac{(\ln dn)^r}{(dn)^s} = d^{-s}\sum_{j=0}^{r}\binom{r}{j}(\ln d)^{r-j}\frac{(\ln n)^j}{n^s}$. In words: the family $\{n^{-s},\ (\ln n)n^{-s},\ (\ln n)^2 n^{-s},\dots\}$ is closed under index scaling, with a triangular matrix of coefficients. That closure is the structural reason Dirichlet series, and not power series, are the correct transform for any sum whose index appears inside a logarithm. The same generalisation solves every problem of the form "periodic weight against $(\ln n)^r/n^s$": decompose the weight along the divisors of the period, then apply the binomial expansion above. See Book 1, field: Number theory, topic: index scaling and Dirichlet convolution.

Telescoping, the alternative engine. $\sum_{n=1}^{N}(g(n)-g(n+1)) = g(1)-g(N+1)$, so $\sum_{n\geq 1}(g(n)-g(n+1)) = g(1)$ whenever $g(n)\to 0$. With $g = b$ this gives $\sum_{k\geq 1}a_k = 0$ for $a_k = b_k - b_{k+1}$, the fact powering the third ranked approach. See Book 1, field: Sequences and recurrences, topic: telescoping.

Circuit Bank

- **WHEN** a series is printed in equal length blocks, or in two forms one of which is bracketed **DO** separate it into (periodic weight) times (term sequence), extend the weight back to $n = 1$ if a term vanishes there, and check whether the weight has mean zero **BECAUSE** the block length is the period of the weight, and mean zero is what makes the thing converge and what makes bracketing matter
- **WHEN** $\ln n$ appears divided by n or by n^s **DO** write down $b_{dn} = d^{-s}(b_n + \frac{\ln d}{n^s})$ for the divisors d of the period before doing anything else **BECAUSE** ln converts multiplication to addition, so index scaling is the only structural operation the expression admits, and it is where every $\ln d$ in the final answer must come from
- **WHEN** a residue class such as $n \equiv 2 \pmod 4$ singles out the special terms **DO** restate it as a condition on $m = n/2$, here "m odd" **BECAUSE** a condition on n/d is what lets you pull the weight back along the scaling map, which is the only way the weight and the terms can be made to talk to each other

- **WHEN** you are tempted to rearrange, regroup or split an infinite sum that is not absolutely convergent **DO** work with the Nth partial sum, keep every identity finite, and take limits exactly once at the end **BECAUSE** Riemann rearrangement means an unjustified regrouping of a conditionally convergent series is not a shortcut, it is an error, and the finite version costs one extra subscript
- **WHEN** an unknown constant appears twice with opposite signs and shifted indices **DO** name it, prove only that it converges, and let it cancel **BECAUSE** evaluating it is usually a much harder problem than the one you were set, and convergence is all the cancellation needs

Backtesting

1. First useful cue: the second, unbracketed display, showing the weights $3, -1, -1, -1$ repeating with period 4 against the single sequence $\ln n/n$; and immediately after it, $3-1-1-1=0$.
2. Method the expression suggested: separate weight from terms; note the only structure $\ln n/n$ has is $b_{2m} = \frac{\ln 2}{2m} + \frac{1}{2}b_m$; therefore decompose the weight along the doubling map.
3. Where you would hesitate, and the thought to reinforce: at the moment of matching ϵ_n to $-(-1)^{n-1} + 2(-1)^{m-1}[n = 2m]$, which can feel like it was pulled from nowhere. Reinforce: it was not guessed, it was forced. $\epsilon = -1 + 4[n \equiv 2 \bmod 4]$, and $n \equiv 2 \pmod 4$ means $n = 2m$ with m odd, and "m odd" carries the sign $(-1)^{m-1}$. Two constants then have to be matched, and there is only one way to match them.
4. The trap the question sets: the problem hands you two displays, one bracketed and one not, so that you will treat them as interchangeable and start regrouping. The unbracketed series is only conditionally convergent. Any solution that shuffles it without a finite partial sum argument is exposed.
5. What makes you recognise it faster next time: the pairing "periodic mean zero weight" with "$\ln n$ over n" is the fingerprint of $-L'(1)$ for a Dirichlet series L. Having once seen that $L(s) = -\zeta(s)(1-2^{1-s})^2$ here, you will recognise the family instantly, and you will know before computing whether the answer is clean (double zero at $s=1$) or contaminated by γ (simple zero).
6. The reusable decision rule: when a sum mixes a periodic weight with a sequence that is self-similar under index scaling, decompose the weight along the scaling maps, keep the algebra finite, and take limits once.

Mirror drills

1. Evaluate $\displaystyle\sum_{k=0}^{\infty}\left(\frac{3}{4k+2} - \frac{1}{4k+3} - \frac{1}{4k+4} - \frac{1}{4k+5}\right)$. *Answer:* 0. Same weight, terms $1/n$. The partial sum is $-\sum_{n\le 4K+1}\frac{(-1)^{n-1}}{n} + 2\sum_{m\le 2K}\frac{(-1)^{m-1}}{2m} \to -\ln 2 + \ln 2 = 0$. This is the statement $L(1)=0$, visible in the factorisation $L(s) = -\zeta(s)(1-2^{1-s})^2$.
2. Let $\epsilon_n = 2$ if $3 \mid n$ and $\epsilon_n = -1$ otherwise. Evaluate $\sum_{n\ge 1}\epsilon_n/n$. *Answer:* $-\ln 3$. Here $\epsilon_n = -1 + 3[3 \mid n]$, so $\sum_{n\le 3N}\epsilon_n/n = -H_{3N} + 3\sum_{m\le N}\frac{1}{3m} = -(H_{3N} - H_N) \to -\ln 3$. Note the contrast with drill 1: the weight is not built from an alternating character, so nothing cancels and the answer is a single ln.
3. Show that $\displaystyle\sum_{m=1}^{\infty}(-1)^{m-1}\left(\frac{\ln(2m)}{2m} - \frac{\ln m}{2m}\right) = \frac{(\ln 2)^2}{2}$. *Answer:* the bracket is $\frac{\ln 2}{2m}$ exactly, so the sum is $\frac{\ln 2}{2}\sum_{m\ge 1}\frac{(-1)^{m-1}}{m} = \frac{(\ln 2)^2}{2}$. This is the engine of B4 in isolation: one $\ln 2$ from splitting the logarithm, one from the alternating harmonic series.

9.11 B5. Triangles with exactly two equalizer lines

A line in the plane of a triangle T is called an equalizer *if it divides T into two regions having equal area and equal perimeter. Find positive integers $a > b > c$, with a as small as possible, such that there exists a triangle with side lengths a, b, c that has exactly two distinct equalizers.*

Read it out loud

A line that cuts a triangle in two cuts across two of its three sides. So every dividing line is described by three pieces of data: which vertex the two cut sides share, and the two distances from that vertex to the cut points. Call the vertex V, call the two sides through it u and v, call the third side w, and call the distances x and y. That is the whole coordinate system, and you should have it on paper within thirty seconds, because the alternative coordinate system (put the triangle in the plane, write the line as $y = mx + k$) has three unknowns instead of two and no relationship to either constraint.

Now read the two constraints in those coordinates. Area: the corner triangle cut off at V has area $\frac{1}{2}xy\sin V$ and the whole triangle has area $\frac{1}{2}uv\sin V$, so equal area says $xy = \frac{1}{2}uv$. A *product*. Perimeter: the corner region has perimeter $x + y + \ell$ where ℓ is the length of the cut, and the other region has perimeter $(u - x) + (v - y) + w + \ell$, so the cut length ℓ cancels and equal perimeter says $x + y = \frac{1}{2}(u + v + w) = s$, the semiperimeter. A *sum*. Two symmetric conditions on the same pair of numbers, one giving their sum and one giving their product. That is a quadratic, and the entire problem is now about how many roots of that quadratic land in the allowed window.

The fields wired in are plane geometry (area from two sides and the included angle; medians; the incentre), elementary convexity (the function $x + k/x$, and AM to GM as its minimum), and at the very end a small piece of number theory (when is $2ab$ a perfect square). The command term "exactly two" is the guiding light: a convex function meets a horizontal level in zero, one or two points, and "exactly one" happens only at tangency. Any problem that asks for an exact count in the presence of a convex function is asking you to find the tangency case.

One more reading, worth the ten seconds it costs. The conditions are symmetric in x and y and involve only the two sides through V and the semiperimeter. Nothing about the angles survives. That is a strong hint that the answer will be a relation purely between the side lengths, and it is: $8ab = (a + b + c)^2$.

Approaches, ranked worst to best

1. **Coordinate bash.** *Cue:* a line in the plane of a triangle. *Circuit:* place the triangle at $(0, 0)$, $(a, 0)$, (p, q) and write the line as $y = mx + k$. Two unknowns for the line, then integrals or shoelace formulae for the area and a distance computation for the perimeter, then a two variable system in m and k with a case split on which sides are cut. This is the zombie route: it throws away the fact that the natural parameters are the two cut distances, and it reintroduces the angles the problem does not care about.

2. **Numerical search over small integer triples.** *Cue:* "find positive integers $a > b > c$ with a as small as possible". *Circuit:* write a loop, count equalizers numerically, report the first hit. It produces $(9, 8, 7)$ and proves nothing, because "exactly two" has not been characterised and no minimality argument exists. Ranks second last: useful only as a check after the real work.

3. **Quadratic with root location conditions.** *Cue:* $x + y = s$ and $xy = uv/2$ make x, y the roots of $t^2 - st + \frac{uv}{2} = 0$. *Circuit:* Vieta, then discriminant plus root location on an interval. Correct, and only a little worse than the winner. The cost is that root location for a quadratic on a closed interval is a six way case analysis (sign of the leading coefficient, sign at each endpoint, position of the vertex, sign of the discriminant), which you must do three times, once per vertex.

4. **Rotate a line about the incentre.** *Cue:* equal area and equal perimeter together are exactly the two conditions that pin a line to the incentre. *Circuit:* prove that every equalizer passes through the incentre I, then parametrise lines through I by direction. The incentre fact is beautiful and you should know it: with $[T] = rs$, decomposing the corner region into triangles from I gives area $= \frac{r}{2}(x+y) \pm \frac{d}{2}\ell$ where d is the distance from I to the cut, so the two conditions force $d\ell = 0$, that is $d = 0$. It gives existence of an equalizer immediately by a continuity argument as the line rotates. It ranks fourth only because counting them still requires the same analysis as the winning route, so you pay for a lovely lemma you do not strictly need.
5. **Level set of a convex function on a geometric window.** *Cue:* a sum and a product, plus a demand for an exact count. *Circuit:* write $y = \frac{uv}{2x}$, so the condition is $\phi(x) := x + \frac{uv}{2x} = s$ on the window $x \in [u/2, u]$, and use convexity. This is the ideal route, and its punchline is that the two endpoint values of the window are forced by a one line geometric fact rather than by algebra: $\phi(u/2) - s = \frac{1}{2}(v-w)$ and $\phi(u) - s = \frac{1}{2}(u-w)$, because the two endpoints of the window are the two medians that cross the side of length u, and a median is an equalizer exactly when the triangle is isosceles at the vertex it comes from. Three sign comparisons of side lengths then give the count at every vertex with no case analysis at all.

The ideal solution

We show the answer is $\boxed{(a,b,c) = (9,8,7)}$.

Throughout, T is a triangle with vertices A, B, C and side lengths $BC = a$, $CA = b$, $AB = c$, with $a > b > c > 0$; write $p = a+b+c$ and $s = p/2$.

Step 1: every dividing line is a pair of distances from a vertex. Let L be a line dividing T into two regions of positive area. Then L meets the interior of T and does not contain a side, so L meets the boundary ∂T in exactly two points $P \neq Q$. We claim there is a vertex $V \notin \{P, Q\}$ such that P and Q lie on the two closed sides through V. If neither P nor Q is a vertex, they lie in the relative interiors of two distinct sides, and those sides share a vertex V, which is neither P nor Q. If exactly one of them, say P, is a vertex, then Q lies in the relative interior of the side opposite P (otherwise L would be a side), and we may take V to be either endpoint of that opposite side. If both are vertices then L is a side, which does not divide T.

So fix such a V, let u and v be the lengths of the two sides through V and w the length of the side opposite V, and let $x \in (0, u]$ and $y \in (0, v]$ be the distances from V to the two boundary points along those sides.

Step 2: the two conditions are a product and a sum. The corner region is the triangle with vertex V and the two cut points, of area $\frac{1}{2}xy \sin V$, while $[T] = \frac{1}{2}uv \sin V$. Hence

$$\text{equal areas} \iff xy = \tfrac{1}{2}uv.$$

If ℓ denotes the length of the cut segment, the corner region has perimeter $x + y + \ell$ and the other region has perimeter $(u-x) + (v-y) + w + \ell$; equating and cancelling ℓ,

$$\text{equal perimeters} \iff 2(x+y) = u+v+w = p \iff x+y = s.$$

Step 3: eliminate y and read the window. Given x, the area condition forces $y = \frac{uv}{2x}$, and the constraint $y \leq v$ is equivalent to $x \geq u/2$. So the equalizers of T whose two boundary points lie on the sides through V correspond exactly to the solutions of

$$\phi(x) := x + \frac{uv}{2x} = s, \qquad x \in \left[\tfrac{u}{2},\, u\right].$$

Step 4: the endpoints of the window are the two medians. At $x = u/2$ we get $y = v$, so the cut passes through the midpoint of the side of length u and through the far endpoint of the side of

length v: it is a median. Its two regions have perimeters $\frac{u}{2}+v+\ell$ and $\frac{u}{2}+w+\ell$, so it is a perimeter bisector exactly when $v=w$. Consistently,

$$\phi\left(\tfrac{u}{2}\right)-s=\tfrac{u}{2}+v-\tfrac{u+v+w}{2}=\frac{v-w}{2}.$$

At $x=u$ we get $y=v/2$, the other median crossing the side of length u, and

$$\phi(u)-s=u+\tfrac{v}{2}-\tfrac{u+v+w}{2}=\frac{u-w}{2}.$$

So the sign of ϕ minus s at each endpoint is decided by comparing an adjacent side with the opposite side. No computation is needed beyond these two lines.

Step 5: convexity converts the count into sign tests. On $(0,\infty)$ we have $\phi''(x)=\frac{uv}{x^3}>0$, so ϕ is strictly convex, with minimum value $\phi(x_*)=\sqrt{2uv}$ attained at $x_*=\sqrt{uv/2}$ (AM to GM). Consequently $\{x:\phi(x)\le s\}$ is a closed interval, and $\phi(x)=s$ has at most two solutions.

Now take the three vertices in turn, using $a>b>c$.

Vertex A, with $u=c$, $v=b$, $w=a$. By Step 4, $\phi(c/2)-s=\frac{b-a}{2}<0$ and $\phi(c)-s=\frac{c-a}{2}<0$. A convex function on a closed interval attains its maximum at an endpoint, so $\phi<s$ throughout $[c/2,c]$: *no equalizer cuts both AB and AC.*

Vertex B, with $u=c$, $v=a$, $w=b$. Now $\phi(c/2)-s=\frac{a-b}{2}>0$ and $\phi(c)-s=\frac{c-b}{2}<0$. The set $\{\phi\le s\}$ is a closed interval $[\alpha,\beta]$ containing c but not $c/2$, so $\alpha\in(c/2,c)$ and, since $\phi(c)<s$, also $\beta>c$. Hence $\phi=s$ has *exactly one* solution in $[c/2,c]$, and it lies strictly inside: *exactly one equalizer cuts both BA and BC, and it meets the relative interiors of those sides.*

Vertex C, with $u=a$, $v=b$, $w=c$. Here $\phi(a/2)-s=\frac{b-c}{2}>0$ and $\phi(a)-s=\frac{a-c}{2}>0$: both endpoints lie strictly above s. Moreover $x_*=\sqrt{ab/2}$ lies in $(a/2,a)$, because $b<a$ gives $x_*<a$ and $a<b+c<2b$ gives $b>a/2$ and hence $x_*>a/2$. Therefore the number of solutions of $\phi=s$ in $[a/2,a]$ is

$$\begin{cases}2 & \text{if } \sqrt{2ab}<s,\\ 1 & \text{if } \sqrt{2ab}=s,\\ 0 & \text{if } \sqrt{2ab}>s,\end{cases}$$

and in the first two cases the solutions again lie strictly inside $(a/2,a)$.

Step 6: no line is counted twice. The equalizer found at B meets the relative interiors of AB and BC; each equalizer found at C meets the relative interiors of CA and CB. Since a dividing line meets ∂T in exactly two points, two such lines can coincide only if their pairs of boundary points coincide, and one pair contains a point of the relative interior of AB while the other does not. So the counts add.

Hence T has 1, 2 or 3 equalizers, and

$$T\text{ has exactly two equalizers}\iff\sqrt{2ab}=s=\frac{a+b+c}{2}\iff 8ab=(a+b+c)^2.$$

Step 7: the integer search. We need integers $a>b>c\ge 1$ with $b+c>a$ and $8ab=(a+b+c)^2$. Since $8ab$ must be a perfect square and $8ab=4\cdot 2ab$, this says $2ab$ is a perfect square, and then $p=2\sqrt{2ab}$, so

$$c=2\sqrt{2ab}-a-b.$$

The triangle inequality $b+c>a$ becomes $2\sqrt{2ab}>2a$, that is $2b>a$. Conversely, suppose $2ab$ is a perfect square with $\frac{a}{2}<b<a$ and $c:=2\sqrt{2ab}-a-b\ge 1$; then $c<b$ holds automatically, because $c<b$ is equivalent to $2\sqrt{2ab}<a+2b$, that is to $8ab<(a+2b)^2$, that is to $0<(a-2b)^2$, which is true since $a\ne 2b$. So the search is: find the least a for which some integer b with $\frac{a}{2}<b<a$ makes $2ab$ a perfect square and $c\ge 1$.

Run it. For $a = 3, \ldots, 8$ the values of $2ab$ over the admissible b are

$$\begin{aligned} a = 3 &: \ 12 \\ a = 4 &: \ 24 \\ a = 5 &: \ 30,\ 40 \\ a = 6 &: \ 48,\ 60 \\ a = 7 &: \ 56,\ 70,\ 84 \\ a = 8 &: \ 80,\ 96,\ 112 \end{aligned}$$

and none is a perfect square. For $a = 9$ the admissible b are $5, 6, 7, 8$, giving $2ab = 90,\ 108,\ 126,\ 144$, and $144 = 12^2$. So $b = 8$, $p = 2 \cdot 12 = 24$, and $c = 24 - 9 - 8 = 7$.

Check the triple $(9, 8, 7)$ directly: $9 > 8 > 7 > 0$, the triangle inequality holds since $8 + 7 = 15 > 9$, and $8ab = 8 \cdot 9 \cdot 8 = 576 = 24^2 = p^2$. So this triangle has exactly two equalizers, and no smaller a admits one.

The two equalizers, explicitly. For $(a, b, c) = (9, 8, 7)$ we have $s = 12$. At C the tangency gives $x = y = \sqrt{ab/2} = \sqrt{36} = 6$: the cut joins the points at distance 6 from C on CB (length 9) and on CA (length 8), splitting the perimeter as $6 + 6$ against $3 + 7 + 2$ and the area in the ratio $\frac{6 \cdot 6}{9 \cdot 8} = \frac{1}{2}$. At B, $\phi(x) = x + \frac{63}{2x} = 12$ gives $x = 6 \pm \frac{3\sqrt{2}}{2}$, of which only $x = 6 - \frac{3\sqrt{2}}{2} \approx 3.879$ lies in $[3.5, 7]$; the corresponding $y \approx 8.121$ lies in $(4.5, 9)$ as it must.

Layer 1: WHAT

Block 1. A dividing line crosses two sides; parametrise it by the two distances x, y from the shared vertex.

Block 2. Equal area gives the product $xy = uv/2$; equal perimeter gives the sum $x + y = s$. The cut length ℓ cancels on its own.

Block 3. Eliminate y: the condition is $\phi(x) = x + \frac{uv}{2x} = s$ on the window $x \in [u/2, u]$, where the window comes from $x \le u$ and $y \le v$.

Block 4. Identify the window endpoints geometrically: they are the two medians crossing the side of length u. Hence $\phi(u/2) - s = \frac{v - w}{2}$ and $\phi(u) - s = \frac{u - w}{2}$.

Block 5. ϕ is strictly convex. Three sign patterns, one per vertex: below and below at A (zero solutions), above and below at B (exactly one), above and above at C (zero, one or two according to the minimum).

Block 6. "Exactly two" therefore means tangency at C: $\sqrt{2ab} = s$, that is $8ab = (a + b + c)^2$.

Block 7. Reduce to "$2ab$ is a perfect square with $a/2 < b < a$" and search $a = 3, 4, 5, \ldots$ until it succeeds, at $a = 9$.

Juxtaposition. The nominal route agrees with you through Block 2 and then reaches for Vieta: x, y are the roots of $t^2 - st + \frac{uv}{2} = 0$, so compute the discriminant $s^2 - 2uv$, then impose $x \le u$, $y \le v$ by root location conditions. That works, but at each vertex you now need the sign of the quadratic at both endpoints, the position of its vertex relative to the interval, and the sign of the discriminant, and you must repeat all of it three times. Worse, the endpoint conditions come out as opaque quadratic expressions rather than as $\frac{v - w}{2}$ and $\frac{u - w}{2}$, so you lose the geometric reading that makes the three cases obvious.

Bifurcation point. The line where the routes split is the one immediately after $x + y = s$ and $xy = uv/2$ are both on the page. The zombie writes the quadratic in t. You write $y = \frac{uv}{2x}$ and treat the problem as counting intersections of a convex curve with a horizontal line, because the question asked "how many", and counting is what convexity is for.

Layer 2: HOW

Block 1. Cue. A line divides a triangle. The triangle has three sides; the line meets exactly two of them. **Why natural.** The definition of a convex region does the work: a line meeting the interior of a

convex set meets its boundary in exactly two points, and the boundary here is three segments, so the two points lie on two sides which share a vertex. Once you know a vertex is distinguished, the two distances from it are the only natural coordinates: the corner region cut off is a triangle with vertex V and legs x, y along the sides. **Link to the next block.** With V, x, y in place, both constraints are about that corner triangle, so write them.

Block 2. **Cue.** "Equal area" next to "equal perimeter". Two constraints on two unknowns; expect a finite answer set. **Why natural.** Area from two sides and the included angle is the only area formula that does not need the third side, and the corner triangle shares its angle with T, so the sin V cancels and the area condition is pure multiplication. For perimeter, notice before computing that the cut segment ℓ belongs to both regions, so it must cancel; that is why the perimeter condition is pure addition, and it is also why the answer will be independent of the angles. **Link to the next block.** A sum and a product of the same two numbers is one of the strongest structures in elementary mathematics. The only question is which of its two readings, quadratic or convex curve, you take.

Block 3. **Cue.** The question asks for a *count*, not for the equalizers themselves. **Why natural.** When you want the number of solutions, put the problem in a form with one unknown and one level: $\phi(x) = s$. Then the count is the number of times a curve crosses a horizontal line, and there is a standard tool for that (monotonicity or convexity) which needs no case analysis. Deriving the window is equally forced: the two constraints $x \le u$ and $y \le v$ are the statement that the cut points really lie on the sides, and under $y = \frac{uv}{2x}$ the second is $x \ge u/2$. **Link to the next block.** The window has two endpoints, so before invoking convexity you need the values of ϕ there.

Block 4. **Cue.** $x = u/2$ says the cut passes through the midpoint of one side; $y = v$ says it passes through a vertex. **Why natural.** Translate before you compute. A cut through a midpoint and the opposite vertex is a median, and for a median the perimeter condition is trivial: the two regions each contain half of the bisected side plus the cut, so they are equal exactly when the other two sides they contain are equal. Hence the sign of $\phi(u/2) - s$ must be the sign of $v - w$, and the algebra confirms it in one line. Computing first and interpreting later would give you the same two numbers with none of the meaning, and meaning is what makes Block 5 free. **Link to the next block.** Now every endpoint sign is a comparison of two side lengths, and the triangle's ordering $a > b > c$ decides all six of them at once.

Block 5. **Cue.** $\phi(x) = x + \frac{k}{x}$ with $k > 0$. That is the standard convex function. **Why natural.** $\phi'' > 0$, so ϕ is strictly convex and its sublevel sets are intervals; a strictly convex function crosses any horizontal level at most twice, and the AM to GM bound $\phi \ge \sqrt{2uv}$ with equality at $x = \sqrt{uv/2}$ is the exact statement of where the minimum sits. Given the endpoint signs, the count is then immediate and identical in structure at all three vertices. The one detail to check is whether the minimiser lies inside the window, and at C it does because $a < b + c < 2b$. **Link to the next block.** At two of the three vertices the count is forced with no further information, so "exactly two" can only be a statement about the third.

Block 6. **Cue.** The count is $1 + (\text{count at } C)$, and $1 + 1 = 2$. **Why natural.** The vertex opposite the largest side contributes nothing and the vertex opposite the middle side contributes exactly one, no matter what the triangle is; so the entire content of "exactly two" is that the vertex opposite the smallest side contributes exactly one, and for a strictly convex function with both endpoints above the level that means tangency. Equality in AM to GM then gives $\sqrt{2ab} = s$, which squares to $8ab = p^2$. **Link to the next block.** A single Diophantine equation in three positive integers, plus the triangle inequality, is now the whole remaining problem.

Block 7. Cue. "Positive integers, a as small as possible." **Why natural.** Minimality demands an exhaustive check, so first make the condition cheap to test: $8ab = p^2$ forces $2ab$ to be a perfect square and determines c from a and b, so for each a there are fewer than $a/2$ candidate values of b, namely those in $(a/2, a)$, and each is a single perfect square test. Establishing that $c < b$ is automatic (it reduces to $(a-2b)^2 > 0$) is what allows you to drop c from the search entirely. **Link to the next block.** Six values of a to eliminate and one to exhibit, and the write up is a finite table.

Layer 3: WHY

Three truths are doing the work.

The first is that both constraints are *symmetric* in the two cut distances, and that they pick out the two elementary symmetric functions: perimeter gives $x + y$, area gives xy. This is not a coincidence of this problem. Perimeter is a first order (length) quantity and area is a second order quantity, so any pair of conditions of the form "bisect a first order quantity" and "bisect a second order quantity" will give you a sum and a product. Once you see a sum and a product, you are on a hyperbola meeting a line, and the intersection count of a line with a convex branch is 0, 1 or 2, with 1 exactly at tangency. Everything in this problem, including the final equation $8ab = p^2$, is the equality case of AM to GM in disguise.

The second is that the window endpoints are medians, so the endpoint comparisons are side comparisons. This is the discernment mechanism worth extracting. Whenever you reduce a geometric family to a parameter running over an interval, ask what the two extreme parameter values *are* geometrically before you evaluate anything. Degenerate configurations are always simpler than generic ones, so their values are always cheaper to obtain by interpretation than by algebra, and the interpretation carries meaning the algebra does not. Here it turns six algebraic evaluations into six comparisons among a, b, c, which the ordering $a > b > c$ settles instantly.

The third is a structural fact worth remembering in its own right: *every triangle has at least one and at most three equalizers*. For a scalene triangle with $a > b > c$, exactly one cuts the two sides adjacent to the middle side, none cuts the two sides adjacent to the longest side, and zero, one or two cut the two sides adjacent to the shortest side according as $8ab$ exceeds, equals or falls short of p^2. The equilateral triangle sits at the other extreme: there every endpoint value equals s, since $u = v = w$, so at each vertex the two roots are exactly the two medians, and the equalizers are precisely the three medians. A very flat triangle, say $(100, 51, 50)$ with $8ab = 40800 > 40401 = p^2$, has exactly one. The problem is asking you to find the boundary between those regimes, in integers.

To train this: whenever a problem asks for an exact count of solutions and there is a convex or monotone function anywhere in sight, immediately look for the tangency or endpoint case, because "exactly n" problems are always boundary problems. And whenever a problem asks for "the smallest integers such that", separate the structural characterisation from the search completely: characterise first, search second, and never search a condition you have not reduced.

How the examiners nudged. They defined a new word, *equalizer*, which is the standard signal that the two halves of the definition are meant to be combined into a single algebraic relation rather than handled separately. They asked for "exactly two distinct", which forces the tangency reading. And they asked for integers with a minimal, which guarantees the structural condition must be clean enough to search by hand; had the characterisation been messy, the question could not have been posed this way. Reading the shape of the question tells you the shape of the answer before you start.

Revolving concepts

Equalizer, splitter, cleaver, median. A line bisecting the area is an area bisector; one bisecting the perimeter is a perimeter bisector; a *splitter* is a perimeter bisector through a vertex, a *cleaver* is a perimeter bisector through the midpoint of a side, and an equalizer is a line that is both an area

and a perimeter bisector. A median is always an area bisector, and is a perimeter bisector exactly when the triangle is isosceles at the vertex it issues from. See Book 1, field: Geometry, topic: triangle cevians.

Every equalizer passes through the incentre. With I the incentre, r the inradius and s the semiperimeter, $[T] = rs$. If a line cuts the sides through V at distances x, y and d is the distance from I to that line, decomposing the corner region into triangles with apex I gives its area as $\frac{r}{2}(x+y) \pm \frac{d}{2}\ell$. Imposing $x + y = s$ and area $= \frac{rs}{2}$ leaves $d\ell = 0$, so $d = 0$. Conversely a line through I bisects the area if and only if it bisects the perimeter. See Book 1, field: Geometry, topic: incentre and $[T] = rs$.

Convexity of $x + k/x$. For $k > 0$ the function is strictly convex on $(0, \infty)$, with $\phi'' = 2k/x^3$ for $\phi = x + k/x$, minimum $2\sqrt{k}$ at $x = \sqrt{k}$; this is AM to GM for two positive numbers, with the equality case being the tangency case. A strictly convex function meets any horizontal line in at most two points, and its sublevel sets are intervals: that is the counting lemma used three times above. See Book 1, field: Inequalities, topic: AM to GM and convexity.

Symmetric functions and Vieta. If $x + y = \sigma$ and $xy = \pi$ then x, y are the roots of $t^2 - \sigma t + \pi = 0$, real precisely when $\sigma^2 \geq 4\pi$. Here $\sigma^2 \geq 4\pi$ reads $s^2 \geq 2uv$, which is the same tangency condition seen through the discriminant instead of through convexity: $s = \sqrt{2uv}$ is where the discriminant vanishes and the two cut distances coincide. Recognising that two different tools (discriminant, AM to GM) give the same boundary is the sign that you have found the real structure. See Book 1, field: Polynomials, topic: Vieta and the discriminant.

A line meeting a convex region. A line that meets the interior of a compact convex set with nonempty interior meets its boundary in exactly two points, unless it contains a boundary segment. This is what licenses the exhaustive case analysis in Step 1 and the "no double counting" argument in Step 6. See Book 1, field: Geometry, topic: convex sets.

Perfect square conditions. $8ab$ is a perfect square if and only if $2ab$ is, since $8ab = 4 \cdot 2ab$ and 4 is a square. In general, mab is a square exactly when the squarefree part of mab is 1, which for a bounded search is faster to test than to characterise. See Book 1, field: Number theory, topic: squarefree parts.

Recursion note. There is no recursion in this problem, and it is worth saying why rather than leaving the box empty: the family of dividing lines is a continuum parametrised by one real number per vertex, not a discrete family built from earlier members. The tool that replaces induction for a continuum family is the intermediate value theorem together with a convexity or monotonicity statement, which is exactly what Step 5 uses. When you meet a one parameter family and want to count members with a property, the standard substitute for induction is: express the property as $\phi(x) =$ constant, establish convexity or monotonicity of ϕ, and evaluate at the endpoints of the parameter window.

Circuit Bank

- **WHEN** a geometric condition gives you the sum of two lengths and another gives you their product **DO** eliminate one of them and write the condition as a level set $\phi(x) = x + \frac{k}{x} = \sigma$ on an explicit window **BECAUSE** a sum and a product is a hyperbola meeting a line, and convexity turns "how many solutions" into three sign tests instead of a six way root location analysis

- **WHEN** your parameter window has endpoints **DO** identify the endpoint configurations geometrically before evaluating anything **BECAUSE** degenerate configurations (here, medians) have one line proofs, so the endpoint values arrive with meaning attached and the case analysis becomes free

- **WHEN** a question asks for "exactly one" or "exactly two" of something governed by a convex function **DO** look immediately for the tangency case, that is the equality case of the relevant

AM to GM **BECAUSE** an exact count is a boundary condition, and for a convex function the only boundary between two solutions and none is the tangent level

- **WHEN** a question ends with "find the smallest positive integers such that" **DO** reduce to a single clean Diophantine condition first, prove which of the side conditions are automatic, and only then run a finite search you can write out **BECAUSE** minimality claims require exhaustion, and exhaustion is only writable once each candidate costs a single test
- **WHEN** a competition defines a new word out of two familiar conditions **DO** expect the two conditions to combine into one algebraic relation, and look for the object they jointly pin down **BECAUSE** setters coin a term when the conjunction is more rigid than either half, here forcing every equalizer through the incentre

Backtesting

1. First useful cue: a dividing line crosses exactly two sides, so it is described by two distances from the vertex they share; and the cut length cancels in the perimeter condition, so no angle can appear in the answer.
2. Method the expression suggested: equal area gives the product $xy = uv/2$, equal perimeter gives the sum $x + y = s$; eliminate y and count solutions of $x + \frac{uv}{2x} = s$ on $[u/2, u]$ by convexity.
3. Where you would hesitate, and the thought to reinforce: at the window $x \in [u/2, u]$, which looks arbitrary. Reinforce: $x \leq u$ says the cut point is on the side; $y \leq v$ says the other cut point is on its side, and under $y = uv/(2x)$ that is $x \geq u/2$. The endpoints are the two medians, and once you know that, the endpoint values $\frac{v-w}{2}$ and $\frac{u-w}{2}$ are predictable before you compute them.
4. The trap the question sets: two of them. First, forgetting the vertex opposite the largest side, or double counting a line at two vertices, so that the count comes out wrong. Second, and more expensive, jumping to the integer search before proving the characterisation $8ab = p^2$, which makes the minimality claim unprovable.
5. What makes you recognise it faster next time: the phrase "exactly two" beside a convex function, and the pairing of a first order and a second order bisection condition. Also worth storing as a fact: every triangle has one, two or three equalizers, exactly one of which cuts the two sides adjacent to the middle length side.
6. The reusable decision rule: turn "count the members of a one parameter geometric family with a property" into "count the solutions of $\phi =$ constant on a closed window", get the endpoints by geometric interpretation, and get the interior behaviour from convexity.

Mirror drills

1. Show that every equalizer of a triangle passes through its incentre. *Answer:* with $[T] = rs$, split the corner region into triangles with apex I: its area is $\frac{r}{2}(x + y) \pm \frac{d}{2}\ell$ where $d = \text{dist}(I, L)$. Equal perimeter gives $x + y = s$ and equal area gives area $= \frac{rs}{2}$, so $d\ell = 0$ and $d = 0$.
2. Prove that every triangle has at least one equalizer and at most three. *Answer:* apply Step 5 at each vertex. At the vertex opposite the longest side both endpoint values lie below s, giving none; at the vertex opposite the middle side the endpoint values straddle s, giving exactly one; at the vertex opposite the shortest side both lie above s, giving at most two. (For an isosceles or equilateral triangle the same argument runs with weak inequalities.)
3. For which triangles is a median an equalizer? *Answer:* the median from a vertex Q bisects the area automatically, and it bisects the perimeter exactly when the two sides through Q are

equal, so precisely when the triangle is isosceles with apex Q. This is the geometric content of the endpoint identities $\phi(u/2) - s = \frac{v-w}{2}$ and $\phi(u) - s = \frac{u-w}{2}$.

9.12 B6. Distinct tuples under a weighted congruence

Find the number of ordered 64*-tuples* $(x_0, x_1, \ldots, x_{63})$ *such that* $x_0, x_1, \ldots, x_{63}$ *are distinct elements of* $\{1, 2, \ldots, 2017\}$ *and*

$$x_0 + x_1 + 2x_2 + 3x_3 + \cdots + 63x_{63}$$

is divisible by 2017*.*

Read it out loud

This is a counting problem with two constraints of completely different natures. One is algebraic and soft: a single linear congruence modulo 2017. The other is combinatorial and hard: the 64 entries must be pairwise distinct. Soft constraints are easy to count against, because a nonzero linear condition modulo a prime cuts a count down by a factor of exactly p. Hard constraints are what makes counting problems difficult, because "all distinct" is a conjunction of $\binom{64}{2}$ negated conditions and is invisible to Fourier analysis. So the entire question is: how do you get the distinctness constraint and the congruence to cooperate rather than fight?

Before answering that, do the arithmetic the setters planted. 2017 is prime, so $\{1, 2, \ldots, 2017\}$ is a complete residue system modulo 2017 (note $2017 \equiv 0$), and you may as well count tuples of distinct elements of $\mathbb{F}_{2017}$. Now add up the coefficients. Writing $c_0 = c_1 = 1$ and $c_i = i$ for $2 \le i \le 63$,

$$\sum_{i=0}^{63} c_i \;=\; 1 + (1 + 2 + \cdots + 63) \;=\; 1 + \frac{63 \cdot 64}{2} \;=\; 1 + 2016 \;=\; 2017.$$

That is the whole problem, and it is why the coefficient of x_0 and the coefficient of x_1 are both 1 instead of the pattern continuing to $0 \cdot x_0$: the setters needed the total to land exactly on the modulus. Sixty seconds spent adding the coefficients is the highest return per unit of ink on the entire paper.

What does it buy? The linear form $\Lambda(x) = \sum_i c_i x_i$ satisfies $\Lambda(x_0 + t, \ldots, x_{63} + t) = \Lambda(x) + t\sum_i c_i = \Lambda(x)$ for every $t \in \mathbb{F}_{2017}$. So Λ is invariant under simultaneous translation of every coordinate. And distinctness is invariant under simultaneous translation too. Two structures, one shared symmetry group: quotient by it. That is the reaction circuit, and everything after it is bookkeeping.

The fields wired in are finite fields (only lightly: what you need is that $\mathbb{F}_p$ is a field so that nonzero elements are invertible), group actions and orbit counting, and the solution of a first order linear recursion with a binomial inhomogeneity. Knowledge gaps are named honestly in Layer 3; this is a B6 and it is not approachable without them.

Approaches, ranked worst to best

1. **Roots of unity filter, applied directly.** *Cue:* a divisibility condition modulo a prime. *Circuit:* write the indicator of $p \mid \Lambda$ as $\frac{1}{p}\sum_{t \in \mathbb{F}_p} \omega^{t\Lambda(x)}$ and sum over the objects. The count becomes $\frac{1}{p}\sum_t \sum_{x \text{ distinct}} \omega^{t\Lambda(x)}$, and the inner sum over *distinct* tuples is a permanent style object with no product formula. The filter is the right instinct for a congruence and the wrong tool here, because distinctness destroys the factorisation that makes the filter useful. It stalls immediately.

2. **Inclusion and exclusion over the pair conditions** $x_i = x_j$**.** *Cue:* "distinct" is the negation of $\binom{64}{2}$ equalities. *Circuit:* subtract the bad events, add back the intersections. The events $\{x_i = x_j\}$ do not intersect independently: their intersections are indexed by set partitions, not by subsets of pairs, so a naive alternating sum over pairs double counts catastrophically. Ranks second last because it is the first thing most candidates write and it cannot be repaired without the next approach.

3. **Inclusion and exclusion over the partition lattice.** *Cue:* intersections of equality events are indexed by partitions of $\{0,\dots,63\}$. *Circuit:* Mobius inversion on the lattice of set partitions. This is correct and complete, and it is the first official solution. For a partition π with $|\pi|$ blocks, the number of tuples constant on blocks and satisfying the congruence is $p^{|\pi|-1}$, unless every block has coefficient sum divisible by p, which here happens only for the single block partition because $0 < \sum_{i\in S} c_i < p$ for every proper nonempty S. You then need the Mobius value of the full partition lattice, $\mu(\hat{0},\hat{1}) = (-1)^{n-1}(n-1)!$, extracted by comparing degree one coefficients in the falling factorial identity. It ranks third because it needs a piece of algebraic combinatorics most candidates do not have installed, and because in the arena the sign bookkeeping is where you lose the paper.

4. **Fix $x_0 = 0$ by translation, then recurse on the number of variables.** *Cue:* the coefficients sum to the modulus, so the form is translation invariant. *Circuit:* quotient by the shared symmetry, then peel off the special value 0. The ideal route. The translation action of $\mathbb{F}_p$ on the set of distinct tuples is free, and Λ is constant on orbits, so the answer is p times the count with x_0 normalised to 0. What remains is a self similar question in the number of variables, which recurses in one line, and the recursion is solved by Pascal's rule. There is no inclusion and exclusion anywhere, no signs to lose, and every step is forced by the previous one.

The ideal solution

Throughout, $p = 2017$, which is prime. Since $2017 \equiv 0 \pmod{2017}$, reduction modulo p is a bijection from $\{1,2,\dots,2017\}$ onto $\mathbb{F}_p$, so we count ordered 64-tuples of pairwise distinct elements of $\mathbb{F}_p$ satisfying the congruence.

Lemma 9.12.1. *Let p be prime and let $a_1,\dots,a_k$ be positive integers with $a_1+\cdots+a_k < p$. Then the number of solutions of*

$$a_1x_1 + \cdots + a_kx_k = 0 \quad \text{in } \mathbb{F}_p$$

with $x_1,\dots,x_k$ pairwise distinct and all nonzero depends only on k and p. Calling it $f(k)$, we have $f(1) = 0$ and

$$f(k) \;=\; \frac{(p-1)!}{(p-k)!} \;-\; k\,f(k-1) \qquad (k \ge 2).$$

Proof. Induct on k. For $k = 1$ we have $0 < a_1 < p$, so a_1 is invertible and $a_1x_1 = 0$ forces $x_1 = 0$, which is excluded: $f(1) = 0$.

Let $k \ge 2$ and assume the statement for $k-1$. Let S be the set of k-tuples of pairwise distinct elements of $\mathbb{F}_p$, with zero allowed; then $|S| = \frac{p!}{(p-k)!}$, which is nonzero since $k \le \sum a_i < p$. The additive group $\mathbb{F}_p$ acts on S by simultaneous translation $x \mapsto x + t\mathbf{1}$, and this action is free (if $t \neq 0$ then every coordinate moves), so S splits into $\frac{1}{p}\cdot\frac{p!}{(p-k)!} = \frac{(p-1)!}{(p-k)!}$ orbits of size p.

Write $\Lambda(x) = \sum_i a_ix_i$ and $\sigma = \sum_i a_i$. Then $\Lambda(x+t\mathbf{1}) = \Lambda(x) + t\sigma$, and σ is invertible because $0 < \sigma < p$. Hence along each orbit Λ takes every value of $\mathbb{F}_p$ exactly once, so *exactly one* member of each orbit is a solution of $\Lambda = 0$. The number of solutions in S is therefore exactly $\frac{(p-1)!}{(p-k)!}$.

Split those solutions according to whether some coordinate is 0. By distinctness at most one coordinate can be 0. If none is, the solution is counted by $f(k)$. If $x_i = 0$ for exactly one index i, delete that coordinate: the remaining x_j are pairwise distinct and nonzero and satisfy $\sum_{j\neq i} a_jx_j = 0$, where the coefficients $(a_j)_{j\neq i}$ are again positive integers with sum less than p, so by the induction hypothesis there are $f(k-1)$ such solutions, for each of the k choices of i. Adding,

$$\frac{(p-1)!}{(p-k)!} \;=\; f(k) + kf(k-1),$$

which is the claimed recursion; in particular $f(k)$ depends only on k and p. □

Lemma 9.12.2. *With f as above, $f(k) = \dfrac{k!}{p}\left(\binom{p-1}{k} + (-1)^k(p-1)\right)$ for $1 \le k < p$.*

Proof. Put $g(k) = \dfrac{p\,f(k)}{k!}$. Dividing the recursion of Lemma 1 by $k!/p$ gives

$$g(k) \;=\; \frac{p}{k!}\cdot\frac{(p-1)!}{(p-k)!} \;-\; \frac{pf(k-1)}{(k-1)!} \;=\; \binom{p}{k} - g(k-1), \qquad g(1) = 0.$$

Pascal's rule $\binom{p-1}{k} + \binom{p-1}{k-1} = \binom{p}{k}$ says that $G(k) = \binom{p-1}{k}$ is a particular solution of the same relation. Hence $h := g - G$ satisfies the homogeneous relation $h(k) = -h(k-1)$, so $h(k) = (-1)^{k-1}h(1) = (-1)^{k-1}\big(0 - (p-1)\big) = (-1)^k(p-1)$. Therefore $g(k) = \binom{p-1}{k} + (-1)^k(p-1)$, and $f(k) = \frac{k!}{p}g(k)$. □

Applying it. Set $c_0 = c_1 = 1$ and $c_i = i$ for $2 \le i \le 63$, so that the given expression is $\Lambda(x) = \sum_{i=0}^{63} c_i x_i$ and

$$\sum_{i=0}^{63} c_i = 1 + \frac{63\cdot 64}{2} = 2017 = p \equiv 0.$$

Let D be the set of 64-tuples of pairwise distinct elements of $\mathbb{F}_p$ and let $M = \#\{x \in D : \Lambda(x) = 0\}$ be the required count. The group $\mathbb{F}_p$ acts on D by simultaneous translation; the action is free, so every orbit has size p and contains exactly one tuple with $x_0 = 0$. Since $\sum_i c_i \equiv 0$, the form Λ is constant on orbits, so the solutions of $\Lambda = 0$ in D are a union of whole orbits and

$$M \;=\; p \cdot \#\big\{x \in D : x_0 = 0,\ \Lambda(x) = 0\big\}.$$

If $x_0 = 0$ then $\Lambda(x) = x_1 + 2x_2 + \cdots + 63x_{63}$ and $x_1, \ldots, x_{63}$ are pairwise distinct and nonzero. The coefficients here are $1, 2, \ldots, 63$, positive integers with sum $2016 < 2017 = p$, so Lemma 1 applies with $k = 63$ and

$$M \;=\; p\,f(63) \;=\; 63!\left(\binom{2016}{63} + (-1)^{63}\cdot 2016\right) \;=\; 63!\binom{2016}{63} - 2016\cdot 63!.$$

Since $63!\binom{2016}{63} = \frac{2016!}{1953!}$, the answer is

$$\boxed{\frac{2016!}{1953!} \;-\; 2016\cdot 63!}.$$

Two checks before you put the pen down. First, the shape: $\frac{2016!}{1953!} = \frac{1}{2017}\cdot\frac{2017!}{1953!}$ is exactly one 2017th of the total number of 64-tuples of distinct residues, so the answer is the equidistribution prediction minus a correction of $2016\cdot 63!$, which is positive and astronomically smaller. A count that came out larger than the equidistribution prediction, or negative, would be wrong on sight. Second, run the whole argument at a miniature prime. Take $p = 7$, where $1 + 1 + 2 + 3 = 7$, so the analogous question asks for ordered 4-tuples of distinct elements of $\{1, \ldots, 7\}$ with $x_0 + x_1 + 2x_2 + 3x_3 \equiv 0 \pmod 7$; the formula predicts $3!\big(\binom{6}{3} - 6\big) = 6(20 - 6) = 84$, and enumerating the $7\cdot 6\cdot 5\cdot 4 = 840$ tuples confirms it. Building the miniature is five minutes and it catches every sign error.

Layer 1: WHAT

Block 1. 2017 is prime and $\{1, \ldots, 2017\}$ is a complete residue system, so work in $\mathbb{F}_p$ and forget the interval.

Block 2. Add the coefficients: $1 + 1 + 2 + \cdots + 63 = 2017$. Therefore Λ is invariant under simultaneous translation of all coordinates.

Block 3. Distinctness is also translation invariant. Quotient: normalise $x_0 = 0$ and multiply the count by p.

Block 4. The normalised question is "distinct nonzero solutions of a positive integer linear form with coefficient sum less than p". That shape is stable under deleting a variable, so it recurses.

Block 5. Count with zeros allowed, exactly, by the same orbit argument: $\frac{(p-1)!}{(p-k)!}$. Then peel off the solutions using 0 to get $f(k) + kf(k-1) = \frac{(p-1)!}{(p-k)!}$.

Block 6. Normalise the recursion to $g(k) + g(k-1) = \binom{p}{k}$ and read Pascal's rule as a particular solution; the homogeneous part is $(-1)^k$.

Block 7. Substitute $p = 2017$, $k = 63$, and convert $63!\binom{2016}{63}$ to $\frac{2016!}{1953!}$.

Juxtaposition. The nominal route treats distinctness as the enemy: it writes the count as (all tuples satisfying the congruence) minus (those with a repeat), and then spends the rest of the attempt trying to organise the repeats. That is a real solution, but it requires the partition lattice and its Mobius function, and it produces the correction term $-2016 \cdot 63!$ as the value $\mu_{\hat{1}} = (-1)^{63}63!$ of an object most candidates cannot even name.

Bifurcation point. The exact line is the one where you have just written $\sum c_i = 2017$. The zombie writes it down as a curiosity and carries on with inclusion and exclusion. You stop, ask what a coefficient sum divisible by p *means* for the linear form, and answer: it means the all ones vector is in the kernel, that is, Λ descends to the quotient by the diagonal. Everything after that is forced.

Layer 2: HOW

Block 1. **Cue.** A specific four digit modulus, and a ground set $\{1, \ldots, 2017\}$ of exactly that many elements. **Why natural.** 2017 is prime (check: it is not divisible by any prime up to 43), and the ground set has exactly p elements with $2017 \equiv 0$, so it is a complete residue system. A setter who wanted you to work in $\mathbb{Z}$ would have given you $\{1, \ldots, 2000\}$; giving you exactly p consecutive integers is an instruction to move to $\mathbb{F}_p$ and stop thinking about size. **Link to the next block.** In $\mathbb{F}_p$ the natural things to look for are symmetries, because a field has plenty and an interval of integers has none.

Block 2. **Cue.** The coefficients are $1, 1, 2, 3, \ldots, 63$, which is a triangular number plus one, and the modulus is 2017. **Why natural.** $1 + 2 + \cdots + 63 = 2016$ is a number you know on sight, and $2016 + 1 = 2017$. The doubled coefficient on x_0 and x_1 is unnatural unless it is there to make the total exact, so the setter has told you what to compute. And the meaning of $\sum c_i \equiv 0$ is not arithmetic but structural: the vector c is orthogonal to $\mathbf{1}$, so the linear functional Λ kills the diagonal line $\{t\mathbf{1}\}$ and therefore factors through the quotient $\mathbb{F}_p^{64}/\mathbb{F}_p\mathbf{1}$. **Link to the next block.** A symmetry of the algebraic constraint is only useful if the combinatorial constraint shares it, so check that next.

Block 3. **Cue.** "Distinct" is a statement about differences $x_i - x_j$, and translation does not change differences. **Why natural.** This is the crux and it is a definition level observation, not an insight. Distinctness is invariant under any map that preserves differences; translation is exactly such a map. So the diagonal $\mathbb{F}_p$ acts on the solution set, freely (a nonzero translation moves every coordinate) and with Λ constant on orbits, which means the solution set is a union of full orbits of size p. Counting orbits is the same as counting normalised representatives, and the cheapest normalisation is $x_0 = 0$. **Link to the next block.** Setting $x_0 = 0$ removes the coefficient c_0 from the form and forces the remaining variables to be nonzero, which is a new and simpler shape.

Block 4. **Cue.** After normalisation the coefficients are $1, 2, \ldots, 63$ with sum 2016, one less than the modulus. **Why natural.** The condition $0 < \sum a_i < p$ is exactly what the orbit argument needs, and it is inherited when you delete any variable, because all the a_i are positive. That inheritance is the signal to set up a recursion in the number of variables, and it is also the signal to state the lemma for *arbitrary* positive coefficients rather than for $1, 2, \ldots, k$: the induction step deletes a variable and

must appeal to the hypothesis for a different coefficient list, so the hypothesis has to be general enough to cover it. Strengthening the statement is what makes the induction possible. **Link to the next block.** To recurse you need the count with zeros allowed, which is the case the orbit argument handles exactly.

Block 5. **Cue.** The obstruction to recursion is the single value 0, which is distinguished only because it kills its coefficient. **Why natural.** If you allow zeros, the orbit argument gives an exact count with no error term: the translation action is free on tuples of distinct elements, and Λ is injective along each orbit precisely because $\sigma = \sum a_i \not\equiv 0$. So the count is (number of orbits) $= \frac{(p-1)!}{(p-k)!}$, independent of the coefficients. Then distinctness does you a favour: at most one coordinate can equal 0, so the correction is a clean $kf(k-1)$ with no overlaps to subtract. **Link to the next block.** You now have a first order recursion with a factorial style inhomogeneity, so normalise it.

Block 6. **Cue.** $f(k) = \frac{(p-1)!}{(p-k)!} - kf(k-1)$: the multiplier is $-k$, which is the derivative of a factorial. **Why natural.** Whenever the multiplier in $f(k) = \lambda_k f(k-1) + R(k)$ is $-k$, dividing by $k!$ turns it into constant multiplier -1; that is the summation factor method, and $k!$ is the summation factor here. Multiplying additionally by p is cosmetic, chosen so that the inhomogeneity becomes exactly $\binom{p}{k}$. At that point Pascal's rule hands you a particular solution for free, because Pascal's rule is literally the identity "a binomial coefficient is the sum of two consecutive binomial coefficients from the row above". The homogeneous solution of $h(k) = -h(k-1)$ is $C(-1)^k$, pinned by $g(1) = 0$. **Link to the next block.** Substitute the numbers.

Block 7. **Cue.** $63!\binom{2016}{63}$. **Why natural.** $\binom{n}{k}k! = \frac{n!}{(n-k)!}$ is the count of ordered selections, so the leading term is $\frac{2016!}{1953!}$, and writing it that way makes the equidistribution reading visible: it is $\frac{1}{2017}$ of the number of ordered 64-tuples of distinct residues. Presenting the answer as (expected count) minus (bias) is not decoration, it is the sanity check. **Link to the next block.** Verify on a miniature prime, then stop.

Layer 3: WHY

The headline truth is this. The congruence and the distinctness condition share a symmetry group, and when two constraints share a symmetry you should quotient by it rather than count against each one separately. Here the group is the diagonal copy of $\mathbb{F}_p$ acting by simultaneous translation. It preserves distinctness because distinctness only sees differences, and it preserves Λ because $\sum c_i \equiv 0$, that is, because $c \perp \mathbf{1}$. Quotienting costs nothing (the action is free, so you just divide by p) and it converts a symmetric problem into a rigid one in which 0 is now a special value, which is exactly the kind of asymmetry a recursion can eat.

The second truth explains the shape of the answer. Over all $\frac{2017!}{1953!}$ ordered 64-tuples of distinct residues, the value $\Lambda(x)$ is very nearly equidistributed among the 2017 residues: the naive prediction is $\frac{2016!}{1953!}$ per residue and the truth is that prediction minus $2016 \cdot 63!$, a relative error of order $63!$ / $\binom{2016}{63}$, which is minuscule. The bias is not noise: it is exactly $(p-1)$ times the number of orderings of 63 objects, and it is the same $(-1)^k(p-1)$ that fell out of the homogeneous part of the recursion. Through the other route it appears as the Mobius number of the partition lattice, $\mu(\hat{0}, \hat{1}) = (-1)^{n-1}(n-1)!$ with $n = 64$, that is $-63!$. Two derivations, one number: the alternating sign of inclusion and exclusion and the alternating homogeneous solution $(-1)^k$ of the recursion are the same phenomenon wearing different clothes.

The third truth is about why only the one block partition matters in the inclusion and exclusion route, and it is worth internalising even though the ideal route avoids it. For a partition π of the index set, the tuples constant on blocks satisfy a linear congruence whose coefficients are the block sums $\sum_{i \in S} c_i$; the congruence is degenerate exactly when every block sum vanishes modulo p. Since

all c_i are positive and their total is exactly p, every proper nonempty subset has block sum strictly between 0 and p, hence nonzero. So only the one block partition is degenerate. The condition $\sum c_i = p$ with all $c_i > 0$ is therefore doing double duty: it creates the translation symmetry and it simultaneously guarantees that no proper subsum is degenerate. That is a designed problem, and recognising design is part of the discernment.

Knowledge gaps, stated bluntly. You cannot approach this problem without the following installed, and none of it is in Further Mathematics.

(i) *Finite fields.* $\mathbb{F}_p = \mathbb{Z}/p\mathbb{Z}$ for p prime is a field: every nonzero element is invertible. You use this twice, and only twice: $\sigma = \sum a_i$ is invertible, and each individual a_i with $0 < a_i < p$ is nonzero. Without primality you would need $\gcd(\sigma, p) = 1$ instead; the problem does not need any deeper property of $\mathbb{F}_p$.

(ii) *Group actions, orbits, and free actions.* A group G acting on a finite set X partitions it into orbits; if the action is free every orbit has size $|G|$ and $|X| = |G| \cdot$ (number of orbits). A function constant on orbits, such as Λ here, descends to the orbit set, so counting its zeros is counting orbits. This is the single most transferable idea in the solution.

(iii) *Recursions with a strengthened hypothesis.* The lemma is stated for arbitrary positive coefficients with sum below p, not for $1, 2, \ldots, k$, because the induction step deletes a variable and lands on a different coefficient list. Learning to notice, before starting an induction, that the natural statement is too weak to induct on is a skill, not luck.

(iv) *First order linear recursions with a variable multiplier.* For $f(k) = \lambda_k f(k-1) + R(k)$, multiply by a summation factor $\prod \lambda_j^{-1}$ to reduce to constant multiplier, then find any particular solution and add $C\lambda^k$. Here $\lambda_k = -k$ and the summation factor is $k!$.

(v) *Optional but valuable: the partition lattice and its Mobius function.* $\mu(\hat{0}, \hat{1}) = (-1)^{n-1}(n-1)!$ on Π_n, equivalently the degree one coefficient of the falling factorial $x(x-1)\cdots(x-n+1)$ is the Stirling number $s(n,1) = (-1)^{n-1}(n-1)!$. You do not need it for the ideal route, but it is what explains the correction term.

Revolving concepts

Complete residue systems. A set of p consecutive integers is a complete residue system modulo p. Here $\{1, \ldots, 2017\}$ works because $2017 \equiv 0$; had the problem said $\{1, \ldots, 2016\}$ the reduction map would have missed 0 and the whole translation argument would collapse. See Book 1, field: Number theory, topic: residue systems.

Translation invariance of a linear form. $\Lambda(x + t\mathbf{1}) = \Lambda(x) + t\sum c_i$. So Λ is translation invariant if and only if $\sum c_i \equiv 0$, and if $\sum c_i \not\equiv 0$ then Λ is a bijection along every translation orbit. Both halves are used: the first to normalise $x_0 = 0$ in the main count, the second to count solutions inside the lemma. See Book 1, field: Linear algebra, topic: linear functionals and kernels.

Free group actions and orbit counting. If a finite group G acts freely on X, then $|X| = |G| \cdot |X/G|$, and any G-invariant subset is a union of orbits. Compare Burnside's lemma, which is the general tool when the action is not free; here freeness makes it unnecessary. See Book 1, field: Abstract algebra, topic: group actions.

Falling factorials, ordered selections and Pascal. $\frac{n!}{(n-k)!} = n(n-1)\cdots(n-k+1)$ counts ordered k-tuples of distinct elements from an n-set; $\binom{n}{k}k! = \frac{n!}{(n-k)!}$; Pascal's rule $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$. See Book 1, field: Combinatorics, topic: binomial identities.

The recursion, fully generalised. This is the recursion in the problem, so generalise it properly. The general first order linear recursion is

$$f(k) \;=\; \lambda_k f(k-1) + R(k), \qquad f(k_0) \text{ given.}$$

Set $\Pi_k = \prod_{j=k_0+1}^{k} \lambda_j$ and divide through by Π_k: with $F(k) = f(k)/\Pi_k$ the relation becomes

$F(k) = F(k-1) + R(k)/\Pi_k$, so

$$f(k) \;=\; \Pi_k \left(f(k_0) + \sum_{j=k_0+1}^{k} \frac{R(j)}{\Pi_j} \right).$$

The quantity $1/\Pi_k$ is the *summation factor*. In this problem $\lambda_k = -k$ and $\Pi_k = (-1)^{k-1}k!$, which is why dividing by $k!$ and flipping signs is the right normalisation; the relation becomes $g(k) = -g(k-1) + \binom{p}{k}$, and for a constant multiplier the alternative and usually faster method is (general solution) = (any particular solution) + $C\lambda^k$. Recognising Pascal's rule as a supply of particular solutions for the relation $g(k) + g(k-1) = \binom{p}{k}$ is the specific move worth banking: whenever the inhomogeneity is a binomial coefficient and the multiplier is -1, try a binomial coefficient from the row above. The same trick solves $g(k) + g(k-1) = \binom{n}{k}$ for every n, and, via $\binom{n}{k} = \binom{n-1}{k} + \binom{n-1}{k-1}$ iterated, any inhomogeneity that is a fixed linear combination of binomial coefficients. See Book 1, field: Sequences and recurrences, topic: first order linear recursions and summation factors.

Inclusion and exclusion over set partitions. The intersections of the events $\{x_i = x_j\}$ are indexed by the partition lattice Π_n, and the correct alternating coefficients are the Mobius numbers μ_π, characterised by $\sum_{\pi|\sigma} \mu_\pi = [\sigma = \hat{0}]$. Comparing degree one coefficients in $|A|(|A|-1)\cdots(|A|-n+1) = \sum_\pi \mu_\pi |A|^{|\pi|}$ gives $\mu_{\hat{1}} = (-1)^{n-1}(n-1)!$, which for $n = 64$ is $-63!$: exactly the correction term in the answer. See Book 1, field: Combinatorics, topic: Mobius inversion on posets.

Circuit Bank

- **WHEN** a congruence has explicit numerical coefficients and an explicit prime modulus **DO** add the coefficients before doing anything else, and compare the total with the modulus **BECAUSE** setters hide the entire structure in that one arithmetic fact, and a coefficient sum divisible by p means the form is invariant under simultaneous translation
- **WHEN** a count is over objects with pairwise distinct entries and you are tempted by inclusion and exclusion **DO** first look for a group action preserving both the constraint and the distinctness, and quotient by it **BECAUSE** quotienting by a free action is exact and sign free, whereas inclusion and exclusion over distinctness lands you in the partition lattice
- **WHEN** a normalisation makes one value special (here 0 after setting $x_0 = 0$) **DO** set up a recursion in the number of variables by peeling off the tuples that use the special value **BECAUSE** distinctness guarantees at most one coordinate takes that value, so the correction term is a clean multiple of the previous count with nothing to subtract twice
- **WHEN** you are about to induct and the statement mentions specific constants **DO** restate the claim for arbitrary data of the same shape, choosing the hypotheses so they are inherited when you delete an object **BECAUSE** the induction step almost never lands on the original constants, and a statement you cannot apply to the reduced case is a statement you cannot induct on
- **WHEN** a recursion has the form $f(k) = -kf(k-1) + R(k)$ **DO** divide by $k!$ to get a constant multiplier, then look for a particular solution of the shape of R **BECAUSE** $k!$ is the summation factor for the multiplier $-k$, and for R a binomial coefficient Pascal's rule supplies the particular solution immediately

Backtesting

1. First useful cue: the coefficient list $1, 1, 2, 3, \ldots, 63$, whose sum is $1 + \frac{63\cdot 64}{2} = 2017$, the modulus itself; and the ground set $\{1, \ldots, 2017\}$ being a complete residue system.

2. Method the expression suggested: the linear form is invariant under simultaneous translation, and so is distinctness, so quotient by the diagonal action, normalise $x_0 = 0$, and recurse on the number of variables.
3. Where you would hesitate, and the thought to reinforce: at the decision to state the lemma for arbitrary positive coefficients with sum below p rather than for $1, 2, \ldots, 63$. Reinforce: the induction deletes a variable, so the hypothesis must survive deletion; positivity and "sum below p" are precisely the two properties that do.
4. The trap the question sets: the distinctness condition, which invites inclusion and exclusion and then punishes it. The secondary trap is arithmetic: the coefficients are $1, 1, 2, 3, \ldots$, not $0, 1, 2, 3, \ldots$, and if you misread the first two you get a coefficient sum of 2016 instead of 2017 and the whole method evaporates.
5. What makes you recognise it faster next time: the pairing of a prime modulus with a coefficient sum equal to that prime is a designed signature. Any time a competition gives an explicit numerical linear form modulo a prime, add the coefficients; if they sum to 0 modulo p, the diagonal translation is your first move.
6. The reusable decision rule: when two constraints in a counting problem share a symmetry group, quotient by the group instead of counting against the constraints separately.

Mirror drills

1. Let $p > 3$ be prime. How many ordered pairs (x_1, x_2) of distinct nonzero elements of $\mathbb{F}_p$ satisfy $x_1 + 2x_2 = 0$? *Answer:* $p - 1$. Directly: $x_1 = -2x_2$ with $x_2 \neq 0$ forces $x_1 \neq 0$, and $x_1 \neq x_2$ fails only if $3x_2 = 0$, impossible for $p > 3$. Via the formula, $f(2) = \frac{2}{p}\left(\binom{p-1}{2} + (p-1)\right) = p - 1$.
2. Prove that for positive integers $a_1, \ldots, a_k$ with $a_1 + \cdots + a_k < p$, the number of k-tuples of distinct elements of $\mathbb{F}_p$ (zeros allowed) with $\sum a_i x_i = 0$ is exactly $\frac{(p-1)!}{(p-k)!}$, independent of the a_i. *Answer:* the diagonal translation acts freely with $\frac{(p-1)!}{(p-k)!}$ orbits, and along each orbit $\Lambda(x + t\mathbf{1}) = \Lambda(x) + t\sigma$ runs over all of $\mathbb{F}_p$ bijectively because σ is invertible, so exactly one member of each orbit is a solution.
3. Compute $\mu(\hat{0}, \hat{1})$ for the partition lattice Π_n. *Answer:* $(-1)^{n-1}(n-1)!$, by comparing the coefficient of x^1 on both sides of $x(x-1)\cdots(x-n+1) = \sum_\pi \mu_\pi x^{|\pi|}$: only $\pi = \hat{1}$ has one block, and the left side has x^1 coefficient $s(n, 1) = (-1)^{n-1}(n-1)!$. For $n = 64$ this is $-63!$, the bias in the answer to B6.

Chapter 10

Putnam 2016

10.1 A1. derivative order forcing divisibility by 2016

Find the smallest positive integer j such that for every polynomial $p(x)$ with integer coefficients and for every integer k, the integer

$$p^{(j)}(k) = \frac{d^j}{dx^j} p(x)\Big|_{x=k}$$

(the j-th derivative of $p(x)$ at k) is divisible by 2016.

Read it out loud

Two universal quantifiers sit inside this question before you ever reach the number 2016: "for every polynomial with integer coefficients" and "for every integer k". Whenever a divisibility claim is dressed up with "for every" over a structured family like this, the claim cannot secretly depend on which member of the family you picked, so your first fifteen seconds should be spent asking what structure the whole family shares. Every integer-coefficient polynomial is built from the monomials x^m by integer linear combinations, and differentiation is linear, so whatever the j-th derivative does to 2016-divisibility, it does uniformly across that whole basis. That is the field this problem lives in: calculus supplies the mechanism, the power rule iterated j times, and number theory supplies the target, divisibility by a fixed composite integer.

The second thing worth writing down inside that first half minute is the prime factorisation of 2016, because a fixed integer appearing as a divisibility target is never meant to be used as 2016; it is meant to be used as $2^5 \cdot 3^2 \cdot 7$. Any claim "N divides something, for every choice of parameters" reduces, by unique factorisation, to a set of independent prime power claims, and those are always the real content.

The final cue is the shape of the question itself: "find the smallest j" is an extremal claim, and an extremal claim always splits into two separate jobs, a proof that some j always works (an upper bound, true for every p and k) and a proof that no smaller j always works (a lower bound, witnessed by one specific failing case). Reading the question out loud, you should already suspect, before writing a single equation, that the eventual argument finds a formula for $p^{(j)}(k)$ that visibly contains a factorial, then turns the whole problem into finding the least j with $2016 \mid j!$.

Approaches, ranked worst to best

1. **numerical trial on sample polynomials.** *Cue:* "find the smallest j" invites experimenting with $p(x) = x^2, x^3, \ldots$ at a handful of integers k. *Circuit:* guess and check. (this closes neither half of a minimality argument: no finite experiment proves a property for every integer-coefficient polynomial and every integer k, and no finite experiment reliably finds the true worst case, since a badly chosen sample polynomial can look fine by accident. It costs unlimited ink and delivers no proof.)

2. **route the derivative through finite differences.** *Cue:* integer-valued sequences and divisibility patterns often go through the finite difference operator Δ, whose iterates $\Delta^j p(k)$ are always integer combinations of values of p. *Circuit:* Polya-Mahler style integer-valued polynomial reasoning. (a genuine technique elsewhere, but D^j and Δ^j are different operators that only agree in the limit as a step size shrinks to zero, so this route needs a translation lemma between the two operators before it can say anything about $p^{(j)}(k)$. That lemma costs more ink than differentiating x^m directly, and lands on the same term regardless.)

3. **expand in the monomial basis and differentiate term by term.** *Cue:* "for every polynomial with integer coefficients" is precisely the hypothesis that licenses working with the general term $a_m x^m$ and invoking linearity. *Circuit:* linearity reduction to the generating basis, closed by

one extremal test polynomial. (this reaches the factor $j!\binom{m}{j}$ in one line, hands you sufficiency for every p and k at once, and needs only $p(x) = x^j$ for the converse. No detour, no unproven translation lemma, and the remaining work is arithmetic on $2016 = 2^5 \cdot 3^2 \cdot 7$.)

The ideal solution

Write $p(x) = \sum_{m=0}^{n} a_m x^m$ with every $a_m \in \mathbb{Z}$. For $m \geq j$, the power rule applied j times gives

$$\frac{d^j}{dx^j} x^m = m(m-1)(m-2)\cdots(m-j+1)\, x^{m-j} = j!\binom{m}{j} x^{m-j},$$

and for $m < j$ this derivative is 0. Summing over m and using linearity of differentiation,

$$p^{(j)}(x) = j! \sum_{m=j}^{n} \binom{m}{j} a_m x^{m-j},$$

so for every integer k,

$$p^{(j)}(k) = j! \sum_{m=j}^{n} \binom{m}{j} a_m k^{m-j}.$$

Every factor inside that sum is an integer: $\binom{m}{j}$ is an integer because it counts subsets, a_m is an integer by hypothesis, and k^{m-j} is an integer power of an integer. So $p^{(j)}(k)$ is $j!$ times an integer, for every integer-coefficient p and every integer k, with no further hypothesis on j at all. This proves that whenever $2016 \mid j!$, the required divisibility holds universally.

For the converse, apply the hypothesis to the single polynomial $p(x) = x^j$, whose j-th derivative is the constant $j!$ for every x, so $p^{(j)}(k) = j!$ for every integer k. If j is to satisfy the problem's condition, this constant must in particular be divisible by 2016. So the condition "$p^{(j)}(k)$ is divisible by 2016 for every integer-coefficient p and every integer k" holds if and only if $2016 \mid j!$: the paragraph above gave sufficiency, this paragraph gives necessity, and together they say the answer is exactly the smallest j with $2016 \mid j!$.

Factorise $2016 = 2^5 \cdot 3^2 \cdot 7$ (since $2016 = 32 \cdot 63 = 2^5 \cdot 9 \cdot 7$). The prime 7 divides none of $1, 2, \ldots, 6$, so $v_7(j!) = 0$ for every $j \leq 6$, and already $j \geq 7$ is forced. At $j = 7$: $7! = 5040$, and $5040/2016 = 2.5$, not an integer, so $2016 \nmid 7!$; equivalently $v_2(7!) = \lfloor 7/2 \rfloor + \lfloor 7/4 \rfloor = 3 + 1 = 4$, one short of the required 5. At $j = 8$: $8! = 40320 = 20 \cdot 2016$ exactly, so $2016 \mid 8!$; checking each prime confirms this is not an accident: $v_2(8!) = 4 + 2 + 1 = 7 \geq 5$, $v_3(8!) = 2 + 0 = 2 \geq 2$, $v_7(8!) = 1 \geq 1$. Since $j = 7$ fails and $j = 8$ succeeds,

$$j = \boxed{8}.$$

Layer 1: WHAT

Block 1: expand p in the monomial basis and push the derivative through the sum by linearity, reducing to the single computation $D^j x^m$. Block 2: recognise the falling factorial $m(m-1)\cdots(m-j+1)$ as $j!\binom{m}{j}$, which strips p and k out of the argument and leaves a bare statement about $j!$. Block 3: realise the converse needs only one polynomial, the one whose j-th derivative collapses to a constant with nothing left over, $p(x) = x^j$. Block 4: convert to pure number theory, factorise 2016, and hunt the rarest prime first.

The nominal route tries to reason about $p^{(j)}(k)$ mod 2016 directly, case by case on small numerical polynomials, without ever writing the general monomial expansion; it never produces a clean sufficient condition and can never rule out smaller j with certainty. The bifurcation point is the decision, at the very first line, to write $p(x) = \sum a_m x^m$ symbolically and differentiate the generic term x^m, rather than picking numbers to differentiate. That one choice, a generic monomial instead of a numerical example, turns an open-ended search into a two-line identity.

Layer 2: HOW

Block 1. **Cue.** The hypothesis quantifies over every integer-coefficient polynomial, and every such polynomial is an integer linear combination of monomials. **Why natural.** Differentiation is linear, so once you know what D^j does to every monomial x^m, superposition tells you what it does to every polynomial at once; there is never a need to reason about an unspecified p again. **Link to the next block.** Applying the power rule j times to x^m produces the falling factorial $m(m-1)\cdots(m-j+1)$, a raw product of j consecutive integers that begs to be simplified.

Block 2. **Cue.** A product of j consecutive integers, $m(m-1)\cdots(m-j+1)$. **Why natural.** This product equals $j!\binom{m}{j}$, and $\binom{m}{j}$ is an integer for every integer m because it counts subsets, so the product is automatically $j!$ times an integer with no extra argument required. **Link to the next block.** With p and k both absorbed into "some integer", the remaining problem is a statement about $j!$ alone, so the next move is to stop thinking about polynomials and start thinking about which j make 2016 divide $j!$.

Block 3. **Cue.** The problem is a minimality claim, so sufficiency alone is not enough; some specific p and k must force the bound to be tight. **Why natural.** The cleanest test case is the one where $p^{(j)}$ collapses to a bare constant with nothing left over: $p(x) = x^j$ gives $p^{(j)}(x) = j!$ identically, independent of k, so there is no residual freedom for the divisibility to slip through on a technicality. **Link to the next block.** Plugging this single polynomial in shows $2016 \mid j!$ is not merely sufficient but forced, so the problem is now, precisely and only, "find the smallest j with $2016 \mid j!$".

Block 4. **Cue.** A fixed composite integer, 2016, must divide a quantity that grows with a parameter j. **Why natural.** Unique factorisation says $2016 \mid j!$ exactly when every prime power in $2016 = 2^5 \cdot 3^2 \cdot 7$ divides $j!$, so one hard question splits into three easy ones, and the rarest prime, here 7, which does not appear at all until $j = 7$, is always worth checking first, since it can never be satisfied early. **Link to the next block.** Checking $j = 7$ shows the count of 2s is one short; checking $j = 8$, the very next candidate, shows all three prime powers are now present, closing the problem.

Layer 3: WHY

The fact beneath the whole solution is that a product of j consecutive integers is always divisible by $j!$, because that product is, by definition, $j!$ times a binomial coefficient, and binomial coefficients count things; every other feature of the statement, which polynomial, which integer k, only disguises that the true content is one arithmetic fact about factorials. The discernment mechanism is to treat "for every polynomial with integer coefficients" as a claim about a $\mathbb{Z}$-module with a known generating set, so that any linear operator on it, derivative, finite difference, or evaluation, is fixed entirely by its action on the monomials. You train this by asking, on first reading such a phrase, whether the final claim could possibly depend on which polynomial was chosen; here it plainly cannot, since the target is the fixed number 2016, and that tells you the reduction to generators is forced before you write a further line. The examiners nudge you the same way by naming a specific composite, 2016, rather than an abstract N: naming it is an invitation to factorise it and treat its three primes as three small separate puzzles, the strategy that finishes the problem in four lines of arithmetic.

Revolving concepts

Falling factorials and binomial coefficients. For integers $m \geq j \geq 0$, $m(m-1)\cdots(m-j+1) = j!\binom{m}{j}$, and $\binom{m}{j}$ extends to every integer m, including negative m, via $\binom{-n}{j} = (-1)^j\binom{n+j-1}{j}$, remaining an integer throughout. This identity underlies the binomial theorem, Vandermonde's identity, and

every argument that a product of consecutive integers is divisible by a factorial. See Book 1, field: Polynomials, topic: falling factorial and binomial coefficient identities.

Legendre's formula. The exponent of a prime p in $n!$ is $v_p(n!) = \sum_{i \geq 1} \lfloor n/p^i \rfloor$, equivalently $\frac{n - s_p(n)}{p-1}$, where $s_p(n)$ is the base p digit sum of n. This is the standard tool for any question asking for the smallest factorial divisible by a given number, and it generalises to Kummer's theorem, which counts the carries produced when adding two numbers in base p and relates that count to the prime valuation of a binomial coefficient. See Book 1, field: Number theory, topic: prime valuations of factorials.

Integer-valued polynomials and Polya's basis theorem. A polynomial with rational coefficients that takes an integer value at every integer, such as $\binom{x}{2} = \frac{x(x-1)}{2}$, need not have integer coefficients itself; Polya's theorem says every such polynomial is a $\mathbb{Z}$-linear combination of the binomial polynomials $\binom{x}{n}$, with coefficients equal to the finite differences of p at 0. Bhargava extended this from $\mathbb{Z}$ to arbitrary subsets of Dedekind domains using a recursively built sequence called a P-ordering, which replaces the ordinary factorial $n!$ by a generalised factorial $k!_P$: at each step you adjoin the element minimising a certain valuation against the elements already chosen, and ordinary Legendre's formula is recovered when the domain is $\mathbb{Z}$ with its usual valuation. This is genuinely beyond the syllabus, and is the honest answer to why the official remark mentions Dedekind domains: it is not a trick, it is a named research-level generalisation.

Circuit Bank

- **WHEN** a hypothesis quantifies over every integer-coefficient polynomial **DO** expand in the monomial basis and use linearity of whichever operator is in play, derivative, finite difference, or evaluation **BECAUSE** a linear operator on a module is determined entirely by its action on a generating set, so testing every monomial once replaces testing every polynomial forever
- **WHEN** a product of consecutive integers appears, in a derivative, a permutation count, or a combinatorial identity **DO** rewrite it immediately as $j!$ times a binomial coefficient **BECAUSE** binomial coefficients are always integers, so this rewriting produces a factorial divisor for free, with no case analysis
- **WHEN** a fixed composite number N must divide a quantity that grows with a parameter **DO** factorise N into prime powers, find the least parameter value meeting each prime power's threshold separately using Legendre's formula, then take the maximum of those thresholds **BECAUSE** divisibility by N splits, by unique factorisation, into independent divisibility by each prime power in N, and each of those is a short, fully determined calculation
- **WHEN** a problem asks for the smallest parameter value that works for every case in some family **DO** budget two separate arguments immediately, a general proof valid for every case, and one explicit failing case at the next value down **BECAUSE** a minimality claim is an equality between a proven bound and an achieved value, and only an explicit example can prove that a smaller value genuinely fails

Backtesting

1. First useful cue: "for every polynomial with integer coefficients", a universally quantified claim over a family with a known generating basis.
2. Method the expression suggested: expand in the monomial basis, differentiate term by term, and watch a factorial fall out of the falling factorial by linearity.
3. Where you would hesitate, and the thought to reinforce: after seeing $j! \mid p^{(j)}(k)$ always, it is tempting to stop, forgetting that minimality demands an explicit failing polynomial one value down, not merely the absence of a proof for it.

4. The trap the question sets: assuming the largest prime power in 2016's factorisation, 2^5, must be the bottleneck, when the prime that appears latest, here 7, absent from $j!$ until $j = 7$, is what actually forces the lower bound, with the power of 2 only barely failing at that same value.
5. What makes you recognise it faster next time: seeing a specific four digit constant like 2016 inside a divisibility claim and reflexively factorising it before doing anything else, since such a constant is chosen for its factorisation, never left arbitrary.
6. The reusable decision rule: WHEN a fixed integer N must divide a quantity that only grows with a parameter j, such as $j!$, DO factorise N and test candidate values of j against every prime power in turn, starting with the prime that first appears latest inside $j!$, BECAUSE that prime, not the one with highest exponent, usually sets the true threshold.

Mirror drills

1. Find the smallest positive integer j such that $p^{(j)}(k)$ is divisible by 720 for every integer-coefficient p and every integer k. *Answer:* $720 = 2^4 \cdot 3^2 \cdot 5$; $j = 5$ fails since $v_2(5!) = 3 < 4$; $j = 6$ gives $6! = 720$ exactly, so $j = 6$.
2. Redo the original problem with 2016 replaced by $4032 = 2^6 \cdot 3^2 \cdot 7$. *Answer:* still $j = 8$, since $v_2(8!) = 7 \geq 6$, $v_3(8!) = 2$ and $v_7(8!) = 1$ already hold, while $j = 7$ still fails on the power of 2 alone.
3. Prove, without quoting Legendre's formula, that $j!$ divides any product of j consecutive integers, even when some of them are negative. *Answer:* if the product includes 0 it is trivially divisible; otherwise the product equals $j!\binom{m}{j}$ for the largest integer m in the range, and $\binom{m}{j}$ remains an integer for negative m by $\binom{-n}{j} = (-1)^j\binom{n+j-1}{j}$.

10.2 A2. limiting growth rate of a binomial threshold

Given a positive integer n, let $M(n)$ be the largest integer m such that

$$\binom{m}{n-1} > \binom{m-1}{n}.$$

Evaluate

$$\lim_{n\to\infty} \frac{M(n)}{n}.$$

Read it out loud

The two objects being compared are binomial coefficients whose indices are each offset from one another by a small fixed amount: the top index moves from m to $m-1$, and the bottom index moves from $n-1$ to n. Whenever two binomial coefficients sit this close to one another, the correct reflex is to form their ratio rather than to expand each one and subtract. A ratio is cheap here because the factorial numerators and denominators of $\binom{m}{n-1}$ and $\binom{m-1}{n}$ overlap almost completely; a difference has no such overlap to exploit and generally produces an expression with no simplification at all. That is the first field this problem lives in: combinatorics supplying the object, algebra supplying the simplification.

The second cue sits in the closing instruction: you are asked to evaluate $\lim_{n\to\infty} M(n)/n$, not to find a closed form for $M(n)$ itself. That single word, limit, tells you before any algebra that $M(n)$ grows linearly in n and that only the proportionality constant is wanted. This licenses a strategy that would be unacceptably lossy if an exact formula for $M(n)$ were required: bounding $M(n)$ between two quantities that share the same behaviour as $n \to \infty$, without ever pinning $M(n)$ down to the nearest integer. Recognising this early saves you from the much harder task of computing a floor function exactly.

The third cue is what happens after the ratio simplifies: a comparison of two combinatorial quantities becomes, in one line, a polynomial inequality in the single variable m. Once a discrete combinatorial condition has been converted into an algebraic inequality with no residual number theoretic content (no coprimality, no parity, nothing but a quadratic), the rest of the problem is pure single variable algebra: find the roots, read off the sign, and pass to a limit. The moment you see the ratio become a rational function of m and n, you should already expect the final constant to be a root of some quadratic, and quadratics with a 3 and a 5 in them on the Putnam are worth being alert to, since $x^2 - 3x + 1 = 0$ is the characteristic equation you get from squaring the roots of the Fibonacci equation $x^2 - x - 1 = 0$, and its larger root is $(3+\sqrt{5})/2$.

Approaches, ranked worst to best

1. **numerical search and pattern spotting.** *Cue:* compute $M(n)$ directly for $n = 1, 2, 3, \ldots$ by testing integers m one at a time. *Circuit:* small case exploration. (this never produces a proof, and it converges slowly: because the limiting constant $(3+\sqrt{5})/2$ is irrational, decimal approximations to $M(n)/n$ from small n settle only gradually, and nothing about the numerical data alone tells you the constant is a root of a quadratic rather than, say, a rational approximation or a different surd entirely.)

2. **Stirling asymptotics on the binomial coefficients directly.** *Cue:* "evaluate the limit" invites approximating both binomial coefficients for $m = cn$ and large n using a Stirling type expansion, then solving for the crossing value of c. *Circuit:* continuous approximation of a discrete extremal problem. (this can locate the constant $c = (3+\sqrt{5})/2$ reasonably quickly as a heuristic, but turning "the approximate crossing point of two Stirling expansions" into a genuine proof that $M(n)/n$ converges to exactly that value requires controlling the error

terms in the approximation to the same precision that the exact algebraic route gives for free. You end up redoing the rigorous part of the argument anyway, after having spent ink on an approximation that was never going to be the final proof.)

3. **exact ratio, then quadratic, then sandwich.** *Cue:* the offset between the two binomial coefficients is small and fixed, which is precisely the situation where the ratio collapses by factorial cancellation to a low degree rational function. *Circuit:* ratio reduction followed by exact root extraction and a squeeze. (this is exact at every step: no approximation is introduced anywhere, the quadratic's roots are found in closed form via the quadratic formula, and the final limit is read off the roots directly. The only extra care needed, confirming that the region where the ratio identity is valid actually contains the relevant root, costs a single short paragraph and closes a gap the fastest write ups sometimes leave open.)

The ideal solution

Fix a positive integer n and restrict attention first to integers $m > n+1$, so that $m-1 \geq n$ and both $\binom{m}{n-1}$ and $\binom{m-1}{n}$ are genuine positive binomial coefficients. For such m,

$$\frac{\binom{m}{n-1}}{\binom{m-1}{n}} = \frac{\dfrac{m!}{(n-1)!\,(m-n+1)!}}{\dfrac{(m-1)!}{n!\,(m-n-1)!}} = \frac{m!}{(m-1)!}\cdot\frac{n!}{(n-1)!}\cdot\frac{(m-n-1)!}{(m-n+1)!} = m\cdot n\cdot\frac{1}{(m-n+1)(m-n)} = \frac{mn}{(m-n+1)}$$

Every factorial term cancelled except the single outermost factor at each of the four ends, because the two binomial coefficients differ by exactly one in each index; that is what makes this ratio worth computing in the first place.

Since $m > n+1$ makes $m-n+1$ and $m-n$ both positive, the ratio above is a ratio of two positive quantities, so $\binom{m}{n-1} > \binom{m-1}{n}$ holds exactly when this ratio exceeds 1, that is when $mn > (m-n+1)(m-n)$. Expanding the right side, $(m-n+1)(m-n) = (m-n)^2 + (m-n) = m^2 - 2mn + n^2 + m - n$, so the condition rearranges to the single quadratic inequality

$$m^2 - (3n-1)m + (n^2 - n) < 0.$$

The coefficient of m^2 is positive, so this holds exactly on the open interval strictly between the two real roots of $m^2 - (3n-1)m + (n^2-n) = 0$, provided those roots are real. By the quadratic formula, the roots are

$$\alpha(n) = \frac{(3n-1) + \sqrt{5n^2 - 2n + 1}}{2}, \qquad \beta(n) = \frac{(3n-1) - \sqrt{5n^2 - 2n + 1}}{2},$$

where the discriminant simplifies as $(3n-1)^2 - 4(n^2-n) = 9n^2 - 6n + 1 - 4n^2 + 4n = 5n^2 - 2n + 1$. This discriminant is positive for every real n, since as a quadratic in n its own discriminant is $4 - 20 < 0$, so $\alpha(n)$ and $\beta(n)$ are always real and distinct, and for $m > n+1$,

$$\binom{m}{n-1} > \binom{m-1}{n} \iff \beta(n) < m < \alpha(n).$$

Two facts confirm this description genuinely captures $M(n)$. First, $\alpha(n) - \beta(n) = \sqrt{5n^2 - 2n + 1}$, and $5n^2 - 2n + 1 - n^2 = 4n^2 - 2n + 1 > 0$ for every $n \geq 1$, so $\alpha(n) - \beta(n) > n \geq 1$: the open interval always contains at least one integer, so the set of valid m is never empty and $M(n)$ is genuinely well defined. Second, the region $m > n+1$ used to derive the ratio identity must actually contain the top of that interval, or the argument would only be describing part of the picture. Since $5n^2 - 2n + 1 > (2n-1)^2 = 4n^2 - 4n + 1$ whenever $n^2 + 2n > 0$, true for every $n \geq 1$, we get $\sqrt{5n^2 - 2n + 1} > 2n - 1$ and hence $\alpha(n) > \frac{(3n-1)+(2n-1)}{2} = \frac{5n-2}{2}$, which exceeds $n+1$ once $n \geq 2$.

So for every $n \geq 2$, $\alpha(n)$ sits well inside the region $m > n+1$ where the ratio identity is exact, and because a single real quadratic's sign pattern (negative strictly between its roots, positive elsewhere) is a global statement about every real m, not merely those in one sub-range, the largest integer anywhere satisfying the original inequality is exactly the largest integer strictly below $\alpha(n)$, with nothing missed by having derived the identity only for $m > n+1$.

Consequently $M(n)$ is the greatest integer strictly less than $\alpha(n)$, so

$$\alpha(n) - 1 \leq M(n) < \alpha(n).$$

To find $\lim \alpha(n)/n$, pull n out from under the square root: for $n > 0$,

$$\frac{\alpha(n)}{n} = \frac{3 - \frac{1}{n} + \sqrt{5 - \frac{2}{n} + \frac{1}{n^2}}}{2} \longrightarrow \frac{3+\sqrt{5}}{2} \qquad (n \to \infty).$$

Dividing the sandwich $\alpha(n) - 1 \leq M(n) < \alpha(n)$ through by n, both outer bounds converge to $\frac{3+\sqrt{5}}{2}$ as $n \to \infty$, so by the sandwich theorem,

$$\lim_{n\to\infty} \frac{M(n)}{n} = \boxed{\frac{3+\sqrt{5}}{2}}.$$

Layer 1: WHAT

Block 1: recognise the two binomial coefficients differ by a fixed small shift in both indices, and choose the ratio over the difference. Block 2: cancel the ratio by factorials down to the rational function $mn/[(m-n+1)(m-n)]$. Block 3: convert "ratio exceeds 1" into a single quadratic inequality in m, and extract its two real roots via the quadratic formula. Block 4: confirm the region of validity actually contains the controlling root, then sandwich $M(n)$ between the larger root and that root minus one, divide by n, and take the limit.

The nominal route tries to manipulate the difference $\binom{m}{n-1} - \binom{m-1}{n}$ directly, perhaps via Pascal's rule or a combinatorial bijection, hunting for a sign without ever forming a ratio; this produces a higher degree, less transparent expression that does not obviously reduce to a quadratic at all. The bifurcation point is the decision, at the very first line, to divide the two binomial coefficients rather than subtract them: that single choice is what turns a combinatorial identity search into a two line factorial cancellation.

Layer 2: HOW

Block 1. **Cue.** The two binomial coefficients have upper indices $m, m-1$ and lower indices $n-1, n$, each pair differing by exactly one. **Why natural.** Nearby binomial coefficients share almost all of their factorial content, so their ratio, unlike their difference, collapses to a low degree rational function; a comparison phrased with a strict inequality between two positive quantities is always more cheaply attacked through a ratio to 1 than through a difference to 0. **Link to the next block.** Writing out the ratio and cancelling the shared factorial terms leaves only the outermost factors, m, n, and the two linear terms $m-n+1$, $m-n$.

Block 2. **Cue.** The simplified ratio $\dfrac{mn}{(m-n+1)(m-n)}$, compared against the threshold 1. **Why natural.** Since $m > n+1$ makes both $m-n+1$ and $m-n$ positive, comparing the ratio to 1 is safe to cross multiply without flipping any inequality, and the resulting condition $mn > (m-n+1)(m-n)$ is a single polynomial inequality with no combinatorics left in it at all. **Link to the next block.** Expanding the right hand side collects everything into one quadratic in m, whose sign is fully governed by its two real roots.

Block 3. **Cue.** A quadratic inequality $m^2 - (3n-1)m + (n^2 - n) < 0$ with a positive leading coefficient. **Why natural.** An upward opening quadratic is negative exactly between its two real roots and positive outside, so the entire solution set for m, for every fixed n, is the open interval between $\beta(n)$ and $\alpha(n)$ found from the quadratic formula, with no further case analysis needed. **Link to the next block.** Since the problem only wants a limiting ratio, the next step is not to solve for $M(n)$ but to bound it tightly enough between quantities whose ratio to n has a known limit.

Block 4. **Cue.** The closing instruction asks for $\lim_{n\to\infty} M(n)/n$, a normalised limit, not an exact value of $M(n)$. **Why natural.** $M(n)$ is the greatest integer below $\alpha(n)$, so it is trapped between $\alpha(n) - 1$ and $\alpha(n)$ regardless of how $\alpha(n)$'s fractional part behaves; dividing by n and pulling n out from under the square root inside $\alpha(n)$ makes both bounds converge to the same value. **Link to the next block.** The sandwich theorem then hands over the limit directly, with no need ever to know $M(n)$ itself.

Layer 3: WHY

The truth beneath the whole solution is that a ratio of two binomial coefficients whose indices are offset by a bounded amount is always a low degree rational function, because the bulk of each factorial is shared and cancels; this is the same cancellation that drives Pascal's rule and the ratio test used to find the largest term of a binomial expansion. The discernment mechanism is to notice, on first reading, whether a comparison of combinatorial quantities is phrased through nearby indices (ratio territory) or through unrelated indices (where no such cancellation exists and a direct combinatorial argument is genuinely needed); training this means asking, before any algebra, "how many index steps separate these two binomial coefficients", since a small answer is the signal that a ratio will collapse cleanly.

A second, independent discernment is needed for the closing instruction: recognising that "evaluate the limit of $X(n)/n$" is a standing licence to bound $X(n)$ rather than compute it, since any two bounds sharing the same leading order behaviour already prove the limit by the sandwich theorem, however badly they might disagree about $X(n)$'s exact value or its fractional part. The examiners nudge you toward the whole route by giving you a problem where the exact value of $M(n)$ genuinely has no clean closed form (it depends on the floor of an irrational number), while the limiting ratio does; asking for the limit rather than the exact value is itself the hint that a closed form for $M(n)$ was never the intended target.

Revolving concepts

Ratios of nearby binomial coefficients. When the upper and lower indices of two binomial coefficients each differ from one another by a bounded fixed amount, their ratio reduces after factorial cancellation to a rational function of bounded degree; here $\binom{m}{n-1}/\binom{m-1}{n} = mn/[(m - n + 1)(m - n)]$ for $m > n + 1$. The same cancellation is what proves the ratio test for the largest binomial coefficient in a row of Pascal's triangle, and what justifies the standard recursive definition $\binom{m}{k} = \binom{m-1}{k-1} \cdot \frac{m}{k}$. See Book 1, field: Combinatorics, topic: consecutive binomial coefficient ratios.

Sign of an upward opening quadratic. For $a > 0$, $am^2 + bm + c$ is negative exactly on the open interval strictly between its two real roots when the discriminant is positive, and nonnegative everywhere else; this is the entire content needed to pass from a quadratic inequality to a description of the solution set, with no further case work. See Book 1, field: Polynomials, topic: quadratic sign charts.

The sandwich theorem for a normalised extremal quantity. If $a_n \le X(n) < b_n$ for all large n and a_n/n, b_n/n both converge to the same limit L, then $X(n)/n \to L$ as well, even when $X(n)$ itself, here an integer defined by a largest-satisfying condition, has no closed form and its distance to a_n or

b_n does not itself converge to anything in particular. See Book 1, field: Real analysis, topic: squeeze arguments for discrete extrema.

The golden ratio's square and the general root squaring transform. The limiting constant $(3+\sqrt{5})/2$ is a root of $x^2-3x+1=0$, and this is exactly ϕ^2, the square of the golden ratio $\phi=(1+\sqrt{5})/2$, itself a root of the Fibonacci characteristic equation $x^2-x-1=0$. This is not a coincidence but an instance of a general transform: if r_1, r_2 are the roots of $x^2-px+q=0$, Vieta's formulas give $r_1+r_2=p$ and $r_1r_2=q$, so $r_1^2+r_2^2=(r_1+r_2)^2-2r_1r_2=p^2-2q$ and $(r_1r_2)^2=q^2$, meaning r_1^2, r_2^2 are the roots of $x^2-(p^2-2q)x+q^2=0$. With $p=1, q=-1$ this gives exactly $x^2-3x+1=0$. The recursion this generalises: a linear recursion $a_{n+1}=pa_n-qa_{n-1}$ has characteristic roots r_1, r_2, and the subsequence obtained by taking every second term satisfies the recursion with characteristic roots r_1^2, r_2^2; concretely the even indexed Fibonacci numbers satisfy $F_{2n+2}=3F_{2n}-F_{2n-2}$ (check: $F_6=8$ and $3F_4-F_2=3\cdot 3-1=8$), and the even indexed Lucas numbers obey the same recursion. The same idea extends to any power t via Newton's identities on the power sums of r_1, r_2, giving, for instance, r_1^3, r_2^3 as the roots of $x^2-(p^3-3pq)x+q^3=0$, which governs every third term of the original sequence. See Book 1, field: Sequences and recurrences, topic: characteristic root transforms under taking subsequences.

Circuit Bank

- **WHEN** two binomial coefficients are compared whose upper and lower indices are each shifted from one another by a small fixed amount **DO** form the ratio rather than the difference **BECAUSE** the factorial terms overlap almost entirely and cancel to a low degree rational function, while a difference of two such terms generally has no simplification at all
- **WHEN** a strict inequality is given between two quantities known to be positive **DO** convert it to a ratio compared against 1 and cross multiply, after explicitly confirming positivity **BECAUSE** cross multiplying an inequality is only valid once both sides are known to be positive, and stating that positivity range explicitly is what turns a heuristic step into a rigorous one
- **WHEN** the final request is a limit of $X(n)/n$ rather than an exact formula for $X(n)$ **DO** bound $X(n)$ between two quantities sharing the same leading asymptotic behaviour and invoke the sandwich theorem **BECAUSE** any tighter computation of $X(n)$ itself, such as pinning down a floor function exactly, is strictly more work than a normalised limit ever requires
- **WHEN** a limiting constant emerges as a root of a quadratic x^2-px+q **DO** check whether p, q relate to a familiar recursion's characteristic equation, in particular whether the quadratic is the "square the roots" transform of a smaller, more recognisable equation such as $x^2-x-1=0$ **BECAUSE** such constants recur across the Putnam, and recognising the family lets you cite a known closed form and known integer sequence identities instead of re-deriving decimal approximations from scratch

Backtesting

1. First useful cue: two binomial coefficients with indices offset from one another by exactly one on top and one on the bottom, compared by a strict inequality.
2. Method the expression suggested: form the ratio of the two binomial coefficients and simplify by factorial cancellation, rather than expanding a difference or a generating function.
3. Where you would hesitate, and the thought to reinforce: after reducing to a quadratic and finding its larger root $\alpha(n)$, it is tempting to declare $M(n)$ equal to the floor of $\alpha(n)$ without checking that this root actually lies inside the region $m>n+1$ where the ratio identity was derived; that check is one short paragraph and closes a real logical gap.

4. The trap the question sets: the phrase "evaluate the limit" invites solving for $M(n)$ exactly via floor functions and case work on whether $\alpha(n)$ is ever an integer, far more work than a normalised limit requires, when a two sided sandwich of $\alpha(n)$ is already exact and sufficient.
5. What makes you recognise it faster next time: seeing binomial coefficients whose indices differ by one step on each side as an immediate ratio cue, and seeing "evaluate $\lim X(n)/n$" as a signal that bounding, not solving exactly, is the intended route.
6. The reusable decision rule: WHEN two binomial coefficients with nearby indices are compared and only an asymptotic constant is required, DO reduce to a ratio, then a quadratic, then sandwich the extremal integer between the quadratic's larger root and that root minus one, BECAUSE this produces the exact limiting constant without ever computing the discrete extremal value itself.

Mirror drills

1. Let $N(n)$ be the largest integer m with $\binom{m}{n} > \binom{m-1}{n+1}$. Find $\lim N(n)/n$. *Answer:* the same cancellation gives $\binom{m}{n}/\binom{m-1}{n+1} = m(n+1)/[(m-n)(m-n-1)]$, leading to $m^2 - (3n+2)m + (n^2+n) < 0$; dividing by n^2 and letting $n \to \infty$, the lower order shifts vanish and the limit is again $(3+\sqrt{5})/2$.
2. Explain why the sandwich argument for $\lim M(n)/n$ succeeds even though the gap $\alpha(n) - M(n)$ need not converge to any particular value as n varies. *Answer:* a bounded (indeed $O(1)$) discrepancy divided by n tends to 0 regardless of whether the discrepancy itself converges, oscillates, or merely stays bounded, so boundedness alone is all the sandwich theorem for a normalised limit ever needs.
3. A quadratic $x^2 - px + q = 0$ has roots r_1, r_2. Find, in terms of p, q, the quadratic satisfied by r_1^3, r_2^3. *Answer:* using $r_1^3 + r_2^3 = (r_1+r_2)^3 - 3r_1r_2(r_1+r_2) = p^3 - 3pq$ and $(r_1r_2)^3 = q^3$, the equation is $x^2 - (p^3 - 3pq)x + q^3 = 0$.

10.3 A3. cyclic functional equation and its arctan integral

Suppose that f is a function from $\mathbb{R}$ to $\mathbb{R}$ such that

$$f(x) + f\left(1 - \frac{1}{x}\right) = \arctan x$$

for all real $x \neq 0$. (As usual, $y = \arctan x$ means $-\pi/2 < y < \pi/2$ and $\tan y = x$.) Find

$$\int_0^1 f(x)\,dx.$$

Read it out loud

There is exactly one structural device in this equation, the substitution $x \mapsto 1 - 1/x$, and it is reused, not varied, so the first thing to do with it is compose it with itself and see where it goes. A functional equation built from a single repeated substitution is secretly a linear system in the values of f at the points of that substitution's orbit, and you cannot know how to solve the system until you know how large the orbit is. Writing $\sigma(x) = 1 - 1/x$ and computing $\sigma(\sigma(x))$ and $\sigma(\sigma(\sigma(x)))$ by hand, within the first thirty seconds, is the single most useful thing you can do with this problem.

The second cue sits on the right hand side: $\arctan x$. Inverse tangent has a standard reflection property linking $\arctan t$ and $\arctan(1/t)$, and whenever a problem's functional equation pairs a variable with a reciprocal-flavoured transform of itself and the right hand side is arctan, that reflection identity is very likely the mechanism that turns several transcendental terms into a single constant. Fields in play: functional equations and the group structure of the substitution (algebra), inverse trigonometry (a short calculus lemma), and a definite integral at the end (calculus again).

The third cue is the target itself: not "find f", but $\int_0^1 f(x)\,dx$, over an interval that is symmetric under $x \mapsto 1 - x$. Whenever an integral is requested over an interval invariant under some involution, and the integrand is only known through a functional equation, the fast route is almost never to solve for f pointwise first; it is to compute $f(x) + f(1-x)$, since the substitution $u = 1 - x$ shows $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$ automatically, so only that sum, not f itself, is ever required.

Approaches, ranked worst to best

1. **solve for f explicitly everywhere, then integrate.** *Cue:* the three term system below determines $2f(x)$ completely, so it is tempting to push through to a genuine closed form for $f(x)$ for every real x. *Circuit:* full inversion of the functional equation before specialising. (this works, since the official remark notes it can be done "with some effort", but it forces you to track the sign of x, the sign of $(x-1)/x$, and the sign of $1/(1-x)$ as three separate case splits inside the reflection identity, merge them into one piecewise formula for f alone, and only then integrate termwise. It reaches the same number after considerably more ink and considerably more room for a sign error.)

2. **substitute the integral representation of arctan and change variables in the resulting double integral.** *Cue:* $\arctan x$ has the integral representation $\int_0^x \frac{dt}{1+t^2}$, and integrating both sides of the functional equation over $[0,1]$ turns the right hand side into a double integral. *Circuit:* representation plus a two dimensional substitution. (a genuine alternate route, credited to an AoPS contributor in the official solutions, but it trades the finite linear system you already have for a new piece of machinery, a change of variables inside a double integral, that has to be set up and justified from scratch. It is not obviously less ink, only differently shaped ink.)

3. **use the substitution's finite order to build a three term linear system, then pair with the interval's own symmetry.** *Cue:* the substitution has order 3, and the target interval $[0,1]$ is itself symmetric under $x \mapsto 1 - x$. *Circuit:* cyclic elimination followed by a reflection identity

and a symmetric pairing. (this needs only linear algebra on three equations, one short calculus lemma proved by differentiating a sum of two arctans, and one substitution to exploit the interval's symmetry. Every step is forced by the structure already present in the statement, and the final integral is the integral of a constant.)

The ideal solution

Write $\sigma(x) = 1 - \frac{1}{x} = \frac{x-1}{x}$. Composing,

$$\sigma(\sigma(x)) = 1 - \frac{1}{\sigma(x)} = 1 - \frac{x}{x-1} = \frac{(x-1)-x}{x-1} = \frac{-1}{x-1} = \frac{1}{1-x}, \qquad \sigma(\sigma(\sigma(x))) = 1-(1-x) = x,$$

so σ has order exactly 3: the three points x, $\sigma(x) = \frac{x-1}{x}$, $\sigma^2(x) = \frac{1}{1-x}$ cycle back to x. For $x \neq 0, 1$ (which keeps x, $\sigma(x)$, $\sigma^2(x)$ all defined and all nonzero, so the given equation may be applied at each), writing the hypothesis at x, at $\sigma(x)$, and at $\sigma^2(x)$ gives

$$\begin{aligned} f(x) + f(\sigma(x)) &= \arctan x, \\ f(\sigma(x)) + f(\sigma^2(x)) &= \arctan\left(\frac{x-1}{x}\right), \\ f(\sigma^2(x)) + f(x) &= \arctan\left(\frac{1}{1-x}\right). \end{aligned}$$

Adding the first and third equations and subtracting the second cancels $f(\sigma(x))$ and $f(\sigma^2(x))$ completely, since each appears once with a plus sign and once with a minus sign, while $f(x)$ appears twice with a plus sign, leaving

$$2f(x) = \arctan x - \arctan\left(\frac{x-1}{x}\right) + \arctan\left(\frac{1}{1-x}\right), \qquad x \neq 0, 1.$$

Before combining terms, record the following short lemma, since it is the only fact doing any real work in what follows. For $t \neq 0$, let $g(t) = \arctan t + \arctan(1/t)$. On $(0, \infty)$, $g'(t) = \frac{1}{1+t^2} + \frac{-1/t^2}{1+1/t^2} = \frac{1}{1+t^2} - \frac{1}{t^2+1} = 0$, so g is constant there, and $g(1) = \arctan 1 + \arctan 1 = \pi/2$; the identical computation on $(-\infty, 0)$ shows g is constant there too, and $g(-1) = -\pi/4 - \pi/4 = -\pi/2$. So

$$\arctan t + \arctan \frac{1}{t} = \begin{cases} \pi/2 & t > 0 \\ -\pi/2 & t < 0 \end{cases} \qquad (t \neq 0).$$

Now substitute $x \to 1 - x$ in the boxed formula for $2f(x)$ (valid since $1 - x \neq 0, 1$ exactly when $x \neq 1, 0$, the same domain):

$$2f(1-x) = \arctan(1-x) - \arctan\left(\frac{-x}{1-x}\right) + \arctan\left(\frac{1}{x}\right).$$

Add this to the formula for $2f(x)$ and regroup into three reciprocal pairs, using $\frac{-x}{1-x} = \frac{x}{x-1}$, the reciprocal of $\frac{x-1}{x}$:

$$2\big[f(x)+f(1-x)\big] = \left[\arctan x + \arctan\frac{1}{x}\right] + \left[\arctan(1-x) + \arctan\frac{1}{1-x}\right] - \left[\arctan\left(\frac{x-1}{x}\right) + \arctan\left(\frac{x}{x-1}\right)\right]$$

For $x \in (0, 1)$: the first bracket has $t = x > 0$, giving $\pi/2$; the second has $t = 1 - x > 0$, giving $\pi/2$; the third has $t = \frac{x-1}{x} < 0$ (numerator negative, denominator positive), giving $-\pi/2$. So

$$2\big[f(x)+f(1-x)\big] = \frac{\pi}{2} + \frac{\pi}{2} - \left(-\frac{\pi}{2}\right) = \frac{3\pi}{2}, \qquad \text{that is,} \qquad f(x)+f(1-x) = \frac{3\pi}{4} \quad \text{for } x \in (0,1).$$

Finally, the substitution $u = 1 - x$ gives $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$, so

$$2\int_0^1 f(x)\,dx = \int_0^1 \big[f(x) + f(1-x)\big]\,dx = \int_0^1 \frac{3\pi}{4}\,dx = \frac{3\pi}{4},$$

and therefore

$$\int_0^1 f(x)\,dx = \boxed{\frac{3\pi}{8}}.$$

(The problem's hypothesis presumes such an f exists; nothing above needs that presumption verified, since the question only asks for the value the integral is forced to take given the constraint, not for a construction of f itself.)

Layer 1: WHAT

Block 1: identify $\sigma(x) = 1 - 1/x$ as the only structural device and compose it with itself to discover it has order 3. Block 2: write the hypothesis at x, $\sigma(x)$, $\sigma^2(x)$ and take the alternating sum that cancels the two unwanted values, isolating $2f(x)$ as a sum of three arctan terms. Block 3: prove the short reflection lemma $\arctan t + \arctan(1/t) = \pm\pi/2$. Block 4: instead of simplifying $2f(x)$ alone, form $2f(1-x)$ by the same formula and add the two, so the six arctan terms regroup into three reciprocal pairs. Block 5: use the symmetry of $[0,1]$ under $x \mapsto 1-x$ to convert the target integral into the integral of the now-constant sum $f(x) + f(1-x)$.

The nominal route stops after Block 2 and tries to resolve the three individual arctan terms in $2f(x)$ into an explicit closed form for f alone, tracking the sign of x, of $(x-1)/x$, and of $1/(1-x)$ as three separate, only loosely related case splits, then integrating that explicit piecewise formula term by term. The bifurcation point is the decision, immediately after obtaining the formula for $2f(x)$, not to simplify it in isolation but to write down $2f(1-x)$ and add: that one choice regroups the six arctan terms into exactly the three reciprocal pairs the lemma was built for, and it is the reason the whole computation ends at a constant rather than at a genuine piecewise function.

Layer 2: HOW

Block 1. **Cue.** The equation reuses a single substitution, $x \mapsto 1 - 1/x$, rather than a family of different ones. **Why natural.** A repeatedly applied substitution generates an orbit under composition, and a functional equation is only as easy to solve as that orbit is small; composing the substitution with itself is the cheapest possible way to discover the orbit's size before committing to any algebra. **Link to the next block.** Composing twice gives $1/(1-x)$ and composing three times returns x, so the orbit has exactly three points, and the equation can be written down at all three.

Block 2. **Cue.** Three points in a cycle, each related to the next by the same substitution. **Why natural.** Writing the hypothesis at each of the three points produces three linear equations in the three unknowns $f(x)$, $f(\sigma x)$, $f(\sigma^2 x)$; because the cycle length is odd, each of the two unwanted unknowns appears in exactly two of the three equations, so giving those equations alternating signs cancels both unwanted unknowns while the target unknown, appearing in two equations with the same sign, survives doubled. **Link to the next block.** The result is a formula for $2f(x)$ as a sum of three arctan terms with arguments x, $(x-1)/x$, and $1/(1-x)$, none of which is obviously simple on its own.

Block 3. **Cue.** Two arctan terms with reciprocal arguments would collapse instantly if such an identity existed; here it does. **Why natural.** Differentiating $\arctan t + \arctan(1/t)$ gives 0 away from $t = 0$, so the sum is constant on each half-line, and evaluating at $t = \pm 1$ pins the two constants at $\pm\pi/2$; this is exactly the shape needed to simplify a sum of arctans with reciprocal arguments, but

the three arguments in the formula for $2f(x)$ do not pair up with each other as reciprocals. **Link to the next block.** Since the terms will not pair up within $2f(x)$ alone, the next move is to bring in a second copy of the same formula, evaluated at a related point, so that cross pairing becomes possible.

Block 4. Cue. The eventual target is $\int_0^1 f(x)\,dx$, an integral over an interval symmetric under $x \mapsto 1-x$. **Why natural.** Substituting $x \to 1-x$ into the formula for $2f(x)$ produces $2f(1-x)$, and adding the two formulas together regroups the six arctan terms into exactly three reciprocal pairs, $(x, 1/x)$, $(1-x, 1/(1-x))$, and $((x-1)/x,\ x/(x-1))$, each ready for the lemma from Block 3. **Link to the next block.** Applying the lemma to each pair, with the sign of each argument checked for $x \in (0,1)$, collapses the whole right hand side to the single constant $3\pi/2$.

Block 5. Cue. $f(x) + f(1-x)$ is now a known constant, and the requested integral runs exactly over $[0,1]$. **Why natural.** The substitution $u = 1-x$ shows $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$, so their sum integrates to twice the target integral, and since the sum is the constant $3\pi/4$, the target integral is immediate. **Link to the next block.** No further work is needed; the answer follows by dividing by two.

Layer 3: WHY

The truth underneath the whole solution is that a functional equation built from one substitution repeated is a statement about the orbit of that substitution, and an orbit of odd length is always solvable for a single value by an alternating sum, because each unwanted term necessarily appears an even number of times (here twice) across the equations written around the cycle, letting a plus and minus assignment cancel it exactly. The second truth is that $\arctan t + \arctan(1/t)$ is piecewise constant, a fact special to inverse tangent among the inverse trigonometric functions in this clean a form, and it is this bounded, sign-dependent constancy, not anything about the substitution, that turns a sum of transcendental terms into a rational multiple of π. The discernment mechanism is to separate these two truths in your head: the algebra of the substitution's orbit tells you how many equations you need and how to combine them, while the specific right hand side (arctan here) tells you what identity will finish the job once the combination is found, and conflating the two, trying to find the identity before finding the orbit, or trying to guess the orbit's size from the right hand side, wastes time on both.

Training this discernment means practising the reflex of composing a functional equation's substitution with itself before reading the right hand side at all, since the orbit size is a purely algebraic fact about σ, independent of what the equation equals. The examiners nudge you toward the paired-integral shortcut specifically by asking for $\int_0^1 f(x)\,dx$ rather than a general antiderivative or a value at a single point: an interval matching a natural symmetry of the substitution's territory ($[0,1]$ is exactly the fixed set of the induced involution $x \mapsto 1-x$, itself expressible via σ since σ^2 fixes the same reflection structure) is the standing invitation to pair rather than solve pointwise.

Revolving concepts

Cyclic substitutions and alternating elimination. If a substitution σ satisfies $\sigma^k = \mathrm{id}$, writing $f(x) + f(\sigma x) = c(x)$ at $x, \sigma(x), \ldots, \sigma^{k-1}(x)$ gives k linear equations in k unknowns. When k is odd, the alternating sum around the whole cycle isolates $2f(x)$ exactly as it did here, because every other unknown occurs an even number of times with cancelling signs. When k is even, the same alternating sum instead telescopes to zero, giving only a consistency condition rather than a value, so individual values of f generally cannot be recovered this way without extra information; this is why the odd order of σ in this problem is not incidental but load bearing. See Book 1, field: Functional equations, topic: solving cyclic functional equations by elimination.

The order of a Mobius substitution. A map $x \mapsto \frac{ax+b}{cx+d}$ with $ad - bc \neq 0$ composes according to 2×2 matrix multiplication up to an overall scalar, so its order as a function equals the order of the associated matrix in the projective general linear group. Here $\sigma(x) = \frac{x-1}{x}$ corresponds to a matrix of trace 1 and determinant 1, whose characteristic equation $\lambda^2 - \lambda + 1 = 0$ has roots that are primitive sixth roots of unity, which is the structural reason σ has order 3 rather than some other value; the six functions x, $1/x$, $1 - x$, $1/(1 - x)$, $(x - 1)/x$, $x/(x - 1)$ together form the full order six group generated by $x \mapsto 1/x$ and $x \mapsto 1 - x$ (the group permuting a cross ratio under its six standard values), and the three used in this problem are its unique subgroup of order three. See Book 1, field: Geometry, topic: Mobius transformations and cross ratio symmetry.

The arctan reflection identity. For $t \neq 0$, $\arctan t + \arctan(1/t)$ equals $\pi/2$ for $t > 0$ and $-\pi/2$ for $t < 0$, proved by showing the derivative vanishes on each half line and evaluating at $t = \pm 1$. It is the $s = -1/t$ special case of the general subtraction formula $\arctan t - \arctan s = \arctan \frac{t-s}{1+ts}$ read modulo π, and it is the standard mechanism whenever a problem pairs a variable with its negative reciprocal inside two arctan terms. See Book 1, field: Real analysis, topic: inverse trigonometric identities via differentiation.

Pairing $f(x)$ with $f(a - x)$ on a symmetric interval. For any integral $\int_0^a f(x)\,dx$, the substitution $u = a - x$ shows it equals $\int_0^a f(a - x)\,dx$, so twice the integral equals $\int_0^a [f(x) + f(a - x)]\,dx$; whenever that sum simplifies, the original integral is found without ever isolating f pointwise. See Book 1, field: Calculus and integrals, topic: symmetry substitutions in definite integrals.

Circuit Bank

- **WHEN** a functional equation reuses a single substitution rather than a family of different ones **DO** compose the substitution with itself to find its order before doing any further algebra **BECAUSE** the order tells you exactly how many equations and unknowns the problem reduces to, and whether a clean alternating sum exists, which happens precisely when the order is odd
- **WHEN** two arctan (or other inverse trigonometric) terms appear with reciprocal or negative reciprocal arguments **DO** apply the reflection identity $\arctan t + \arctan(1/t) = \pm\pi/2$, fixing the sign from the sign of t **BECAUSE** this is the standard mechanism collapsing a sum of two inverse trig terms to a constant, and the sign of t is the only case work it requires
- **WHEN** a definite integral of an unknown function defined by a functional equation is requested over an interval symmetric under some involution **DO** pair $f(x)$ with f of the involution's image of x and integrate their sum, rather than solving for f pointwise first **BECAUSE** the substitution defining the involution shows both orderings of the integral are equal, so only the sum needs to be known, and that sum is often forced to a constant even when f itself is not
- **WHEN** a problem only asks for a definite integral of a function pinned down by a functional equation **DO** resist finding a full closed form for f before integrating **BECAUSE** the target combination, an orbit sum or an involution pair sum, is frequently far simpler than f itself, and finding f completely solves a harder problem than the one asked

Backtesting

1. First useful cue: a single substitution, $1 - 1/x$, reused inside the equation, feeding a function whose right hand side is arctan.
2. Method the expression suggested: compose the substitution with itself to find its order, then write the given equation at every point of the resulting orbit.

3. Where you would hesitate, and the thought to reinforce: after finding the order is three and writing three equations, the right sign pattern for the alternating sum is not automatic; the reinforcing thought is that each unwanted unknown occurs in exactly two of the three equations, so opposite signs on those two equations cancel it while the target unknown, occurring in two equations with the same natural sign, survives doubled.
4. The trap the question sets: it looks as though f must be found explicitly, sign case by sign case, before the integral can be evaluated, when pairing $f(x)$ with $f(1-x)$ using the interval's own symmetry bypasses ever finding f pointwise.
5. What makes you recognise it faster next time: seeing a request for $\int_0^1 f(x)\,dx$ specifically, an interval symmetric under $x \mapsto 1-x$, as an immediate invitation to compute $f(x)+f(1-x)$ rather than $f(x)$ alone.
6. The reusable decision rule: WHEN a functional equation's substitution has finite odd order k, DO write the equation around the whole cycle and take the alternating sum to isolate $2f(x)$, and WHEN the actual target is a symmetric definite integral, DO additionally pair that result with its image under the interval's own symmetry before integrating, BECAUSE the pairing frequently collapses transcendental terms to a constant via a standard reflection identity that the isolated formula alone does not display.

Mirror drills

1. Verify directly that composing $\sigma(x) = 1 - 1/x$ with itself three times returns x, and write the two other nontrivial maps in the cycle in simplest form. *Answer:* $\sigma^2(x) = \dfrac{1}{1-x}$ and $\sigma^3(x) = x$; the cycle is x, $\dfrac{x-1}{x}$, $\dfrac{1}{1-x}$.
2. Suppose $g(x) + g\left(\dfrac{1}{1-x}\right) = x$ for all $x \neq 0, 1$. Using the same cyclic elimination (note $\tau(x) = 1/(1-x)$ is σ^2 from the main solution, so it also has order 3), find the analogous formula for $2g(x)$. *Answer:* $2g(x) = x - \dfrac{1}{1-x} + \left(1 - \dfrac{1}{x}\right)$, obtained by adding the equations at x and at $\tau^2(x) = \sigma(x)$ and subtracting the equation at $\tau(x)$.
3. If the right hand side of the original equation were an even function $E(x) = E(-x)$ instead of $\arctan x$, would $f(x) + f(1-x)$ still be forced to a constant on $(0,1)$ by the same method. *Answer:* not in general; the collapse relied specifically on $\arctan t + \arctan(1/t)$ being piecewise constant, a special reflection property of arctan, and a generic even function satisfies no such reciprocal-sum identity, so the paired sum would typically remain a genuine, non-constant function of x.

10.4 A4. Minimum tile count for an odd rectangle

Consider a $(2m-1) \times (2n-1)$ rectangular region, where m and n are integers such that $m, n \geq 4$. This region is to be tiled using tiles of the two types shown:



(the dotted lines of the printed figure divide the tiles into 1×1 squares: the first tile is a 2×2 square with one unit cell deleted, the second is a 2×3 rectangle with two opposite corner cells deleted). The tiles may be rotated and reflected, as long as their sides are parallel to the sides of the rectangular region. They must all fit within the region, and they must cover it completely without overlapping.

What is the minimum number of tiles required to tile the region?

Read it out loud

This is a tiling question with two polyominoes: an L tromino of area 3 and an S tetromino of area 4, on a board whose two side lengths are both odd, written for you as $2m-1$ and $2n-1$. It asks for a minimum, so it is an extremal question, which means two separate jobs: a lower bound that no tiling can beat, and one explicit tiling that meets it. Neither half alone scores.

The fields wired into it are combinatorial tiling, colouring and weighting invariants, and constructive induction on a board. There is no algebra here and no geometry beyond incidence of unit cells. What is really being tested is whether you know that the minimum tile count of a tiling is bounded below by a sparse witness set, and whether you can then run a routine strip decomposition to build the matching construction. It is a two-sided extremal problem dressed up as a puzzle.

The first cue is the parametrisation itself. The examiners did not write "an odd by odd rectangle"; they wrote $(2m-1) \times (2n-1)$, and they did not write the answer in terms of side lengths. They have handed you m and n and are waiting for an answer in m and n. The set of cells with both coordinates odd has exactly m choices of column and n choices of row, so it has exactly mn points. Within thirty seconds you should have written down "odd cells: mn of them" and the guess "answer mn", because that is the only clean function of m and n anywhere near the area count.

The second cue is the total area, $(2m-1)(2n-1) = 4mn - 2m - 2n + 1$. Divide by the largest tile: you get $mn - \frac{2m+2n-1}{4}$, strictly less than mn. So if the answer really is mn, the area bound is not the obstruction and you must find a combinatorial one. That single comparison is what tells you to go hunting for a colouring instead of grinding on constructions.

Approaches, ranked worst to best

1. **Area accounting only.** *Cue:* the words "minimum number of tiles" with tiles of areas 3 and 4. *Circuit:* divide the area by the largest tile and round up. This gives at least $\lceil (4mn - 2m - 2n + 1)/4 \rceil = mn - \lfloor (2m + 2n - 1)/4 \rfloor$ tiles, which for $m = n = 4$ reads 13. It is true and useless: you would then burn the rest of the session trying to build a 13 tile or 14 tile covering of the 7×7 square that does not exist. Worse than useless, in fact, because it makes you confident in a wrong target.

2. **Checkerboard colouring.** *Cue:* it is a tiling problem, so colour it. *Circuit:* count the imbalance between the two colours. The tromino covers two of one colour and one of the other, the tetromino covers two of each. The board has one more black cell than white. You get an inequality with no bite, because both tiles are nearly balanced. This is the reflex answer and it fails; knowing why it fails is what points you at the right colouring.

3. **Strip induction for a construction, with the bound assumed.** *Cue:* the side lengths $2m-1$ and $2n-1$ peel naturally into strips of width 2. *Circuit:* reduce a large board to a fixed base block plus strips. This half is genuinely necessary and you will use it, but on its own it proves only that mn tiles suffice. A Putnam grader gives you at most half the marks for one side of an extremal problem.

4. **Sparse witness colouring plus strip construction.** *Cue:* the tromino fits inside a 2×2 box, and the board is given as $2m-1$ by $2n-1$. *Circuit:* find the largest set of cells that no single tile can cover twice, then match it with a construction. Best. The lower bound is three lines, the construction is a decomposition into a 7×7 base case and strips of width 2, and the two meet exactly at mn. Every step is finite and checkable in the arena.

The ideal solution

The answer is $\boxed{mn}$.

Index the cells of the region by (i, j) with $1 \le i \le 2m-1$ the column and $1 \le j \le 2n-1$ the row. Call the two tiles the *tromino* (a 2×2 block of cells with one cell deleted) and the *tetromino* (a 2×3 block of cells with two opposite corner cells deleted). Call a cell *red* if i and j are both odd.

Lemma 10.4.1. *There are exactly mn red cells, and every legal placement of either tile covers at most one red cell. Consequently every tiling of the region uses at least mn tiles.*

Proof. The odd values of i in $\{1, \dots, 2m-1\}$ are $1, 3, \dots, 2m-1$, so there are m of them, and likewise n odd values of j; hence mn red cells.

A tromino occupies three of the four cells of some 2×2 block, with columns $\{i, i+1\}$ and rows $\{j, j+1\}$. Exactly one of $i, i+1$ is odd and exactly one of $j, j+1$ is odd, so that block contains exactly one red cell, and the tromino covers at most one.

A tetromino in horizontal orientation occupies columns $\{i, i+1, i+2\}$ and rows $\{j, j+1\}$, taking two horizontally adjacent cells in one row and the two horizontally adjacent cells shifted by one column in the other row. Exactly one of the rows $j, j+1$ is odd, so every red cell it covers lies in that one row; in that row it occupies exactly two consecutive columns, exactly one of which is odd. So it covers exactly one red cell. The vertical orientation is the same statement with rows and columns interchanged.

Since each tile covers at most one red cell and all mn red cells must be covered, a tiling has at least mn tiles. □

Lemma 10.4.2. *For every integer $k \ge 2$, a $2 \times (2k-1)$ rectangle can be tiled by exactly k tiles.*

Proof. Take the two rows to be a top row and a bottom row, and number the columns $1, \dots, 2k-1$.

For $k = 2$ the rectangle is 2×3 and two trominoes tile it: one covering the whole of column 1 together with the bottom cell of column 2, the other covering the top cell of column 2 together with the whole of column 3.

For $k \ge 3$, place a tromino on the whole of column 1 together with the bottom cell of column 2. Then for $t = 1, 2, \dots, k-2$ place a tetromino on the top cells of columns $2t$ and $2t+1$ together with the bottom cells of columns $2t+1$ and $2t+2$; this is a legal placement since $2t+2 \le 2k-2 < 2k-1$. An immediate induction on t shows that after the tth tetromino, columns 1 through $2t+1$ are completely covered and exactly the bottom cell of column $2t+2$ is covered beyond them. Taking $t = k-2$: columns 1 through $2k-3$ are full and the bottom cell of column $2k-2$ is covered. The three cells that remain are the top cell of column $2k-2$ and both cells of column $2k-1$, which form a tromino. The count is $1 + (k-2) + 1 = k$. □

Lemma 10.4.3. *The 7×7 square can be tiled by 16 tiles.*

Proof. Here is such a tiling, written as a 7×7 array in which the cells carrying the same label form one tile (rows listed from top to bottom, columns from left to right):

1	1	3	3	5	6	6
2	1	4	3	5	5	6
2	2	4	4	7	8	8
13	14	14	4	7	7	8
13	13	14	9	9	11	11
15	16	16	10	9	12	11
15	15	16	10	10	12	12

Tile 4 occupies the cells in rows 2,3,3,4 and columns 3,3,4,4, that is, a 3×2 block with two opposite corners removed, so it is a tetromino in vertical orientation. Each of the other fifteen labels occupies three cells of a 2×2 block, so each is a tromino. The labels exhaust all 49 cells with no repetition, and $15 \cdot 3 + 4 = 49$ confirms the areas. □

To see where that array comes from rather than merely that it works: delete the centre cell of the 7×7 square and the remaining 48 cells split into four 3×4 rectangles arranged as a pinwheel, namely rows 1 to 3 with columns 1 to 4, rows 1 to 4 with columns 5 to 7, rows 5 to 7 with columns 4 to 7, and rows 4 to 7 with columns 1 to 3. Each 3×4 rectangle splits into two 2×3 rectangles and each 2×3 rectangle is two trominoes, which would give 16 trominoes and area 48. The one cell left over is the centre, so exactly one of those 2×3 rectangles must be enlarged by the centre cell into a 7 cell region, and 7 cells is one tromino plus one tetromino. The array above does this with the 2×3 block in rows 1 to 3, columns 3 to 4.

Now assemble. Let $m, n \geq 4$ and take the region to have $2m - 1$ columns and $2n - 1$ rows. Split it into

1. the 7×7 square formed by columns 1 to 7 and rows 1 to 7;
2. for each $s = 1, \ldots, m - 4$, the block of columns $\{6 + 2s, 7 + 2s\}$ and rows 1 to 7, a 2×7 rectangle;
3. for each $t = 1, \ldots, n - 4$, the block of all $2m - 1$ columns and rows $\{6 + 2t, 7 + 2t\}$, a $(2m - 1) \times 2$ rectangle.

These pieces are pairwise disjoint and their union is the whole region, because $7 + 2(m - 4) = 2m - 1$ and $7 + 2(n - 4) = 2n - 1$. By Lemma 2 with $k = 4$, each 2×7 rectangle takes 4 tiles; by Lemma 2 with $k = m$, each $(2m - 1) \times 2$ rectangle takes m tiles; by Lemma 3 the square takes 16. The total is

$$16 + 4(m - 4) + m(n - 4) = 16 + 4m - 16 + mn - 4m = mn.$$

So mn tiles suffice, and by Lemma 1 no tiling uses fewer. The minimum is mn.

Remark 10.4.4. The count of each tile type in a minimum tiling is forced. If a tiling uses a trominoes and b tetrominoes then $3a + 4b = 4mn - 2m - 2n + 1$, and $a + b = mn$ gives $b = mn - 2m - 2n + 1 = (m - 2)(n - 2) - 3$. For $m, n \geq 4$ this is at least 1, so every minimum tiling must contain at least one tetromino; at $m = n = 4$ it equals 1, which is why the 7×7 base case uses exactly one.

Layer 1: WHAT

Block 1. Name the tiles: L tromino, area 3, bounding box 2×2; S tetromino, area 4, bounding box 2×3.

Block 2. Compute the area $4mn - 2m - 2n + 1$ and the area bound $mn - \frac{2m+2n-1}{4}$. Note that the natural guess mn exceeds it, so the binding constraint is not area.

Block 3. Ask the only question that produces a better bound: what is the largest set of cells that no single tile can cover twice? Because the tromino lives in a 2×2 box, such a set may not contain

two cells of one 2×2 box, so its cells must differ by at least 2 in one coordinate. The densest such set is the cells with both coordinates odd.

Block 4. Verify the tetromino also covers at most one of them. Count them: mn. Lower bound complete.

Block 5. Build. Strips of width 2 are the natural unit because $2m - 1 = 7 + 2(m - 4)$ and $2n - 1 = 7 + 2(n - 4)$: peel strips until only a 7×7 square is left.

Block 6. Tile a $2 \times (2k - 1)$ strip with k tiles (tromino, then tetrominoes climbing the staircase, then tromino), and tile the 7×7 with 16 by pinwheeling four 3×4 rectangles around the centre cell and paying for the centre cell with the single tetromino.

Block 7. Add: $16 + 4(m - 4) + m(n - 4) = mn$. Bound and construction meet, so the minimum is mn.

The nominal route goes: compute the area, divide by 4, start drawing tilings of the 7×7 square on squared paper, discover after twenty minutes that 15 tiles seems to be the floor, fail to prove it, and hand in a construction with a hand wave. The ideal route never draws a single tiling until the lower bound is in hand.

The bifurcation point is the line immediately after the area computation. Zombie: "the bound is about $mn - 2$, so let me try to hit it". Ideal: "the bound is not an integer of the right shape, so area is the wrong invariant; which cells can a single tile not hit twice?" Everything downstream follows from that one question.

Layer 2: HOW

Block 1. Cue. Two tiles are drawn with dotted subdivision lines, and the question asks only for a count. **Why natural.** Before anything else you convert a picture into two numbers and two boxes: area 3 in a 2×2 box, area 4 in a 2×3 box. The bounding box is not decoration; for tiling bounds the bounding box is the whole content of the tile, because any invariant that is constant on 2×2 boxes automatically controls the tile. **Link to the next block.** Once you have the areas you can price the board, and pricing the board is the only cheap thing available.

Block 2. Cue. "Minimum number of tiles" with tiles of different areas. **Why natural.** The area bound is free and takes one line, so you always take it, but you always then ask whether it is tight. Here the total area $4mn - 2m - 2n + 1$ is odd while the tiles have areas 3 and 4, and the bound $mn - \frac{2m+2n-1}{4}$ has a denominator that does not belong in an answer expressed in m and n. The examiners gave you m and n; they expect an answer built from them. **Link to the next block.** A bound you do not believe is a bound you must replace, and the standard replacement in tiling is a colouring.

Block 3. Cue. The tromino fits inside a 2×2 block of cells. **Why natural.** This is the definition-level step. A tile covers at most one cell of a set R precisely when R is sparse relative to the tile's footprint. For a tile confined to a 2×2 box, "no two cells of R inside one 2×2 box" is exactly "any two cells of R differ by at least 2 in the row index or by at least 2 in the column index". The densest set with that property is the sublattice of cells whose row and column indices are both odd, and it is the densest because every 2×2 box must contain at most one point of R while the odd sublattice contains exactly one. So the colouring is derived, not guessed. **Link to the next block.** You now have a candidate witness set; before using it you must confirm the second tile obeys the same rule, since one bad tile destroys the bound.

Block 4. Cue. The tetromino has a 2×3 box, which contains two red cells, so the check is not automatic. **Why natural.** You look at where the tetromino's cells sit inside its box. Only one of its two rows is odd, and in that row it holds exactly two consecutive columns, of which exactly one

is odd. So it covers exactly one red cell, not two. That "exactly", stronger than what you need, is the sign that this colouring is the right one: the tetromino cannot even waste a red cell. **Link to the next block.** With the invariant verified, the count of red cells on a $(2m-1)\times(2n-1)$ board is immediate and gives the lower bound mn.

Block 5. **Cue.** The side lengths are $2m-1$ and $2n-1$ with $m,n\geq 4$, and $2m-1-7=2(m-4)$. **Why natural.** The hypothesis $m,n\geq 4$ is the examiner telling you the base case is 7×7 and the increment is 2. When a construction must work for all large parameters, you never build the general board directly; you build one fixed block and one repeating unit, and check that the repeating unit is cheap enough. Here the repeating unit is a strip of width 2. **Link to the next block.** Strips of width 2 and odd length are small enough to tile by hand, so the whole construction reduces to two finite checks.

Block 6. **Cue.** A $2\times(2k-1)$ strip has odd length, so it cannot be cut into 2×2 blocks; something must absorb the odd cell. **Why natural.** The tetromino is exactly a staircase step: it advances the boundary between covered and uncovered by two columns while flipping which row is ahead. So a chain of tetrominoes moves along the strip, and the two ends, where the staircase must start and stop flush, are exactly where the trominoes go. That is why the pattern is tromino, then tetrominoes, then tromino, and why the count is $1+(k-2)+1=k$, which matches the red cell count k of the strip. The strip is itself extremal. **Link to the next block.** The same accounting applied to the 7×7 base case says you need 16 tiles there, and area forces 15 trominoes and 1 tetromino.

Block 7. **Cue.** $49=15\cdot 3+4$, and $48=8\cdot 6$. **Why natural.** Since fifteen of the sixteen tiles are trominoes and two trominoes tile a 2×3 rectangle, aim to cover the board by 2×3 rectangles and see what is left over. Deleting the centre cell of an odd square leaves a region that pinwheels into four congruent rectangles, which is the standard way to handle a square with a distinguished centre. Then one 2×3 rectangle adopts the centre cell to make a 7 cell region, and $7=3+4$ tells you that region is one tromino plus one tetromino. **Link to the next block.** Adding the base case to the strips gives $16+4(m-4)+m(n-4)=mn$, which meets the lower bound, so both halves are done and the problem is closed.

Layer 3: WHY

The mathematical truth underneath the lower bound is this. Let R be any set of cells such that every legal placement of every tile covers at most one cell of R. Then in any tiling, the map sending each cell of R to the tile covering it is injective, so the number of tiles is at least $|R|$. That is the whole of the colouring method, and it is a special case of linear programming duality: the minimum number of tiles is the value of a covering programme, and any assignment of weights to cells whose sum over each legal tile placement is at most 1 gives a lower bound equal to the total weight. Colourings are the 0 and 1 weightings. When 0 and 1 weightings are too crude, fractional weightings sometimes do better; here they do not need to, because the 0 and 1 witness already meets a construction.

The discernment mechanism, the thing you must actually train, is the choice of R. You do not guess colourings. You compute the bounding boxes of the tiles, take the smallest one, and ask for the maximum set of cells meeting every translate of that box at most once. For a 2×2 box the answer is forced: the four cosets of the sublattice $2\mathbb{Z}\times 2\mathbb{Z}$ are the only candidates of maximum density, and on a board with both side lengths odd the coset "odd, odd" is strictly the largest, with mn points against $m(n-1)$, $(m-1)n$ and $(m-1)(n-1)$ for the other three. That is the entire reason the problem says $2m-1$ by $2n-1$ instead of giving numbers: odd side lengths maximise the witness set, and the answer mn is just its size. Read the parametrisation as a hint about what to count and you have solved the problem before you have started it.

The second truth is that in an extremal tiling problem the area bound and the colouring bound measure different kinds of waste. The area bound measures how much space the tiles occupy; the colouring bound measures how much of a scarce resource they consume. Whenever the tile set contains pieces of unequal area, the area bound is achieved only if the largest tile is used almost everywhere, and you can test that instantly: here $b = (m-2)(n-2) - 3$ tetrominoes in a minimum tiling against $a = mn - b$ trominoes, so the tetromino is rare and the area bound is hopeless. Do that arithmetic before choosing which bound to chase.

The examiners nudged in three places. They wrote the sides as $2m-1$ and $2n-1$ rather than as odd numbers, which is a direct instruction to count an odd sublattice. They imposed $m, n \geq 4$, which is exactly the condition making at least one tetromino necessary and making the 7×7 plus strips decomposition legal, so it names the base case for you. And they drew a tromino whose bounding box is 2×2, which is the smallest useful footprint and therefore the one that dictates the colouring. Three cues, one route.

To train this: for every tiling problem you meet, before attempting anything, write the bounding box of each tile, list the sublattices that meet each box at most once, and count them on the given board. If the largest count exceeds the area bound, you have your answer and you only need a construction. That takes ninety seconds and it decides the problem.

Revolving concepts

Polyomino and tiling. A polyomino is a finite edge connected union of unit cells; a tiling of a region is a partition of its cells into congruent copies (here under rotation and reflection) of the allowed pieces. Packing is the same with "cover at most once" instead of "exactly once"; every lower bound proved by a witness set is really a packing bound and therefore survives if you weaken tiling to packing. See Book 1, field: Combinatorics, topic: polyomino tilings.

Colouring and weighting invariants. The general statement is the one proved above: a weight function w on cells with $\sum_{c \in T} w(c) \leq 1$ for every legal placement T gives #tiles $\geq \sum_c w(c)$. The classical instances you must have installed are the mutilated chessboard (dominoes and a two colouring), the 8×8 board and the T tetromino (a four colouring), and de Bruijn's theorem that an $a \times b$ brick tiles a box only if the box is "harmonic" in the same sense. The deep generalisation, which is beyond A level and worth naming as a gap, is the Conway and Lagarias tiling group, where you replace colourings with a boundary word in a finitely presented group and get obstructions no colouring can see. See Book 1, field: Combinatorics, topic: colouring invariants for tilings.

Maximum sparse set. The witness set is the maximum independent set of the graph whose vertices are cells and whose edges join two cells lying in a common bounding box of some tile. For a 2×2 box this graph is the king graph, and its independent sets are exactly the sets with no two cells within Chebyshev distance 1, so the maximum on a $(2m-1) \times (2n-1)$ board has mn elements, realised uniquely by the odd sublattice. See Book 1, field: Discrete structures and graphs, topic: independent sets.

The recursions, generalised. Two recursions appear and both are of the "add a fixed slab, pay a fixed cost" type. The strip recursion is $T(k) = T(k-1) + 1$ with $T(2) = 2$, whose solution is $T(k) = k$: adding two columns to a strip of height 2 costs one more tile. The board recursion is $N(m,n) = N(m-1,n) + n$ and $N(m,n) = N(m,n-1) + m$ with base $N(4,4) = 16$, whose solution is $N(m,n) = mn$. The general form is $N(m,n) = N(m-1,n) + f(n)$ with f independent of m, solved by $N(m,n) = N(m_0,n) + (m-m_0)f(n)$; whenever a construction is built by appending congruent slabs, write this recursion down explicitly rather than describing the picture, because the closed form is what you compare to the lower bound. The same pattern governs every "tile a large board by peeling strips" argument. See Book 1, field: Sequences and recurrences, topic: linear recursions with constant increment.

Extremal two sidedness. Any answer to "find the minimum" is a pair (bound, example). The write-up convention is to state the answer first, prove the bound, then exhibit the example and

verify it explicitly. See Book 1, field: Combinatorics, topic: extremal arguments.

Circuit Bank

- **WHEN** a tiling question asks for the minimum number of tiles **DO** look for the largest set of cells that no single tile can cover twice, and count it **BECAUSE** each tile then consumes at most one witness cell, so the tile count is at least the witness count, and this beats the area bound whenever the tiles have unequal areas
- **WHEN** a tile family contains a piece whose bounding box is 2×2 **DO** test the four sublattices given by the parity of the two coordinates **BECAUSE** every 2×2 box contains exactly one cell of each parity class, so each class is automatically a valid witness set, and on an odd by odd board the "odd, odd" class is the biggest
- **WHEN** a board is described as $(2m-1) \times (2n-1)$ rather than by explicit numbers **DO** read it as an instruction to count something of size mn **BECAUSE** the examiner chose the parametrisation so that the answer is a clean function of m and n, and the odd sublattice is the only natural object of size mn
- **WHEN** you must construct a tiling for all large m and n **DO** fix one base block and one repeating slab of width 2, and write the slab recursion **BECAUSE** the construction then reduces to two finite checks and one arithmetic identity, instead of an unbounded family of pictures
- **WHEN** your lower bound and your construction agree **DO** stop, and state the answer **BECAUSE** extremal problems are two sided, and further optimisation of either half is wasted ink

Backtesting

1. First useful cue: the side lengths are written as $2m-1$ and $2n-1$, and the tromino's bounding box is 2×2. Together these say "count the cells with both coordinates odd", and there are exactly mn of them.
2. Method the expression suggested: a witness colouring for the lower bound, then a decomposition of the board into one 7×7 base square and strips of width 2 for the construction.
3. Where you would hesitate, and the thought to reinforce: at the tetromino check, because its 2×3 bounding box does contain two red cells and for a moment the whole colouring looks dead. The thought to reinforce is that a tile is not its bounding box; you must look at which cells of the box the tile actually occupies. Here the tetromino uses only one of the two rows that could be odd, and only two consecutive columns of that row, so it covers exactly one red cell.
4. The trap the question sets: the area bound. It is easy, it is wrong as a target, and it points you at 13 or 14 tiles for the 7×7 square, which sends you hunting for a tiling that does not exist. The second trap is stopping after the construction, which is worth about half the marks.
5. What makes you recognise it faster next time: the reflex "unequal tile areas plus a request for a minimum equals colouring, not area". Also the reflex that a hypothesis like $m, n \geq 4$ is naming your base case.
6. The reusable decision rule: for any tiling minimum, compute the tile bounding boxes first, take the maximum set of cells meeting every translate of the smallest box at most once, count it on the given board, and compare with the area bound before writing anything else.

Mirror drills

1. Show that in any tiling of the $(2m-1)\times(2n-1)$ board that uses exactly mn tiles, the number of tetrominoes is exactly $(m-2)(n-2)-3$, and deduce that at least one tetromino is used when $m,n\geq 4$. *Answer:* with a trominoes and b tetrominoes, $3a+4b=4mn-2m-2n+1$ and $a+b=mn$; subtract three times the second from the first to get $b=mn-2m-2n+1=(m-2)(n-2)-3\geq 4-3=1$.
2. Prove that no $(2m-1)\times(2n-1)$ board can be tiled by tetrominoes alone, for any $m,n\geq 1$. *Answer:* the area $(2m-1)(2n-1)$ is odd and every tetromino has area 4, so the area is not divisible by 4; in fact it is not even.
3. A $2\times(2k-1)$ strip: prove that k tiles are needed as well as sufficient. *Answer:* the strip has k cells with both coordinates odd (one row is odd, and k of the $2k-1$ columns are odd), each tile covers at most one of them, and Lemma 2 achieves k.

10.5 A5. Bounded alternating words in two generators

Suppose that G is a finite group generated by the two elements g and h, where the order of g is odd. Show that every element of G can be written in the form

$$g^{m_1}h^{n_1}g^{m_2}h^{n_2}\cdots g^{m_r}h^{n_r}$$

with $1 \leq r \leq |G|$ and $m_1, n_1, m_2, n_2, \ldots, m_r, n_r \in \{-1, 1\}$. (Here $|G|$ is the number of elements of G.)

Read it out loud

This is a finite group theory question with two separate demands hidden inside one sentence. The first demand is qualitative: every element must be reachable at all by words of this restricted shape. The second is quantitative: reachable in at most $|G|$ blocks. Read the word out loud and you hear the structure immediately, because $g^{m_1}h^{n_1}g^{m_2}h^{n_2}\cdots$ is not a word in g and h at all. It is a word in the four products

$$gh, \quad gh^{-1}, \quad g^{-1}h, \quad g^{-1}h^{-1},$$

and r is its length. That single re-reading is the whole problem: you were handed two generators with unrestricted exponents and you must work instead with an alphabet of four fixed letters and no inverses allowed.

The fields wired in are elementary finite group theory (generated subgroups, order of an element, Lagrange and index 2 subgroups) and a counting principle that is really graph reachability: how many steps does it take for a growing set to fill a finite set. There is nothing analytic and nothing about the specific structure of G; every step must work for an arbitrary finite group.

Within thirty seconds you should have written down $T = \{g^{\pm 1}h^{\pm 1}\}$ and the two questions "does T generate G?" and "why $|G|$?". You should also have noticed where the odd order hypothesis must live. It cannot be about the bound, because the bound $|G|$ is a counting statement that has nothing to do with parity. So it must be about reachability, that is, about whether T generates. And it takes one line to see how: the four letters let you form

$$(gh)(g^{-1}h)^{-1} = ghh^{-1}g = g^2,$$

so the subgroup generated by T contains g^2, and g^2 is a generator of $\langle g \rangle$ precisely when the order of g is odd. The hypothesis is not decoration; it is the hinge.

Approaches, ranked worst to best

1. **Direct word rewriting.** *Cue:* every element is some word in $g^{\pm 1}, h^{\pm 1}$, so massage that word into the required shape. *Circuit:* normal form manipulation, inserting hh^{-1} and gg^{-1} to force alternation. You can indeed pad a general word into alternating shape, but g^k for k large needs k blocks, and nothing in the process bounds r by $|G|$. It also never uses the odd order hypothesis, which is proof that it is the wrong route: a proof that does not consume the hypothesis is a proof of something false.

2. **Cayley digraph and shortest paths.** *Cue:* the phrase "with $1 \leq r \leq |G|$" is a diameter bound. *Circuit:* build the digraph on G with an edge $s \to st$ for each $t \in T$, then take a shortest path. Correct and intuitive: a shortest path cannot repeat a vertex, so it has at most $|G|$ edges. But you still have to prove the digraph is strongly connected, which is exactly the reachability half of the problem, so the graph language buys you only the bound, and it buys it at the cost of a paragraph setting up vocabulary.

3. **Shortest representation plus pigeonhole.** *Cue:* a minimum exists because lengths are positive integers. *Circuit:* take a shortest word, look at its partial products, and delete a repeated one. Clean and short. You prove reachability separately (a nonempty subset of a finite group closed under multiplication is a subgroup), then bound the length by pigeonhole on the partial products. Two arguments, both standard.
4. **Grow or stabilise.** *Cue:* the reachable sets $S_1 \subseteq S_2 \subseteq \cdots$ live inside a finite set. *Circuit:* a monotone chain in a finite set either strictly grows or is stable, and stability upgrades to invariance under the whole group. Best, because it does both jobs at once. You never need a separate reachability argument and never need minimality: each step either adds an element or ends the proof, so after $|G|$ steps you are done. It is also the argument that generalises without change to any finite state reachability problem.

The ideal solution

Write e for the identity of G and set

$$T = \{\, gh,\ gh^{-1},\ g^{-1}h,\ g^{-1}h^{-1} \,\} \subseteq G.$$

An expression $g^{m_1}h^{n_1}\cdots g^{m_r}h^{n_r}$ with all $m_i, n_i \in \{-1, 1\}$ is precisely a product of r elements of T, so the assertion to be proved is: every element of G is a product of at most $|G|$ elements of T, using at least one factor. For $r \geq 1$ let

$$S_r = \{\, t_1 t_2 \cdots t_k \ :\ 1 \leq k \leq r,\ t_1, \ldots, t_k \in T \,\},$$

so that $S_1 \subseteq S_2 \subseteq \cdots \subseteq G$ and $S_r \neq \varnothing$. The goal is $S_{|G|} = G$.

Lemma 10.5.1. *The subgroup $\langle T \rangle$ generated by T is the whole of G.*

Proof. Since $\langle T \rangle$ is a subgroup containing gh and $g^{-1}h$, it contains

$$(gh)(g^{-1}h)^{-1} = (gh)(h^{-1}g) = g^2.$$

Let d be the order of g; by hypothesis d is odd, say $d = 2e + 1$. Then

$$(g^2)^{e+1} = g^{2e+2} = g^{d+1} = g,$$

so $g \in \langle T \rangle$. Hence $h = g^{-1}(gh) \in \langle T \rangle$ as well, and therefore $\langle T \rangle \supseteq \langle g, h \rangle = G$. □

Lemma 10.5.2. *If $S_r = S_{r+1}$ for some $r \geq 1$, then $S_r = G$.*

Proof. Let $t \in T$. Appending the letter t to a product of at most r letters gives a product of at most $r + 1$ letters, so $S_r t \subseteq S_{r+1} = S_r$. Right multiplication by t is a bijection of G, so $|S_r t| = |S_r|$, and a subset of S_r of the same finite size as S_r equals S_r; thus $S_r t = S_r$, and applying t^{-1} on the right of both sides gives $S_r t^{-1} = S_r$ too.

Therefore the set $H = \{x \in G : S_r x = S_r\}$ contains T, and it is a subgroup of G because $S_r(xy) = (S_r x)y$ and $S_r x^{-1} = S_r$ whenever $S_r x = S_r$. By Lemma 1, $H \supseteq \langle T \rangle = G$, so $S_r x = S_r$ for every $x \in G$. Pick any $u \in S_r$ (possible since $gh \in S_1 \subseteq S_r$). Then $ux \in S_r$ for every $x \in G$, and $\{ux : x \in G\} = G$, so $S_r = G$. □

Proof of the required statement. We show by induction on $r \geq 1$ that $|S_r| \geq \min\{r, |G|\}$.

For $r = 1$ this says $|S_1| \geq 1$, which holds since $gh \in S_1$. Suppose it holds for some $r \geq 1$. If $S_r = G$ then $S_{r+1} = G$ and the claim for $r + 1$ is immediate. If $S_r \neq G$ then Lemma 2 forces $S_r \neq S_{r+1}$, so $S_{r+1} \supsetneq S_r$ and $|S_{r+1}| \geq |S_r| + 1 \geq \min\{r, |G|\} + 1 \geq \min\{r + 1, |G|\}$, where the last inequality uses that $|S_r| \geq |G|$ would give $S_r = G$. This completes the induction.

Taking $r = |G|$ gives $|S_{|G|}| \geq |G|$, hence $S_{|G|} = G$. So every $s \in G$ is a product of r elements of T for some r with $1 \leq r \leq |G|$, that is,

$$s = g^{m_1} h^{n_1} \cdots g^{m_r} h^{n_r}$$

for some $m_1, n_1, \ldots, m_r, n_r \in \{-1, 1\}$, as required. □

Remark 10.5.3. The hypothesis on the order of g cannot be dropped. Take $G = \mathbb{Z}/2\mathbb{Z}$ written additively and $g = h = 1$. Every allowed word is a sum of $2r$ terms each equal to ± 1, hence even, hence 0 in G; the element 1 has no representation of the required form. Here g has order 2.

Layer 1: WHAT

Block 1. Re-read the target word: it is a product of r letters drawn from the four element set $T = \{g^{\pm 1} h^{\pm 1}\}$, with no inverses of letters permitted.

Block 2. Split the task in two: reachability (is every element such a product at all?) and the length bound $r \leq |G|$.

Block 3. For reachability, compute inside $\langle T \rangle$: $(gh)(g^{-1}h)^{-1} = g^2$. So $g^2 \in \langle T \rangle$.

Block 4. Odd order upgrades g^2 to g, because $g = (g^2)^{(d+1)/2}$ when $d = \mathrm{ord}(g)$ is odd. Then $h = g^{-1}(gh)$ and so $\langle T \rangle = G$.

Block 5. For the bound, define the reachable sets S_r (products of at most r letters) and observe that the chain $S_1 \subseteq S_2 \subseteq \cdots$ lives in a finite set, so it either strictly grows or repeats.

Block 6. If it repeats at step r, then S_r is stable under right multiplication by every letter, hence by every inverse letter, hence by all of $\langle T \rangle = G$, which forces $S_r = G$.

Block 7. So each step before termination adds at least one element: $|S_r| \geq \min\{r, |G|\}$, giving $S_{|G|} = G$. Done.

The nominal route reads the word as "a word in g and h with restricted exponents" and tries to rewrite an arbitrary product $g^{a_1} h^{b_1} \cdots$ into that shape by inserting cancelling pairs. It produces pages of algebra, never controls r, and never uses the odd order hypothesis.

The bifurcation point is the very first line of work. Zombie: treat g and h as the letters and the exponents as the problem. Ideal: treat $gh, gh^{-1}, g^{-1}h, g^{-1}h^{-1}$ as the letters and ask what semigroup they generate. Everything else in the solution is forced once you make that substitution.

Layer 2: HOW

Block 1. Cue. The exponents are confined to $\{-1, 1\}$ and the pattern alternates g then h rigidly. **Why natural.** A restriction on exponents is never about exponents. It means the problem has secretly changed alphabet: the atoms are no longer g and h but the four syllables that the pattern permits. You always perform this substitution explicitly, because the restriction that looked awkward ("only ± 1") becomes the natural statement "a word of length r in a four letter alphabet". **Link to the next block.** With the alphabet named, the target statement splits cleanly into reachability and length.

Block 2. Cue. The problem asks for existence of a representation and a bound on its length in one sentence. **Why natural.** Two quantifiers, two proofs. Separating them tells you where each hypothesis lives: a bound of $|G|$ smells of counting the elements of G, while the parity condition on $\mathrm{ord}(g)$ cannot possibly affect a counting bound. So the odd order must be spent on reachability. Locating a hypothesis before using it is the fastest way to find the intended route. **Link to the next block.** Reachability means asking what the four letters generate.

Block 3. **Cue.** You have four products of $g^{\pm 1}$ and $h^{\pm 1}$ and you want to eliminate h. **Why natural.** To isolate g you must cancel the h on the right, and the only cancellation available is to multiply by the inverse of another letter with the same h exponent. There are exactly two letters ending in h, namely gh and $g^{-1}h$, so $(gh)(g^{-1}h)^{-1}$ is the only nontrivial product of that shape, and it collapses to g^2. There is no search here: the computation is the unique one available. **Link to the next block.** You have g^2 but you want g, and that is exactly where a hypothesis about the order of g is needed.

Block 4. **Cue.** "the order of g is odd". **Why natural.** The cyclic group $\langle g \rangle$ has order d, and $\langle g^k \rangle = \langle g \rangle$ exactly when $\gcd(k, d) = 1$. With $k = 2$ this is exactly the condition that d is odd. Explicitly $g = (g^2)^{(d+1)/2}$, since $g^d = e$. So the hypothesis is precisely the statement "g^2 is as good a generator as g", which is what the previous block needs. **Link to the next block.** Once $g \in \langle T \rangle$, then $h = g^{-1}(gh) \in \langle T \rangle$, and generation of G by g and h finishes reachability.

Block 5. **Cue.** The requested bound is $|G|$, the size of the group, not a power or a factorial. **Why natural.** A bound equal to the number of states is the signature of a reachability argument: you are filling a finite set one layer at a time and each layer that fails to grow ends the process. So you define S_r, the set reachable in at most r letters, and study the chain. This is the same object as the ball of radius r in the Cayley digraph of T. **Link to the next block.** A monotone chain in a finite set has only two behaviours, so you now analyse the stable case.

Block 6. **Cue.** $S_r = S_{r+1}$ means appending a letter never escapes S_r. **Why natural.** Right multiplication by a fixed element is a bijection of G, so $S_r t \subseteq S_r$ upgrades from an inclusion to an equality by counting, and equality is what lets you invert: $S_r t^{-1} = S_r$. That is the crucial move, and it is the reason finiteness is doing real work. Once S_r is right invariant under T and under T^{-1}, it is right invariant under the group they generate, which Lemma 1 says is all of G; a nonempty right G invariant set is G. **Link to the next block.** So stability implies completion, and hence growth is compulsory until completion.

Block 7. **Cue.** Each step adds at least one element until the process ends. **Why natural.** This converts the dichotomy into an inequality: $|S_r| \geq \min\{r, |G|\}$, proved by a two line induction. The bound $r \leq |G|$ falls out with no further work, and note that the hypothesis $r \geq 1$ in the problem is automatically respected because S_r was defined with $k \geq 1$, so no empty product ever appears. **Link to the next block.** Nothing remains: reachability and the bound are both proved, and the odd order hypothesis has been used exactly once, in Lemma 1.

Layer 3: WHY

Three mathematical truths carry this problem.

The first is that in a finite group, closure under multiplication is already closure under inverses. If $\varnothing \neq S \subseteq G$ satisfies $SS \subseteq S$, then for $a \in S$ the map $x \mapsto ax$ is injective on S, hence surjective on S since S is finite, so a has a left inverse inside S and $e \in S$, and repeating gives $a^{-1} \in S$. The same finiteness is what upgrades $S_r t \subseteq S_r$ to $S_r t = S_r$ in Lemma 2. Every time you see a semigroup condition inside a finite group, this upgrade is available and you should take it without comment.

The second is the grow or stabilise dichotomy. A chain $S_1 \subseteq S_2 \subseteq \cdots$ of subsets of an N element set, in which stability at any step implies completion, must be complete by step N. This is the abstract skeleton shared by breadth first search, by transitive closure algorithms, by the proof that a connected Cayley graph has diameter less than $|G|$, and by the proof that a finite state machine's reachable set is found in at most (number of states) rounds. When a problem asks for a bound equal to the size of the ambient finite set, this is almost always the intended engine.

The third is the obstruction theory that explains the odd order hypothesis, and this is the part that turns the problem from a puzzle into a piece of structure. Let $N = \langle T \rangle$. As computed, $g^2 \in N$, and also $h^2 = (gh^{-1})^{-1}(gh) \in N$. A short check shows N is normal: for $t = g^{\varepsilon}h^{\delta} \in T$ one has $gtg^{-1} = g^2(gh^{-\delta})^{-1}$ if $\varepsilon = 1$ and $gtg^{-1} = (gh^{-\delta})^{-1}$ if $\varepsilon = -1$, while $hth^{-1} = (g^{-\varepsilon}h^{-1})^{-1}h^{\delta-1}$ with $h^{\delta-1} \in \{e, h^{-2}\}$; all of these lie in N. In the quotient G/N the images satisfy $\overline{g}\,\overline{h} = \overline{e}$ and $\overline{g}^2 = \overline{e}$, so G/N is generated by $\overline{g}$ with $\overline{g}^2 = \overline{e}$ and therefore has order 1 or 2. In other words, the only possible obstruction to T generating G is a homomorphism $\varphi\colon G \to \mathbb{Z}/2\mathbb{Z}$ with $\varphi(g) = \varphi(h) = 1$, which sends every letter of T to 0. The order of $\varphi(g)$ divides the order of g; if that order is odd then $\varphi(g) = 0$ and the obstruction cannot exist. So the hypothesis is exactly, and only, the death of the mod 2 obstruction, and the counterexample $G = \mathbb{Z}/2\mathbb{Z}$ with $g = h = 1$ is the obstruction in its smallest form.

The discernment mechanism is therefore: when a hypothesis looks arbitrary, construct the counterexample it forbids, and you will have found both the place it is used and the shape of the argument. Here "odd order" forbids a parity, parity comes from index 2 subgroups, index 2 subgroups are kernels of maps to $\mathbb{Z}/2\mathbb{Z}$, and maps to $\mathbb{Z}/2\mathbb{Z}$ kill the alternating word because every letter has an even total exponent. That chain of thought takes fifteen seconds once trained and it hands you the entire solution plan.

Train it by making the substitution reflex automatic: whenever exponents in a word are constrained to a finite set, rename the permitted syllables as an alphabet and ask two questions, does the alphabet generate, and how fast does the reachable set grow. The examiners nudged you toward both: they wrote the exponent set $\{-1, 1\}$ explicitly rather than saying "nonzero", and they wrote the bound as $|G|$ rather than as some function of the order of g.

Revolving concepts

Generated subgroup and word length. For $T \subseteq G$, $\langle T \rangle$ is the smallest subgroup containing T, equivalently the set of finite products of elements of T and their inverses. If G is finite, then the set of finite products of elements of T alone (no inverses) is already a subgroup, so the two notions coincide. See Book 1, field: Abstract algebra, topic: generators and relations.

Order of an element and powers that still generate. If $\operatorname{ord}(g) = d$ then $\langle g^k \rangle = \langle g^{\gcd(k,d)} \rangle$, so $\langle g^k \rangle = \langle g \rangle$ if and only if $\gcd(k, d) = 1$. The case used here is $k = 2$: squaring loses nothing exactly when d is odd. The same fact powers many problems about squares in groups, for instance that in a group of odd order every element is a square, uniquely. See Book 1, field: Abstract algebra, topic: cyclic groups.

Cayley digraph and diameter. Given $T \subseteq G$, the Cayley digraph has vertex set G and an edge $s \to st$ for each $t \in T$. Reachability in this digraph is exactly representability by words in T, and the sets S_r are the balls of radius r about e. If the digraph is strongly connected then its diameter is at most $|G| - 1$, since a shortest path repeats no vertex. See Book 1, field: Discrete structures and graphs, topic: Cayley graphs.

Index 2 subgroups and the mod 2 obstruction. A subgroup of index 2 is normal, and quotients of order 2 correspond to surjections $G \to \mathbb{Z}/2\mathbb{Z}$, which factor through $G/[G, G]$ tensored with $\mathbb{Z}/2\mathbb{Z}$. Any expression in which every generator appears an even number of times in total lands in the kernel of every such map, which is precisely why alternating ± 1 words are trapped. See Book 1, field: Abstract algebra, topic: normal subgroups and quotients.

The recursion, generalised. The reachable sets satisfy the recursion

$$S_1 = T, \qquad S_{r+1} = S_r \cup S_r T,$$

a monotone recursion in the finite lattice of subsets of G. The general theory: any monotone map F on the subsets of an N element set, applied repeatedly from a starting set, gives an increasing chain that becomes constant after at most N steps, and the limit is the least fixed point containing the start. In this problem the extra structural input is that the only fixed points here are $\varnothing$ and G,

which turns "becomes constant" into "equals G". The same generalised recursion governs transitive closure, reachable states of an automaton, the set of positions in a combinatorial game reachable from the start, and closure under a set of moves in an invariants problem: write $S_{r+1} = S_r \cup F(S_r)$, prove that a fixed point must be everything, and read off a bound equal to the size of the ambient set. See Book 1, field: Discrete structures and graphs, topic: reachability and closure.

Circuit Bank

- **WHEN** a word has its exponents restricted to a finite set, or a rigid alternating pattern **DO** rename the permitted syllables as a new alphabet T and restate the problem as "words in T" **BECAUSE** the restriction is a change of generating set in disguise, and once renamed the question becomes the standard pair "does T generate?" and "how long are the words?"
- **WHEN** a bound equal to the size of a finite set is requested **DO** run the chain $S_{r+1} = S_r \cup S_r T$ and prove that stability forces completion **BECAUSE** each round before termination adds at least one new element, so the process finishes within the size of the set, which is exactly the requested bound
- **WHEN** a nonempty subset of a finite group is closed under multiplication **DO** call it a subgroup immediately, and if it also satisfies $St \subseteq S$ then upgrade to $St = S$ **BECAUSE** right multiplication is a bijection of the group, so an inclusion between finite sets of equal size is an equality
- **WHEN** a hypothesis in the problem looks arbitrary, for instance a parity or coprimality condition **DO** build the counterexample it forbids before you start proving anything **BECAUSE** the counterexample names the obstruction, and the obstruction names the single line of the proof where the hypothesis must be spent
- **WHEN** you need g but only have g^2 inside a subgroup **DO** use $g = (g^2)^{(d+1)/2}$ where $d = \operatorname{ord}(g)$ is odd **BECAUSE** $\langle g^k \rangle = \langle g \rangle$ exactly when $\gcd(k, \operatorname{ord}(g)) = 1$, so odd order makes squaring lossless

Backtesting

1. First useful cue: the exponent set $\{-1, 1\}$ with the rigid g then h alternation, which says the real alphabet is $T = \{gh, gh^{-1}, g^{-1}h, g^{-1}h^{-1}\}$ and r is a word length in that alphabet.
2. Method the expression suggested: prove $\langle T \rangle = G$ using the odd order, then bound the word length by the grow or stabilise argument on the reachable sets S_r.
3. Where you would hesitate, and the thought to reinforce: at the moment when you have $g^2 \in \langle T \rangle$ but want g. Reinforce that $\langle g^k \rangle = \langle g \rangle$ whenever $\gcd(k, \operatorname{ord} g) = 1$, and that the explicit exponent is $(d+1)/2$. Also hesitate at the step $S_r t \subseteq S_r \Rightarrow S_r t = S_r$; reinforce that finiteness plus injectivity is what makes this legal.
4. The trap the question sets: never using the odd order hypothesis. Any argument that "clearly you can write everything as such a word" is false, as $G = \mathbb{Z}/2\mathbb{Z}$ with $g = h = 1$ shows. The second trap is producing a representation of length 0 for the identity, which the condition $r \geq 1$ forbids; define your sets with at least one letter from the start.
5. What makes you recognise it faster next time: the pairing of a bound equal to $|G|$ with a generation question is the fingerprint of a reachability or Cayley graph argument, and a parity hypothesis on the order of a generator is the fingerprint of an index 2 obstruction.
6. The reusable decision rule: when constrained words are involved, change alphabet first, then ask generation and growth separately, and locate every hypothesis before writing a line.

Mirror drills

1. Show that the statement fails without the odd order hypothesis. *Answer:* take $G = \mathbb{Z}/2\mathbb{Z}$ additively with $g = h = 1$; each block $g^m h^n$ contributes $m + n \in \{-2, 0, 2\}$, which is 0 modulo 2, so only the identity is representable.

2. Prove the length bound instead by taking a shortest representation $s = t_1 \cdots t_r$ and applying pigeonhole to the partial products, and explain why you must use $s_0 = e, s_1, \ldots, s_{r-1}$ rather than $s_1, \ldots, s_r$. *Answer:* if $r > |G|$ two of the r partial products $s_0, \ldots, s_{r-1}$ coincide, say $s_i = s_j$ with $i < j$, and deleting $t_{i+1} \cdots t_j$ gives a representation of length $i + (r - j) \geq 1$ and less than r; using $s_1, \ldots, s_r$ instead could delete every letter and produce the forbidden empty word when $s = e$.

3. Prove that if a finite group G is generated by a set T that is closed under inverses, every element is a product of at most $|G| - 1$ elements of T. *Answer:* let B_r be the ball of radius r about e; if $B_r = B_{r+1}$ then $B_r T = B_r$ so B_r is right invariant under $\langle T \rangle = G$ and equals G; otherwise $|B_r| \geq r + 1$, so $B_{|G|-1} = G$.

10.6 A6. Sharp integral bound for cubics with a root

Find the smallest constant C such that for every real polynomial $P(x)$ of degree 3 that has a root in the interval $[0,1]$,

$$\int_0^1 |P(x)|\, dx \le C \max_{x\in[0,1]} |P(x)|\,.$$

Read it out loud

This compares two norms of the same function on the same interval: the L^1 norm $\int_0^1 |P|$ and the sup norm $\max_{[0,1]} |P|$, over the class of real cubics with a root in $[0,1]$. Both sides are homogeneous of degree 1 in P, so the class being a cone (closed under multiplication by scalars) means you may normalise: set $\max_{[0,1]} |P| = 1$ and the question becomes "how large can $\int_0^1 |P|$ be?".

The fields wired in are approximation theory (extremal problems for polynomials under a sup norm constraint), quadrature (exact integration formulas for low degree polynomials), and finite dimensional convexity (a linear functional on a compact convex set attains its maximum at an extreme point). What is really being tested is whether you know that integration against a fixed interval is a linear functional on a four dimensional space, and that such a functional can be written exactly in terms of finitely many point values.

Before touching the cubic, ask what the root is for. Without it the answer is 1: take P with $\max|P| = 1$ and almost constant on $[0,1]$, say $P(x) = 1 - \varepsilon x^3$, and the ratio tends to 1. So the root is the entire source of the inequality: it forces $|P|$ to touch zero somewhere in $[0,1]$, which wastes integral, and the question is exactly how much waste is forced. That reframing tells you the extremal P will hug the value ± 1 everywhere except near its root.

Within thirty seconds you should have written three things. First, normalise $\max|P| = 1$. Second, "the absolute value must go", which means splitting $[0,1]$ at the roots of P so that $|P|$ is a genuine polynomial on each piece. Third, and this is the cue that decides the problem, "$\int_0^1$ is a linear functional on a four dimensional space, so it equals a fixed weighted sum of a few point values": Simpson's rule with weights $\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$ at $0, \frac{1}{2}, 1$ is exact on cubics, and if the root sits at the node 0 then that whole node contributes nothing, leaving at most $\frac{4}{6} + \frac{1}{6} = \frac{5}{6}$. The answer is visible before any algebra.

Approaches, ranked worst to best

1. **Coefficient bashing with multipliers.** *Cue:* "find the smallest constant" over a family of polynomials. *Circuit:* parametrise $P = ax^3 + bx^2 + cx + d$ and optimise. You must maximise $\int_0^1 |P|$, a piecewise defined function of four parameters, subject to a root condition and to a sup norm constraint that is an infinite family of inequalities. Lagrange multipliers do not apply cleanly to max constraints, the absolute value breaks differentiability, and you will not finish. Worst by a distance.

2. **Guess the extremal polynomial and verify.** *Cue:* extremal polynomials under a sup norm constraint touch the constraint repeatedly, as with Chebyshev. *Circuit:* equioscillation. This produces $4x^3 - 8x^2 + 5x$ and the value $5/6$, so you get the lower bound $C \ge 5/6$ and the correct answer to write in a box. It proves nothing about the upper bound, and the upper bound is where the marks are.

3. **Convexity and extreme points.** *Cue:* the constraint set is convex and compact and the objective is linear. *Circuit:* maximise a linear functional at an extreme point, or exhibit a supporting hyperplane. Genuinely the discovery route: the set of cubics with $P(0) = 0$ and $0 \le P \le 1$ on $[0,1]$ is a compact convex body in a three dimensional space, and $\int_0^1$ is linear on it, so the

maximum is at an extreme point, and extreme points are pinned by three active constraints. Correct and illuminating, but classifying extreme points rigorously in a write up costs a page.

4. **Localise at the roots, then integrate exactly by quadrature.** *Cue:* a root in $[0,1]$, an integral of $|P|$, and degree 3. *Circuit:* split at the roots so $|P|$ is a polynomial on each piece, rescale each piece to $[0,1]$ so its root sits at an endpoint, then use a quadrature rule that is exact on cubics and has that endpoint as a node. Best. The upper bound becomes five lines, everything is elementary, and the equality analysis hands you the extremal polynomial without guessing.

The ideal solution

The answer is $C = \boxed{\dfrac{5}{6}}$.

Lemma 10.6.1 (Simpson's rule is exact on cubics)**.** *For every real polynomial Q of degree at most 3,*

$$\int_0^1 Q(t)\,dt = \frac{1}{6}\Big(Q(0) + 4Q\big(\tfrac{1}{2}\big) + Q(1)\Big).$$

Proof. Both sides are linear in Q, so it suffices to check them on the basis $1, t, t^2, t^3$ of the space of polynomials of degree at most 3. The left sides are $1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}$; the right sides are $\frac{1}{6}(1+4+1) = 1$, $\frac{1}{6}(0+2+1) = \frac{1}{2}$, $\frac{1}{6}(0+1+1) = \frac{1}{3}$ and $\frac{1}{6}(0+\frac{1}{2}+1) = \frac{1}{4}$. □

Lemma 10.6.2 (the normalised case)**.** *Let Q be a real polynomial of degree at most 3 with $Q(0) = 0$ and $Q(x) \geq 0$ for all $x \in [0,1]$, and put $\mu = \max_{x\in[0,1]} Q(x)$. Then $\int_0^1 Q(x)\,dx \leq \frac{5}{6}\mu$.*

Proof. By Lemma 1 and $Q(0) = 0$,

$$\int_0^1 Q(x)\,dx = \frac{1}{6}\Big(4Q\big(\tfrac{1}{2}\big) + Q(1)\Big) \leq \frac{1}{6}(4\mu + \mu) = \frac{5}{6}\mu.$$ □

Proposition 10.6.3. *If P is a real polynomial of degree 3 with a root in $[0,1]$, then $\int_0^1 |P(x)|\,dx \leq \frac{5}{6}M$ where $M = \max_{x\in[0,1]} |P(x)|$.*

Proof. List the distinct roots of P lying in $[0,1]$ as $r_1 < r_2 < \cdots < r_j$ with $j \geq 1$, and set $r_0 = 0$, $r_{j+1} = 1$. The intervals $[r_0, r_1], [r_1, r_2], \ldots, [r_j, r_{j+1}]$ cover $[0,1]$, overlap only at endpoints, and their lengths sum to 1; discard any of length 0. Fix one of them, $I = [\alpha, \beta]$ with $\alpha < \beta$. At least one endpoint of I is a root of P: for $[r_0, r_1]$ it is β, for $[r_j, r_{j+1}]$ it is α, and for the others both. Also P has no root in the open interval (α, β), so by the intermediate value theorem P has constant sign there; choose $\varepsilon \in \{-1, 1\}$ with $\varepsilon P \geq 0$ on I.

Let $\phi\colon [0,1] \to I$ be the affine bijection with $\phi(0)$ equal to a root endpoint of I and $\phi(1)$ the other endpoint, and set $Q = \varepsilon\,(P \circ \phi)$. Then Q is a polynomial of degree at most 3, $Q(0) = \varepsilon P(\phi(0)) = 0$, and $Q \geq 0$ on $[0,1]$ with $\max_{[0,1]} Q = \max_I |P| \leq M$. Substituting $x = \phi(t)$, so that $dx = \pm(\beta - \alpha)\,dt$, gives

$$\int_I |P(x)|\,dx = (\beta - \alpha)\int_0^1 Q(t)\,dt \leq (\beta - \alpha)\cdot\frac{5}{6}\max_I |P| \leq \frac{5}{6}M(\beta - \alpha),$$

using Lemma 2. Summing over all the intervals and using that their lengths sum to 1,

$$\int_0^1 |P(x)|\,dx \leq \frac{5}{6}M\sum(\beta - \alpha) = \frac{5}{6}M.$$ □

It remains to prove that no smaller constant works, that is, to exhibit a cubic attaining equality. The equality analysis tells you what it must be rather than leaving you to guess. Equality in Lemma 2 forces $Q(\frac{1}{2}) = Q(1) = \mu$; normalise $\mu = 1$. Since $\frac{1}{2}$ lies in the interior of $[0,1]$ and $Q \leq 1$ there,

$x = \frac{1}{2}$ is an interior maximum, so $1 - Q$ vanishes to order at least 2 at $\frac{1}{2}$, and it also vanishes at 1. As $1 - Q$ has degree at most 3 and is not identically 0 (because $Q(0) = 0$), it must be exactly

$$1 - Q(x) = k\big(x - \tfrac{1}{2}\big)^2(x - 1)$$

for some constant $k \neq 0$, and $Q(0) = 0$ gives $1 = k \cdot \frac{1}{4} \cdot (-1)$, so $k = -4$ and

$$Q(x) = 1 + 4\big(x - \tfrac{1}{2}\big)^2(x - 1) = 1 + (2x - 1)^2(x - 1) = 4x^3 - 8x^2 + 5x.$$

Now verify that this Q really is admissible. It has degree 3 and $Q(0) = 0$, so it has a root in $[0, 1]$. Its derivative is

$$Q'(x) = 12x^2 - 16x + 5 = (2x - 1)(6x - 5),$$

so on $[0, 1]$ the polynomial increases on $[0, \frac{1}{2}]$, decreases on $[\frac{1}{2}, \frac{5}{6}]$ and increases on $[\frac{5}{6}, 1]$, with

$$Q(0) = 0, \qquad Q\big(\tfrac{1}{2}\big) = 1, \qquad Q\big(\tfrac{5}{6}\big) = \tfrac{25}{27}, \qquad Q(1) = 1.$$

Hence $0 \leq Q \leq 1$ on $[0, 1]$ and $\max_{[0,1]} |Q| = 1$. Finally

$$\int_0^1 Q(x)\,dx = 1 - \frac{8}{3} + \frac{5}{2} = \frac{12 - 32 + 30}{12} = \frac{5}{6},$$

so $\int_0^1 |Q| = \frac{5}{6} \max_{[0,1]} |Q|$ and no constant smaller than $\frac{5}{6}$ can work. Combined with the Proposition, the smallest constant is $C = \frac{5}{6}$.

Remark 10.6.4. The choice of quadrature rule is not free, and the comparison is instructive. Any rule $\int_0^1 Q = \sum_i w_i Q(x_i)$ with nodes $x_i \in [0, 1]$, nonnegative weights and exactness on cubics satisfies $\sum_i w_i = 1$, so for Q with $0 \leq Q \leq \mu$ and $Q(x_{i_0}) = 0$ at a node x_{i_0} it gives $\int_0^1 Q \leq (1 - w_{i_0})\mu$. You therefore want the rule, with the root as a node, whose weight at that node is as large as possible. The two point Gauss rule is exact on cubics but its nodes are $\frac{1}{2} \pm \frac{1}{2\sqrt{3}}$ and neither is a root, so it gives nothing. The closed four node rule $\frac{1}{8}(Q(0) + 3Q(\frac{1}{3}) + 3Q(\frac{2}{3}) + Q(1))$ has $w_0 = \frac{1}{8}$ and gives only $\frac{7}{8}$. Simpson has $w_0 = \frac{1}{6}$ and gives $\frac{5}{6}$, which the example above shows is optimal.

Layer 1: WHAT

Block 1. Both sides are homogeneous of degree 1, so normalise $\max_{[0,1]} |P| = 1$ and maximise $\int_0^1 |P|$.

Block 2. The modulus sign is the obstruction. Split $[0, 1]$ at the roots of P; on each piece P has constant sign, so $|P| = \pm P$ is a polynomial there.

Block 3. Every piece has a root of P at one of its endpoints. Rescale each piece affinely onto $[0, 1]$ with that root sent to 0: the maximum is unchanged, the integral scales by the length, and the lengths sum to 1. So it is enough to prove the inequality once, for a nonnegative cubic on $[0, 1]$ vanishing at 0.

Block 4. Now you need $\int_0^1 Q$ in terms of a few values of Q, one of which is the known zero. Integration is a linear functional on the four dimensional space of cubics; Simpson's rule reproduces it exactly from the three values at $0, \frac{1}{2}, 1$ with weights $\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$.

Block 5. Since $Q(0) = 0$, the node at 0 contributes nothing and $\int_0^1 Q = \frac{1}{6}(4Q(\frac{1}{2}) + Q(1)) \leq \frac{5}{6}\mu$. The upper bound is done.

Block 6. Equality forces $Q(\frac{1}{2}) = Q(1) = \mu$ with $\frac{1}{2}$ an interior maximum, so $\mu - Q$ has a double root at $\frac{1}{2}$ and a root at 1; together with $Q(0) = 0$ this pins Q uniquely as $4x^3 - 8x^2 + 5x$ after normalising $\mu = 1$.

Block 7. Verify that this Q actually satisfies $0 \leq Q \leq 1$ on $[0, 1]$, which the equality conditions do not guarantee, and compute $\int_0^1 Q = \frac{5}{6}$.

The nominal route writes $P = ax^3 + bx^2 + cx + d$, imposes $P(r) = 0$, and tries to optimise a four parameter family against an infinite family of sup norm constraints. It never terminates.

The bifurcation point is the line after the normalisation. Zombie: introduce coefficients. Ideal: ask what kind of object $\int_0^1$ is on the space of cubics, and answer "a linear functional, hence a fixed weighted sum of point values". Everything after that is forced.

Layer 2: HOW

Block 1. **Cue.** The inequality has P on both sides, once inside an integral and once inside a maximum, and both are homogeneous of degree 1. **Why natural.** Scaling P by λ scales both sides by $|\lambda|$, and the hypothesis (degree 3, a root in $[0,1]$) is preserved, so the ratio is what matters and you may normalise the maximum to 1. Always strip the homogeneity of an extremal problem first: it removes one dimension for free and turns "find the smallest C" into "maximise one quantity over a bounded set". **Link to the next block.** With the maximum fixed at 1, the only remaining enemy is the modulus sign.

Block 2. **Cue.** $\int_0^1 |P|$, with P a polynomial whose roots you are told something about. **Why natural.** A modulus of a polynomial is not a polynomial, but it is one on each interval between consecutive roots, and the problem has handed you a root to split at. This is the standard move: when an absolute value sits inside an integral, cut the domain at the sign changes of the integrand. Here it is doubly natural because the hypothesis is about roots, so the splitting points are exactly the objects the problem gives you. **Link to the next block.** After splitting, each piece carries a root at one of its endpoints, which is a much more rigid situation than "a root somewhere in the interval".

Block 3. **Cue.** Each piece is an interval with a root at an endpoint, and the desired inequality is invariant under affine changes of variable. **Why natural.** Under $x = \phi(t)$ with ϕ affine from $[0,1]$ onto $[\alpha,\beta]$, the maximum of $|P|$ over the piece is unchanged and the integral picks up the factor $\beta - \alpha$. So the inequality $\int_I |P| \le C\,\ell(I) \max_I |P|$ says the same thing on every interval, and proving it once on $[0,1]$ with the root at 0 proves it everywhere. Then summing over the pieces works only because the lengths add to 1, which is why the factor $\ell(I)$ had to be there. **Link to the next block.** The problem is now the single clean question: how large can $\int_0^1 Q$ be for a nonnegative cubic with $Q(0) = 0$ and $Q \le 1$?

Block 4. **Cue.** Degree at most 3 means a four dimensional vector space; $Q \mapsto \int_0^1 Q$ is linear. **Why natural.** Any linear functional on a finite dimensional space of polynomials is a finite weighted sum of point evaluations, and for integration those are exactly the quadrature rules. Since you know one value already ($Q(0) = 0$), you want a rule with 0 as a node, so that the unknown part of the estimate involves as little weight as possible. Simpson's rule is the standard rule with node set $\{0, \frac{1}{2}, 1\}$, and it is exact on cubics, not merely on quadratics, because the node set is symmetric about $\frac{1}{2}$ and both sides annihilate the odd function $(x - \frac{1}{2})^3$. **Link to the next block.** Substituting the known zero into the rule immediately produces a numerical bound.

Block 5. **Cue.** $Q(0) = 0$ sits precisely on a node of the rule. **Why natural.** The weights of a quadrature rule exact on constants sum to 1, so knocking out one node removes its weight from the estimate: $\int_0^1 Q \le (1 - w_0)\mu$. For Simpson $w_0 = \frac{1}{6}$ and the bound is $\frac{5}{6}\mu$. The whole content of the answer $5/6$ is "one minus the Simpson endpoint weight", and that is why the constant looks the way it does. **Link to the next block.** An inequality proved by a weighted sum of values has an obvious equality condition, so you get the extremal polynomial for free.

Block 6. **Cue.** Equality needs every term that was bounded by μ to equal μ. **Why natural.** That means $Q(\frac{1}{2}) = Q(1) = \mu$. Since $\frac{1}{2}$ is interior and $Q \le \mu$ throughout, $\mu - Q$ has a minimum of 0 at an interior point, hence a double root there, plus a root at 1. A cubic with a prescribed double root and

a prescribed simple root is determined up to a scalar, and the remaining condition $Q(0) = 0$ fixes the scalar. No guessing, no equioscillation folklore: three linear conditions on a three dimensional space. **Link to the next block.** The candidate is forced, but forced candidates still have to be checked.

Block 7. Cue. The equality conditions used only $Q(\frac{1}{2})$ and $Q(1)$, never the constraint $0 \le Q \le \mu$ on the whole interval. **Why natural.** A necessary condition can produce a polynomial that violates the constraints it was derived under, so you must verify admissibility. Factor $Q' = (2x-1)(6x-5)$, read the shape (up, down, up), and evaluate at the four relevant points $0, \frac{1}{2}, \frac{5}{6}, 1$. The local minimum value $\frac{25}{27} > 0$ and the maxima both equal 1, so $0 \le Q \le 1$ and the example is legitimate. **Link to the next block.** Upper bound plus attained example gives the smallest constant, so the problem is closed.

Layer 3: WHY

The first truth is that integration over a fixed interval, restricted to polynomials of degree at most n, is a linear functional on an $(n+1)$ dimensional space, and every such functional is a weighted sum of at most $n+1$ point evaluations. Quadrature rules are not numerical analysis trivia here; they are the coordinates in which the problem is linear. Once you accept that, the entire difficulty of the problem collapses to choosing the right nodes.

The second truth is the exactness statement itself, and it deserves its honest reason rather than a citation. Write $u = x - \frac{1}{2}$. A rule with nodes symmetric about $\frac{1}{2}$ and weights symmetric to match sends the odd function u^3 to 0, and $\int_0^1 u^3\,dx = 0$ as well. So a symmetric rule that is exact on degree 2 is automatically exact on degree 3: the extra degree is a gift of symmetry, not an accident. That is why Simpson, a three node rule, integrates cubics exactly, and it is the reason a solver who knows only "Simpson is a numerical approximation" will not see the route.

The third truth is the localisation principle. An inequality of the form $\int_I |P| \le C\,\ell(I)\max_I |P|$ is invariant under affine reparametrisation of I, so its truth on one interval implies its truth on all, and a quantity that is additive over a partition (the integral) can then be bounded by summing the local bounds, provided the local bound carries the length factor. Whenever both sides of an inequality scale the same way under $x \mapsto \alpha + (\beta - \alpha)x$, the correct move is always the same: localise to the worst configuration, prove the normalised statement, and sum. Here the localisation also upgrades the hypothesis "a root somewhere in $[0,1]$" into the far stronger "a root at an endpoint", which is what makes the bound sharp.

The discernment mechanism is the choice of node set, and it is a genuinely finite search. The bound produced by a rule with nonnegative weights and nodes in $[0,1]$, exact on cubics, and with the root at a node, is $1 - w_0$. So you want the largest possible weight at the root node. The two point Gauss rule is exact on cubics but has no node at 0, so it is worthless here. The closed four node rule has $w_0 = \frac{1}{8}$ and gives $\frac{7}{8}$. Simpson has $w_0 = \frac{1}{6}$ and gives $\frac{5}{6}$. Since $\frac{5}{6}$ is attained, the search stops. This same reasoning solves the analogous problem for other degrees: for degree at most 1 the trapezium rule $\int_0^1 Q = \frac{1}{2}(Q(0) + Q(1))$ gives the sharp constant $\frac{1}{2}$, attained by $Q(x) = x$; for degree at most 2 the correct rule is the two node Radau rule $\int_0^1 Q = \frac{1}{4}Q(0) + \frac{3}{4}Q(\frac{2}{3})$, which is exact on quadratics and gives the sharp constant $\frac{3}{4}$, attained by $Q(x) = 3x - \frac{9}{4}x^2$. Notice that for degree 2 the Newton to Cotes rule is not the right one; the rule must be chosen for the problem, not from memory.

How to train this: whenever you meet an inequality between $\int |P|$ and $\max |P|$ on a finite dimensional space of functions, write the objective as a linear functional, write the constraint set, and check whether the functional can be expressed through point values at the constrained points. If yes, the problem is a weighing exercise and it is over in five lines.

Now the knowledge gaps, stated plainly, because this problem is not approachable without them. First, quadrature: you must know what "a rule exact on polynomials of degree at most n" means, that the weights of such a rule sum to 1 on $[0,1]$, that Simpson's rule has weights $\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$ and

degree of exactness 3, and the symmetry argument for that extra degree. A level numerical methods teaches the trapezium rule as an approximation and stops; exactness is the whole point here and it is a gap. Second, extremal theory of polynomials: the fact that a linear functional on a compact convex set attains its maximum at an extreme point, and the Chebyshev style heuristic that extremal polynomials touch their constraints at as many points as the dimension allows. You do not need this to write the final proof, but without it you will not believe that an extremal polynomial exists or know how many contact conditions to impose. Third, and mildest, the localisation principle for affine invariant inequalities, which is standard practice but rarely stated as a principle.

Revolving concepts

Norms on a space of polynomials. On $[0,1]$ the quantities $\|P\|_1 = \int_0^1 |P|$ and $\|P\|_\infty = \max_{[0,1]} |P|$ are both norms on the four dimensional space of cubics, and all norms on a finite dimensional space are equivalent, so constants like the one asked for necessarily exist. Always $\|P\|_1 \le \|P\|_\infty$; the reverse comparison $\|P\|_\infty \le c_n \|P\|_1$ also holds with a constant depending only on the degree, and inequalities of that shape are called Nikolskii inequalities. This problem asks for the exact constant in the easy direction after the class has been cut down by a root condition. See Book 1, field: Real analysis, topic: norms on finite dimensional spaces.

Quadrature and degree of exactness. A rule $\int_0^1 Q \approx \sum_{i=0}^{k} w_i Q(x_i)$ is exact on degree n when it holds with equality for $Q = 1, x, \dots, x^n$, that is, when the moment equations $\sum_i w_i x_i^m = \frac{1}{m+1}$ hold for $m = 0, 1, \dots, n$. With $k+1$ prescribed nodes you can always achieve $n = k$ by solving a Vandermonde system; symmetry of the node set buys one extra degree when k is even, which gives Simpson its exactness on cubics; and choosing the nodes optimally (Gauss to Legendre) buys $k+1$ extra degrees. Radau rules are the intermediate case where one node is fixed at an endpoint and the rest are optimised, which is exactly the family this problem wants. See Book 1, field: Calculus and integrals, topic: quadrature rules.

Extreme points and linear objectives. If K is a compact convex subset of $\mathbb{R}^d$ and f is linear, then $\max_K f$ is attained at an extreme point of K. Here K is the set of cubics with $Q(0) = 0$ and $0 \le Q \le 1$ on $[0,1]$, sitting inside $\mathbb{R}^3$, and its extreme points are pinned by three active constraints; the extremal polynomial found above has exactly three, namely $Q(\frac{1}{2}) = 1$, $Q'(\frac{1}{2}) = 0$ and $Q(1) = 1$. Counting active constraints against the dimension is the reliable way to predict the shape of an extremiser before computing it. See Book 1, field: Inequalities, topic: convexity and extreme points.

Chebyshev polynomials and equioscillation. The classical extremal problems for polynomials under a sup norm constraint are solved by Chebyshev polynomials, whose defining feature is that they attain their extreme value at the maximum possible number of points with alternating signs. The polynomial $4x^3 - 8x^2 + 5x = 1 + (2x-1)^2(x-1)$ is the analogue for this constrained problem: it hits its maximum twice and its minimum once inside $[0,1]$. See Book 1, field: Polynomials, topic: Chebyshev extremal problems.

Structure of the argument, generalised. The scheme is: normalise by homogeneity, localise by affine invariance, linearise by quadrature, and close by equality analysis. The generalised statement it proves is this. Let $n \ge 1$ and let C_n be the smallest constant with $\int_0^1 |P| \le C_n \max_{[0,1]} |P|$ for all real P of degree at most n with a root in $[0,1]$. If there is a quadrature rule with nodes in $[0,1]$, nonnegative weights, exactness on degree n, and node 0 carrying weight w_0, then $C_n \le 1 - w_0$; and if some admissible P attains that value, the bound is sharp. This gives $C_1 = \frac{1}{2}$, $C_2 = \frac{3}{4}$ and $C_3 = \frac{5}{6}$. There is no recursion in n here; each degree needs its own optimal rule, which is why you must know the Radau construction rather than one memorised formula. See Book 1, field: Calculus and integrals, topic: quadrature rules.

Circuit Bank

- **WHEN** an inequality is homogeneous of the same degree in the unknown function on both sides **DO** normalise one side to 1 and maximise the other over the resulting bounded set **BECAUSE** the class is a cone, so the ratio is the only content, and normalising removes a dimension for free
- **WHEN** an absolute value sits inside an integral of a polynomial **DO** split the domain at the roots, where the sign changes, so the integrand is a polynomial on each piece **BECAUSE** a modulus is only piecewise smooth, and the pieces are exactly the intervals between roots, which the hypothesis is usually about
- **WHEN** both sides of an inequality scale identically under an affine change of variable **DO** prove the normalised version on $[0, 1]$ with the length factor included, then sum over a partition **BECAUSE** affine invariance means one proof covers every interval, and the length factor is what makes the sum telescope to the whole domain
- **WHEN** you need the integral of a polynomial of degree at most n and you know its value at some specific point **DO** write the integral as a quadrature rule exact on degree n whose nodes include that point **BECAUSE** integration is a linear functional on an $(n + 1)$ dimensional space, so it is a fixed weighted sum of point values, and a known zero deletes a whole weight
- **WHEN** you have proved an upper bound by bounding several terms by the same quantity **DO** read off the equality conditions and use them to construct the extremal example **BECAUSE** the equality conditions are linear conditions on a finite dimensional space and usually determine the extremiser uniquely, which converts guessing into solving

Backtesting

1. First useful cue: degree 3 with an integral and a maximum over the same interval, plus a prescribed root. That says linear functional on a four dimensional space, and a known value at a point, so quadrature.
2. Method the expression suggested: normalise, split at the roots, rescale each piece so the root sits at an endpoint, apply Simpson's rule, and read off $1 - \frac{1}{6}$.
3. Where you would hesitate, and the thought to reinforce: at the reduction. It feels like a lot of bookkeeping for what looks like a special case. Reinforce that the reduction is not bookkeeping but the step that upgrades "a root somewhere" into "a root at an endpoint", which is what makes the constant sharp; and that it is legitimate only because the local inequality carries the factor $\ell(I)$ and the lengths sum to 1.
4. The trap the question sets: producing only the example. It is the more enjoyable half, it gives the right numerical answer, and it is worth about half the credit. The second trap is verifying the example incompletely: the equality conditions do not imply $0 \leq Q \leq 1$, so the sign analysis via $Q' = (2x - 1)(6x - 5)$ is compulsory.
5. What makes you recognise it faster next time: the pairing of $\int$ and max over one interval on a finite dimensional function space is the fingerprint of a quadrature or extreme point argument, and a constraint of the form "vanishes at a point" is telling you to put that point at a node.
6. The reusable decision rule: to bound a linear functional on polynomials of bounded degree subject to sup norm and interpolation constraints, express the functional exactly as a nonnegative weighted sum of values at points where you have information, and the bound is the total weight at the points where the value is unconstrained.

Mirror drills

1. Find the smallest C such that $\int_0^1 |P| \le C \max_{[0,1]} |P|$ for every real polynomial of degree 1 with a root in $[0,1]$. *Answer:* $\frac{1}{2}$; the trapezium rule $\int_0^1 Q = \frac{1}{2}(Q(0)+Q(1))$ is exact on degree 1, and $Q(x) = x$ attains it.
2. Do the same for degree 2. *Answer:* $\frac{3}{4}$; use the two node rule $\int_0^1 Q = \frac{1}{4}Q(0) + \frac{3}{4}Q(\frac{2}{3})$, which is exact on quadratics (check on $1, x, x^2$), and note $Q(x) = 3x - \frac{9}{4}x^2$ has maximum 1 at $x = \frac{2}{3}$ and integral $\frac{3}{4}$.
3. Prove that Simpson's rule is exact on cubics using symmetry rather than by checking four basis elements, and explain why the closed four node rule with weights $\frac{1}{8}, \frac{3}{8}, \frac{3}{8}, \frac{1}{8}$ gives only the weaker constant $\frac{7}{8}$. *Answer:* write $u = x - \frac{1}{2}$; the rule is symmetric so it annihilates u^3, as does $\int_0^1$, and exactness on degree 2 is a direct check, so exactness on degree 3 follows. The four node rule has endpoint weight $\frac{1}{8}$, and the bound produced is $1 - w_0 = \frac{7}{8}$: more nodes means less weight sitting on the known zero.

10.7 B1. exponentiate the log to expose telescoping

Let $x_0, x_1, x_2, \ldots$ be the sequence such that $x_0 = 1$ and, for $n \geq 0$,

$$x_{n+1} = \ln\left(e^{x_n} - x_n\right)$$

(as usual, $\ln$ denotes the natural logarithm). Show that the infinite series $x_0 + x_1 + x_2 + \cdots$ converges and find its sum.

Read it out loud

This is a sequences and series question wearing a transcendental disguise. Strip the disguise and what is actually being tested is completely standard: a recursively defined sequence, a request to prove convergence, and a request for the exact sum. The only unusual feature is that the recursion is defined through a composition of ln and exp, and that composition is itself the first command term. Whenever the rule for x_{n+1} places e^{x_n} inside a ln, the two functions are begging to be cancelled, because they are inverses of each other and cancelling them is the one manipulation guaranteed to make progress.

The second thing to read out loud is the phrase find its sum, not merely does it converge. The general term x_n has no elementary closed form: check this within thirty seconds by noting $x_1 = \ln(e-1)$ is already not a nice number. So an exact sum can only exist if the partial sums collapse, not the individual terms. That is the second cue: it tells you, before you write a line of proof, that you are hunting for a telescoping identity, not a formula for x_n.

The field embedded here is real analysis dressed as calculus: monotone bounded sequences, continuity of exp, and the inequality $e^x \geq x + 1$. None of this needs anything beyond what you already carry from Further Maths calculus; what is new is the discipline of reading the composition $\ln(e^{(\cdot)} - \cdot)$ as an instruction to exponentiate, rather than as an expression to simplify by brute force.

Approaches, ranked worst to best

1. **hunt for a closed form of x_n.** *Cue:* a recursively defined sequence always tempts a pattern search. *Circuit:* small case pattern matching. (worst: $x_1 = \ln(e-1)$ and $x_2 = \ln(e^{x_1} - x_1)$ are not recognisable numbers, so this stalls before it starts and costs ink for nothing).
2. **numerical iteration to guess the target value.** *Cue:* you want a number to aim at before committing to a proof. *Circuit:* decimal iteration as a sanity check. (a fine backtest, but it is not a proof and earns nothing on its own; keep it in your head, not on the page).
3. **exponentiate the recursion to expose a telescoping identity.** *Cue:* ln applied to an expression built from e^{x_n}. *Circuit:* cancel inverse functions on sight. (best: converts an opaque transcendental recursion into the algebraic identity $x_n = e^{x_n} - e^{x_{n+1}}$, which is exactly the shape of a telescoping summand).

The ideal solution

First certify that the recursion is always well defined. For real x let $\varphi(x) = e^x - x$; then $\varphi'(x) = e^x - 1$, negative for $x < 0$, zero at $x = 0$, positive for $x > 0$, so φ has a strict global minimum at $x = 0$ with $\varphi(0) = 1$. Hence

$$e^x - x \geq 1 \quad \text{for every real } x, \text{ with equality only at } x = 0.$$

So $e^{x_n} - x_n$ is always at least 1, meaning $\ln(e^{x_n} - x_n)$ is always defined, and moreover $x_{n+1} = \ln(e^{x_n} - x_n) \geq \ln 1 = 0$, with equality only if $x_n = 0$.

Since $x_0 = 1 \neq 0$, induction gives $x_n > 0$ for every n: the base case is immediate, and if $x_n > 0$ then $x_n \neq 0$, so the inequality above is strict and $x_{n+1} > 0$.

Now exponentiate the defining relation. Since $x_{n+1} = \ln(e^{x_n} - x_n)$,

$$e^{x_{n+1}} = e^{x_n} - x_n, \qquad \text{that is,} \qquad x_n = e^{x_n} - e^{x_{n+1}}.$$

This single identity carries the whole problem. Because $x_n > 0$, it gives $e^{x_{n+1}} < e^{x_n}$, and since exp is strictly increasing, $x_{n+1} < x_n$. So (x_n) is strictly decreasing and bounded below by 0, and by the Monotone Convergence Theorem it converges to some limit $L \geq 0$. Taking $n \to \infty$ in $x_n = e^{x_n} - e^{x_{n+1}}$, using continuity of exp on both sides, gives $L = e^L - e^L = 0$. So $x_n \to 0$.

Finally, sum the identity for $n = 0, \ldots, N$:

$$\sum_{n=0}^{N} x_n = \sum_{n=0}^{N} \left(e^{x_n} - e^{x_{n+1}}\right) = e^{x_0} - e^{x_{N+1}} = e - e^{x_{N+1}},$$

the middle terms cancelling in pairs. Since $x_{N+1} \to 0$ as $N \to \infty$ and exp is continuous, $e^{x_{N+1}} \to e^0 = 1$. Hence the partial sums converge, and

$$\sum_{n=0}^{\infty} x_n = \boxed{e - 1}.$$

Layer 1: WHAT

Block 1: certify positivity of the argument of every ln in the recursion via $e^x \geq x + 1$, and deduce $x_n > 0$ for all n by induction. Block 2: exponentiate the recursion to trade the opaque log identity for the algebraic identity $x_n = e^{x_n} - e^{x_{n+1}}$. Block 3: read monotonicity straight off that same identity. Block 4: invoke the Monotone Convergence Theorem, then find the limit by substituting it back into the identity. Block 5: sum the identity from $n = 0$ to N, watch it telescope, and let $N \to \infty$.

The nominal route computes $x_1, x_2, x_3, \ldots$ directly and searches for a pattern in the numbers themselves; it dies at $x_1 = \ln(e - 1)$, a number with no visible structure. The ideal route never looks at x_n as a number at all; it manipulates the relation between consecutive terms. The bifurcation point sits at the very first line you write after reading the problem: do you compute x_1 numerically, or do you exponentiate the defining equation symbolically before computing anything? Everything downstream follows from that choice.

Layer 2: HOW

Block 1. **Cue.** The argument of ln is $e^{x_n} - x_n$, an exponential minus a linear image of its own exponent. **Why natural.** A ln recursion is not trustworthy until you know its argument is positive for every term it will ever be fed; checking this costs one calculus minimisation and, as a bonus, hands you the sign of every x_n for free. **Link to the next block.** Knowing $x_n > 0$ for all n is exactly the hypothesis Block 3 needs to turn an identity into a strict inequality.

Block 2. **Cue.** ln sits directly outside an expression built from e^{x_n}. **Why natural.** ln and exp are mutual inverses, so applying exp to both sides is the one operation guaranteed to simplify rather than complicate; and since the problem demands an exact sum rather than an exact term, you should already be looking for the result to have an additive, telescoping shape. **Link to the next block.** The identity $x_n = e^{x_n} - e^{x_{n+1}}$ literally reads term equals a difference of consecutive values of a function of the sequence, which is the definition of a telescoping summand: summing is now free.

Block 3. **Cue.** The identity from Block 2, combined with the positivity from Block 1. **Why natural.** A positive difference $e^{x_n} - e^{x_{n+1}}$ forces $e^{x_{n+1}} < e^{x_n}$, and since exp is injective and increasing, the ordering pulls straight back onto $x_{n+1} < x_n$; no separate derivative argument on the sequence itself is needed. **Link to the next block.** A decreasing sequence bounded below is exactly the hypothesis of the Monotone Convergence Theorem, so existence of a limit is automatic and only its value remains open.

Block 4. **Cue.** A convergent sequence defined by a continuous recursion. **Why natural.** If $x_n \to L$ and $x_{n+1} = g(x_n)$ for continuous g, then $L = g(L)$; it is easier here to pass the limit through the un-logged identity of Block 2 than to reintroduce ln. **Link to the next block.** Knowing $x_n \to 0$ pins down the one unknown quantity left in the telescoped partial sum, the tail term $e^{x_{N+1}}$.

Block 5. **Cue.** The telescoped finite sum $e - e^{x_{N+1}}$ together with $x_{N+1} \to 0$. **Why natural.** Continuity of exp at 0 converts a statement about the sequence into a statement about numbers: $e^{x_{N+1}} \to e^0 = 1$. **Link to the next block.** None; this closes the problem.

Layer 3: WHY

The truth underneath this problem is that a recursively defined sequence can hide an exact potential function even when it has no closed form: if the defining map, once you undo its outer wrapper, expresses the increment x_n as $F(x_n) - F(x_{n+1})$ for some function F, then the sum is exact and equals $F(x_0)$ minus the limit of F along the sequence. The wrapper here is ln composed with $\exp(\cdot) - \mathrm{id}$, and undoing it is what exposes $F = \exp$.

The discernment mechanism is this: the moment a Putnam problem asks for an exact sum of a sequence you cannot solve explicitly, refuse the instinct to search for the general term. Assume instead that the sum itself is the accessible object, and go looking for a telescoping identity by undoing whichever transcendental function wraps the recursion. You train this discernment by noticing, across many problems, that find the exact sum and no closed form for the term are never a contradiction: together they are a promise that a telescoping identity exists.

What the examiners nudged toward is visible in the exact placement of x_n inside the log: it is not $e^{x_n} - 1$ or $e^{x_n} - 2x_n$, it is $e^{x_n} - x_n$ specifically, the one combination whose exponentiated form returns x_n as a clean difference. A quieter nudge is the word converges: by demanding a proof of convergence rather than asserting it, the setters force you to extract monotonicity and boundedness from the same identity that gives the telescoping, rather than treating the two halves of the problem as separate tasks.

Revolving concepts

The inequality $e^x \geq x + 1$. This is the tangent line to the convex function e^x at $x = 0$; convexity guarantees the tangent lies below the curve everywhere, with equality only at the point of tangency. It underlies AM-GM style bounds and any problem needing a cheap lower bound on an exponential. See Book 1, field: Inequalities, topic: convexity and tangent line bounds.

Monotone Convergence Theorem. A monotonic bounded sequence of real numbers converges; this is equivalent to completeness of $\mathbb{R}$ and is the standard first move whenever a recursively defined sequence needs a convergence proof before its limit can be computed. See Book 1, field: Real analysis, topic: monotone sequences and completeness.

Fixed points of continuous iteration. If $x_{n+1} = g(x_n)$ for continuous g and $x_n \to L$, then $L = g(L)$. Whenever a recursion is already known to converge, its limit is always a root of $L = g(L)$, which trades an infinite process for one algebraic equation. See Book 1, field: Sequences and recurrences, topic: fixed points of iterated maps.

Telescoping through an invertible wrapper, generalised. Suppose a recursion, after undoing an invertible continuous function h, takes the form $h(x_{n+1}) = h(x_n) - x_n$. Rearranging gives $x_n = h(x_n) - h(x_{n+1})$ for every n, so

$$\sum_{n=0}^{N} x_n = h(x_0) - h(x_{N+1}),$$

and if $x_n \to L$, the infinite sum equals $h(x_0) - h(L)$. This is the general shape behind this problem with $h = \exp$, and it applies verbatim whenever a recursion reads h of the next term equals h of the current term minus the current term, for any invertible continuous h: the same telescoping falls out with $h = \sinh$, $h = \tan$, or any other cancellable wrapper. See Book 1, field: Sequences and recurrences, topic: telescoping sums and generating potentials.

Circuit Bank

- **WHEN** a recursion is written as an invertible function applied to an expression built from that function's inverse of the previous term, for example $x_{n+1} = \ln(\text{stuff with } e^{x_n})$ **DO** apply the invertible function to both sides immediately, before doing anything else **BECAUSE** the two functions cancel by definition, and the resulting identity is almost always the algebraic relation the problem actually wants you to sum or bound
- **WHEN** a problem asks for the exact sum of a sequence with no closed form for its general term **DO** search for a telescoping identity $x_n = F(x_n) - F(x_{n+1})$ hidden inside the recursion, rather than solving for x_n **BECAUSE** an exact closed form sum of an opaque sequence is only possible when the partial sums collapse, and telescoping is the mechanism that makes them collapse
- **WHEN** you need the limit of a sequence defined by a continuous recursion and already know it converges **DO** substitute the limit directly into the recursion and solve $L = g(L)$ **BECAUSE** continuity lets you pass to the limit termwise, converting an infinite process into one algebraic equation
- **WHEN** you meet the expression $e^x - x$ in any context **DO** reach immediately for $e^x \geq x + 1$ **BECAUSE** it is the tangent line to a convex function at its unique point of contact, and it pins the minimum of $e^x - x$ at exactly 1

Backtesting

1. First useful cue: the argument of ln contains e^{x_n}, an exponential of the very term being logged.
2. Method the expression suggested: exponentiate both sides to cancel the outer ln.
3. Where you would hesitate, and the thought to reinforce: whether you have the right identity to sum; reinforce by noticing $x_n = e^{x_n} - e^{x_{n+1}}$ is shaped exactly like a telescoping summand.
4. The trap the question sets: assuming you must find a closed form for x_n before you can find $\sum x_n$.
5. What makes you recognise it faster next time: treating $\ln(e^{(\cdot)} - \cdots)$ as a standing visual signature that means exponentiate now.
6. The reusable decision rule: never attempt to solve a transcendental recursion for its general term; test first whether undoing the outer function converts the increment into a telescoping difference.

Mirror drills

1. Let $y_0 = 2$ and $y_{n+1} = \ln(e^{y_n} + y_n)$ for $n \geq 0$, note the sign flip. Does $\sum y_n$ converge? *Answer:* No. Exponentiating gives $e^{y_{n+1}} = e^{y_n} + y_n > e^{y_n}$ since $y_0 > 0$, so (y_n) increases, cannot tend to 0, and the series diverges; the sign in the recursion was load bearing, not decorative.
2. Same recursion as the main problem but started at $x_0 = 2$. Find $\sum x_n$. *Answer:* $e^2 - 1$, by the identical telescoping argument, since it only ever used $x_0 > 0$, never the specific value 1.
3. Let $v_0 = 1$ and $v_{n+1} = \operatorname{arcsinh}(\sinh v_n - v_n)$ for $n \geq 0$. Find $\sum v_n$. *Answer:* $\sinh 1$, since $\sinh x > x$ for $x > 0$ gives $v_n > 0$ for all n; applying sinh to both sides gives $v_n = \sinh v_n - \sinh v_{n+1}$, the same wrapper trick with $h = \sinh$ in place of exp, and the limit is again 0 by the same monotone argument.

10.8 B2. counting squarish numbers via basin sizes

Define a positive integer n to be squarish *if either n is itself a perfect square or the distance from n to the nearest perfect square is a perfect square. For example,* 2016 *is squarish, because the nearest perfect square to* 2016 *is* $45^2 = 2025$ *and* $2025 - 2016 = 9$ *is a perfect square. (Of the positive integers between* 1 *and* 10*, only* 6 *and* 7 *are not squarish.)*

For a positive integer N, let $S(N)$ be the number of squarish integers between 1 *and N, inclusive. Find positive constants α and β such that*

$$\lim_{N\to\infty} \frac{S(N)}{N^\alpha} = \beta,$$

or show that no such constants exist.

Read it out loud

Read past the word squarish immediately; the definition is not the content of the problem, it is packaging. What is actually being tested is a counting function built from a nearest neighbour condition, and the command term is the phrase distance to the nearest perfect square. That phrase is not computable in isolation: distance to the nearest something only means anything once you have fixed which two somethings are competing to be nearest. So the first thirty seconds of this problem should produce one picture, the number line with the perfect squares $1, 4, 9, 16, 25, \ldots$ marked on it, and a mental note that every integer sits between two consecutive marks and is closer to one of them.

That picture is a partition. Around each perfect square n^2 there is a block of consecutive integers that are closer to n^2 than to $(n-1)^2$ or $(n+1)^2$, and squarishness, for any integer inside that block, becomes a purely local question about the block. The problem has silently converted a global counting question into a sum of local counting questions, one per block, and it did this the moment it said nearest.

The second cue sits in the answer format, $S(N)/N^\alpha \to \beta$. This is the signature of a counting function with power law growth, and whenever such a limit is demanded for a function built by summing a slowly growing local count across a range of blocks, the growth rate is recovered by comparing a sum to an integral, not by finding an exact formula. The fields embedded here are number theory (perfect squares, floor of the square root), elementary combinatorics (partition into blocks and local counting), and real analysis (the discrete integral test). Nothing here is beyond Further Maths calculus except the discipline of using an integral to control a sum you cannot evaluate exactly, which is exactly the tool this problem is built to make you reach for.

Approaches, ranked worst to best

1. **tabulate small values and curve fit α, β numerically.** *Cue:* explicit small examples are given in the statement. *Circuit:* numerical extrapolation. (worst: might land on the right numbers, but supplies no proof at all, and the problem explicitly demands either a proof of existence or a proof of nonexistence).

2. **search for an exact closed form for $S(N)$ valid for every N.** *Cue:* the small cases look almost formulaic. *Circuit:* direct formula hunting. (costly: an exact global formula does not exist cleanly because of the floor functions involved, and none is needed since only the asymptotic power law is asked for; this route burns ink chasing more precision than the question wants).

3. **partition into basins around each perfect square, count squarish numbers locally via floor of the square root, then estimate the resulting sum by comparison with an integral.** *Cue:* the phrase nearest perfect square, and the shape of the target limit. *Circuit:* basin decomposition

plus the discrete integral test. (best: turns an intractable global count into an exact local formula plus one standard asymptotic estimate, and answers both branches of the question at once).

The ideal solution

Step 1: the basin of n^2. Fix a positive integer n. The point equidistant from $(n-1)^2$ and n^2 is their average, $n^2 - n + \frac{1}{2}$, which is never an integer, so every integer is strictly closer to one side; an integer k satisfies $k > n^2 - n + \frac{1}{2}$, that is $k \geq n^2 - n + 1$, exactly when k is closer to n^2 than to $(n-1)^2$. Likewise the point equidistant from n^2 and $(n+1)^2$ is $n^2 + n + \frac{1}{2}$, so $k \leq n^2 + n$ exactly when k is closer to n^2 than to $(n+1)^2$. Hence the integers for which n^2 is the nearest perfect square form exactly the interval

$$\mathrm{Basin}(n) = \{\, n^2 - n + 1,\ n^2 - n + 2,\ \ldots,\ n^2 + n \,\},$$

a block of exactly $2n$ consecutive integers. Since $\mathrm{Basin}(n)$ ends at $n^2 + n$ and $\mathrm{Basin}(n+1)$ begins at $(n+1)^2 - (n+1) + 1 = n^2 + n + 1$, these blocks are pairwise disjoint and their union over $n = 1, 2, 3, \ldots$ is all of $\{1, 2, 3, \ldots\}$ with no gaps and no ties, so every positive integer lies in exactly one basin.

Step 2: the local count. Inside $\mathrm{Basin}(n)$, an integer $k = n^2 - d$ with $1 \leq d \leq n-1$ is squarish exactly when d is a perfect square, an integer $k = n^2 + d$ with $1 \leq d \leq n$ is squarish exactly when d is a perfect square, and $k = n^2$ itself is squarish outright. So the number of squarish integers in $\mathrm{Basin}(n)$ is

$$c(n) = 1 + \#\{\text{perfect squares in } [1, n-1]\} + \#\{\text{perfect squares in } [1, n]\} = 1 + \lfloor \sqrt{n-1} \rfloor + \lfloor \sqrt{n} \rfloor.$$

(Check this against the statement's own data: $\mathrm{Basin}(1) = \{1, 2\}$, $c(1) = 1 + 0 + 1 = 2$, both squarish; $\mathrm{Basin}(2) = \{3, 4, 5, 6\}$, $c(2) = 1 + 1 + 1 = 3$, and indeed $3, 4, 5$ are squarish while 6 is not.)

Step 3: sandwiching $S(N)$. Let $T(M) = \sum_{n=1}^{M} c(n)$, the number of squarish integers in $\{1, \ldots, M^2 + M\}$, and let $m = \lfloor \sqrt{N} \rfloor$, so $m^2 \leq N < (m+1)^2$. Since $(m-1)^2 + (m-1) = m^2 - m < N$, the whole of $\mathrm{Basin}(1)$ through $\mathrm{Basin}(m-1)$ lies inside $[1, N]$, so $T(m-1) \leq S(N)$. Since $N < (m+1)^2 \leq (m+1)^2 + (m+1)$, the interval $[1, N]$ lies inside $\mathrm{Basin}(1)$ through $\mathrm{Basin}(m+1)$, so $S(N) \leq T(m+1)$. Thus

$$T(m-1) \ \leq\ S(N) \ \leq\ T(m+1), \qquad m = \lfloor \sqrt{N} \rfloor. \tag{$*$}$$

Step 4: the growth rate of $T(M)$. From $\sqrt{n} - 1 < \lfloor \sqrt{n} \rfloor \leq \sqrt{n}$ and the same bound for $\sqrt{n-1}$, $c(n) = 2\sqrt{n} + O(1)$ uniformly in n, so $T(M) = 2\sum_{n=1}^{M} \sqrt{n} + O(M)$. Since $x \mapsto \sqrt{x}$ is increasing, $\int_{n-1}^{n} \sqrt{x}\,dx \leq \sqrt{n} \leq \int_{n}^{n+1} \sqrt{x}\,dx$ for every $n \geq 1$; summing over $n = 1, \ldots, M$ gives

$$\frac{2}{3}M^{3/2} = \int_0^M \sqrt{x}\,dx \ \leq\ \sum_{n=1}^{M} \sqrt{n} \ \leq\ \int_1^{M+1} \sqrt{x}\,dx = \frac{2}{3}\left((M+1)^{3/2} - 1\right).$$

Expanding $(M+1)^{3/2} = M^{3/2} + \frac{3}{2}M^{1/2} + O(M^{-1/2})$ shows both bounds equal $\frac{2}{3}M^{3/2} + O(\sqrt{M})$, so

$$T(M) = \frac{4}{3}M^{3/2} + O(M).$$

Step 5: translating back to N. Since $m \leq \sqrt{N} < m + 1$, $m = \sqrt{N} + O(1)$, and expanding $(m \pm 1)^{3/2} = (\sqrt{N} + O(1))^{3/2} = N^{3/4} + O(N^{1/4})$ shows both $T(m-1)$ and $T(m+1)$ equal $\frac{4}{3}N^{3/4} + O(N^{1/2})$. By $(*)$,

$$S(N) = \frac{4}{3}N^{3/4} + O(N^{1/2}), \qquad \text{hence} \qquad \frac{S(N)}{N^{3/4}} = \frac{4}{3} + O(N^{-1/4}) \ \longrightarrow\ \frac{4}{3}.$$

So $\alpha = \dfrac{3}{4}$, $\beta = \dfrac{4}{3}$ work. This pair is forced: if such a limit exists for some α with finite nonzero β, then for any other positive $\alpha' > \alpha$, $S(N)/N^{\alpha'} = (S(N)/N^{\alpha}) \cdot N^{\alpha-\alpha'} \to 0$, and for $\alpha' < \alpha$ the same ratio tends to infinity, so no other pair can work. Therefore

$$\boxed{\alpha = \tfrac{3}{4},\ \beta = \tfrac{4}{3}}$$

are the unique constants with $\lim_{N\to\infty} S(N)/N^{\alpha} = \beta$.

Layer 1: WHAT

Block 1: convert nearest perfect square into a partition of the positive integers into basins around each n^2, each of exactly $2n$ consecutive integers. Block 2: count squarish integers inside one basin exactly, via the number of perfect square offsets on each side, giving $c(n) = 1 + \lfloor\sqrt{n-1}\rfloor + \lfloor\sqrt{n}\rfloor$. Block 3: sandwich $S(N)$ between $T(m-1)$ and $T(m+1)$, where $m = \lfloor\sqrt{N}\rfloor$ counts how many basins fit inside $[1, N]$. Block 4: replace the floor functions by $\sqrt{n}$ up to bounded error, then estimate $\sum \sqrt{n}$ by comparison with $\int \sqrt{x}\,dx$. Block 5: translate the resulting power of m back into a power of N, read off $\alpha = 3/4$, $\beta = 4/3$, and note the pair is unique by a scaling argument.

The nominal route tries to reason about squarishness globally, perhaps by hunting for an exact formula for $S(N)$ valid for every N; it produces nothing usable, because nearest perfect square has no meaning until you have committed to a partition. The ideal route commits to that partition in the very first line. The bifurcation point is exactly there: does your first written sentence try to describe $S(N)$ directly, or does it draw the basins first and only then ask what squarish means inside one of them? Only the second choice gives you an object, the local count $c(n)$, that is simple enough to sum.

Layer 2: HOW

Block 1. **Cue.** The phrase distance to the nearest perfect square. **Why natural.** Nearest is only well defined relative to two competitors, so the natural first move is to find, for each n, exactly which integers have n^2 as their nearest perfect square; the midpoints to the two neighbouring squares are never integers, so this partition has no ambiguous boundary cases to worry about. **Link to the next block.** Once every integer has a well defined home basin, squarishness of any integer can be decided by looking only inside its own basin.

Block 2. **Cue.** Inside one basin, the offset from n^2 ranges over a bounded set of integers on each side. **Why natural.** Squarishness of $k = n^2 \pm d$ reduces to whether d is a perfect square, and counting perfect squares up to a bound is exactly what floor of a square root computes; this converts a verbal condition into an arithmetic formula with no further case analysis needed. **Link to the next block.** The local formula $c(n)$ is now a function of n alone, which is precisely what you need to sum over basins to build $S(N)$.

Block 3. **Cue.** $S(N)$ counts squarish integers up to N, while $c(n)$ counts them per whole basin; the two do not line up exactly because N generally falls in the middle of a basin. **Why natural.** Rather than compute the exact partial contribution of the boundary basin, sandwich it: every basin fully below N contributes exactly $c(n)$, and N can never lie before the start of Basin$(\lfloor\sqrt{N}\rfloor)$ nor after the end of Basin$(\lceil\sqrt{N}\rceil)$, which pins the error to at most one basin at each end. **Link to the next block.** The bound $(*)$ replaces the awkward function $S(N)$ with the much tamer partial sum $T(M)$, evaluated at $M = m \pm 1$.

Block 4. **Cue.** $T(M)$ is a sum of the slowly growing function $c(n) \approx 2\sqrt{n}$. **Why natural.** There is no exact closed form for $\sum \sqrt{n}$, but $\sqrt{x}$ is increasing and continuous, so its sum over integers is trapped between two integrals of $\sqrt{x}$ shifted by one step; the trapping error is always one power lower than the integral itself, which is exactly harmless once you only need the leading power law. **Link to the next block.** This produces $T(M) = \frac{4}{3}M^{3/2} + O(M)$, a clean power of M, but the problem is stated in terms of N, not M.

Block 5. **Cue.** $m = \lfloor \sqrt{N} \rfloor$ links M back to N. **Why natural.** Substituting $m = \sqrt{N} + O(1)$ into $M^{3/2}$ and expanding converts a power of m into a power of N plus a controlled lower order error, so the $O(1)$ discrepancy between m and $\sqrt{N}$, and between $S(N)$ and $T(m)$, both wash out once divided by the leading term $N^{3/4}$. **Link to the next block.** None; dividing through by $N^{3/4}$ and letting $N \to \infty$ finishes the problem, and the uniqueness of α, β follows immediately from how a power law limit scales.

Layer 3: WHY

The truth underneath this problem is that any condition phrased as nearest element of a fixed increasing sequence secretly defines a partition of the integers into blocks, a one dimensional Voronoi decomposition, and any counting problem stated in those terms is really a sum of local counts, one per block. You cannot see the local count until you have drawn the blocks, and you cannot draw the blocks until you take the phrase nearest literally enough to compute its boundary as a midpoint.

The discernment mechanism is this: whenever a definition compares an object to the nearest member of a sparse reference sequence, sketch the reference sequence on a line first, before touching the actual condition. This turns an amorphous global statement into finitely many local ones, and it works regardless of what the reference sequence is, perfect squares here, but perfect cubes, primes, or any other increasing sequence in other problems. Separately, whenever an asymptotic growth rate rather than an exact value is requested, train yourself to stop as soon as you can write the summand as a power of n up to a bounded error, because from that point the integral comparison is mechanical and adds no further insight, only bookkeeping.

What the examiners nudged toward is encoded in the exact choice of perfect squares as the reference sequence: consecutive squares are $2n$ apart near n^2, a gap that itself grows, and it is exactly this growing gap that produces a genuine power law with a non-trivial exponent instead of a fixed size window; had they chosen a reference sequence with constant gaps, the whole problem would collapse to elementary counting with $\alpha = 1$. The specific exponent $3/4$ is the fingerprint of composing two square root scales, once in the gap size, once in the fraction of that gap taken up by squarish offsets, a composition worth recognising as its own pattern rather than as an isolated coincidence.

Revolving concepts

Floor of the square root as a counting device. $\lfloor \sqrt{x} \rfloor$ is the unique integer k with $k^2 \le x < (k+1)^2$, and it counts exactly the perfect squares in $[1, x]$; more generally the perfect squares in an interval $[a, b]$ number $\lfloor \sqrt{b} \rfloor - \lfloor \sqrt{a-1} \rfloor$. This is the single arithmetic tool that turns nearest perfect square into an explicit count. See Book 1, field: Number theory, topic: floor functions and perfect power counting.

Basin decomposition by nearest reference point. For any strictly increasing sequence $a_1 < a_2 < \cdots$, the integers nearer to a_n than to any other term form an interval bounded by the (generically non integer) midpoints to the two neighbouring terms, and these intervals partition all sufficiently large integers. The generalisation worth installing: this is a one dimensional Voronoi diagram, and it applies to any nearest to a known sequence condition, not only perfect squares; for perfect k-th

powers the basin of n^k has size of order $k\,n^{k-1}$, matching the derivative of n^k. See Book 1, field: Discrete structures and graphs, topic: nearest neighbour partitions of the integers.

The discrete integral test, generalised. If f is increasing and continuous, $\int_0^M f(x)\,dx \le \sum_{n=1}^{M} f(n) \le \int_1^{M+1} f(x)\,dx$, so $\sum_{n=1}^{M} f(n) = \int_0^M f(x)\,dx + O(f(M))$. Viewing $T(M) = T(M-1) + c(M)$ as a running total, a first order recurrence with increment $c(M)$, this is the general solution method whenever the increment grows like a power: if $c(n) \sim Cn^p$ for $p > -1$, then $T(M) \sim \dfrac{C}{p+1} M^{p+1}$, recovering $T(M) \sim \frac{4}{3}M^{3/2}$ here with $C = 2$, $p = \frac{1}{2}$. The same method solves every power sum $\sum n^p$ and every running total whose per step growth is a power law. See Book 1, field: Real analysis, topic: sum to integral comparison and power sum asymptotics.

Uniqueness of a power law exponent. If $f(N)/N^\alpha \to \beta$ for some finite $\beta > 0$, the pair (α, β) is unique among positive α: raising or lowering α multiplies the ratio by a vanishing or exploding factor $N^{\alpha-\alpha'}$. See Book 1, field: Real analysis, topic: regular variation and power law scaling.

Circuit Bank

- **WHEN** a definition compares an integer to the nearest member of a known increasing sequence, such as nearest perfect square **DO** partition the integers into basins around each term of the sequence, using the midpoints to the two neighbours as boundaries **BECAUSE** nearest has no computable meaning until the competitors are fixed, and the basins turn a global condition into finitely many local ones
- **WHEN** you need the exact count of a local property, such as an offset being a perfect square, inside a bounded window **DO** express the count via floor of the appropriate root, not by searching for a global closed form **BECAUSE** floor functions are the exact bookkeeping device for how many terms of a sparse sequence land inside a window of known size
- **WHEN** you need the growth rate, not the exact value, of a sum of a smooth increasing function over a growing range **DO** sandwich the sum between two shifted integrals of that function **BECAUSE** the resulting error is always one power lower than the leading term, so it vanishes after dividing by the leading term and taking the limit
- **WHEN** a problem asks you to find α, β with $f(N)/N^\alpha \to \beta$ finite and positive **DO** treat any working pair you find as the unique answer, without searching for alternatives **BECAUSE** scaling by any other power sends the ratio to zero or to infinity, never to a second finite positive constant

Backtesting

1. First useful cue: distance to the nearest perfect square is undefined until you fix a partition of the integers by nearest square.
2. Method the expression suggested: draw the perfect squares on the number line and read off the basin of each n^2 as the midpoint boundaries to its neighbours.
3. Where you would hesitate, and the thought to reinforce: matching basins to a specific N that falls mid basin; reinforce by remembering $m = \lfloor \sqrt{N} \rfloor$ counts exactly how many full basins fit inside $[1, N]$, with error confined to at most one basin at either end.
4. The trap the question sets: chasing an exact formula for $S(N)$ for every N, which is unnecessary and does not exist cleanly, when only the leading power law is required.
5. What makes you recognise it faster next time: treating nearest [sparse increasing sequence] as an instant cue for a basin decomposition, and $f(N)/N^\alpha \to \beta$ as an instant cue for a sum to integral comparison.

6. The reusable decision rule: reduce a nearest neighbour counting problem to a sum, over basins, of a local count depending only on the basin index, then find that sum's growth rate against an integral of the same order.

Mirror drills

1. Call n *cubeish* if n is a perfect cube or its distance to the nearest perfect cube is a perfect cube. Find α with $S(N)/N^{\alpha} \to \beta$ for some finite $\beta > 0$. *Answer/hint:* the basin of n^3 has size of order $3n^2$ (the derivative of n^3), so the local cubeish count is of order $n^{2/3}$; summing that over n up to $N^{1/3}$ and comparing with an integral gives $\alpha = 5/9$. In general, for nearest perfect k-th power, $\alpha = (2k-1)/k^2$, which recovers $3/4$ at $k = 2$.

2. Compute $S(100)$ directly from the basin method. *Answer:* $T(9) = \sum_{n=1}^{9} c(n) = 38$, and the truncated basin of $n = 10$ contributes only its left half and centre, since $N = 10^2$ exactly: the squarish numbers $91, 96, 99, 100$, four more; so $S(100) = 42$.

3. Suppose instead n is called squarish-prime if n is a perfect square or its distance to the nearest perfect square is prime. Does $S(N)/N^{\alpha}$ still converge to a finite positive limit for some α? *Answer/hint:* no clean power law survives; the count of primes up to a bound x grows like $x/\ln x$ by the prime number theorem, not like a power of x, so the local count per basin is $\Theta(n/\ln n)$ rather than $\Theta(\sqrt{n})$, and the resulting sum carries an extra logarithmic factor that no single power N^{α} can absorb.

10.9 B3. double the extremal triangle to cover the set

Suppose that S is a finite set of points in the plane such that the area of triangle $\triangle ABC$ is at most 1 whenever A, B, and C are in S. Show that there exists a triangle of area 4 that (together with its interior) covers the set S.

Read it out loud

Two words carry this whole problem: finite and 4. Finite is a command term for the extremal principle, choose the object in S that maximises some quantity, because a finite set of real numbers always attains its maximum, with no limiting or compactness argument required. The moment you see finite set together with a bounded quantity attached to triples of that set, your first written line should be let $A, B, C \in S$ maximise the area of triangle ABC.

The number 4 is the second cue, and it is not an arbitrary bound, it is 2^2. Whenever a covering object is asked to have k^2 times the area of an extremal object found inside the same configuration, the natural construction is to scale that extremal object outward by a linear factor of k. Here $k = 2$, and the triangle scaled by 2 outward from an extremal triangle ABC is a very concrete classical object, the triangle through each vertex of ABC parallel to the opposite side, which has ABC sitting inside it as its own medial triangle.

The field embedded here is combinatorial geometry: the extremal principle, similarity and area scaling, and the fact that a triangle is exactly the intersection of the three half planes bounded by its own sides. None of this is hidden, the two numbers in the statement, finite and 4, hand you the whole strategy before you have drawn a single picture; what remains is to certify, rigorously and for every point of S at once, that the doubled triangle really does contain the original set, and to notice the one configuration the argument does not immediately cover: a set S with no three points forming a genuine triangle at all.

Approaches, ranked worst to best

1. **bound the convex hull of S directly via its diameter or width.** *Cue:* S is a bounded planar set of points. *Circuit:* diameter and width estimates. (worst: there is no clean route from a diameter bound to the exact constant 4, and this approach does not use the specific hypothesis, area at most 1 for every triple, in any sharp way).
2. **normalise via an affine transformation before arguing.** *Cue:* area ratios are affine invariant, so a convenient triangle shape can be assumed. *Circuit:* affine reduction to a standard triangle. (a legitimate simplification of bookkeeping, but it removes none of the real content; the covering argument itself still has to be supplied afterwards, so this only ever ranks as a helper, never as the main idea).
3. **extremal principle plus doubling the maximal triangle outward, then certify containment side by side using half planes.** *Cue:* the words finite and the target area $4 = 2^2$. *Circuit:* extremal choice, medial/anticomplementary triangle duality, half plane intersection. (best: the two cues in the statement point directly at this construction, and it settles every point of S at once with three short area comparisons).

The ideal solution

Reduction to a genuine triangle. If no three points of S are non collinear, then S lies on a single line ℓ (this covers $|S| \leq 2$ as a special case). If $|S| \leq 1$, any triangle of area 4 containing that single point works. Otherwise let $B, C \in S$ be two points at the greatest distance $d = |BC| > 0$ apart; since S lies on ℓ and B, C are its extreme points along ℓ, every point of S lies on segment BC. Take the triangle with base BC and apex at perpendicular height $h = 8/d$ above the midpoint of BC: its area

is $\frac{1}{2} \cdot d \cdot h = 4$, and since $S \subseteq BC$, which is one full edge of this triangle, S is covered. This disposes of every degenerate case, so from now on assume S contains three non collinear points.

The extremal triangle. Since S is finite, choose $A, B, C \in S$ maximising the area of triangle ABC among all triples of points of S; write $T = \text{area}(ABC)$, so $0 < T \leq 1$.

Doubling outward. Let A', B', C' be the points such that A, B, C are the midpoints of $B'C', C'A', A'B'$ respectively; equivalently, draw through each of A, B, C the line parallel to the opposite side of ABC, and let these three lines meet in A', B', C'. Then ABC is the medial triangle of $A'B'C'$, similar to it with ratio $\frac{1}{2}$, so

$$\text{area}(A'B'C') = 4T \leq 4.$$

Every point of S lies inside $A'B'C'$. Fix $P \in S$. For a triangle with a fixed base QR, area equals $\frac{1}{2}|QR| \cdot d(X, \text{line } QR)$ for the apex X, so two such triangles sharing the base QR have areas in the same order as the perpendicular distances of their apexes from line QR. Applying this with base BC: since $\text{area}(PBC) \leq T = \text{area}(ABC)$ by maximality of T, we get $d(P, \text{line } BC) \leq d(A, \text{line } BC)$. But line $B'C'$ is, by construction, exactly the line through A parallel to BC, at perpendicular distance $d(A, \text{line } BC)$ from it, on the same side as A (hence the same side as B, C, since B, C lie on line BC itself, distance 0 from it). So the inequality $d(P, \text{line } BC) \leq d(A, \text{line } BC)$ says exactly that P lies in the closed half plane bounded by line $B'C'$ that contains B and C.

The identical argument with base CA (using $\text{area}(PCA) \leq T$) places P in the half plane bounded by line $C'A'$ containing C, A; and with base AB (using $\text{area}(PAB) \leq T$) places P in the half plane bounded by line $A'B'$ containing A, B. A triangle is exactly the intersection of the three closed half planes determined by its own three sides, the ones containing the triangle's interior, so P lying in all three such half planes for $A'B'C'$ means precisely $P \in \triangle A'B'C'$ (closed region). Since $P \in S$ was arbitrary,

$$S \subseteq \triangle A'B'C'.$$

Fixing the area at exactly 4. If $4T = 4$, the triangle $A'B'C'$ already has area 4 and covers S: done. If $4T < 4$, apply a homothety centred at the centroid of $A'B'C'$ with ratio $1/\sqrt{T} > 1$; this maps $A'B'C'$ to a similar triangle of area exactly $\left(1/\sqrt{T}\right)^2 \cdot 4T = 4$, and since a homothety with ratio greater than 1 centred at an interior point of a convex region maps that region into a strictly larger region containing it, the image triangle still contains $A'B'C'$ and hence still covers S. Either way, a triangle of area exactly 4 covering S has been produced.

such a triangle of area 4 always exists.

Layer 1: WHAT

Block 1: rule out the degenerate case first, covering a fully collinear S directly with a thin triangle of area 4 built on its two extreme points. Block 2: invoke the extremal principle, choosing $A, B, C \in S$ of maximal triangle area $T \leq 1$, now guaranteed non degenerate. Block 3: double ABC outward into $A'B'C'$, the triangle of which ABC is the medial triangle, giving area exactly $4T$. Block 4: for an arbitrary $P \in S$, turn each of the three maximality inequalities into a half plane membership via a shared base distance comparison, then intersect the three half planes to place P inside $A'B'C'$. Block 5: patch the area up to exactly 4 by a homothety when $T < 1$.

The nominal route reasons about the convex hull of S directly, perhaps via its diameter or width, hoping some general covering inequality falls out; nothing forces the constant 4 to appear this way, and the specific hypothesis (area at most 1 for every triple) is never used sharply. The ideal route never looks at S as a shape at all until the very last step; it manufactures one fixed triangle from S and tests every other point of S against that one triangle, three inequalities at a time. The bifurcation point is the first line of the proof: do you start describing S's global shape, or do you start by naming an extremal triple? Only the second choice produces an object rigid enough to argue against.

Layer 2: HOW

Block 1. **Cue.** The set S is finite, and the quantity in question is an area attached to triples of S. **Why natural.** A finite set of real numbers always attains its maximum, so naming a maximiser costs nothing and immediately produces a single fixed triangle to build the whole proof around, in place of an unwieldy quantifier over all triples. **Link to the next block.** But this maximiser might be degenerate if S has no non collinear triple at all, so that possibility has to be checked and dispatched before relying on it.

Block 2. **Cue.** The target covering area is 4, and the bound on every triple is 1; $4 = 2^2$. **Why natural.** A covering region with k^2 times the extremal area is exactly what you get by scaling the extremal triangle outward by linear factor k; for $k = 2$ that scaled triangle is the classical one obtained by drawing, through each vertex of ABC, the line parallel to the opposite side, of which ABC is then the medial triangle. **Link to the next block.** This produces a specific candidate triangle $A'B'C'$ of area exactly $4T \leq 4$, but being a candidate is not the same as covering S; that still needs proof, point by point.

Block 3. **Cue.** Maximality of $T = \text{area}(ABC)$ applied to the triples (P, B, C), (P, C, A), (P, A, B) for an arbitrary $P \in S$. **Why natural.** Each inequality $\text{area}(PBC) \leq T$ shares the base BC with $\text{area}(ABC)$, and two triangles on a common base have areas in the same order as their apexes' distances from that base, so the inequality converts directly into a statement about how far P can be from line BC. **Link to the next block.** The line at that maximal permitted distance from BC, on the side containing A, is exactly line $B'C'$ by construction, so the distance bound becomes a half plane membership bounded by one side of $A'B'C'$.

Block 4. **Cue.** Three half plane memberships, one per side of $A'B'C'$, each containing the whole of ABC. **Why natural.** A triangle is nothing more than the intersection of the three half planes cut out by its own side lines on the side of its interior, so three simultaneous half plane memberships are exactly membership in the triangle $A'B'C'$ itself; nothing more needs to be checked. **Link to the next block.** Since $P \in S$ was arbitrary, this holds for every point of S at once, so $S \subseteq \triangle A'B'C'$, a triangle of area at most 4.

Block 5. **Cue.** The problem asks for a triangle of area exactly 4, while the construction so far only guarantees area at most 4. **Why natural.** Growing a convex region outward from an interior centre by a homothety of ratio greater than 1 always produces a similar, strictly larger region containing the original, so a shortfall in area can always be made up without losing the covering property. **Link to the next block.** None; this closes the problem for the non degenerate case, and Block 1 already closed the degenerate one.

Layer 3: WHY

The truth underneath this problem is that maximality of an extremal object, tested one replaced vertex at a time, is exactly strong enough to certify a half plane containment for every other point of the configuration, one half plane per side of the extremal object. This is not special to triangles: it is a general fact about extremal simplices inside a bounded family, that doubling the extremal simplex outward via parallel supporting hyperplanes through its own vertices always produces a region containing every competitor, because the maximality condition is precisely the statement that no competitor can push past any one of those supporting hyperplanes.

The discernment mechanism is this: whenever a problem gives you finite together with a quantity bounded over all tuples drawn from that finite set, name the extremiser immediately, before attempting any other manipulation, because maximality is a hypothesis that only pays off

once you have something concrete to test other elements against. Separately, whenever a covering object is asked to be some fixed multiple of an extremal object's size, translate that multiple into a linear scale factor and look for the classical similarity construction realising it, rather than trying to bound the covering region's size by direct computation.

What the examiners nudged toward is the exact pairing of finite (licensing the free extremal choice) with the numerical target 4 (pointing at a linear doubling). A quieter nudge is the phrase together with its interior, which tells you in advance that the covering claim is about a closed region, not merely about the three final vertices, preparing you for the half plane argument rather than a vertex counting one.

Revolving concepts

The extremal principle. A nonempty finite set of real numbers attains its maximum and minimum; applied to a numerical quantity indexed by tuples from a finite set, this licenses fixing an extremal tuple as a distinguished object with no limiting argument. See Book 1, field: Combinatorics, topic: extremal principle.

Area via base and height, as a distance comparison. For a fixed base QR, $\text{area}(XQR) = \frac{1}{2}|QR|\, d(X, \text{line } QR)$, so comparing two such areas is exactly comparing two perpendicular distances to the same line. This is the bridge, used three times here, between an area inequality and a half plane membership. See Book 1, field: Geometry, topic: area as base times height.

Medial triangle and its inverse, the anticomplementary triangle. If A, B, C are the midpoints of the sides of $A'B'C'$, then $A'B'C'$ is similar to ABC with ratio 2 and area $4\,\text{area}(ABC)$; equivalently each side of $A'B'C'$ passes through one vertex of ABC parallel to the opposite side. See Book 1, field: Geometry, topic: medial triangle and area scaling under similarity.

A convex polygon as an intersection of half planes. Any triangle, and more generally any convex polygon, equals the intersection of the closed half planes bounded by its own edge lines and containing its interior; membership in the polygon is equivalent to simultaneous membership in every one of these half planes. See Book 1, field: Geometry, topic: half plane (H) representations of convex regions.

The construction, generalised to k dimensions. The same three ingredients, extremal principle, doubling outward via hyperplanes through each vertex parallel to the opposite facet, and half space intersection, generalise verbatim to a finite set of points in k-dimensional space where every k-simplex among them has volume at most 1. Carrying out the construction on the standard simplex with vertices at the origin and the standard basis vectors shows the resulting covering simplex has volume exactly k^k times the extremal simplex's volume: this recovers the factor $4 = 2^2$ used here for $k = 2$, but gives $27 = 3^3$, not 8, for tetrahedra in three dimensions, since a tetrahedron's six edge midpoints form an octahedron rather than a smaller tetrahedron, so the clean medial triangle picture is special to two dimensions even though the extremal argument itself is not. See Book 1, field: Geometry, topic: simplices and extremal covering constructions.

Circuit Bank

- **WHEN** a finite set has an associated numerical quantity, such as an area or a length, defined over tuples drawn from it, and that quantity is bounded **DO** name the specific tuple that attains the maximum of that quantity **BECAUSE** a nonempty finite set of real numbers always attains its maximum, handing you a fixed, concrete object with no limiting argument required
- **WHEN** a covering object is required to have a fixed multiple k^2 (or k^3, and so on) of an extremal object's size found inside the same configuration **DO** construct the object scaled outward by linear factor k from the extremal one, via the parallel supporting hyperplane through each of its vertices **BECAUSE** similarity multiplies size by the appropriate power of the linear scale factor, and this is the classical construction realising exactly that scaling

- **WHEN** two triangles (or simplices) share a common base (or facet) and you need to compare their areas (or volumes) **DO** convert the comparison into a comparison of perpendicular distances from the apex to that shared base **BECAUSE** area or volume with a fixed base is proportional to height alone, so an order on areas is exactly an order on heights
- **WHEN** you have established that a point lies in several half planes, one bounded by each side of a given triangle and containing that triangle's interior **DO** conclude the point lies inside the closed triangular region itself **BECAUSE** a triangle is precisely the intersection of the half planes determined by its own three sides
- **WHEN** an extremal argument is built on a nondegenerate object, such as a genuine triangle, that might fail to exist for some inputs, such as a collinear point set **DO** isolate and dispatch that degenerate case with a direct, separate construction before relying on the extremal object **BECAUSE** an extremal principle argument silently assumes the extremum is attained by a nondegenerate instance, and a full proof must check that assumption rather than leave it implicit

Backtesting

1. First useful cue: the word finite licenses choosing a maximal area triangle among triples of S with no further justification.
2. Method the expression suggested: since the target area 4 is 2^2 times the bound 1, double the extremal triangle outward via the parallel line through each of its vertices.
3. Where you would hesitate, and the thought to reinforce: proving containment for an arbitrary point of S, not just the three chosen vertices; reinforce by converting each of the three area inequalities (against a shared side of the extremal triangle) into a distance inequality, then into a half plane membership bounded by the matching side of the doubled triangle.
4. The trap the question sets: assuming a direct convex hull or diameter argument can reach the exact constant 4, or forgetting that a fully collinear S has no genuine triangle to extremise at all.
5. What makes you recognise it faster next time: reading finite plus bounded quantity as an instant cue for the extremal principle, and covering by k^2 times an extremal size as an instant cue for a linear-k doubling construction.
6. The reusable decision rule: extremise, double outward by the required linear scale factor, certify containment side by side using shared-base distance comparisons converted to half planes, then separately confirm the extremal object cannot degenerate for the inputs the problem allows.

Mirror drills

1. If S consists of exactly the three vertices of a single triangle of area 1, does the construction still give a covering triangle of area 4, and is that the smallest triangle covering S? *Answer:* the construction still produces the anticomplementary triangle of area 4, but the triangle ABC itself, area 1, already covers S since S is exactly its vertex set; the bound of 4 is far from tight when S is small, and only accumulates real content once many points of S press against the extremal triangle from several directions.
2. Suppose instead every triangle from S has area at most T for some constant $T > 0$, not necessarily 1. What area of covering triangle does the identical construction produce? *Answer:* $4T$, by the same argument verbatim; the scaling factor 4 is forced by the doubling construction alone and does not depend on the specific numerical bound.

3. Generalise the construction, not the numeric answer, to a finite set of points in three dimensional space where every tetrahedron among them has volume at most 1: what volume results from drawing, through each vertex of the maximal volume tetrahedron, the plane parallel to the opposite face? *Answer:* 27, not 8. Checking on the standard tetrahedron with vertices at the origin and the three unit basis points confirms the scaling factor in k dimensions is k^k, here $3^3 = 27$, because a tetrahedron's edge midpoints form an octahedron rather than a smaller tetrahedron, so the two dimensional picture of ABC as a medial triangle does not carry over unchanged.

10.10 B4. Expected determinant of a random skew matrix

Let A be a $2n \times 2n$ matrix, with entries chosen independently at random. Every entry is chosen to be 0 or 1, each with probability $1/2$. Find the expected value of $\det(A - A^t)$ (as a function of n), where A^t is the transpose of A.

Read it out loud

This is a probability question wearing a linear algebra coat. Read the object before you read the words: you are asked for the expected value of a determinant, and the matrix inside the determinant is not A but $A - A^t$. The moment you see $A - A^t$ you say the word out loud: skew-symmetric. That is not decoration, that is the entire question. A skew-symmetric matrix of even size has determinant equal to the square of its Pfaffian, so the answer will be a positive number, and the combinatorial objects that eventually appear will be perfect matchings rather than general permutations. You do not need the Pfaffian to solve this, but knowing it tells you in advance what the shape of the answer must be, and that is worth ten seconds of your time.

The fields wired in are three. Matrices and determinants, specifically the Leibniz expansion, which is the only description of a determinant that plays well with expectation. Probability, specifically linearity of expectation, independence, and the fact that a mean zero factor annihilates a monomial. Combinatorics, specifically counting fixed point free involutions of a $2n$ element set.

What is really being tested is whether you know which of your tools survives contact with a determinant. Expectation is linear. The determinant is not. So $\mathrm{E}[\det M] \neq \det \mathrm{E}[M]$ in general, and the candidate who writes $\mathrm{E}[A - A^t] = 0$ and concludes that the answer is 0 has thrown the paper away in a single line. That is the trap, and it is a good one, because $\mathrm{E}[A - A^t] = 0$ is true. It is the next step that is false.

Inside thirty seconds you should have written three lines. Set $M := A - A^t$. For $i < j$, $M_{ij} = A_{ij} - A_{ji}$ takes the values $-1, 0, 1$ with probabilities $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$, and the $\binom{2n}{2}$ variables M_{ij} with $i < j$ are independent because they are built from pairwise disjoint pairs of entries of A. Finally $M_{ii} = 0$ and $M_{ji} = -M_{ij}$. The $4n^2$ coin flips have collapsed into $n(2n-1)$ independent three valued variables, and everything after that is bookkeeping.

Approaches, ranked worst to best

1. **Expectation of the determinant equals the determinant of the expectation.** *Cue:* the matrix $A - A^t$ has mean zero. *Circuit:* linearity, misapplied. This is not an approach, it is the trap, and it returns the wrong answer 0. The determinant is a polynomial of degree $2n$ in the entries, and expectation commutes with polynomials only through their monomials, never through the polynomial as a whole. Named here so that the reflex to reject it is trained rather than assumed.

2. **Compute small cases and guess.** *Cue:* the answer is requested as a function of n. *Circuit:* pattern completion. You get $1/2$ for $n = 1$ and, after some grinding, $3/4$ for $n = 2$. Two data points do not identify $(2n)!/(4^n n!)$, and a guess earns nothing. Useful only as a check at the end, which is exactly where you should use it.

3. **Cofactor expansion or eigenvalue structure.** *Cue:* a determinant of a structured matrix. *Circuit:* expand along a row, or diagonalise. A skew-symmetric real matrix has eigenvalues in conjugate pairs $\pm i\lambda$, so $\det M = \prod \lambda_k^2 \geq 0$, which is genuine structure. But the eigenvalues of a random integer matrix have no tractable joint distribution, and cofactor expansion produces a recursion in objects that are not independent. Both stall.

4. **Pfaffian.** *Cue:* even sized skew-symmetric matrix. *Circuit:* $\det M = \mathrm{Pf}(M)^2$, and Pf is a signed sum over perfect matchings. Fast and structurally honest: square the Pfaffian, note that a cross

term between two distinct matchings contains some M_{ij} to the first power and therefore has mean zero, and the diagonal terms each contribute $\mathrm{E}[M_{ij}^2]^n = 2^{-n}$. The only cost is that you must cite the identity $\det = \mathrm{Pf}^2$, which is a real theorem and not something you can prove in the margin.

5. **Leibniz expansion with cycle decomposition.** *Cue:* expected value of a determinant with independent entries. *Circuit:* write det as a sum of monomials, take expectations monomial by monomial, and kill every monomial containing a mean zero factor to an odd power. Best. It is completely self-contained, it needs no theorem beyond the definition of the determinant and independence, and it produces the matching count as an output rather than requiring it as an input. This is the route below.

The ideal solution

Write $M = A - A^t$, so $M_{ii} = 0$ and $M_{ji} = -M_{ij}$ for all i, j. For $i < j$ the variable $M_{ij} = A_{ij} - A_{ji}$ is a function of the two entries A_{ij} and A_{ji} alone, and distinct unordered pairs $\{i, j\}$ use disjoint pairs of entries of A, so

$$\{M_{ij} : 1 \le i < j \le 2n\}$$

is an independent family. Each has the distribution

$$\mathrm{P}(M_{ij} = -1) = \tfrac{1}{4}, \quad \mathrm{P}(M_{ij} = 0) = \tfrac{1}{2}, \quad \mathrm{P}(M_{ij} = 1) = \tfrac{1}{4},$$

so that

$$\mathrm{E}[M_{ij}] = 0, \qquad \mathrm{E}[M_{ij}^2] = \tfrac{1}{4} + \tfrac{1}{4} = \tfrac{1}{2}.$$

By the Leibniz formula,

$$\det M = \sum_{\sigma \in S_{2n}} \operatorname{sgn}(\sigma) \prod_{i=1}^{2n} M_{i\sigma(i)}.$$

This is a finite sum of random variables, each bounded, so expectation passes through it:

$$\mathrm{E}[\det M] = \sum_{\sigma \in S_{2n}} \operatorname{sgn}(\sigma) E_\sigma, \qquad E_\sigma := \mathrm{E}\left[\prod_{i=1}^{2n} M_{i\sigma(i)}\right].$$

Fix σ and decompose $\{1, \dots, 2n\}$ into the orbits (cycles) of σ. The product splits accordingly, one factor group per cycle, and cycles occupy disjoint index sets, so the unordered pairs $\{i, \sigma(i)\}$ arising in different cycles are distinct.

Claim 10.10.1. *$E_\sigma = 0$ unless σ is a product of n disjoint transpositions.*

Proof. If σ has a fixed point i then the product contains the factor $M_{ii} = 0$, so the product is identically zero and $E_\sigma = 0$.

Suppose σ has a cycle $C = (c_1\, c_2 \cdots c_m)$ of length $m \ge 3$. The factors contributed by C are $M_{c_1c_2}, M_{c_2c_3}, \dots, M_{c_mc_1}$, and the m unordered pairs $\{c_1, c_2\}, \dots, \{c_m, c_1\}$ are distinct: if $\{i, \sigma(i)\} = \{j, \sigma(j)\}$ with $i \ne j$ then $\sigma(i) = j$ and $\sigma(j) = i$, which puts i and j in a cycle of length 2. Hence every one of the $2n$ factors of the whole product is $\pm M_{ij}$ for some $i < j$, and the variable $M_{c_1c_2}$ occurs in exactly one factor overall. Grouping the product as $\pm M_{c_1c_2}$ times a product of factors involving other pairs, independence gives

$$E_\sigma = \pm\, \mathrm{E}[M_{c_1c_2}] \cdot \mathrm{E}[\text{rest}] = 0,$$

since $\mathrm{E}[M_{c_1c_2}] = 0$. □

So only the fixed point free involutions survive. Let σ be such a permutation, with transpositions $\{a_1, b_1\}, \ldots, \{a_n, b_n\}$ partitioning $\{1, \ldots, 2n\}$. Then

$$\prod_{i=1}^{2n} M_{i\sigma(i)} = \prod_{\ell=1}^{n} M_{a_\ell b_\ell} M_{b_\ell a_\ell} = \prod_{\ell=1}^{n} \left(-M_{a_\ell b_\ell}^2\right) = (-1)^n \prod_{\ell=1}^{n} M_{a_\ell b_\ell}^2,$$

and $\operatorname{sgn}(\sigma) = (-1)^n$ because σ is a product of n transpositions. The two signs cancel, so

$$\operatorname{sgn}(\sigma) \prod_{i=1}^{2n} M_{i\sigma(i)} = \prod_{\ell=1}^{n} M_{a_\ell b_\ell}^2,$$

which is nonnegative, and by independence of the n distinct variables involved,

$$\operatorname{sgn}(\sigma) E_\sigma = \prod_{\ell=1}^{n} \mathrm{E}\left[M_{a_\ell b_\ell}^2\right] = \left(\tfrac{1}{2}\right)^n = 2^{-n}.$$

Every surviving permutation contributes the same amount, and no cancellation occurs.

It remains to count the fixed point free involutions of a $2n$ element set, equivalently the perfect matchings. Choose the partner of the smallest unused element in $2n-1$ ways, then repeat on the remaining $2n-2$ elements: the count satisfies $P_{2n} = (2n-1)P_{2n-2}$ with $P_0 = 1$, hence

$$P_{2n} = (2n-1)!! = \frac{(2n)!}{2^n\, n!}.$$

Therefore

$$\mathrm{E}[\det(A - A^t)] = \frac{(2n)!}{2^n n!} \cdot 2^{-n} = \boxed{\frac{(2n)!}{4^n\, n!}}.$$

Two checks. For $n = 1$, $M = \begin{pmatrix} 0 & M_{12} \\ -M_{12} & 0 \end{pmatrix}$ has $\det M = M_{12}^2$ with expectation $\frac{1}{2}$, and the formula gives $2!/(4 \cdot 1) = \frac{1}{2}$. For $n = 2$ the formula gives $24/32 = 3/4$, which is 3 matchings each worth $\frac{1}{4}$. Finally, the answer is positive, as it must be: $\det M = \operatorname{Pf}(M)^2 \geq 0$ pointwise, so its expectation cannot be zero unless $\det M$ vanishes almost surely, which it does not.

Layer 1: WHAT

Block 1. Name the object: $M = A - A^t$ is skew-symmetric, its strictly upper triangular entries are independent, each taking $-1, 0, 1$ with probabilities $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$; the rest of the matrix is determined by them.

Block 2. Determinant into monomials: use the Leibniz formula and push expectation inside the finite sum, so the problem becomes one expectation per permutation.

Block 3. Decompose each permutation into cycles, because that is what tells you which variables appear and to what power.

Block 4. Kill everything: fixed points give a zero factor; cycles of length at least 3 give a mean zero factor to the first power. Only fixed point free involutions survive.

Block 5. Evaluate a survivor: the sign $(-1)^n$ of the involution cancels the sign from $M_{ba} = -M_{ab}$, leaving a product of squares with expectation 2^{-n}.

Block 6. Count the survivors: $(2n-1)!! = (2n)!/(2^n n!)$, multiply, done.

Juxtaposition. The nominal route reads "expected value of a determinant" and reaches for the matrix as a whole: take expectations of the entries first, or find the distribution of $\det M$, or expand along a row. All three treat det as a function of a matrix. The ideal route refuses to see a matrix at all: it sees a degree $2n$ polynomial in $n(2n-1)$ independent scalars, and it knows that expectation is a linear functional on polynomials. The bifurcation point is the line where you write the Leibniz sum instead of writing $\mathrm{E}[M]$. Everything downstream is forced.

Layer 2: HOW

Block 1. **Cue.** The expression $A - A^t$, and nothing else in the problem. A difference of a matrix and its transpose has exactly one name. **Why natural.** Because the definition of skew-symmetry is $M^t = -M$, and $(A - A^t)^t = A^t - A = -(A - A^t)$ is a one line verification. Naming the structure is worth doing before any computation because it fixes the diagonal at zero and halves the number of random variables. The independence claim is also definition level: M_{ij} for $i < j$ is a function of A_{ij}, A_{ji} only, and those pairs are disjoint across distinct $\{i, j\}$, so functions of disjoint independent families are independent. **Link to the next block.** Once you know the entries are independent with mean zero, the only question left is which monomials in those entries survive the expectation. That question is asked and answered by the determinant expansion.

Block 2. **Cue.** "Expected value of det". Expectation is linear; the determinant is not. **Why natural.** The Leibniz formula is the unique description of a determinant as a linear combination of monomials in the entries. Since expectation is linear and the sum is finite, E passes through the sum with no convergence concerns at all. This is the only structurally available move: any route that keeps det intact has to know the distribution of $\det M$, which is far harder than the question asked. **Link to the next block.** Now every term is E of a product of $2n$ entries, so you must understand which entries appear in a given term, and how often.

Block 3. **Cue.** A product indexed by a permutation. Permutations decompose into cycles, always, and that decomposition is what controls repetitions. **Why natural.** The pair $\{i, \sigma(i)\}$ repeats only when σ swaps i and $\sigma(i)$. That is a statement about cycles of length 2, so the cycle type is precisely the right invariant to sort the terms by. This is not a guess: you are looking for repeated variables, and repetition of the unordered pair $\{i, \sigma(i)\}$ is equivalent to a transposition. **Link to the next block.** Sorting by cycle type immediately splits the permutations into three classes, and two of them die.

Block 4. **Cue.** A factor with mean zero multiplied by independent stuff. **Why natural.** If X is independent of Y and $\mathrm{E}[X] = 0$ then $\mathrm{E}[XY] = \mathrm{E}[X]\,\mathrm{E}[Y] = 0$. In a cycle of length at least 3, the variable attached to any one edge of that cycle appears exactly once in the entire $2n$ fold product, so it can be split off as such an X. A fixed point is even cheaper: the factor is the diagonal entry, which is identically zero. **Link to the next block.** Only fixed point free involutions remain, and they are exactly the perfect matchings of $\{1, \ldots, 2n\}$, which is the combinatorial object you predicted the instant you said "skew-symmetric".

Block 5. **Cue.** Two sources of sign in the same term, one from $\mathrm{sgn}(\sigma)$ and one from $M_{ba} = -M_{ab}$. **Why natural.** Never leave two signs unresolved: multiply them out and see what happens. An involution with n transpositions has sign $(-1)^n$, and the product of the n pairs $M_{ab}M_{ba}$ carries $(-1)^n$. They cancel, so every surviving term is a nonnegative random variable with the same expectation. That is the structural reason there is no cancellation in the final count, and it is also a consistency check against $\det M = \mathrm{Pf}(M)^2 \geq 0$. **Link to the next block.** With all survivors equal to 2^{-n} in expectation, the answer is 2^{-n} times the number of survivors.

Block 6. **Cue.** "Count the perfect matchings of a $2n$ element set." **Why natural.** Match the smallest unused element to any of the remaining $2n - 1$ elements and recurse. That is a recursion because the object it produces is the same object one size down, which is the definition level reason the double factorial appears. **Link to the next block.** Multiply, then check $n = 1$ by hand.

Layer 3: WHY

Three truths make this route ideal, and they are the same three truths behind almost every expected value of a polynomial statistic of a random matrix.

First, expectation is a linear functional on the space of polynomials in the entries, so any polynomial statistic is computable as soon as you know the joint moments of the entries. The determinant is a polynomial of degree $2n$, homogeneous, multilinear in the rows. Therefore the correct representation is the monomial one, and the Leibniz expansion is not an implementation detail, it is the entire method.

Second, independence converts E of a product into a product of Es, but only when the factors involve disjoint sets of variables. So the work in any such problem is combinatorial bookkeeping: which monomials use which variables, and to what power. Here the bookkeeping is the cycle decomposition.

Third, a mean zero variable annihilates any monomial in which it appears to an odd power. This is why the answer is supported on perfect matchings and nowhere else, and it is why the symmetry of the distribution of M_{ij} never actually mattered: all you used was $\mathrm{E}[M_{ij}] = 0$ and $\mathrm{E}[M_{ij}^2] = \frac{1}{2}$.

That last observation is the discernment mechanism, and it generalises with no extra work. If M is any $2n \times 2n$ skew-symmetric random matrix whose entries M_{ij}, $i < j$, are independent with $\mathrm{E}[M_{ij}] = 0$ and $\mathrm{E}[M_{ij}^2] = s_{ij}$, then exactly the same argument gives

$$\mathrm{E}[\det M] = \sum_{\text{perfect matchings } \{a_\ell, b_\ell\}} \prod_{\ell=1}^{n} s_{a_\ell b_\ell},$$

the hafnian of the second moment matrix. The problem as set is the special case $s_{ij} \equiv \frac{1}{2}$. Recognising that your solution never used the specific distribution is how you convert one Putnam problem into a permanent piece of machinery.

How to train the discernment: whenever a random object is fed into an algebraic function, ask which representation of that function is linear in the randomness. For determinants it is the Leibniz sum. For traces of powers it is the sum over closed walks. For characteristic polynomials it is the sum over principal minors. In each case the expectation collapses onto the sub-family of index patterns in which every variable appears to an even power, and that sub-family is always some kind of matching or pairing. That principle, that mean zero variables force pairings, is the same principle that produces Wick's theorem and the Catalan numbers in random matrix theory, and the Putnam examiners were nudging at it by choosing $A - A^t$ rather than A, and by choosing the even size $2n$ rather than a general size.

Revolving concepts

Skew-symmetric matrix. Definition: $M^t = -M$, forcing $M_{ii} = 0$ over a field of characteristic not 2. Orbiting facts: over $\mathbb{R}$ the eigenvalues come in pairs $\pm i\lambda$; the rank is always even; an odd sized skew-symmetric matrix has determinant 0, since $\det M = \det(-M^t) = (-1)^{2m+1} \det M$. See Book 1, field: Matrices and determinants, topic: symmetry classes of matrices.

Pfaffian. For M skew-symmetric of size $2n$, $\mathrm{Pf}(M) = \sum_\alpha \mathrm{sgn}(\alpha) \prod_{\ell=1}^{n} M_{i_\ell j_\ell}$, summed over perfect matchings α, and $\det M = \mathrm{Pf}(M)^2$. This is the reason the answer is positive and the reason matchings, not permutations, index the surviving terms. You did not need it, but it predicts the shape of the answer before you compute anything. See Book 1, field: Matrices and determinants, topic: Pfaffian.

Leibniz expansion. $\det M = \sum_\sigma \mathrm{sgn}(\sigma) \prod_i M_{i\sigma(i)}$, a sum of $(2n)!$ signed monomials. Orbiting facts: sgn is a homomorphism to $\{\pm 1\}$; a permutation with cycle type having c cycles on N letters has sign $(-1)^{N-c}$; a product of n disjoint transpositions has sign $(-1)^n$. See Book 1, field: Matrices and determinants, topic: Leibniz formula.

Mean zero kills odd powers. If $X_1, \dots, X_r$ are independent and $\mathrm{E}[X_1] = 0$, then $\mathrm{E}[X_1^{e_1} \cdots X_r^{e_r}] = 0$ whenever $e_1 = 1$. More generally the monomial survives only if each variable appears with an exponent whose moment is nonzero. See Book 1, field: Probability, topic: moments and independence.

The matching recursion, generalised. Let P_{2n} be the number of ways to partition $2n$ labelled objects into unordered pairs. Pair the smallest element with one of the other $2n-1$, then recurse: $P_{2n} = (2n-1)P_{2n-2}$, $P_0 = 1$, hence $P_{2n} = (2n-1)!! = (2n)!/(2^n n!)$. The general form of this recursion is $P_N = (N-1)P_{N-2}$, and the general solution method is the same every time: iterate and read off a product, then convert the product of odd numbers into $(2n)!/(2^n n!)$ by inserting the missing even factors. The same recursion counts: the number of ways to pair up $2n$ people for a chess round, the number of chord diagrams on $2n$ points, the $2n$-th moment of a standard normal divided by nothing at all (indeed $\mathrm{E}[Z^{2n}] = (2n-1)!!$, which is the same pairing count, for exactly the same mean zero reason). Whenever you see $(2n-1)!!$ in an answer, look for the pairing that produced it. See Book 1, field: Combinatorics, topic: perfect matchings and double factorials.

Cycle decomposition. Every permutation is a product of disjoint cycles, uniquely up to order; the orbits partition the index set. Load bearing here because disjoint orbits mean disjoint variable sets, which is what licences the factorisation of the expectation. See Book 1, field: Abstract algebra, topic: symmetric group.

Circuit Bank

- **WHEN** you are asked for the expected value of a determinant, a permanent, or any polynomial in random entries **DO** expand into monomials (Leibniz for det) and take expectations term by term **BECAUSE** expectation is linear on monomials but does not commute with the polynomial as a whole; $\mathrm{E}[\det M] \neq \det \mathrm{E}[M]$.
- **WHEN** a monomial contains an independent factor with mean zero to the first power **DO** delete the whole monomial **BECAUSE** $\mathrm{E}[XY] = \mathrm{E}[X]\,\mathrm{E}[Y] = 0$ when X and Y are independent and $\mathrm{E}[X] = 0$; this is what forces surviving terms to be pairings.
- **WHEN** you see $A - A^t$, or any even sized skew-symmetric matrix **DO** write down that the diagonal is zero, that the strictly upper entries determine everything, and that $\det = \mathrm{Pf}^2 \geq 0$ **BECAUSE** skew-symmetry is a definition level constraint that halves the number of free variables and predicts a nonnegative answer.
- **WHEN** a permutation indexes a product and you need to know which variables repeat **DO** decompose into cycles first **BECAUSE** repetition of the unordered pair $\{i, \sigma(i)\}$ happens exactly on 2-cycles, so cycle type is the sorting invariant.
- **WHEN** your final answer is a formula in n **DO** test $n = 1$ by direct computation before you write the last line **BECAUSE** a single arithmetic slip in a factorial count is invisible in the algebra and instantly visible at $n = 1$.

Backtesting

1. First useful cue: the expression $A - A^t$, which names the matrix skew-symmetric and immediately gives a zero diagonal, an independent family of $\binom{2n}{2}$ entries, and the prediction that perfect matchings will index the answer.
2. Method the expression suggested: Leibniz expansion of the determinant, expectation term by term, cycle decomposition of each permutation to see which terms survive.
3. Where you would hesitate, and the thought to reinforce: at the claim that a cycle of length at least 3 contributes zero. Reinforce it by asking not "is this zero" but "does some variable appear here exactly once", because that is the only question the mean zero property can answer.

4. The trap the question sets: $\mathrm{E}[A - A^t] = 0$ tempting you into $\mathrm{E}[\det(A - A^t)] = \det(0) = 0$. The determinant is degree $2n$, not degree 1; the mean of a square is not the square of the mean.
5. What makes you recognise it faster next time: the pairing signature. If a problem gives you independent mean zero randomness and asks for the expectation of a polynomial, the answer is always a sum over pairings of the variables, with weights the second moments.
6. The reusable decision rule: expectation of a polynomial statistic equals the sum over monomials whose every variable has even exponent; find that sub-family combinatorially, weight it by the moments, count it.

Mirror drills

1. Let B be an $(2n+1) \times (2n+1)$ matrix with entries chosen as in the problem. Find $\mathrm{E}[\det(B - B^t)]$. *Answer:* 0, and not because of any expectation argument: an odd sized skew-symmetric matrix satisfies $\det M = \det(-M^t) = -\det M$, so $\det M = 0$ pointwise.
2. Let the entries of A be independent and uniform on $\{0, 1, \ldots, q-1\}$ instead of on $\{0, 1\}$. Find $\mathrm{E}[\det(A - A^t)]$. *Answer:* the proof used only $\mathrm{E}[M_{ij}] = 0$ and $\mathrm{E}[M_{ij}^2] = s$, and here $s = \mathrm{E}[(A_{ij} - A_{ji})^2] = 2\,\mathrm{Var}(A_{ij}) = (q^2-1)/6$, so the answer is $\frac{(2n)!}{2^n n!}\left(\frac{q^2-1}{6}\right)^n$.
3. Let A be $n \times n$ with independent entries of mean 0 and variance 1. Show that $\mathrm{E}[\det A] = 0$ and $\mathrm{E}[(\det A)^2] = n!$. *Answer:* in the single expansion every monomial has some entry to the first power; in the squared expansion the pair (σ, τ) survives only when $\sigma = \tau$, contributing $\mathrm{sgn}(\sigma)^2 \prod \mathrm{E}[A_{i\sigma(i)}^2] = 1$, and there are $n!$ such pairs.

10.11 B5. Multiplicative sandwich forces a pure power

Find all functions f from the interval $(1,\infty)$ to $(1,\infty)$ with the following property: if $x, y \in (1,\infty)$ and $x^2 \le y \le x^3$, then $(f(x))^2 \le f(y) \le (f(x))^3$.

Read it out loud

This is a functional equation that is not an equation. It is a two sided inequality, a sandwich: the hypothesis sandwiches y between two powers of x, and the conclusion sandwiches $f(y)$ between the same two powers of $f(x)$. Say that out loud, because it is the whole structure: the numbers 2 and 3 appear on the input side and on the output side, unchanged. The map f is being asked to transport a multiplicative window of a fixed shape to a multiplicative window of the same shape.

Now notice what is missing. There is no continuity, no monotonicity, no measurability, nothing. So you cannot reach for any theorem about regular solutions of Cauchy's equation; the rigidity has to be manufactured out of the inequality itself. That is the real content of the problem, and it is why this sits at B5 and not at B2.

The fields wired in: functional equations (the normal form game), real analysis (chaining inequalities, no regularity assumed), number theory (the multiplicative independence of 2 and 3), and equidistribution (density of $\{a\theta - b\}$ for irrational θ). The answer will be $f(x) = x^c$ for $c > 0$, and you should write that down in the first ten seconds after checking it works: for $x > 1$ and $x^2 \le y \le x^3$, $x^{2c} \le y^c \le x^{3c}$, which is exactly $f(x)^2 \le f(y) \le f(x)^3$.

Within thirty seconds you should also have written the substitution $x = e^u$. Powers of x become multiples of u, so with $g(u) = \log f(e^u)$ the condition reads: $2u \le v \le 3u$ implies $2g(u) \le g(v) \le 3g(u)$. Look at what that did. It did not solve the problem; it reproduced it. The output side is still multiplicative. So do it again. That doubling of the same move is the cue the examiners planted by using the same two constants on both sides.

Approaches, ranked worst to best

1. **Guess the family and verify it.** *Cue:* "find all functions" with an obvious candidate. *Circuit:* pattern recognition. Verifying that x^c works takes one line and earns close to nothing, because the entire difficulty is the word "all". Do it anyway, but only as target practice: it tells you the exact statement you must eventually prove.

2. **Assume continuity or monotonicity, then use Cauchy's equation.** *Cue:* a functional condition on $\mathbb{R}$. *Circuit:* regularity plus additivity gives linearity. This is the standard reflex and it is wrong here, because no regularity is given and none is derivable cheaply. Writing "we may assume f is continuous" is an unearned hypothesis and a grader will take the marks.

3. **Iterate the sandwich on powers of the argument.** *Cue:* $y = x^2$ and $y = x^3$ are the extreme allowed choices. *Circuit:* apply the hypothesis repeatedly and stack the inequalities. You learn a great deal about f on the sparse set of exponents $2^a 3^b$, and nothing at all about f between them. It stalls because the information lives on a multiplicative semigroup with no obvious way to interpolate.

4. **One logarithm, then attack g directly.** *Cue:* powers on both sides. *Circuit:* log turns exponents into coefficients. A real improvement: $2u \le v \le 3u \Rightarrow 2g(u) \le g(v) \le 3g(u)$. But the constraint is still "multiply by something between 2 and 3", and multiplications do not chain into a useful additive bookkeeping. Most solvers stop here and grind.

5. **Two logarithms, then chain displacements.** *Cue:* the constraint is a sandwich of the same shape on both sides. *Circuit:* convert the constraint into "a shift of size in $[\log 2, \log 3]$ maps to a shift of size in $[\log 2, \log 3]$", so that steps compose additively. Best. With $h(s) = \log g(e^s)$ the condition becomes a translation statement, forward and backward steps are both available,

and the reachable displacements after a forward and b backward steps form the interval $[a\log 2 - b\log 3,\ a\log 3 - b\log 2]$, whose endpoints are dense in $\mathbb{R}$ because $\log 3/\log 2$ is irrational. That density is the rigidity, and it is where the numbers 2 and 3 finally earn their keep.

The ideal solution

The answer: exactly the functions $f(x) = x^c$ with $c > 0$.

These work: if $c > 0$ then $x^c > 1$ for $x > 1$, and $x^2 \le y \le x^3$ gives $x^{2c} \le y^c \le x^{3c}$, that is, $f(x)^2 \le f(y) \le f(x)^3$. Conversely, suppose f has the stated property. Throughout write $L_2 = \log 2$ and $L_3 = \log 3$.

Step 1: two changes of variable. Define $g : (0,\infty) \to (0,\infty)$ by $g(u) = \log f(e^u)$; this is positive because $f > 1$. Putting $x = e^u$ and $y = e^v$ with $u > 0$, the hypothesis $x^2 \le y \le x^3$ is exactly $2u \le v \le 3u$ (and then $v > 0$, so $y > 1$ automatically), while the conclusion is exactly $2g(u) \le g(v) \le 3g(u)$.

Define $h : \mathbb{R} \to \mathbb{R}$ by $h(s) = \log g(e^s)$; this is defined because $g > 0$. Putting $u = e^s$ and $v = e^t$, the condition $2u \le v \le 3u$ is exactly $s + L_2 \le t \le s + L_3$, and the conclusion $2g(u) \le g(v) \le 3g(u)$ is exactly $h(s) + L_2 \le h(t) \le h(s) + L_3$. So the hypothesis has become:

$$\text{(P)} \qquad t - s \in [L_2, L_3] \implies h(t) - h(s) \in [L_2, L_3].$$

Interchanging the roles of s and t in (P) gives its mirror image for free:

$$\text{(P}') \qquad t - s \in [-L_3, -L_2] \implies h(t) - h(s) \in [-L_3, -L_2].$$

It suffices to prove that $h(t) - h(s) = t - s$ for all s, t. Indeed, then $h(s) = s + \gamma$ for a constant γ, so $g(e^s) = e^{s+\gamma}$, that is $g(u) = cu$ for all $u > 0$ with $c = e^\gamma > 0$; then $\log f(e^u) = cu$, so $f(e^u) = (e^u)^c$ and $f(x) = x^c$ for all $x > 1$.

Step 2: chaining steps. For nonnegative integers a, b put

$$I_{a,b} = [\,aL_2 - bL_3,\ aL_3 - bL_2\,],$$

a nonempty interval since $(a+b)L_2 \le (a+b)L_3$, with $I_{0,0} = \{0\}$, $I_{1,0} = [L_2, L_3]$ and $I_{0,1} = [-L_3, -L_2]$. Note $I_{a-1,b} + I_{1,0} = I_{a,b}$ and $I_{a,b-1} + I_{0,1} = I_{a,b}$, where the sum of intervals is taken endpoint by endpoint.

Claim 10.11.1. *For all nonnegative integers a, b and all $s, t \in \mathbb{R}$: if $t - s \in I_{a,b}$ then $h(t) - h(s) \in I_{a,b}$.*

Proof. Induct on $a + b$. If $a + b = 0$ then $t = s$ and the statement is trivial. If $a + b = 1$ the statement is (P) or (P′).

Let $a + b = N \ge 2$ and assume the claim for all smaller values of $a + b$. Suppose first that $a > 0$, and let $\tau = t - s \in I_{a,b}$. At the left endpoint $\tau = aL_2 - bL_3$ the decomposition of τ into a forward steps and b backward steps is forced: every forward step is exactly L_2 and every backward step exactly L_3, so deleting one forward step leaves $(a-1)L_2 - bL_3$, the left endpoint of $I_{a-1,b}$. At the right endpoint the same deletion leaves the right endpoint of $I_{a-1,b}$. So interpolate affinely between these two forced choices: define

$$j(\tau) = \frac{(N-1)\tau - b(L_2 + L_3)}{N}.$$

A direct computation gives $j(aL_2 - bL_3) = (a-1)L_2 - bL_3$ and $j(aL_3 - bL_2) = (a-1)L_3 - bL_2$, and j is affine with nonnegative slope $(N-1)/N$, so j maps $I_{a,b}$ into $I_{a-1,b}$. Moreover

$$\tau - j(\tau) = \frac{\tau + b(L_2 + L_3)}{N}$$

is affine and increasing in τ, with value L_2 at the left endpoint of $I_{a,b}$ and L_3 at the right endpoint, so $\tau - j(\tau) \in [L_2, L_3]$ for every $\tau \in I_{a,b}$.

Now set $r = s + j(\tau)$. Then $r - s = j(\tau) \in I_{a-1,b}$, so the inductive hypothesis gives $h(r) - h(s) \in I_{a-1,b}$; and $t - r = \tau - j(\tau) \in [L_2, L_3]$, so (P) gives $h(t) - h(r) \in I_{1,0}$. Adding,

$$h(t) - h(s) \in I_{a-1,b} + I_{1,0} = I_{a,b}.$$

If instead $b > 0$, use $\tilde{j}(\tau) = \big((N-1)\tau + a(L_2 + L_3)\big)/N$, which sends the endpoints of $I_{a,b}$ to the corresponding endpoints of $I_{a,b-1}$ and satisfies $\tau - \tilde{j}(\tau) = (\tau - a(L_2 + L_3))/N \in [-L_3, -L_2]$ for $\tau \in I_{a,b}$; the same two step argument with (P′) in place of (P) finishes the induction. □

Step 3: the arithmetic input.

Lemma 10.11.2. *$\theta = L_3/L_2$ is irrational, and for every $T \in \mathbb{R}$, every $\varepsilon > 0$ and every N, there exist integers $a, b > N$ with $T < aL_3 - bL_2 < T + \varepsilon$. The same holds with the roles of L_2 and L_3 exchanged, and also with the approximation from below, $T - \varepsilon < aL_2 - bL_3 < T$.*

Proof. If $L_3/L_2 = p/q$ with positive integers p, q then $3^q = 2^p$, contradicting unique factorisation. For the density, let $\|\cdot\|$ denote distance to the nearest integer. Given M, the $M+1$ fractional parts of $0, \theta, \dots, M\theta$ lie in $[0,1)$, so two of them lie within $1/M$ of each other, giving a positive integer $q \le M$ with $\|q\theta\| < 1/M$, and $\|q\theta\| \neq 0$ by irrationality. The points $\{mq\theta\}$, $m = 1, 2, \dots$, therefore step around the circle $\mathbb{R}/\mathbb{Z}$ in increments of size less than $1/M$ and so meet every interval of length $1/M$; taking m as large as we please, we may take $a = mq$ as large as we please. Writing $aL_3 - bL_2 = L_2(a\theta - b)$ and choosing b to be the appropriate integer near $a\theta$, we obtain $a\theta - b$ in any prescribed interval of length ε/L_2, with $b \to \infty$ as $a \to \infty$; since M was arbitrary the approximation can be made from above or from below as required. The exchanged version is the same argument with $\theta' = L_2/L_3$, also irrational. □

Step 4: squeeze. Fix $s < t$ and set $T = t - s > 0$. Let $\varepsilon > 0$.

Upper bound. Choose $a, b > N$ with $T < aL_3 - bL_2 < T + \varepsilon$. Since

$$aL_2 - bL_3 = \frac{L_2}{L_3}\,(aL_3 - bL_2) \;-\; b\,\frac{L_3^2 - L_2^2}{L_3},$$

the first term is bounded while the second tends to $-\infty$ as $b \to \infty$, so for N large enough $aL_2 - bL_3 < T$. Hence $T \in I_{a,b}$, and the Claim gives $h(t) - h(s) \le aL_3 - bL_2 < T + \varepsilon$.

Lower bound. Choose $a, b > N$ with $T - \varepsilon < aL_2 - bL_3 < T$. Since

$$aL_3 - bL_2 = \frac{L_3}{L_2}\,(aL_2 - bL_3) \;+\; b\,\frac{L_3^2 - L_2^2}{L_2} \longrightarrow +\infty,$$

for N large enough $aL_3 - bL_2 > T$, so again $T \in I_{a,b}$ and the Claim gives $h(t) - h(s) \ge aL_2 - bL_3 > T - \varepsilon$.

As $\varepsilon > 0$ was arbitrary, $h(t) - h(s) = T = t - s$ whenever $s < t$; the cases $s = t$ and $s > t$ follow trivially and by symmetry. By Step 1 this forces $f(x) = x^c$ with $c = e^\gamma > 0$, and all such f do satisfy the condition. ■

Layer 1: WHAT

Block 1. Verify that x^c, $c > 0$, works, so that the target is pinned: you must prove $h(s) - s$ is constant.

Block 2. First logarithm, $g(u) = \log f(e^u)$: powers of the argument become multiples.

Block 3. Second logarithm, $h(s) = \log g(e^s)$: multiples become shifts, and the hypothesis becomes "a displacement in $[L_2, L_3]$ maps to a displacement in $[L_2, L_3]$".

Block 4. Read the constraint backwards by swapping s and t: backward steps of size in $[-L_3, -L_2]$ are free.

Block 5. Chain a forward and b backward steps. Reachable displacements form $I_{a,b} = [aL_2 - bL_3,\ aL_3 - bL_2]$, and h maps $I_{a,b}$ displacements into $I_{a,b}$ displacements. Prove it by induction on $a + b$, with the intermediate point chosen by affine interpolation between the two forced endpoint decompositions.

Block 6. $I_{a,b}$ is long (length $(a + b)(L_3 - L_2)$), so it cannot trap anything; but its endpoints are individually dense in $\mathbb{R}$, because L_3/L_2 is irrational.

Block 7. Squeeze: tune the right endpoint just above T for the upper bound and the left endpoint just below T for the lower bound, letting the other endpoint run off harmlessly.

Block 8. Unwind both logarithms.

Juxtaposition. The nominal route is: prove f is increasing, prove f is continuous, reduce to Cauchy. It fails at the first step, since nothing forces monotonicity locally. The bifurcation point is Block 3, the second logarithm. Almost everyone takes one logarithm, because "powers become multiples" is a reflex from A level; almost nobody takes the second, because after the first the condition looks like homogeneity and invites scaling arguments. The ideal route takes the second logarithm for a stated reason: the only constraints that chain are constraints on displacements, and only after two logarithms is the constraint a statement about displacements.

Layer 2: HOW

Block 1. **Cue.** "Find all functions", with 2 and 3 appearing symmetrically in the hypothesis and the conclusion. **Why natural.** The symmetry of the exponents means the condition is invariant under conjugating f by any power map, which is the fingerprint of the family $x \mapsto x^c$. Verifying the candidate first is not optimism, it is target setting: it converts a vague "find all" into the sharp statement "prove $\log f(e^{e^s}) - s$ is constant". **Link to the next block.** With the target known, every step afterwards is judged by whether it moves the constraint closer to a statement about a single unknown function of a real variable with an additive constraint.

Block 2. **Cue.** x^2 and x^3: powers of the variable. **Why natural.** The definition of x^k as $e^{k \log x}$ says that the logarithm is the exact translation device from exponents to coefficients. This is forced, not chosen. **Link to the next block.** After one logarithm the hypothesis is additive in the input ($2u \le v \le 3u$) but still multiplicative in the output ($2g(u) \le g(v) \le 3g(u)$). The asymmetry is the signal that the job is half done.

Block 3. **Cue.** A constraint of the form "multiply by a factor in $[2, 3]$", appearing on both sides. **Why natural.** Multiplicative constraints do not compose into anything you can count; additive ones do, because displacements add. So apply the same device again: $u = e^s$ turns "multiply by a factor in $[2, 3]$" into "add a shift in $[L_2, L_3]$". The reason this is natural is not that it works, it is that a second logarithm is the only operation that turns a scaling into a translation. **Link to the next block.** Now the hypothesis says h transports $[L_2, L_3]$ displacements to $[L_2, L_3]$ displacements, and you can start walking.

Block 4. **Cue.** A statement about an unordered pair of points, phrased with s on the left. **Why natural.** Any condition of the form "if $t - s$ lies in J then $h(t) - h(s)$ lies in J" automatically yields the reflected condition on $-J$, because you may relabel which point you call s. Free information should always be extracted before you start working. **Link to the next block.** Forward and backward steps together generate far more displacements than forward steps alone, and this is what will eventually reach every real number.

Block 5. **Cue.** "I have steps of two types and I want to reach a general displacement": a reachability problem. **Why natural.** The set of displacements achievable with a forward and b backward steps

is the Minkowski sum $aI_{1,0} + bI_{0,1} = I_{a,b}$, and induction on $a + b$ is the natural bookkeeping. The only real question is where to put the intermediate point, and here the choice is forced rather than inspired: at the two endpoints of $I_{a,b}$ the decomposition into steps is unique, so define the intermediate point by affinely interpolating those two forced choices. That is exactly the map j, and the required property $\tau - j(\tau) \in [L_2, L_3]$ then holds automatically because $\tau - j(\tau)$ is affine with the correct endpoint values. **Link to the next block.** You now know the exact rigidity available: for every (a, b), h compresses $I_{a,b}$ into itself.

Block 6. Cue. The length of $I_{a,b}$ is $(a + b)(L_3 - L_2)$, which grows. **Why natural.** So the interval never shrinks around a target, and a naive squeeze is impossible. What can be controlled is one endpoint at a time. This is where the arithmetic enters: $aL_3 - bL_2 = L_2(a\theta - b)$ with $\theta = L_3/L_2$ irrational, so by Kronecker the endpoint values are dense, and by taking a, b large the other endpoint is pushed to infinity in the harmless direction. **Link to the next block.** Two applications of the Claim, one tuned at each endpoint, give the two sided bound.

Block 7. Cue. Two inequalities, each with an arbitrary ε. **Why natural.** $h(t) - h(s) \le T + \varepsilon$ and $h(t) - h(s) \ge T - \varepsilon$ for all $\varepsilon > 0$ force equality. Nothing subtler is available, and nothing subtler is needed. **Link to the next block.** Equality means $h(s) - s$ is constant, and Step 1 was set up so that this single fact unwinds to $f(x) = x^c$ in three lines.

Layer 3: WHY

The mathematical truth beneath this problem is that the condition is a statement about the additive group generated by L_2 and L_3, and the rigidity of f is exactly the density of that group in $\mathbb{R}$.

Look at what the two logarithms achieve. They put the problem in the normal form: an unknown $h : \mathbb{R} \to \mathbb{R}$ such that a displacement lying in the step set $J = [L_2, L_3]$ is mapped to a displacement lying in J. Iterating with signs, a displacement in $aJ - bJ$ maps into $aJ - bJ$. If the numbers $\log 2$ and $\log 3$ were commensurable, the endpoints $aL_3 - bL_2$ would live in a discrete subgroup of $\mathbb{R}$ and could not approach an arbitrary target, and the conclusion would be false. It is false: take exponents 2 and 4 instead of 2 and 3, so that $L_4/L_2 = 2$. Then for any δ with $0 < \delta < L_2/(2\pi)$ the function

$$h(s) = s + \delta \sin\left(\frac{2\pi s}{L_2}\right)$$

satisfies the corresponding condition, since $h(t) - h(s) = \tau + 2\delta \cos\big(2\pi(s + \tau/2)/L_2\big) \sin(\pi\tau/L_2)$ for $\tau = t - s$, and the oscillating term vanishes at $\tau = L_2$ and $\tau = 2L_2$ and is small enough in between. Unwinding gives

$$f(x) = x^{c(x)}, \qquad c(x) = \exp\left(\delta \sin\left(\frac{2\pi \log\log x}{\log 2}\right)\right),$$

a solution of the $(2, 4)$ problem that is not a pure power. So the hypothesis is not decoration: the problem is true for $(2, 3)$ and false for $(2, 4)$, and the dividing line is whether $\log 3/\log 2$ is irrational, which is whether 2 and 3 are multiplicatively independent, which is unique factorisation.

The discernment mechanism, and the thing to train: when a functional condition is stated with explicit constants, ask what group or semigroup those constants generate after you have normalised the condition into additive form. If the generated set is dense, expect total rigidity and go hunting for the density theorem. If it is discrete, expect a family of oscillating solutions and go hunting for a counterexample. That single question decides the shape of the whole solution before any technical work begins, and it is the reason to normalise first and compute second.

The second reusable truth is the interpolation device. Whenever an inductive step needs an intermediate point inside a constraint interval, look at the extreme cases: at an endpoint of the

interval the decomposition is usually forced, and the affine map interpolating the forced choices is then the correct intermediate point everywhere in between, with the verification reducing to checking two endpoint values because everything in sight is affine. That is not a lucky formula, it is the unique affine map with the required endpoint behaviour.

What the examiners nudged with: the identical constants on both sides (telling you the answer is a power map and the normal form is a translation condition), the choice of 2 and 3 rather than 2 and 4 (telling you the arithmetic of the exponents is load bearing), and the codomain $(1,\infty)$ rather than $(0,\infty)$ (making both logarithms legal, since $f>1$ gives $g>0$ gives h real, and forcing $c>0$).

Revolving concepts

Normal form of a functional condition. Definitions: conjugating f by $x=e^u$ replaces power maps by linear maps; conjugating again by $u=e^s$ replaces linear maps by translations. The pair of substitutions is the standard route from the multiplicative Cauchy family $f(xy)=f(x)f(y)$ to the additive one, and from power maps to shifts. Orbiting facts: $x\mapsto x^c$ is the conjugate of $s\mapsto s+\log c$; the group of power maps on $(1,\infty)$ is isomorphic to $(\mathbb{R}_{>0},\times)$ and to $(\mathbb{R},+)$. See Book 1, field: Functional equations, topic: conjugation to normal form.

Kronecker density and multiplicative independence. If θ is irrational then $\{a\theta-b : a,b\in\mathbb{Z},\ a,b>N\}$ is dense in $\mathbb{R}$ for every N; the proof is pigeonhole to find $\|q\theta\|$ small, then walk. Weyl's theorem strengthens this to equidistribution. Irrationality of $\log 3/\log 2$ is unique factorisation: $3^q=2^p$ forces $p=q=0$. See Book 1, field: Number theory, topic: irrationality and equidistribution.

Minkowski sums of intervals, and the reachability recursion. For intervals, $[p,q]+[p',q']=[p+p',q+q']$. The recursion here is $I_{a,b}=I_{a-1,b}+I_{1,0}=I_{a,b-1}+I_{0,1}$ with $I_{0,0}=\{0\}$, whose solution is $I_{a,b}=aI_{1,0}+bI_{0,1}=[aL_2-bL_3,\ aL_3-bL_2]$. Generalised: if the available steps are the intervals $J_1,\dots,J_r$ used $a_1,\dots,a_r$ times, the reachable set is $\sum a_iJ_i$, an interval whose endpoints are the corresponding integer combinations of the endpoints of the J_i, and whose length is $\sum a_i|J_i|$. That single formula tells you in advance that the reachable set widens, hence that rigidity must come from the endpoints rather than from the width. The same recursion governs every "steps of bounded size" argument: random walk ranges, interval exchange bookkeeping, and the reachable set of a discrete control system. See Book 1, field: Real analysis, topic: reachable sets and interval arithmetic.

Induction on a total weight $a+b$. When two parameters can each be decremented, induct on their sum and treat the two decrements as symmetric cases. Orbiting technique: the intermediate point should be produced by an explicit affine formula so that verification reduces to endpoint checks. See Book 1, field: Logic and set theory, topic: induction on a well founded measure.

Sandwich conditions. A two sided constraint $\alpha\le\Phi\le\beta$ is much stronger than either half; the standard exploitation is to run the two halves against each other in opposite directions, which is exactly what the upper and lower bound in Step 4 do. See Book 1, field: Inequalities, topic: two sided estimates and squeezing.

Circuit Bank

- **WHEN** a functional condition involves powers of the argument on both sides with the same exponents **DO** take logarithms twice, so that powers become multiples and multiples become shifts **BECAUSE** only displacement conditions chain additively, and chaining is the only way to move information from a sparse set of arguments to all arguments.
- **WHEN** a condition reads "if $t-s\in J$ then $h(t)-h(s)\in J$" **DO** immediately write down the reflected version by swapping s and t, and then study the semigroup of reachable displacements **BECAUSE** the condition is symmetric in the two points, so backward steps are free, and reachability is the whole content.
- **WHEN** after normalising, the constants generate a subgroup of $\mathbb{R}$ **DO** ask whether that

subgroup is dense or discrete **BECAUSE** dense means the answer is rigid and you need a density theorem, discrete means there is an oscillating counterexample and the problem as stated must have excluded it.

- **WHEN** an induction needs an intermediate point inside a constraint interval **DO** take the affine interpolation of the decompositions forced at the two endpoints **BECAUSE** everything in sight is affine, so the verification collapses to checking two endpoint values.
- **WHEN** a reachable interval grows in length instead of shrinking **DO** tune one endpoint against the target and let the other endpoint escape to infinity, then repeat with the other endpoint **BECAUSE** a two sided bound does not need one interval that traps the target, only two intervals each of which pins one side.

Backtesting

1. First useful cue: the exponents 2 and 3 occur identically in the hypothesis and the conclusion, which says the condition is invariant under conjugation by power maps and therefore that power maps are the expected answer.
2. Method the expression suggested: two logarithmic substitutions to reach the translation normal form, then chaining of forward and backward steps, then a density argument on $\mathbb{Z}L_2 + \mathbb{Z}L_3$.
3. Where you would hesitate, and the thought to reinforce: at the induction, wondering where the intermediate point comes from. Reinforce this: at the endpoints of the constraint interval the step decomposition is forced, and the affine interpolation of forced endpoint data is always the right guess for the interior.
4. The trap the question sets: assuming regularity. There is no continuity hypothesis, and the problem is false in the analogous $(2, 4)$ case, so any argument that does not use the irrationality of $\log 3/\log 2$ is wrong.
5. What makes you recognise it faster next time: the shape "$x^2 \le y \le x^3$ implies $f(x)^2 \le f(y) \le f(x)^3$" is a transport of a window, and windows transported by a map are always studied in the normal form where the window is a translation.
6. The reusable decision rule: normalise the constraint until it is additive, identify the group generated by the constants, and let the arithmetic of that group decide whether the answer is rigid.

Mirror drills

1. Find all $f : (1, \infty) \to (1, \infty)$ such that $x^3 \le y \le x^5$ implies $f(x)^3 \le f(y) \le f(x)^5$. *Answer:* again exactly $f(x) = x^c$, $c > 0$; the proof is unchanged with L_2, L_3 replaced by $\log 3, \log 5$, and the only fact used is that $\log 5/\log 3$ is irrational.
2. Show that the analogous problem with exponents 2 and 4 has solutions that are not pure powers. *Answer:* $\log 4/\log 2 = 2$ is rational, so the reachable displacements lie in $(\log 2)\mathbb{Z}$; take $h(s) = s + \delta \sin(2\pi s/\log 2)$ with $0 < \delta < \log 2/(2\pi)$ and unwind two logarithms to get $f(x) = x^{c(x)}$ with $c(x) = \exp(\delta \sin(2\pi \log\log x/\log 2))$.
3. Prove that $\{a \log 3 - b \log 2 : a, b \in \mathbb{Z}_{\ge 0}\}$ is dense in $\mathbb{R}$. *Answer:* $\theta = \log 3/\log 2$ is irrational by unique factorisation; pigeonhole on the fractional parts of $0, \theta, \ldots, M\theta$ gives q with $\|q\theta\| < 1/M$, and the multiples of $q\theta$ then step around $\mathbb{R}/\mathbb{Z}$ in increments below $1/M$, hitting every interval of that length.

10.12 B6. A dyadic double sum that telescopes to one

Evaluate

$$\sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k} \sum_{n=0}^{\infty} \frac{1}{k2^n+1}.$$

Read it out loud

An alternating outer sum in k, and an inner sum over n in which the index n appears only inside $k2^n$. Read the interaction of the two indices, because that is the only structural feature here: the summand is

$$\frac{(-1)^{k-1}}{k} \cdot \frac{1}{k2^n+1},$$

so k appears twice and n appears once, and the two indices meet in exactly one place, the product $k2^n$. Everything else is a function of k alone. A double sum whose indices meet in a single combined quantity is a double sum you should reindex by that quantity.

The fields wired in: real analysis (absolute convergence, and the theorem that an absolutely summable family may be summed in any order or grouping), number theory (every positive integer is uniquely an odd number times a power of two, which is the base 2 valuation), and series technique (geometric collapse and telescoping). There is no calculus in the ideal route at all, which is worth noticing given that three of the published solutions use integrals.

What is really being tested is restraint. The inner sum $\sum_{n\geq 0} 1/(k2^n+1)$ is a Lambert type series and has no elementary closed form; for $k=1$ its value is $0.5+1/3+1/5+1/9+\cdots \approx 1.2645$, a constant with no known expression in elementary functions. A candidate who spends the first twenty minutes trying to evaluate it has lost the problem to consumption rather than analysis. The correct first act is to declare the inner sum unevaluatable and change what you are summing over.

In the first thirty seconds you should have written the convergence audit, because it is what licences everything else:

$$0 < \sum_{n=0}^{\infty} \frac{1}{k2^n+1} < \sum_{n=0}^{\infty} \frac{1}{k2^n} = \frac{2}{k},$$

so the outer terms are bounded by $2/k^2$ and the whole double family is absolutely summable, with total absolute mass at most $2\zeta(2) = \pi^2/3$. Once that is on the page you may treat the expression as an unordered sum over the set of pairs (k,n), and you are free.

Approaches, ranked worst to best

1. **Evaluate the inner sum first.** *Cue:* a nested sum invites you to do the inside first. *Circuit:* compute inward out. This is the zombie route and it never terminates: $\sum_n 1/(k2^n+1)$ has no elementary closed form. The cue that should stop you is that k also sits outside, so an exact closed form in k would be needed, not merely a number.

2. **Numerical evaluation and recognition.** *Cue:* "Evaluate", with no free parameter in the answer. *Circuit:* compute enough terms to recognise the constant. Summing to $k \approx 10^5$ gives $1.0000\ldots$, which tells you the target is 1. Worth sixty seconds as target setting; worth zero marks as a solution.

3. **Expand $1/(k2^n+1)$ as a geometric series in $1/(k2^n)$.** *Cue:* a denominator of the form $t+1$ with t large. *Circuit:* $1/(1+t) = \sum(-1)^m t^{-m-1}$ for $t>1$. This works and is the published first solution: peel off the $n=0$ term, expand the rest, and sum a triple series to get $2\log 2 - 1$ and $2-2\log 2$. It costs the most ink, needs the $n=0$ term treated separately because the expansion fails there, and needs three separate convergence checks. The logarithms cancelling at the end is a loud signal that a route exists in which they never appear.

4. **Integral representation.** *Cue:* a denominator $t+1$ with t a nonnegative exponent. *Circuit:* $\frac{1}{t+1} = \int_0^1 x^t\,dx$, which converts denominators into exponents. Very pretty: the k sum becomes $\log(1+x^{2^n})$, the n sum becomes $\log\prod_{n\ge0}(1+x^{2^n}) = -\log(1-x)$ by binary expansion, and $-\int_0^1 \log(1-x)\,dx = 1$. Ranks below the last two only because it needs term by term integration, monotone convergence, an infinite product, and an improper integral, which is four separate things a grader can ask you to justify.

5. **Parity folding.** *Cue:* the index set $\{k2^n\}$ is invariant under $k \mapsto 2k$. *Circuit:* split the sum by the parity of k, substitute $k = 2\ell$ in the even part, and let the sum fold onto itself. Split $S = S_0 - S_1$ into odd and even k with all signs removed. Substituting $k = 2\ell$ in S_1 reproduces $S_0 + S_1$ with the $n = 0$ term removed, giving $S_1 = \frac{1}{2}(S_0 + S_1) - \frac{1}{2}$ and hence $S = S_0 - S_1 = 1$ in four lines. Short and honest, but it computes the answer without explaining it.

6. **Reindex by the level sets of $k2^n + 1$.** *Cue:* the two indices meet only in the product $k2^n$. *Circuit:* sum along the fibres of the map $(k,n) \mapsto k2^n$, which are exactly the doubling chains of an odd number. Best. The fibre sum collapses geometrically to $1/(m-1)$, the outer sum becomes $\sum_{m\ge2} 1/(m(m-1))$, and the answer 1 appears as a telescoping sum. It is the shortest complete route, it uses one convergence theorem and nothing else, and it explains why the answer is 1.

The ideal solution

Write $a_{k,n} = \dfrac{(-1)^{k-1}}{k(k2^n+1)}$ for integers $k \ge 1$ and $n \ge 0$, so that the given expression is the iterated sum $\sum_{k\ge1}\sum_{n\ge0} a_{k,n}$.

Step 1: absolute summability. Since $k2^n + 1 > k2^n$,

$$\sum_{k=1}^{\infty}\sum_{n=0}^{\infty}|a_{k,n}| < \sum_{k=1}^{\infty}\sum_{n=0}^{\infty}\frac{1}{k^2 2^n} = 2\sum_{k=1}^{\infty}\frac{1}{k^2} = \frac{\pi^2}{3} < \infty.$$

Hence the family $(a_{k,n})$ indexed by the countable set $P = \mathbb{Z}_{\ge1}\times\mathbb{Z}_{\ge0}$ is absolutely summable. Its unordered sum S therefore exists, equals the given iterated sum, and may be computed by partitioning P into any collection of disjoint blocks, summing each block, and adding the block totals.

Step 2: the right partition. The pair (k,n) enters $a_{k,n}$ through k and through the single quantity $k2^n$. Partition P by the value of

$$m = k2^n + 1.$$

As (k,n) ranges over P, m ranges over the integers $m \ge 2$, and the block over a given m is

$$P_m = \{(k,n) \in P : k2^n = m-1\}.$$

Step 3: describe a block. Fix $m \ge 2$ and write $m - 1 = 2^e j$ with j odd and $e = v_2(m-1) \ge 0$; this decomposition into an odd part and a power of two exists and is unique. If $k2^n = 2^e j$ with $k \ge 1$ and $n \ge 0$, then comparing the powers of 2 on both sides forces $n \le e$ and $k = 2^{e-n}j$. So, writing $i = e - n$,

$$P_m = \{(2^i j,\ e-i) : i = 0,1,\dots,e\},$$

a chain of $e+1$ pairs whose k values are $j, 2j, 4j, \dots, 2^e j$.

Step 4: sum a block. On P_m the factor $1/(k2^n+1)$ is the constant $1/m$, so

$$\sum_{(k,n)\in P_m} a_{k,n} = \frac{1}{m}\sum_{i=0}^{e}\frac{(-1)^{2^i j-1}}{2^i j}.$$

The parity along the chain is decided at once: j is odd, so $(-1)^{j-1} = 1$; every later term $2^i j$ with $i \geq 1$ is even, so $(-1)^{2^i j - 1} = -1$. Hence

$$\sum_{i=0}^{e} \frac{(-1)^{2^i j - 1}}{2^i j} = \frac{1}{j}\left(1 - \sum_{i=1}^{e} \frac{1}{2^i}\right) = \frac{1}{j}\left(1 - \left(1 - \frac{1}{2^e}\right)\right) = \frac{1}{2^e j} = \frac{1}{m-1},$$

where the sum over $i \geq 1$ is empty when $e = 0$, in which case the identity reads $1/j = 1/(m-1)$ and still holds. Therefore

$$\sum_{(k,n)\in P_m} a_{k,n} = \frac{1}{m(m-1)}.$$

Step 5: add the blocks. By Step 1 we may sum the block totals in any order:

$$S = \sum_{m=2}^{\infty} \frac{1}{m(m-1)} = \sum_{m=2}^{\infty} \left(\frac{1}{m-1} - \frac{1}{m}\right) = \lim_{M\to\infty}\left(1 - \frac{1}{M}\right) = 1.$$

So

$$\sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k} \sum_{n=0}^{\infty} \frac{1}{k2^n + 1} = \boxed{1}.$$

Layer 1: WHAT

Block 1. Convergence audit: $\sum_n 1/(k2^n+1) < 2/k$, so the terms are $O(1/k^2)$ and the double family is absolutely summable with mass at most $\pi^2/3$.

Block 2. Refuse to evaluate the inner sum: it is a Lambert type series with no elementary closed form, and in any case a number would not help since k sits outside as well.

Block 3. Flatten: stop reading the expression as a sum of sums, and read it as one sum over the lattice of pairs (k, n).

Block 4. Reindex: the indices meet only in $k2^n$, so partition the lattice into the level sets $m = k2^n + 1$.

Block 5. Identify a level set: unique factorisation into odd part times power of two makes the fibre the doubling chain $j, 2j, \ldots, 2^e j$.

Block 6. Sum the chain: one odd term with sign $+$ and e even terms with sign $-$, and the geometric identity $1 - (2^{-1} + \cdots + 2^{-e}) = 2^{-e}$ collapses the chain sum to $1/(m-1)$.

Block 7. Telescope $\sum_{m\geq 2} 1/(m(m-1)) = 1$.

Juxtaposition. The nominal route is: evaluate the inner sum, then substitute back. That is what nesting invites, and it is a dead end, because the inner object is a genuinely transcendental thing with no closed form. The ideal route never evaluates the inner sum at all; it dissolves the nesting and re-partitions the index set along a different axis. The bifurcation point is Block 3: the line at which you replace two nested indices by one flat index set. Everything after that is forced, because on a flat index set the only decision left is which partition to use, and the summand answers that question by telling you where k and n meet.

Layer 2: HOW

Block 1. Cue. Any manipulation of an infinite double sum. Before you rearrange anything, you owe a bound. **Why natural.** Rearranging a conditionally convergent series can produce any value at all, and the outer sum here is alternating, which is the classic signature of conditional convergence. So the audit is not bureaucracy, it is the licence for the entire method. The bound itself is immediate: dropping the $+1$ in the denominator gives a geometric series in n and the harmless factor $2/k$. **Link to the next block.** Once absolute summability is on the page, the expression is an unordered sum over pairs and any partition is legal. That is a much freer object than a nested sum.

Block 2. **Cue.** The inner sum has k in it and no closed form. **Why natural.** Ask what a closed form would have to look like: to be useful it would have to be an explicit function of k that can then be summed against $(-1)^{k-1}/k$. Series of the shape $\sum_n 1/(q^n + c)$ are Lambert type and are known not to have elementary closed forms. Recognising the family costs nothing and saves the whole problem. **Link to the next block.** If the inner sum cannot be evaluated, the only remaining lever is the index set itself.

Block 3. **Cue.** A nested sum with independent index ranges: $k \geq 1$ and $n \geq 0$ with no coupling in the limits. **Why natural.** Independent ranges mean the index set is a product, so the sum is a sum over a lattice, and a sum over a lattice can be organised along any family of disjoint slices. The nesting given in the problem is just one such family (vertical lines), and there is no reason it is the best one. **Link to the next block.** So look at the summand and ask which slicing makes it simple.

Block 4. **Cue.** The summand is $F(k) \cdot G(k2^n)$ with $F(k) = (-1)^{k-1}/k$ and $G(t) = 1/(t+1)$. The two indices meet only in $k2^n$. **Why natural.** On the level set $k2^n =$ constant, the factor G is constant and comes out, leaving a pure sum of F over the fibre. That is the definition level reason to slice this way: it turns a two variable sum into a constant times a one variable sum, per slice. Whenever a summand factors through a single combined quantity, the level sets of that quantity are the natural slices. **Link to the next block.** So describe the fibres of the map $(k, n) \mapsto k2^n$.

Block 5. **Cue.** Powers of 2 in a product equal to a fixed integer. **Why natural.** Every positive integer is uniquely 2^e times an odd number, so $k2^n = 2^e j$ forces $n \leq e$ and $k = 2^{e-n} j$. The fibre is therefore the doubling chain of the odd part, of length $e + 1 = v_2(m-1) + 1$. This is arithmetic, not cleverness: the map $(k, n) \mapsto k2^n$ is exactly the multiplication that unique factorisation controls. **Link to the next block.** Now sum F along a chain, where the terms are $\pm 1/(2^i j)$.

Block 6. **Cue.** A geometric chain, with an alternating sign that depends on parity, and parity is constant after the first step. **Why natural.** $2^i j$ is odd exactly when $i = 0$. So the sign is $+$ once and $-$ forever after, and the chain sum is $\frac{1}{j}(1 - \sum_{i=1}^{e} 2^{-i})$. The partial geometric sum $\sum_{i=1}^{e} 2^{-i} = 1 - 2^{-e}$ is the identity that makes the bracket exactly 2^{-e}, and $2^e j = m - 1$. This is why the base 2 in the problem is not decorative: with 3^n in place of 2^n the sign would be constant along the whole chain and the partial sums would not complement to a single reciprocal. **Link to the next block.** Each block contributes $1/(m(m-1))$, and that is a telescoping term.

Block 7. **Cue.** $\frac{1}{m(m-1)}$, a product of consecutive integers in a denominator. **Why natural.** Partial fractions give $\frac{1}{m-1} - \frac{1}{m}$, and consecutive difference means telescoping. The sum from $m = 2$ is 1. **Link to the next block.** Answer 1, matching the numerical check made at the start.

Layer 3: WHY

One arithmetic fact runs underneath every route to this problem: every positive integer is uniquely an odd number times a power of two. Watch it appear in four disguises.

In the ideal route it is explicit: the fibre of $(k, n) \mapsto k2^n$ over $m - 1$ is the doubling chain of the odd part of $m - 1$.

In the integral route it is the identity $\prod_{n=0}^{\infty}(1 + x^{2^n}) = \frac{1}{1-x}$ for $|x| < 1$, which one proves by telescoping, since $(1 - x)\prod_{n=0}^{N}(1 + x^{2^n}) = 1 - x^{2^{N+1}}$, and which says in generating function language that every nonnegative integer has exactly one binary expansion. That is the same fact, dressed as a product.

In the parity folding route it is the invariance of the index set $\{k2^n\}$ under $k \mapsto 2k$, which is the self-similarity of the dyadic chains; the substitution $k = 2\ell$ maps the even part of the sum back onto

the whole sum with the $n = 0$ layer stripped, and that produces a linear equation for the unknown. Same fact, dressed as a symmetry.

In the published triple series route it is the Euler factor: the alternating zeta value $\sum_k (-1)^{k-1} k^{-s}$ equals $(1 - 2^{1-s})\zeta(s)$ precisely because splitting $\mathbb{Z}_{\geq 1}$ by parity is splitting off the prime 2. Same fact, dressed as a factorisation of a Dirichlet series.

So the discernment mechanism is this: when powers of a fixed base appear in an index, the arithmetic of that base is the engine of the problem, and your job is to choose the disguise in which the engine is most visible. The reindexing route makes it most visible, which is why it is also the shortest.

There is a second, more portable truth here, and it is about double sums in general. If a summand can be written as $F(k)\, G(\phi(k, n))$, then slicing the index set by the level sets of ϕ turns the double sum into

$$\sum_{N \in \operatorname{im} \phi} G(N) \sum_{(k,n) \in \phi^{-1}(N)} F(k),$$

and the whole problem becomes: describe the fibres, and sum F over one. That is a completely mechanical procedure, and it is the correct default for any double sum in which the indices interact through a single expression. Dirichlet's hyperbola method, the convolution identity $\sum_{d|N}$, and the standard proof that $\sum_{k,n} 1/(k^2 n^2) = \zeta(2)^2$ are all instances. The only skill it demands is being able to describe the fibre, and for $\phi(k, n) = k2^n$ the fibre is handed to you by unique factorisation.

How to train it: every time you meet a double sum, write down ϕ, the quantity through which the two indices actually interact, before you write anything else. If ϕ is one of the indices alone, the sum factorises and you are done. If ϕ is a product, a sum or a maximum of the two indices, slice by ϕ. The examiners nudged at exactly this by putting k both outside and inside: they made sure that no evaluation of the inner sum could be useful, so that the only surviving move is the structural one.

Knowledge gaps, stated plainly, because this is a B6 and the theory must be installed before the problem is even approachable. First, unordered summation: the theorem that an absolutely summable family over a countable index set has a well defined sum, invariant under every rearrangement and every partition into blocks. This is the whole licence for Steps 2 to 5, and it is not part of A level or Further Maths. Second, the base 2 valuation v_2 and the odd part decomposition, which is elementary but must be fluent, including the uniqueness that pins down $n \leq e$. Third, telescoping and partial fractions for $1/(m(m-1))$, which you have, and the value $\pi^2/6$ of $\zeta(2)$, which you need only as a finite bound and could replace by 2. Fourth, if you want the integral route: term by term integration of a uniformly convergent series, the monotone convergence theorem for series of nonnegative functions, infinite products, and the convergent improper integral $\int_0^1 \log(1 - x)\, dx = -1$. Fifth, if you want the published route: the alternating harmonic series $\sum (-1)^{k-1}/k = \log 2$ and the eta to zeta relation. The honest ranking is that the ideal route needs exactly one theorem beyond school mathematics, and every other route needs at least three.

Revolving concepts

Absolute summability of a family. Definition: a family $(a_p)_{p \in P}$ over a countable set P is absolutely summable if $\sup_{F \text{ finite}} \sum_{p \in F} |a_p| < \infty$. Theorem: it then has a unique unordered sum, equal to every iterated sum and to every block decomposition $\sum_{\text{blocks}} \sum_{p \in \text{block}} a_p$. This is Fubini for counting measure, and it is the only analytic input the ideal solution uses. Contrast: Riemann's rearrangement theorem says a conditionally convergent real series can be rearranged to any value, which is why the audit is not optional. See Book 1, field: Real analysis, topic: unordered sums and Fubini for series.

Base 2 *valuation and the doubling chain.* Definition: $v_2(N)$ is the exponent of 2 in N, and $N = 2^{v_2(N)} \cdot \text{odd}$ uniquely. The set of (k, n) with $k2^n = N$ is $\{(2^i \operatorname{odd}(N), v_2(N) - i) : 0 \leq i \leq v_2(N)\}$, a chain of length $v_2(N) + 1$. Orbiting facts: $v_2(ab) = v_2(a) + v_2(b)$; the number of (k, n) with $k2^n = N$

is $v_2(N)+1$, which is also the number of divisors of N that are powers of 2 times a fixed odd part. See Book 1, field: Number theory, topic: p-adic valuation.

Slicing a double sum by a level set. Technique: if the summand is $F(k)G(\phi(k,n))$ then $\sum_{k,n} = \sum_N G(N)\sum_{\phi^{-1}(N)} F$. Generalised recursion form: when $\phi(k,n) = k2^n$, the fibres are orbits of the map $k \mapsto 2k$ truncated at the odd end, so the fibre sums obey the one step recursion $T(2N) = T(N) + F(\text{new element})$, and the closed form is the geometric collapse used above. The same slicing produces Dirichlet convolution when $\phi(k,n) = kn$, and the hyperbola method when the fibres are too long to handle directly. See Book 1, field: Generating functions, topic: rearranging double sums by level sets.

Self-similar sums and linear equations. When an index set is invariant under a substitution, splitting the sum along the orbits of the substitution gives a linear equation in the unknown sum. Here, splitting by the parity of k and substituting $k = 2\ell$ gives $S_1 = \frac{1}{2}(S_0 + S_1) - \frac{1}{2}$, hence $S = S_0 - S_1 = 1$. General method: write $S = \sum_{\text{fixed part}} + (\text{copy of } S \text{ scaled})$, then solve the resulting linear relation. This is the same manoeuvre that evaluates $\sum 2^{-n}$ by $S = 1 + S/2$, and it generalises to any sum with a scaling symmetry. See Book 1, field: Sequences and recurrences, topic: self-similar sums.

Binary expansion as an infinite product. $\prod_{n\geq 0}(1+x^{2^n}) = 1/(1-x)$ for $|x| < 1$, proved by $(1-x)\prod_{n=0}^{N}(1+x^{2^n}) = 1 - x^{2^{N+1}}$. Orbiting facts: the same telescoping proves $\sum_{n\geq 0} \frac{x^{2^n}}{1-x^{2^{n+1}}} = \frac{x}{1-x}$; the general Euler style identity $\prod(1+x^{q^n})$ encodes base q expansions with digits 0 and 1. See Book 1, field: Generating functions, topic: binary expansions and telescoping products.

Lambert type series. $\sum_{n\geq 0} \frac{1}{q^n+c}$ has no elementary closed form in general; the Erdős and Borwein constant $\sum_{n\geq 1} 1/(2^n-1)$ is the standard example and is known to be irrational. Recognising the family is what tells you to stop trying. See Book 1, field: Real analysis, topic: Lambert series.

Circuit Bank

- **WHEN** a double sum has independent index ranges and a summand in which the two indices meet in one combined quantity ϕ **DO** partition the index set by the level sets of ϕ and sum each fibre **BECAUSE** on a fibre the factor depending on ϕ is constant and comes out, so a two variable sum becomes a constant times a one variable sum.
- **WHEN** you are about to rearrange, regroup or interchange any infinite sum **DO** first bound the sum of absolute values, and write the bound down **BECAUSE** unordered summation is legal only for absolutely summable families, and an alternating outer sum is exactly the case where conditional convergence would break everything.
- **WHEN** an inner sum has no closed form but the outer index also appears inside it **DO** stop evaluating and start reindexing **BECAUSE** a closed form would have to be an explicit function of the outer index to be useful, so the absence of one is a signal that the intended route is structural rather than computational.
- **WHEN** powers of a fixed base appear inside an index **DO** invoke the unique factorisation into that base times a coprime part, and expect chains or orbits **BECAUSE** the fibres of $(k,n) \mapsto kq^n$ are exactly the q-multiplication chains of the q-free part.
- **WHEN** a block or fibre sum comes out as $\frac{1}{m(m-1)}$ or any product of consecutive integers **DO** split by partial fractions and telescope **BECAUSE** consecutive denominators are the fingerprint of a difference, and a sum of differences collapses to its endpoints.

Backtesting

1. First useful cue: the two indices meet only in the product $k2^n$, and k appears both inside and outside, so no evaluation of the inner sum can be the intended route.

2. Method the expression suggested: bound everything to get absolute summability, flatten the nested sum into a sum over the lattice of pairs, and reindex by $m = k2^n + 1$.
3. Where you would hesitate, and the thought to reinforce: at the description of the fibre. Reinforce it by writing $m - 1 = 2^e j$ with j odd before you try to solve $k2^n = m - 1$; with the odd part named, the fibre is immediate.
4. The trap the question sets: the nesting itself, which invites you to compute the inner sum first, and the alternating outer sum, which invites an illegal rearrangement if you skip the convergence audit.
5. What makes you recognise it faster next time: powers of 2 inside an index plus a sign depending on the parity of that index means the dyadic chain structure is the whole problem, since parity is constant along the chain after the first step.
6. The reusable decision rule: in a double sum, identify the single quantity through which the indices interact and slice by its level sets; if the fibres are chains under multiplication by a base, sum them geometrically.

Mirror drills

1. Prove that $\prod_{n=0}^{\infty}(1 + x^{2^n}) = \frac{1}{1-x}$ for $|x| < 1$, and deduce $\sum_{n=0}^{\infty} \log(1 + x^{2^n}) = -\log(1 - x)$. *Answer:* $(1 - x)\prod_{n=0}^{N}(1 + x^{2^n}) = 1 - x^{2^{N+1}} \to 1$ by repeated difference of squares; take logarithms, legal since all factors are positive.
2. Evaluate $\sum_{n=0}^{\infty} \frac{x^{2^n}}{1 - x^{2^{n+1}}}$ for $|x| < 1$. *Answer:* $\frac{x}{1-x}$, since $\frac{x^{2^n}}{1 - x^{2^{n+1}}} = \frac{1}{1 - x^{2^n}} - \frac{1}{1 - x^{2^{n+1}}}$ telescopes.
3. Evaluate $\sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k(k+1)}$. *Answer:* $2\log 2 - 1$, by $\frac{1}{k(k+1)} = \frac{1}{k} - \frac{1}{k+1}$ and the alternating harmonic series; this is the $n = 0$ layer of the problem, and comparing it with the total 1 shows that the layers $n \geq 1$ contribute $2 - 2\log 2$.

Chapter 11

Putnam 2015

11.1 A1. maximal triangle on a hyperbola branch

Let A and B be points on the same branch of the hyperbola $xy = 1$. Suppose that P is a point lying between A and B on this hyperbola, such that the area of the triangle APB is as large as possible. Show that the region bounded by the hyperbola and the chord AP has the same area as the region bounded by the hyperbola and the chord PB.

Read it out loud

Two points fixed on one branch of $xy = 1$, a third point free on the arc between them, and a claim about the area of the triangle it cuts out: this is extremal geometry with a fixed base, so the whole first half is really asking where the perpendicular distance from the arc to the chord AB is largest. The second sentence is not a new problem bolted on, it is a demand that whatever description you find for P must be strong enough to feed straight into an area equality, so an ugly formula for P is a sign you solved the first half the wrong way.

The phrase "same branch" is the first useful cue: it licenses one sign convention, $x > 0$, for the whole argument. Within thirty seconds you should have $A = (a, 1/a)$, $B = (b, 1/b)$, $P = (t, 1/t)$, $0 < a < t < b$. The field quietly embedded underneath the coordinate geometry is the group of area preserving affine self maps of $xy = 1$, and finding it is what separates a two page computation from a four line proof.

Approaches, ranked worst to best

1. **Coordinates and two integrals.** *Cue:* "area of the region bounded by the hyperbola" asked twice. *Circuit:* integrate line minus curve over each sub arc separately. (costs two logarithmic integrals and an algebraic reconciliation, and never explains why the two answers should agree, it only checks that they do.)

2. **AM-GM plus one integral, symmetry asserted for the other.** *Cue:* the product $t \cdot (ab/t)$ appearing once the base is fixed. *Circuit:* minimise $t + ab/t$ to find P, integrate once, then assert the second integral mirrors the first. (still one full integral, and the asserted symmetry of the closed form is not yet a proof of anything.)

3. **Tangent parallel to chord, then reflect.** *Cue:* the branch is visibly convex, so "biggest triangle on a fixed base" is a tangent line question. *Circuit:* locate P by matching slopes, rescale by a squeeze map into mirror position, read the equality off a reflection. (no integral is ever computed; the equality is forced, not checked.)

The ideal solution

Without loss of generality take the branch with $x > 0$: the map $(x, y) \mapsto (-x, -y)$ carries the other branch to this one and preserves the hyperbola and every area, so nothing is lost. Write $A = (a, 1/a)$, $B = (b, 1/b)$ with $0 < a < b$, and $P = (t, 1/t)$, $t \in (a, b)$.

Locating P. The area of triangle APB is half of $|AB|$ times the perpendicular distance from P to the line AB, so maximising it over t is maximising that distance. Since $y = 1/x$ has $y'' = 2/x^3 > 0$ on $x > 0$, the branch is strictly convex, so the signed height from the arc to the chord AB is an affine function of t minus the strictly convex function $1/t$, hence strictly concave in t; it vanishes at $t = a, b$ and is positive between, so it has a unique interior maximum, at the zero of its derivative, that is, where the tangent to the curve is parallel to AB. The chord has slope $\dfrac{1/b - 1/a}{b - a} = -\dfrac{1}{ab}$, the tangent at $(t, 1/t)$ has slope $-1/t^2$, and these agree exactly when $t = \sqrt{ab} \in (a, b)$. So $P = (\sqrt{ab}, 1/\sqrt{ab})$, uniquely.

Normalising. For $\lambda = 1/\sqrt{ab}$, the linear map $T(x,y) = (\lambda x, y/\lambda)$ has determinant 1, so it preserves every area; and since $(\lambda x)(y/\lambda) = xy$, it fixes the hyperbola, sending arcs to arcs and, being linear, chords to chords. It sends A, B, P to

$$T(A) = \left(\sqrt{\tfrac{a}{b}}, \sqrt{\tfrac{b}{a}}\right), \qquad T(B) = \left(\sqrt{\tfrac{b}{a}}, \sqrt{\tfrac{a}{b}}\right), \qquad T(P) = (1,1),$$

so with $s = \sqrt{a/b}$, $T(A) = (s, 1/s)$ and $T(B) = (1/s, s)$ are mirror images across the line $y = x$, and $T(P)$ is the fixed point of that mirror lying on the curve.

Reflecting. Let $\sigma(x,y) = (y,x)$. This isometry preserves area exactly, fixes the hyperbola, and satisfies $\sigma(T(A)) = T(B)$, $\sigma(T(B)) = T(A)$, $\sigma(T(P)) = T(P)$. So σ carries the arc and chord running from $T(A)$ to $T(P)$ onto the arc and chord running from $T(B)$ to $T(P)$, hence carries the region bounded by them onto the region bounded by the image pair, and an isometry cannot change area, so the two image regions are equal in area. Because T itself is area preserving and carries the original two regions, bounded by the hyperbola and the chords AP, PB, exactly onto these two image regions, the original two regions are equal in area too, which is the claim.

Layer 1: WHAT

Block 1: coordinatise the branch by one positive real each for A, P, B. Block 2: recognise a fixed base area extremum as a height extremum, and use convexity to know the maximiser is unique, at tangent parallel to chord. Block 3: solve one slope equation, $t = \sqrt{ab}$. Block 4: normalise by the squeeze map into mirror position. Block 5: reflect to swap the two regions and read off the equality, then transport back through the normalising map, which is itself area preserving.

The nominal and ideal routes agree through block 3; the bifurcation point sits exactly after P is found, where the nominal route reaches for two definite integrals of line minus curve, and the ideal route reaches instead for a symmetry that exchanges the two regions outright.

Layer 2: HOW

Block 1. **Cue.** "same branch", "P... on this hyperbola": both statements about position on one connected curve. **Why natural.** one positive parameter per point folds the constraint $xy = 1$ into the coordinate itself, so no multiplier or substitution is needed later. **Link to the next block.** three numbers a, t, b are now ready to feed a fixed base, moving height argument.

Block 2. **Cue.** "area... as large as possible" with A, B fixed. **Why natural.** fixed base area is half base times height, and a strictly convex arc gives a strictly concave height function, hence a unique interior maximiser. **Link to the next block.** a unimodal maximum is one derivative condition, which is a geometric statement: tangent parallel to chord.

Block 3. **Cue.** both slopes have the shape $-1/(\text{something})$, coming from the same function $1/x$. **Why natural.** comparing like with like collapses immediately to $t^2 = ab$. **Link to the next block.** $t = \sqrt{ab}$ is a geometric mean, an invariant that a scaling map can send to a fixed point.

Block 4. **Cue.** a geometric mean is exactly what the squeeze map $(x,y) \mapsto (\lambda x, y/\lambda)$ can send to 1. **Why natural.** this one parameter family is the full determinant one symmetry group of $xy = 1$ fixing the curve, so it is the natural group to normalise by. **Link to the next block.** once P sits at $(1,1)$ and A, B are swapped by $y = x$, the picture has a genuine symmetry, not a coincidence.

Block 5. **Cue.** A, B mirror images, P on the mirror line. **Why natural.** an isometry that swaps two regions cannot leave their areas different. **Link to the next block.** the equality proved after normalising transports back unchanged, because the normalising map was itself area preserving.

Layer 3: WHY

The truth underneath is that the extremal area question and the equal area question are one symmetry, not two computations that happen to agree. The squeeze maps and the single reflection generate enough of an area preserving group on $xy = 1$ that any A, B on one branch can be moved to mirror position at zero cost, and once there, the extremiser P has no choice but to sit on the fixed locus of the reflection, because the optimisation problem is carried to itself when A, B are swapped.

The discernment mechanism: when two fixed points and a free third point on a curve appear together, check for an area or distance preserving symmetry exchanging the two fixed points before reaching for calculus on an explicit formula. Calculus always works and always costs more; the symmetry, when present, replaces two computations with one act of swapping. Train it by asking, for every curve given by an equation invariant under some transformation group ($xy = c$ under $(x, y) \mapsto (\lambda x, y/\lambda)$, a circle under rotation about its centre, $y = x^2$ under $x \mapsto -x$), what that group does to the specific two points on the page.

The examiners' nudge is the request for an equality rather than a number: a question that could have asked you to compute one area and instead asks you to prove two are equal is telling you the intended route never computes either.

Revolving concepts

Tangent parallel to chord. For strictly convex g on $[p, q]$, the secant through the endpoints minus g is strictly concave, zero at the ends, positive between, so uniquely maximised where the tangent to g matches the secant's slope. The same fact locates the point on a parabola or convex arc farthest from a given line. See Book 1, field: Calculus and integrals, topic: convexity and extremal distance to a chord.

Squeeze maps of $xy = c$. The maps $(x, y) \mapsto (\lambda x, y/\lambda)$, $\lambda > 0$, have determinant 1 and fix $xy = c$ setwise; with the reflection $(x, y) \mapsto (y, x)$ they generate the area preserving symmetries of the curve, and can move any two points on one branch into mirror position. Any curve invariant under a one parameter scaling pair with product one repays writing this group down first. See Book 1, field: Geometry, topic: affine symmetries of conics.

Reflections as area preserving involutions used to prove equalities. An isometry that exchanges two regions proves them equal without measuring either. See Book 1, field: Geometry, topic: symmetry arguments for equal area.

No sequence, induction or self similar structure appears in this problem, so there is no recursion to generalise here; the reusable structure is entirely the symmetry group argument above.

Circuit Bank

- **WHEN** two fixed points and a free third point sit on a curve with a visible area preserving symmetry group **DO** find the map or reflection that swaps the two fixed points before setting up any integral **BECAUSE** the extremiser lies on the fixed locus of that symmetry, usually settling both the extremum and any follow up equality at once
- **WHEN** you need the point on a convex or concave arc farthest from a fixed chord **DO** set tangent slope equal to chord slope and solve **BECAUSE** this is the geometric content of the same derivative condition, reached with two slope formulas instead of a full determinant expansion
- **WHEN** a curve is given by an equation invariant under a scaling pair $(x, y) \mapsto (\lambda x, \mu y)$ with $\lambda\mu = 1$ **DO** write that family down explicitly and use it to move any two points on the curve into mirror position **BECAUSE** it is an area preserving bijection fixing the curve, so the symmetric case is reached at zero cost

- **WHEN** a problem asks you to prove two quantities equal rather than to compute either **DO** look for a symmetry of the whole picture that exchanges them before computing anything **BECAUSE** a proof by exchange is shorter and generalises to cases where neither quantity has a closed form

Backtesting

1. First useful cue: "same branch", fixing one sign convention for the whole argument.
2. Method the expression suggested: fixed base, moving height, so compare tangent slope to chord slope rather than differentiate the full triangle area formula.
3. Where you would hesitate, and the thought to reinforce: right after finding $t = \sqrt{ab}$, when integration tempts; reinforce that the question asked for an equality, so look for what exchanges the two regions.
4. The trap the question sets: both areas are honestly computable by a direct integral, which rewards grinding and never announces the shortcut.
5. What makes you recognise it faster next time: seeing $xy = c$ should now fire the squeeze map reflex, the way a circle already fires rotate about the centre.
6. The reusable decision rule: when a symmetry preserving group and two fixed points are both present, normalise by the group before touching calculus.

Mirror drills

1. $A = (-2, 4)$, $B = (3, 9)$ on $y = x^2$; find P on the arc between them with tangent parallel to AB, and say whether the two cut off regions must be equal. *Answer:* $P = (1/2, 1/4)$ from $2t = 1$; no squeeze symmetry moves this A, B into mirror position on a parabola, so the two regions need not be equal here, and are not.
2. Two points lie on the same branch of $xy = 4$. Name the map replacing $(x, y) \mapsto (\lambda x, y/\lambda)$ and locate P. *Answer:* the same family still fixes $xy = 4$; with $A = (a, 4/a)$, $B = (b, 4/b)$, matching slopes gives $P = (\sqrt{ab}, 4/\sqrt{ab})$, since the constant cancels out of the slope comparison exactly as before.
3. On $xy = 1$, redefine P as the point where the arc length from A to P equals the arc length from P to B. Is this still the tangent parallel to chord point? *Answer:* no in general; arc length bisection and area maximisation are different extremal conditions and coincide only once the extra mirror symmetry is imposed on A, B specifically, a caveat worth clarifying before assuming two notions of midpoint agree.

11.2 A2. a Lucas recurrence and an odd prime factor

Let $a_0 = 1$, $a_1 = 2$, and $a_n = 4a_{n-1} - a_{n-2}$ for $n \geq 2$. Find an odd prime factor of a_{2015}.

Read it out loud

A linear recurrence with constant coefficients, order two, asked about one specific enormous index, 2015. The command is "find an odd prime factor," not "find a_{2015}," which is the cue that you are not meant to compute the number itself, you are meant to find structure that survives without ever writing all its digits.

The field embedded is linear recurrences and their closed forms: order two, constant coefficients, means a characteristic quadratic with two roots, and the sequence is a combination of their powers. The number $2015 = 5 \times 13 \times 31$ is not decoration, it is the whole plan: an index that factors is an invitation to relate a_{2015} to a_5 or a_{13} or a_{31}, terms small enough to compute by hand, through some divisibility law connecting a_n and a_{kn}.

Within thirty seconds you should have written down the characteristic equation $x^2 - 4x + 1 = 0$ and noticed $2015 = 5 \times 403$ with 403 odd, because "odd" in the question and an odd cofactor in the index are the same word for a reason: the factorisation identity for a sum of like powers only exists at an odd exponent.

Approaches, ranked worst to best

1. **Iterate the recursion to $n = 2015$.** *Cue:* a recursively defined sequence and one specific index. *Circuit:* just keep applying $a_n = 4a_{n-1} - a_{n-2}$ two thousand times. (the term has hundreds of digits by $n = 2015$, and even granting a machine to reach it, you would then need to factor an enormous number, which the problem cannot intend by hand.)

2. **Extend the sequence to negative indices and argue periodicity mod a candidate prime.** *Cue:* the recursion reads both ways, $a_{n-2} = 4a_{n-1} - a_n$. *Circuit:* show $a_{-n} = a_n$, then chase a prime dividing some a_m through the shifted recursion modulo that prime. (correct and avoids irrational numbers entirely, but needs a careful backward induction and a modular argument before it ever touches 2015.)

3. **Closed form via the characteristic roots, then factor $x^k + y^k$ for odd k.** *Cue:* a constant coefficient order two recurrence, and 2015 has an odd factor 403. *Circuit:* write $a_n = (\alpha^n + \beta^n)/2$, use $x^k + y^k = (x + y)(\cdots)$ for odd k to show $a_n \mid a_{kn}$, then compute the small term a_5 by hand. (five arithmetic steps and one factoring identity; the only overhead is a short integrality argument.)

The ideal solution

The characteristic equation of $a_n = 4a_{n-1} - a_{n-2}$ is $x^2 - 4x + 1 = 0$, with roots $\alpha = 2 + \sqrt{3}$, $\beta = 2 - \sqrt{3}$, satisfying $\alpha + \beta = 4$ and $\alpha\beta = 1$. The general solution is $a_n = A\alpha^n + B\beta^n$; the initial conditions $a_0 = 1 = A + B$ and $a_1 = 2 = A\alpha + B\beta$ are symmetric under swapping $\alpha \leftrightarrow \beta$, and $A = B = \frac{1}{2}$ solves them, since then $\frac{1}{2}(\alpha + \beta) = 2$. Hence

$$a_n = \frac{\alpha^n + \beta^n}{2} \qquad \text{for every } n \geq 0.$$

Key lemma. For every positive integer n and every odd positive integer k, $a_n \mid a_{kn}$.
Proof. Write $x = \alpha^n$, $y = \beta^n$. Since k is odd,

$$x^k + y^k = (x + y)\big(x^{k-1} - x^{k-2}y + \cdots - xy^{k-2} + y^{k-1}\big),$$

an alternating sum of k terms. Dividing both sides by 2 gives

$$\frac{a_{kn}}{a_n} = \frac{\alpha^{kn} + \beta^{kn}}{\alpha^n + \beta^n} = x^{k-1} - x^{k-2}y + \cdots + y^{k-1} =: Q.$$

Now $\alpha, \beta \in \mathbb{Z}[\sqrt{3}] = \{p + q\sqrt{3} : p, q \in \mathbb{Z}\}$, a ring closed under addition and multiplication, so $x = \alpha^n$ and $y = \beta^n$ lie in $\mathbb{Z}[\sqrt{3}]$, and hence so does Q, say $Q = p + q\sqrt{3}$. But $Q = a_{kn}/a_n$ is also a ratio of two positive integers (the sequence is a sequence of positive integers by immediate induction from the recursion and its integer initial data), hence rational; and a rational number of the form $p + q\sqrt{3}$ with $p, q \in \mathbb{Z}$ forces $q = 0$, since $\sqrt{3}$ is irrational. So $Q = p \in \mathbb{Z}$, meaning $a_n \mid a_{kn}$. ■

Applying it. $2015 = 5 \times 403$ and $403 = 13 \times 31$ is odd, so the lemma with $n = 5$, $k = 403$ gives $a_5 \mid a_{2015}$. Compute a_5 directly from the recursion: $a_2 = 4(2) - 1 = 7$, $a_3 = 4(7) - 2 = 26$, $a_4 = 4(26) - 7 = 97$, $a_5 = 4(97) - 26 = 362 = 2 \times 181$. Trial division by every prime up to $\lfloor\sqrt{181}\rfloor = 13$, that is $2, 3, 5, 7, 11, 13$, shows none divides 181, so 181 is prime, and it is odd. Since $362 \mid a_{2015}$ and $362 = 2 \times 181$, in particular $181 \mid a_{2015}$.

$$\boxed{181 \text{ is an odd prime factor of } a_{2015}.}$$

Layer 1: WHAT

Block 1: solve the recurrence in closed form through its characteristic equation. Block 2: notice the target index 2015 factors with an odd cofactor, and look for a divisibility law relating a_n and a_{kn} for odd k. Block 3: prove that law using the factorisation of $x^k + y^k$ for odd k, then an integrality argument inside $\mathbb{Z}[\sqrt{3}]$. Block 4: apply the law once, at the smallest available factor pair $(5, 403)$, and finish by direct arithmetic on a_5.

The nominal route shares block 1, a closed form is the obvious first move for any constant coefficient recurrence, but then tries to reach 2015 directly, either by brute iteration or by chasing the sequence backwards through negative indices and modular periodicity. The bifurcation point is the decision, once the closed form is in hand, to look at how 2015 factors rather than at 2015 itself: factoring the index converts one intractable huge term into an algebraic identity between two small ones.

Layer 2: HOW

Block 1. **Cue.** a linear recurrence with constant coefficients, order two. **Why natural.** such a recurrence always has a two dimensional solution space spanned by the powers of the roots of its characteristic quadratic, so the closed form is the standard first move, not a clever one. **Link to the next block.** having $a_n = (\alpha^n + \beta^n)/2$ turns "does a_n divide a_m" into a question about powers of two fixed algebraic numbers, not about the recursively defined integers directly.

Block 2. **Cue.** the specific index 2015, and the word "odd" in "odd prime factor." **Why natural.** an index that factors as a small number times an odd cofactor invites a law that needs an odd exponent, because the natural factoring identity for a sum of like powers, $x^k + y^k = (x + y)(\cdots)$, exists only when k is odd. **Link to the next block.** this converts the target into: find $a_n \mid a_{2015}$ for a small divisor n of 2015 with odd cofactor, then compute a_n.

Block 3. **Cue.** the ratio a_{kn}/a_n equals $(\alpha^{kn} + \beta^{kn})/(\alpha^n + \beta^n)$, exactly the shape $x^k + y^k$ over $x + y$. **Why natural.** the factoring identity gives an explicit integer polynomial Q in $x = \alpha^n, y = \beta^n$; since $\mathbb{Z}[\sqrt{3}]$ is closed under the ring operations, Q lands back in $\mathbb{Z}[\sqrt{3}]$, and comparing that with Q being manifestly rational forces Q to be a plain integer. **Link to the next block.** a plain integer Q with $a_{kn} = a_n \cdot Q$ is exactly the statement $a_n \mid a_{kn}$, available now for any odd k and any n.

Block 4. **Cue.** $2015 = 5 \times 403$, 403 odd. **Why natural.** the lemma applies directly with $n = 5, k = 403$, and a_5 is five arithmetic steps from the given initial data, cheap enough to compute and factor by hand. **Link to the next block.** $a_5 = 362 = 2 \times 181$ with 181 prime hands you the odd prime factor the question asked for, with no further work.

Layer 3: WHY

The truth underneath is that $a_n = (\alpha^n + \beta^n)/2$ is a close relative of $\cosh(n\theta)$, since α, β are real and $\alpha\beta = 1$, and sequences of this shape obey the same product identities that Chebyshev polynomials obey: whenever the index has an odd factor, the sum of like powers factoring identity is available and hands you a divisor at a smaller index. This is the general fact about Lucas sequences of the second kind: if $V_n = \alpha^n + \beta^n$ for the roots of $x^2 - Px + Q = 0$, then $V_n \mid V_{kn}$ for every odd k, by exactly the argument above, with $\mathbb{Z}[\sqrt{3}]$ replaced by $\mathbb{Z}[\sqrt{P^2 - 4Q}]$ or its full ring of integers in general.

The discernment mechanism: whenever a problem defines a linear recurrence and then asks about one specific, large, composite index, factor that index before doing anything else, and ask which factorisation identity an odd exponent in that factorisation unlocks. Train it by treating "the index is 2015" the same way you would treat "the index is n" in a general olympiad divisibility lemma: write the closed form first, then let the arithmetic of the index, not the recursion, choose your route.

The examiners' nudge is the phrase "find an odd prime factor," not "find the smallest prime factor" or "find all prime factors": it tells you in advance that the intended factor is not unique and is not obtained by full factorisation, which is a green light to stop as soon as one clean factor appears, exactly at a_5.

Revolving concepts

Linear recurrences with constant coefficients. Order two, $a_n = Pa_{n-1} - Qa_{n-2}$, has characteristic equation $x^2 - Px + Q = 0$; distinct real roots α, β give $a_n = A\alpha^n + B\beta^n$, a repeated root α gives $a_n = (A + Bn)\alpha^n$, and complex conjugate roots $re^{\pm i\theta}$ give $a_n = r^n(A\cos n\theta + B\sin n\theta)$. This single tool turns any such recurrence into an explicit formula, and every divisibility question about it starts here. See Book 1, field: Sequences and recurrences, topic: characteristic equation of a linear recurrence.

Lucas sequences and the odd index divisibility law. For α, β roots of $x^2 - Px + Q = 0$, define $U_n = (\alpha^n - \beta^n)/(\alpha - \beta)$ and $V_n = \alpha^n + \beta^n$; both satisfy the same recurrence as a_n. The identity $x^k + y^k = (x + y)(x^{k-1} - \cdots + y^{k-1})$ for odd k gives $V_n \mid V_{kn}$ whenever k is odd, by the integrality argument used above; the companion fact for the U sequence is the stronger $\gcd(U_m, U_n) = U_{\gcd(m,n)}$, needing a different argument, and is the standard generalisation used for the Fibonacci numbers, $P = 1, Q = -1$. See Book 1, field: Number theory, topic: Lucas sequences and divisibility.

Rational algebraic numbers in $\mathbb{Z}[\sqrt{d}]$ are ordinary integers. Any element $p + q\sqrt{d}$, $p, q \in \mathbb{Z}$, d not a perfect square, that happens to be rational must have $q = 0$, since 1 and $\sqrt{d}$ are linearly independent over $\mathbb{Q}$. This one line of algebraic number theory does all the work in the key lemma, and it recurs whenever a sequence is built from two conjugate irrational roots and a divisibility or integrality claim follows. See Book 1, field: Number theory, topic: rings of integers of real quadratic fields.

Factoring $x^n + y^n$ against $x^n - y^n$. The difference always factors, $x^n - y^n = (x - y)(x^{n-1} + \cdots + y^{n-1})$; the sum factors this way only when n is odd. The odd case is the flip side of the fact that $2^n + 1$ can be prime only if n is a power of 2, since otherwise n has an odd factor $k > 1$ and $2^{n/k} + 1$ divides $2^n + 1$. See Book 1, field: Polynomials, topic: factorisations of $x^n \pm y^n$.

The recursion here is order two and constant coefficient throughout, with no further nested or self similar structure, so the generalisation above, arbitrary P, Q and the U_n, V_n pair, is the full

extent of it.

Circuit Bank

- **WHEN** a linear recurrence with constant coefficients is given and you are asked about one specific large index **DO** solve the characteristic equation and write the closed form in the two roots before touching the specific index **BECAUSE** every subsequent divisibility or congruence question becomes an algebraic identity in the roots, not an arithmetic slog through the recursion
- **WHEN** the target index of a recursively defined sequence has an odd factor k **DO** look for a law of the shape $a_n \mid a_{kn}$ coming from $x^k + y^k = (x+y)(\cdots)$ **BECAUSE** a sum of like powers only factors this way for an odd exponent, so an odd cofactor is exactly the signal that the identity is available
- **WHEN** a closed form built from two conjugate irrational roots produces a ratio you can see is rational, both numerator and denominator are terms of an integer sequence **DO** show the ratio also lies in $\mathbb{Z}[\sqrt{d}]$, then conclude it is a plain integer **BECAUSE** a rational element of $\mathbb{Z}[\sqrt{d}]$ has zero irrational part, a one line argument replacing a much longer direct proof of divisibility
- **WHEN** a problem says "find an odd prime factor" rather than "find the prime factorisation" **DO** stop as soon as one clean factor is found, usually at the smallest useful divisor of the index **BECAUSE** the phrasing tells you the intended answer is not unique, so hunting for more is wasted ink

Backtesting

1. First useful cue: the specific index 2015 next to the word "odd," inviting a factorisation of 2015 itself before anything else.
2. Method the expression suggested: a constant coefficient order two recurrence suggested the characteristic equation and closed form immediately.
3. Where you would hesitate, and the thought to reinforce: after finding $a_n = (\alpha^n + \beta^n)/2$, unsure how to connect it to a_{2015}; reinforce that factoring the index first, then hunting for an identity needing an odd exponent, is the route.
4. The trap the question sets: a_{2015} genuinely has hundreds of digits, and a candidate who tries to iterate the recursion by hand is punished with an unfactorable giant number.
5. What makes you recognise it faster next time: any recursively defined sequence with roots multiplying to plus or minus one should now fire "this is a Lucas sequence, check the odd index divisibility law" the way "arithmetic progression" fires "common difference."
6. The reusable decision rule: when a recurrence's closed form has two conjugate roots, test whether a ratio of two of its terms is forced to be an ordinary integer by comparing rationality against membership in the relevant ring of integers.

Mirror drills

1. Let $b_0 = 2, b_1 = 1, b_n = b_{n-1} - b_{n-2}$. Find the period of b_n mod 5 and decide whether 5 divides any term. *Answer:* the roots of $x^2 - x + 1 = 0$ are primitive sixth roots of unity, so b_n is periodic with period 6; the values $b_0, \ldots, b_5$ are $2, 1, -1, -2, -1, 1$, none 0 mod 5, so $5 \nmid b_n$ for any n, a reminder that the odd cofactor circuit only fires once the sequence actually hits a multiple of the target prime.

2. Fibonacci numbers satisfy $F_m \mid F_n$ whenever $m \mid n$, for every m, not only odd cofactors. Explain in one sentence why Fibonacci gets the unrestricted statement while a_n above only gets it for odd cofactors. *Answer:* Fibonacci is a $U_n = (\alpha^n - \beta^n)/(\alpha - \beta)$ type sequence, tied to the identity $x^k - y^k = (x - y)(\cdots)$, valid for every k, whereas a_n is a $V_n = \alpha^n + \beta^n$ type sequence, tied to $x^k + y^k$, which factors through $x + y$ only for odd k.

3. Using the lemma, name three indices $m > 2015$ for which $181 \mid a_m$ follows with no further computation. *Answer:* any $m = 5 \times (\text{odd number})$, for instance $m = 2025 = 5 \times 405$, $m = 2035 = 5 \times 407$, $m = 2045 = 5 \times 409$, since the lemma gives $a_5 \mid a_m$ whenever $m/5$ is odd, and $a_5 = 2 \times 181$.

11.3 A3. Roots of unity collapsing to a gcd sum

Compute

$$\log_2\left(\prod_{a=1}^{2015}\prod_{b=1}^{2015}(1+e^{2\pi iab/2015})\right).$$

Here i is the imaginary unit (that is, $i^2=-1$).

Read it out loud

Say it in plain terms before touching a symbol: you are asked for $\log_2$ of a single complex number, and that number is a double product of 2015^2 factors of the form $1+\zeta^{ab}$, where $\zeta = e^{2\pi i/2015}$ is a primitive 2015th root of unity. Two integer indices a and b each range over a full residue system mod 2015, and they interact only through the product ab in the exponent. Strip away the dressing and this is a problem about the multiplicative structure of exponents modulo 2015, wrapped in a complex analysis shell.

Field tags, out loud: complex numbers (roots of unity, cyclotomic factorisation), number theory (gcd, Euler's totient, multiplicative arithmetic functions), and a sliver of polynomial theory (the identity $x^n-1=\prod(x-\zeta^b)$). There is no calculus here and no geometry; every genuine idea in this problem lives in how a, b and 2015 interact as integers, with the roots of unity acting only as a bookkeeping device for that interaction.

The first cue fires before you write a single equation: $\log_2$ of an unevaluated product, set as a Putnam answer, is a promise. A well posed exam question does not ask you to take $\log_2$ of a number with no reason to be a power of 2, so within the first thirty seconds you should have written two things on your page. First, the product must equal $2^{f(a)}$ for some integer valued f, so the whole answer is $\sum_a f(a)$. Second, $2015 = 5\times 13\times 31$, because any modulus chosen to be the competition year is chosen for its factorisation, and you check that before anything else. A caveat worth clarifying immediately: 2015 is odd. Hold onto that; it turns out to be the one property that keeps every factor in the product away from zero.

The double product is symmetric in a and b only in appearance. Structurally read it as: fix a, let b run over a full residue system mod 2015, and ask what the multiplication by a map does to that residue system. That question, asked of a cyclic group, always has the same answer: it depends on a only through $d=\gcd(a,2015)$. That is the guiding light this problem is built around.

Approaches, ranked worst to best

1. **Direct numerical expansion.** *Cue:* multiply out all 2015^2 complex factors. *Circuit:* no circuit, pure enumeration. (why it ranks here: this is over four million complex multiplications, it produces a decimal approximation rather than a proof, and it uses none of the structure the problem is built on).

2. **Casework on the eight divisors of 2015.** *Cue:* 2015 has divisors $1,5,13,31,65,155,403,2015$, so split a into eight classes by $\gcd(a,2015)$ and treat each separately. *Circuit:* repeated ad hoc use of the same unproven idea, no single lemma. (why it ranks here: it eventually reaches the truth, but forces you to rediscover the same root of unity argument eight times over, at eight times the ink and eight times the chance of a slip on one case).

3. **Solve the coprime case cleanly, patch the rest afterwards.** *Cue:* $\gcd(a,2015)=1$ collapses immediately by a permutation argument, then handle non-coprime a as an afterthought. *Circuit:* half a lemma, stitched rather than unified. (why it ranks here: the coprime case is genuinely the heart of the argument, but treating it in isolation means writing the reduction $n=dn_1$, $a=da_1$ a second time later, instead of building it into the lemma from the first line).

4. **One lemma for every a via $d = \gcd(a, 2015)$, then a multiplicative divisor sum.** *Cue:* the exponent depends on a, b only through ab mod n, and the outer $\log_2$ demands an exponent function. *Circuit:* reduce mod the gcd, evaluate the coprime case by plugging $x = -1$ into $x^{n_1} - 1 = \prod(x - \zeta^c)$, then evaluate $\sum_a \gcd(a, n)$ by Dirichlet convolution and multiplicativity. (this is the route worth taking: it proves $\prod_b (1 + \zeta^{ab}) = 2^{\gcd(a,2015)}$ uniformly, for every a from 1 to 2015 in one pass, then finishes the numerical part in three prime evaluations instead of any casework).

The ideal solution

Lemma 11.3.1. *Let n be a positive odd integer, let $\zeta_n = e^{2\pi i/n}$, and let a be any integer with $d = \gcd(a, n)$. Then*

$$\prod_{b=1}^{n} \left(1 + \zeta_n^{ab}\right) = 2^d.$$

Proof. Write $n = dn_1$ and $a = da_1$ with $\gcd(a_1, n_1) = 1$; since n_1 divides the odd number n, n_1 is itself odd. As b runs over $1, \ldots, n$, the value $ab/n = a_1 b/n_1$ depends on b only through b mod n_1, and each residue class mod n_1 occurs exactly $d = n/n_1$ times among $b = 1, \ldots, n$. Hence

$$\prod_{b=1}^{n} \left(1 + \zeta_n^{ab}\right) = \left(\prod_{c=0}^{n_1 - 1} \left(1 + e^{2\pi i a_1 c/n_1}\right) \right)^d.$$

Because $\gcd(a_1, n_1) = 1$, the map $c \mapsto a_1 c$ mod n_1 permutes $\{0, 1, \ldots, n_1 - 1\}$, so with $\zeta = e^{2\pi i/n_1}$,

$$\prod_{c=0}^{n_1 - 1} \left(1 + e^{2\pi i a_1 c/n_1}\right) = \prod_{c=0}^{n_1 - 1} (1 + \zeta^c).$$

The numbers $\zeta^0, \ldots, \zeta^{n_1 - 1}$ are exactly the roots of $z^{n_1} - 1$, so $z^{n_1} - 1 = \prod_{c=0}^{n_1 - 1}(z - \zeta^c)$. Setting $z = -1$,

$$(-1)^{n_1} - 1 = \prod_{c=0}^{n_1 - 1} (-1 - \zeta^c) = (-1)^{n_1} \prod_{c=0}^{n_1 - 1} (1 + \zeta^c).$$

Since n_1 is odd, $(-1)^{n_1} = -1$, so $-2 = -\prod_{c=0}^{n_1 - 1}(1 + \zeta^c)$, giving $\prod_{c=0}^{n_1 - 1}(1 + \zeta^c) = 2$. Substituting back,

$$\prod_{b=1}^{n} \left(1 + \zeta_n^{ab}\right) = 2^d.$$

□

Apply the lemma with $n = 2015$, which is odd. For every a from 1 to 2015, $\prod_{b=1}^{2015}(1 + e^{2\pi i ab/2015}) = 2^{\gcd(a,2015)}$, so

$$\log_2 \left(\prod_{a=1}^{2015} \prod_{b=1}^{2015} (1 + e^{2\pi i ab/2015}) \right) = \sum_{a=1}^{2015} \gcd(a, 2015).$$

For each divisor d of $n = 2015$, the number of $a \in \{1, \ldots, n\}$ with $\gcd(a, n) = d$ equals the number of $a' = a/d \in \{1, \ldots, n/d\}$ with $\gcd(a', n/d) = 1$, which is $\phi(n/d)$. So

$$\sum_{a=1}^{n} \gcd(a, n) = \sum_{d \mid n} d\, \phi(n/d).$$

This is the Dirichlet convolution of the identity function $n \mapsto n$ with ϕ, and a convolution of two multiplicative functions is multiplicative. In other words, it suffices to know the value on each

prime factor separately. On a single prime p, the divisors are 1 and p, giving $\sum_{d|p} d\,\phi(p/d) = 1\cdot\phi(p) + p\cdot\phi(1) = (p-1)+p = 2p-1$.

Since $2015 = 5\times 13\times 31$ is a product of three distinct primes,

$$\sum_{a=1}^{2015} \gcd(a,2015) = (2\cdot 5-1)(2\cdot 13-1)(2\cdot 31-1) = 9\times 25\times 61.$$

Now $9\times 25 = 225$ and $225\times 61 = 225\times 60 + 225 = 13500 + 225 = 13725$. Hence

$$\log_2\left(\prod_{a=1}^{2015}\prod_{b=1}^{2015}(1+e^{2\pi iab/2015})\right) = \boxed{13725}.$$

Layer 1: WHAT

Block 1: read the outer $\log_2$ as a constraint, not a formatting choice; commit to hunting for an exponent function $f(a)$ with inner product $= 2^{f(a)}$. Block 2: fix a, write $d = \gcd(a,2015)$, $n_1 = 2015/d$, $a_1 = a/d$, and reduce the inner product over b to the dth power of a coprime-index product over n_1 terms. Block 3: use $\gcd(a_1,n_1) = 1$ to permute the coprime-index product into $\prod_c(1+\zeta^c)$ over all n_1th roots of unity, independent of a. Block 4: evaluate that product by plugging $z=-1$ into $z^{n_1}-1 = \prod(z-\zeta^c)$, using that n_1 is odd because 2015 is odd. Block 5: conclude the lemma, take $\log_2$, and turn the double product into $\sum_a \gcd(a,2015)$. Block 6: evaluate that sum by regrouping over divisors, invoking multiplicativity, and multiplying three numbers built from the prime factorisation of 2015.

Juxtapose the zombie route against the ideal one. The nominal route stalls at Block 2: without the gcd substitution, a zombie attack handles $a=1$ (easy, coprime to 2015) and then treats every other a as a separate, unrelated case, either by brute force or by rediscovering the permutation argument from scratch eight times over, once per divisor of 2015. It is worth repeating: the divisor structure of 2015 is not eight separate problems, it is one lemma evaluated at eight different values of d.

The bifurcation point sits exactly between Block 1 and Block 2: do you write $d = \gcd(a,2015)$ before you have solved any single case, or only after solving the coprime case and looking for a way to extend it? Writing it first is the reroute; writing it second is the zombie route wearing the reroute's clothes one step too late, after the ink for the coprime case has already been spent twice.

Layer 2: HOW

Block 1. **Cue.** The final requested quantity is $\log_2$ of an unevaluated product, in a competition that only sets answers that resolve cleanly. **Why natural.** A logarithm of a number the setter expects you to compute exactly is only ever clean when the argument is a power of the log's base; this is a fact about how the problem was constructed, and it correctly predicts that the inner product must equal $2^{\text{something}}$ before you have proved anything about roots of unity. **Link to the next block.** Having fixed the target shape, the only free variable left is which integer exponent depends on a, so you turn to the inner product over b for one fixed a.

Block 2. **Cue.** The exponent $ab/2015$ depends on b only through ab mod 2015, and multiplication by a fixed a on $\mathbb{Z}/2015\mathbb{Z}$ is the acting object. **Why natural.** Multiplication by a on a cyclic group of order n has orbits of size $n/\gcd(a,n)$, with each element of the image hit $\gcd(a,n)$ times; this is bare orbit counting, so writing $d=\gcd(a,n)$, $n = dn_1$, $a = da_1$ is the reflex the moment an index appears only through a product modulo n. **Link to the next block.** Once the inner product is known to be the dth power of a coprime-index product, all remaining work is to evaluate that one coprime-index product.

Block 3. **Cue.** $\gcd(a_1, n_1) = 1$ turns multiplication by a_1 into a genuine bijection of $\mathbb{Z}/n_1\mathbb{Z}$. **Why natural.** A bijection of the index set of a product leaves the product unchanged, since multiplication of complex numbers is commutative; this converts the a-dependent coprime-index product into the a-independent product $\prod_c(1+\zeta^c)$ over every n_1th root of unity, and a product of a linear expression over every root of a known monic polynomial is exactly what the factorisation $x^{n_1}-1=\prod(x-\zeta^c)$ exists to evaluate. **Link to the next block.** Evaluating that factorisation needs a choice of point to substitute, and the natural point is the one that makes the linear factor vanish, namely $x=-1$.

Block 4. **Cue.** n_1 divides the odd number 2015, so n_1 is odd. **Why natural.** Substituting $z=-1$ gives $(-1)^{n_1}-1=(-1)^{n_1}\prod(1+\zeta^c)$; oddness of n_1 forces $(-1)^{n_1}=-1$ and hence $\prod(1+\zeta^c)=2$ rather than 0. Had n_1 been even, -1 would itself be an n_1th root of unity, one factor would vanish, and the original product, and its logarithm, would not exist. This is the exact point where the problem's choice of an odd modulus stops being decoration and starts being load-bearing. **Link to the next block.** With the coprime case pinned at 2, raising to the power d gives the full lemma $\prod_b(1+\zeta^{ab})=2^{\gcd(a,2015)}$ for every a from 1 to 2015.

Block 5. **Cue.** $\log_2$ applied to a product of powers of 2. **Why natural.** $\log_2$ is a homomorphism from the positive reals under multiplication to the reals under addition, so $\log_2\prod_a 2^{\gcd(a,n)}=\sum_a\gcd(a,n)$ without further work; this is the one line where the outer $\log_2$ finally earns its place, converting the multiplicative statement about roots of unity into a purely additive number-theoretic sum. **Link to the next block.** $\sum_{a=1}^{n}\gcd(a,n)$ is a named, well studied object with a standard evaluation route: partition the index set by the value of the gcd.

Block 6. **Cue.** The sum has 2015 terms, but $\gcd(a,2015)$ only ever takes one of eight divisor values. **Why natural.** For each divisor $d \mid n$, exactly $\phi(n/d)$ values of a satisfy $\gcd(a,n)=d$, which regroups the sum as $\sum_{d\mid n} d\,\phi(n/d)$; because $n\mapsto n$ and ϕ are both multiplicative, their convolution is multiplicative, so the whole calculation collapses to one evaluation per prime factor of 2015. **Link to the next block.** With $2015=5\times13\times31$ squarefree, the sum is the product of three single-prime evaluations $2p-1$, and the last step is arithmetic, not mathematics.

Layer 3: WHY

Two mathematical truths carry this entire problem. The first: any product $\prod_b g(ab \bmod n)$, for g a fixed function and a a fixed integer, is governed entirely by the orbit structure of multiplication by a on $\mathbb{Z}/n\mathbb{Z}$, and that orbit structure depends on a only through $\gcd(a,n)$. You should never see an exponent of the form $ab \bmod n$ without immediately asking what $\gcd(a,n)$ is; every other feature of a is invisible to the problem. The second: a product of a fixed linear factor evaluated at every root of a known monic polynomial is not something you multiply out, it is something you read off, because $\prod_b(x_0-\zeta^b)$ is by definition the defining polynomial evaluated at x_0 up to the sign forced by degree. More generally, whenever you must multiply $f(\zeta^b)$ over every nth root of unity for a polynomial f, you are computing a resultant of f with x^n-1; for linear f, as here, that resultant is a single evaluation, and this problem never needs more than that special case.

The jump from a permutation of residues to evaluating a polynomial at a chosen point looks like a quantum leap the first time you meet it, until you notice both steps are answering the same question: which arrangement of the roots makes the product easiest to name. The permutation step tells you the product does not depend on a_1 at all; the evaluation step tells you what the a_1-independent value actually is. Neither is a leap once you see them as two halves of one question.

The discernment mechanism that chooses this route over casework is a habit of asking, before writing a single line of algebra, whether the index set carries algebraic structure. Here the index set is $\mathbb{Z}/2015\mathbb{Z}$ acting on itself by multiplication, a structure you should recognise on sight whenever a

problem's only operations on the indices are addition and multiplication modulo a fixed n. Train this by refusing to write out a sum or product term by term until you have named the group or set acting on the index, and named the orbit sizes. The second habit, layered on top, is checking the modulus for compositeness the moment you see one: 2015 factoring as three distinct primes is the signal that whatever sum survives to the end will be finished by multiplicativity, not by sixteen individual divisor terms.

What the examiners were nudging toward: choosing the modulus to be the competition year is not cosmetic, it is the entire design of the problem. Oddness of 2015 is predicated on nothing more than keeping the product away from zero, which is the reason a solution exists at all; squarefreeness with three prime factors keeps the final arithmetic to three small evaluations instead of a long divisor list. A caveat worth clarifying: nothing about this problem needed 2015 to be prime, or to have any special residue property beyond parity and its factorisation, and a solution that goes looking for deeper number-theoretic properties of 2015 is chasing a red herring the setters did not lay.

Revolving concepts

Roots of unity and the cyclotomic factorisation. Definition: $\zeta_n = e^{2\pi i/n}$, and $\zeta_n^0, \dots, \zeta_n^{n-1}$ are exactly the n solutions of $x^n = 1$, forming a cyclic group of order n under multiplication. Orbiting facts: $x^n - 1 = \prod_{b=0}^{n-1}(x - \zeta_n^b)$; the sum of all nth roots of unity is 0 for $n > 1$; every subgroup of this group has order dividing n and consists of the mth roots of unity for some $m \mid n$. The generalised move: any product of a fixed polynomial evaluated at every nth root of unity is a resultant with $x^n - 1$, collapsing for linear factors to one substitution. See Book 1, field Complex numbers, topic roots of unity and cyclotomic factorisation.

Multiplication orbits modulo n. Definition: for fixed a, the map $b \mapsto ab \bmod n$ on $\mathbb{Z}/n\mathbb{Z}$ has image of size $n/\gcd(a,n)$, with every value in the image attained $\gcd(a,n)$ times. Orbiting facts: Bezout's identity; the subgroup lattice of $\mathbb{Z}/n\mathbb{Z}$ indexed by divisors of n; multiplication by a is a bijection of $\mathbb{Z}/n\mathbb{Z}$ exactly when $\gcd(a,n) = 1$. See Book 1, field Number theory, topic gcd and modular arithmetic.

Euler's totient function. Definition: $\phi(n)$ counts the integers in $\{1, \dots, n\}$ coprime to n. Orbiting facts: ϕ is multiplicative; $\phi(p^k) = p^{k-1}(p-1)$ for a prime p; $\sum_{d \mid n} \phi(d) = n$. See Book 1, field Number theory, topic Euler's totient function.

Multiplicative arithmetic functions and Dirichlet convolution. Definition: f is multiplicative if $f(1) = 1$ and $f(mn) = f(m)f(n)$ whenever $\gcd(m,n) = 1$; the Dirichlet convolution of f and g is $(f * g)(n) = \sum_{d \mid n} f(d)g(n/d)$, and the convolution of two multiplicative functions is multiplicative. This is itself a recursion in disguise: given a multiplicative f's values on prime powers, f on any n is built up by peeling off one prime power at a time and multiplying, base case $f(1) = 1$; install this as the general method for evaluating any multiplicative function, whether it is ϕ, the divisor count, the divisor sum, or the gcd sum below. See Book 1, field Number theory, topic multiplicative functions and Dirichlet convolution.

The gcd sum function. Definition: $P(n) = \sum_{a=1}^{n} \gcd(a,n)$, sometimes called Pillai's arithmetical function. Orbiting facts: $P(n) = \sum_{d \mid n} d\,\phi(n/d) = (\mathrm{id} * \phi)(n)$, multiplicative; on a prime power, $P(p^k) = p^k + k\,p^{k-1}(p-1)$, and in particular $P(p) = 2p - 1$. See Book 1, field Number theory, topic gcd sum functions.

Circuit Bank

- **WHEN** you must multiply a fixed polynomial expression evaluated at every nth root of unity **DO** evaluate the cyclotomic factorisation $x^n - 1 = \prod_b (x - \zeta^b)$ at the point that kills the polynomial's own factor, rather than multiplying the roots out one by one **BECAUSE**

a product over all roots of a known monic polynomial is that polynomial evaluated at the complementary point, up to a sign fixed by the degree

- **WHEN** an index b appears in an exponent only through $ab \bmod n$ for a fixed integer a **DO** write $d = \gcd(a, n)$, $n = dn_1$, $a = da_1$ with $\gcd(a_1, n_1) = 1$, and reduce to the coprime case raised to the dth power **BECAUSE** multiplication by a on $\mathbb{Z}/n\mathbb{Z}$ has orbits of size $n/\gcd(a, n)$, so the gcd is the only feature of a the final answer can depend on
- **WHEN** a sum runs over all $1 \le a \le n$ and the summand depends on a only through $\gcd(a, n)$ **DO** regroup by divisor $d \mid n$, replace each class's size with $\phi(n/d)$, and check whether the resulting divisor sum is a Dirichlet convolution of known multiplicative functions **BECAUSE** grouping by gcd converts a length n sum into a sum over the divisors of n, and multiplicativity finishes it in one evaluation per prime factor
- **WHEN** the modulus or parameter in a Putnam problem is the competition year itself **DO** factorise it before writing anything else, and check its parity and whether it is squarefree **BECAUSE** setters choose the year as the parameter precisely because its factorisation drives the finishing arithmetic, and parity and squarefreeness are the two properties most often load-bearing
- **WHEN** the requested answer is $\log_b$ of an unevaluated expression **DO** treat this as a promise that the expression equals b to an integer power, and hunt for that exponent as a computable arithmetic function rather than evaluating the expression directly **BECAUSE** a well posed exam answer is never the logarithm of a number with no reason to be a clean power of the base, so the outer function tells you the shape of the lemma to look for before you have proved anything

Backtesting

1. First useful cue: $\log_2$ wraps an unevaluated product, and the modulus 2015 is visibly odd.
2. Method the expression suggested: reduce the inner product over b by the orbit of multiplication by a modulo 2015, writing $d = \gcd(a, 2015)$ before attempting any single case.
3. Where you would hesitate, and the thought to reinforce: hesitating over whether coprime and non-coprime a need separate lemmas; reinforce that $n = dn_1$, $a = da_1$ swallows both in one substitution, so never solve the coprime case in isolation first.
4. The trap the question sets: forgetting to justify why oddness of 2015 is necessary, and so presenting a lemma that would silently claim a nonzero product even when n is even.
5. What makes you recognise it faster next time: seeing an exponent of the form $ab \bmod n$ should now instantly summon $d = \gcd(a, n)$, and seeing $\log_b$ of an unevaluated expression should instantly summon a search for an exponent function; that sentence has to be self-generated, in the arena, not read off a memorised template.
6. The reusable decision rule: when an exponent or index depends on two integers only through their product modulo n, substitute $d = \gcd(a, n)$ and reduce to the coprime case, because permutation arguments and cyclotomic evaluations both require coprimality to hold. Install this as the first line you write whenever you meet the pattern.

Mirror drills

1. Compute $\prod_{b=1}^{9} \left(1 + e^{2\pi i b/9}\right)$. *Answer:* 9 is odd, so the same substitution $z = -1$ into $z^9 - 1 = \prod(z - \zeta^b)$ applies directly with $d = 1$, giving 2.

2. For $n = 45 = 3^2 \times 5$, evaluate $\sum_{a=1}^{45} \gcd(a, 45)$. *Answer:* multiplicativity gives $P(9) \times P(5)$, with $P(9) = 9 + 2 \cdot 3 \cdot 2 = 21$ and $P(5) = 2 \cdot 5 - 1 = 9$, so the sum is 189.
3. Explain in one line why $\prod_{b=1}^{n} \left(1 + e^{2\pi i b/n}\right) = 0$ whenever n is even. *Answer:* $b = n/2$ gives $e^{2\pi i b/n} = e^{i\pi} = -1$, so that single factor is $1 + (-1) = 0$.

11.4 A4. Binary weights on an even floor set

For each real number x, let

$$f(x) = \sum_{n \in S_x} \frac{1}{2^n},$$

where S_x is the set of positive integers n for which $\lfloor nx \rfloor$ is even. What is the largest real number L such that $f(x) \geq L$ for all $x \in [0, 1)$? (As usual, $\lfloor z \rfloor$ denotes the greatest integer less than or equal to z.)

Read it out loud

This is a subset sum of the geometric series $\sum 2^{-n}$, where the subset is chosen for you by a parity test on a floor function. There are exactly two moving parts: which n survive the test, and what a surviving n is worth. The second part is already installed in you and must not be re-derived in the arena: with ratio $1/2$, each term equals the sum of the entire tail after it, $2^{-n} = \sum_{m>n} 2^{-m}$. So an early survivor outweighs everything that could follow it. The consequence is immediate and it is the whole strategic read: to make $f(x)$ small you must vacate the early positions, and once you cannot vacate a position any more, the game is over. The question is therefore not analytic at all. It is combinatorial: how sparse, and how late, can S_x be forced to be?

The fields wired in are geometric series (atomic), the floor function used as a slow counter, and the distinction between an infimum and a minimum. Notice the phrasing: "the largest real number L such that $f(x) \geq L$ for all x". That is the definition of an infimum, written out longhand, and the examiners wrote it longhand rather than saying "the minimum value of f" because the value is approached and never attained. Suspecting that costs you nothing and stops you wasting time hunting for a minimiser that does not exist.

Within thirty seconds you should have two lines on the page. First: $\lfloor x \rfloor = 0$ for every $x \in [0, 1)$, so $1 \in S_x$ always, hence $f(x) \geq 1/2$ for every x, and $L \in [1/2, 1)$. Second: for $x < 1$ the step from nx to $(n+1)x$ is less than 1, so $\lfloor (n+1)x \rfloor - \lfloor nx \rfloor \in \{0, 1\}$. That second line is the engine of the whole problem, because a counter that moves by 0 or 1 cannot stay odd for long without standing completely still.

Everything now hangs on one number: the longest run of consecutive integers that can all fail the parity test. If runs of missing integers have length at most g, then from position 2 onward every block of g consecutive integers donates something, and the sum of the worst case is a geometric series you can write down without thinking. So find g, then add up.

Approaches, ranked worst to best

1. **Grid evaluation and pattern spotting.** *Cue:* a completely explicit series and the words "largest real number". *Circuit:* evaluate, tabulate, conjecture. Compute f at a few hundred rationals and watch the minimum settle near 0.5714. This gives you the target $4/7$ and, far more usefully, its address: the small values live just below $x = 2/3$. As a proof it is worth zero marks. As thirty seconds of reconnaissance it is worth a lot, because the second half of the problem asks you to reverse engineer an x, and now you know where to look. Ranks last because it proves nothing.

2. **Continued fraction classification of the parity word.** *Cue:* $\lfloor nx \rfloor$ is a Beatty sequence. *Circuit:* Sturmian words and the three distance theorem. You can describe the parity pattern of $\lfloor nx \rfloor$ completely through the continued fraction of x. It is correct and it is enormous. You would be importing a theory of infinite words to establish a fact about runs of length three. Costs pages of ink for one line of content.

3. **Lexicographic comparison with the extremal set.** *Cue:* the weights 2^{-n} are literally binary place values. *Circuit:* one place value dominates its whole tail. Set $T = \{1, 4, 7, 10, \dots\}$, note $F(T) = 4/7$, and argue that if $f(x) < 4/7$ then the first position where S_x and T disagree must belong to T, which forces three consecutive integers to be absent from S_x, which forces $x < 1/2$, which contradicts. This is the published route and it is sound. It ranks below the next one only because it runs through a contradiction and needs the comparison lemma stated with care, when the same information is available directly.

4. **Run length bound plus block partition.** *Cue:* consecutive values of nx differ by $x < 1$. *Circuit:* a step smaller than 1 cannot skip a parity class three times running. Prove that three consecutive integers can all miss S_x only when $x < 1/2$; kill $x < 1/2$ in one line because then $2 \in S_x$; then chop $\{2, 3, 4, \dots\}$ into blocks of three and charge each block its worst element. No contradiction, no comparison lemma, and the equality pattern falls out of the proof, which is exactly what you need for the sharpness half. This is the ideal route.

The ideal solution

The answer is $\boxed{L = \frac{4}{7}}$.

Throughout, $x \in [0, 1)$ and $S_x = \{n \in \mathbb{Z}_{>0} : \lfloor nx \rfloor \text{ is even}\}$.

Lemma 11.4.1. *For every $x \in [0, 1)$ and every positive integer n, $\lfloor (n+1)x \rfloor - \lfloor nx \rfloor \in \{0, 1\}$.*

Proof. Since $0 \le x < 1$ we have $nx \le (n+1)x < nx + 1$. The floor function is nondecreasing, so $\lfloor nx \rfloor \le \lfloor (n+1)x \rfloor \le \lfloor nx + 1 \rfloor = \lfloor nx \rfloor + 1$. □

Lemma 11.4.2. *If $x \in [1/2, 1)$, then no three consecutive positive integers all lie outside S_x.*

Proof. Suppose $\lfloor nx \rfloor$, $\lfloor (n+1)x \rfloor$ and $\lfloor (n+2)x \rfloor$ are all odd. The difference of two odd integers is even, so by Lemma 1 each of the two consecutive differences is 0. Hence $\lfloor nx \rfloor = \lfloor (n+2)x \rfloor$, and two reals with the same floor differ by less than 1, so $(n+2)x - nx = 2x < 1$. That contradicts $x \ge 1/2$. □

Proposition 11.4.3. *$f(x) \ge 4/7$ for every $x \in [0, 1)$.*

Proof. Since $\lfloor x \rfloor = 0$ is even, $1 \in S_x$, so the term $1/2$ is always present.

If $x < 1/2$ then $\lfloor 2x \rfloor = 0$ is even, so $2 \in S_x$ as well and $f(x) \ge \frac{1}{2} + \frac{1}{4} = \frac{3}{4} > \frac{4}{7}$.

If $x \ge 1/2$, partition the integers from 2 upward into the blocks

$$B_k = \{3k + 2,\ 3k + 3,\ 3k + 4\}, \qquad k = 0, 1, 2, \dots,$$

which are disjoint and cover $\{2, 3, 4, \dots\}$. Each B_k consists of three consecutive integers, so by Lemma 2 it contains some element of S_x, and that element contributes at least $2^{-(3k+4)}$, the value of the largest member of B_k. Adding the guaranteed term from $n = 1$,

$$f(x) \ \ge\ \frac{1}{2} + \sum_{k=0}^{\infty} \frac{1}{2^{3k+4}} = \frac{1}{2} + \frac{1/16}{1 - 1/8} = \frac{1}{2} + \frac{1}{14} = \frac{4}{7}.$$ □

Proposition 11.4.4. *For every $\varepsilon > 0$ there exists $x \in [0, 1)$ with $f(x) < 4/7 + \varepsilon$.*

Proof. Fix a positive integer N and set $x = \frac{2}{3} - \delta$ with $0 < \delta < \frac{1}{9N}$, so that $0 < m\delta < \frac{1}{3}$ for every $m \le 3N$. For $1 \le k \le N$,

$$\begin{aligned}
&(3k-2)x = (2k-2) + \tfrac{2}{3} - (3k-2)\delta, && \text{fractional part in } (\tfrac{1}{3}, \tfrac{2}{3}), && \lfloor (3k-2)x \rfloor = 2k-2,\\
&(3k-1)x = (2k-1) + \tfrac{1}{3} - (3k-1)\delta, && \text{fractional part in } (0, \tfrac{1}{3}), && \lfloor (3k-1)x \rfloor = 2k-1,
\end{aligned}$$

$$(3k)x = 2k - 3k\delta, \qquad \text{value in } (2k - \tfrac{1}{3}, 2k), \qquad \lfloor (3k)x \rfloor = 2k - 1.$$

So among $1, 2, \ldots, 3N$ the elements of S_x are exactly $1, 4, 7, \ldots, 3N - 2$. Bounding the tail beyond $3N$ by the whole of it,

$$f(x) \;\le\; \sum_{k=1}^{N} \frac{1}{2^{3k-2}} + \sum_{n>3N} \frac{1}{2^n} = \frac{4}{7}\left(1 - 8^{-N}\right) + 8^{-N} < \frac{4}{7} + 8^{-N}.$$

Choosing N with $8^{-N} < \varepsilon$ gives the claim. □

The two propositions together say that $4/7$ is a lower bound and that nothing larger is, so $L = 4/7$.

Remark 11.4.5. The value is never attained. Since $4/7 = 0.\overline{100}$ in binary and that expansion neither terminates nor ends in all ones, the only set with $F(S) = 4/7$ is $S = \{1, 4, 7, \ldots\}$. If S_x were that set, then by Lemma 1 the floor must increase by exactly 1 at each of the two steps inside a block and by 0 at the block boundary, so $\lfloor (3k+1)x \rfloor = 2k$ and $\lfloor (3k+3)x \rfloor = 2k+1$ for all $k \ge 0$. The first gives $x \ge 2k/(3k+1)$ for all k, hence $x \ge 2/3$; the second gives $(3k+3)x < 2k+2$, hence $x < 2/3$. So $f(x) > 4/7$ for every x, and the bound is an infimum only.

Layer 1: WHAT

Block 1. $\lfloor x \rfloor = 0$, so $1 \in S_x$ for every x in the range: the term $1/2$ is free and permanent.

Block 2. Because the weights halve, minimising f means vacating positions as early and as long as possible; so ask how long a gap in S_x can be.

Block 3. The step $\lfloor (n+1)x \rfloor - \lfloor nx \rfloor$ is 0 or 1 because $x < 1$; three odd floors in a row therefore force two steps of 0, hence $2x < 1$.

Block 4. So a triple gap forces $x < 1/2$, and $x < 1/2$ puts 2 into S_x, giving $f(x) \ge 3/4$. That branch is dead, well above the target.

Block 5. On the surviving branch every three consecutive integers meet S_x. Chop $\{2, 3, \ldots\}$ into blocks of three, charge each block its worst (largest) element, sum the geometric series: $1/2 + 1/14 = 4/7$.

Block 6. The worst case in Block 5 is the pattern $\{1, 4, 7, \ldots\}$, so engineer an x that realises it for as long as you like: the pattern says $\lfloor nx \rfloor$ is even exactly on $n \equiv 1 \pmod 3$, which is the behaviour of x just below $2/3$. Verify the three floors and let the truncation run to infinity.

Juxtaposition. The nominal route treats f as a function to be evaluated: tabulate f at rationals, try to write a formula for $f(x)$, perhaps try to integrate or to differentiate something. It goes nowhere because f is nowhere continuous and has no formula. The ideal route never evaluates f at all except at the very end; it studies the set S_x and quotes the series. The bifurcation point is the sentence "how long can a run of consecutive integers all fail the test?". Everything before that sentence is arithmetic, everything after it is a two line proof.

Layer 2: HOW

Block 1. Cue. The domain $x \in [0, 1)$ appears before anything else, and the first index is $n = 1$. **Why natural.** Testing the smallest case is not a habit, it is a duty: $\lfloor 1 \cdot x \rfloor = 0$ is even for the entire domain, so one term of the series is unconditional. A term that is present for every x is a floor on the answer, and it also fixes the scale: the answer is between $1/2$ and 1. **Link to the next block.** Having one guaranteed term, you now want to know how much more is guaranteed, which means asking what the parity test can and cannot do.

Block 2. **Cue.** The weights are 2^{-n}, and the question asks for a greatest lower bound. **Why natural.** With ratio $1/2$ a term equals its own tail, so the comparison of two subset sums is decided at the first position where they differ. Minimising is therefore a scheduling problem: push each forced element as late as the constraints allow. This reframing is forced by the ratio, not chosen. **Link to the next block.** Scheduling needs constraints, so the next move is to find the constraint on S_x.

Block 3. **Cue.** $x < 1$ multiplying an integer counter n. **Why natural.** $\lfloor nx \rfloor$ advances by $(n+1)x - nx = x < 1$ per step, so it is a counter that moves by 0 or 1. Parity therefore alternates whenever the counter moves and repeats whenever it stalls, so a long run of one parity is a long stall, and a long stall is a statement that x is small. That is the definition-level reason: the floor of a slowly increasing sequence cannot repeat a parity three times without repeating a value. **Link to the next block.** Turn the stall into an inequality on x and see which x survive.

Block 4. **Cue.** A triple stall gives $2x < 1$. **Why natural.** Small x is exactly the regime where the early integers all have floor 0, which is even, so small x is the worst possible regime for someone trying to make f small. The case split is not a nuisance, it is the mechanism: the enemy cannot both starve the early terms and have a long gap. **Link to the next block.** With the small x branch dead, the remaining branch has a hard structural constraint you can sum.

Block 5. **Cue.** Every window of three consecutive integers meets S_x. **Why natural.** Disjoint blocks are the right partition because the constraint is stated on windows, and a lower bound on a sum of nonnegative terms may be computed block by block. Starting the blocks at 2 rather than at 1 is deliberate: it keeps the free term $1/2$ separate and aligns the blocks so that the worst case is $\{1, 4, 7, \dots\}$ rather than a weaker pattern. Charging each block its largest element is the worst case by construction. **Link to the next block.** The worst case pattern is now explicit, so the natural next question is whether it can be realised.

Block 6. **Cue.** $4/7 = 0.\overline{100}$ in binary, a period of three starting at position 1. **Why natural.** A repeating binary answer is a blueprint for the extremal set, and the extremal set is a blueprint for x: you need $\lfloor nx \rfloor$ even exactly when $n \equiv 1 \pmod 3$, that is, the counter must gain exactly 2 every three steps, so x must be just under $2/3$. Approaching from below rather than sitting at $2/3$ is forced, because at $x = 2/3$ the value $\lfloor 3kx \rfloor = 2k$ is even and the pattern breaks in the wrong direction. **Link to the next block.** Verify the three residue classes, bound the tail crudely, and let N grow.

Layer 3: WHY

Three truths carry this problem. The first is the self-dominating property of the ratio $1/2$: $2^{-n} = \sum_{m>n} 2^{-m}$. It converts a question about a sum into a question about the earliest admissible position, which is why the answer is a repeating binary fraction and why the extremal configuration is periodic rather than exotic. Whenever you see weights r^n with $r \le 1/2$, expect the extremal problem to be solved by a greedy or blockwise argument, never by calculus.

The second is that $n \mapsto \lfloor nx \rfloor$ with $0 \le x < 1$ is a counter with step 0 or 1. All the rigidity in the problem comes from that single fact. Runs of a fixed parity are stalls, stalls are upper bounds on x, and the interaction between "small x gives long stalls" and "small x gives even floors early" is the tension the problem is built on. In general the maximum stall length for $x \in (0, 1)$ is $\lceil 1/x \rceil$, so the enemy is choosing x to trade stall length against how quickly the counter reaches an odd value at all.

The third is the difference between an infimum and a minimum, which the wording flags for you. When the extremal object is periodic but the constraint is only satisfiable for finitely many

periods at a time, the bound is approached and not attained. Recognising this early stops you from trying to prove that some specific x is a minimiser, and tells you the second half of the proof will be a family x_N with a parameter, not a single point.

The discernment mechanism to train: when a sum runs over a set defined by an arithmetic condition, do not try to describe the set. Find the strongest simple combinatorial constraint the condition imposes (here, maximum gap length), then solve the resulting clean extremal problem, then check that the extremal configuration is achievable. Describe less, constrain more. The examiners nudged you toward exactly this by using base 2 (so that place value logic dominates), by starting the domain at 0 (so that $n = 1$ is free), and by choosing the interval $[0, 1)$ (so that the step is less than 1 and the counter is slow).

Revolving concepts

Geometric series with ratio $1/2$. The identity $\sum_{m>n} 2^{-m} = 2^{-n}$ says binary place values are exactly self-dominating, so subset sums of $\{2^{-n}\}$ are ordered lexicographically by the sets, with ties only at the dyadic rationals (the $0.1000\ldots = 0.0111\ldots$ ambiguity). This is why $4/7 = 0.\overline{100}$ has a unique representing set. See Book 1, field: Real analysis, topic: geometric series and binary expansions.

The floor function as a counter. For $x \in [0, 1)$, $\lfloor nx \rfloor$ counts how many multiples of 1 the ray nx has passed. Standard companions: $\lfloor y \rfloor + \lfloor z \rfloor \le \lfloor y + z \rfloor \le \lfloor y \rfloor + \lfloor z \rfloor + 1$; $\lfloor \lfloor y \rfloor / d \rfloor = \lfloor y/d \rfloor$ for a positive integer d; and Hermite's identity $\sum_{j=0}^{d-1} \lfloor y + j/d \rfloor = \lfloor dy \rfloor$. See Book 1, field: Number theory, topic: floor function identities.

Beatty and Sturmian structure. The sequence $\lfloor nx \rfloor$ is a Beatty sequence; its difference word $\lfloor (n+1)x \rfloor - \lfloor nx \rfloor$ is a Sturmian word over $\{0, 1\}$ with density x. The three distance theorem controls the gaps between the fractional parts $\{nx\}$. You do not need this here, but it is the correct home for any harder problem about $\lfloor nx \rfloor$. See Book 1, field: Sequences and recurrences, topic: Beatty sequences.

The recursion, generalised. The parity word $e_n = \lfloor nx \rfloor \bmod 2$ satisfies an exact recursion whenever x is rational. Write $x = p/q$ in lowest terms; then $\lfloor (n+q)x \rfloor = \lfloor nx + p \rfloor = \lfloor nx \rfloor + p$, so

$$e_{n+q} \equiv e_n + p \pmod 2.$$

Unrolling this is the general solution: $e_{n+jq} \equiv e_n + jp$, so the word is periodic with period q when p is even and period $2q$ when p is odd. Every question about parities of $\lfloor nx \rfloor$ at rational x reduces to computing one period. The same unrolling handles $\lfloor nx \rfloor$ modulo any m: the word is periodic with period $q \cdot m / \gcd(p, m)$. Applied here with x near $2/3$, the period is 3 and the pattern $(0, 1, 1)$ is exactly the block structure the proof used.

Infimum versus minimum, and the extremal series. The general block bound the proof used is worth storing in full generality: if S contains 1 and meets every window of g consecutive integers, then for $0 < r < 1$,

$$\sum_{n \in S} r^n \;\ge\; r + \frac{r^{g+1}}{1 - r^g},$$

with the bound approached by $S = \{1,\ 1+g,\ 1+2g, \ldots\}$. For $r = 1/2$, $g = 3$ this is $1/2 + 1/14 = 4/7$; for $r = 1/3$, $g = 3$ it is $1/3 + 1/78 = 9/26$. See Book 1, field: Inequalities, topic: extremal subset sums of geometric series.

Circuit Bank

- **WHEN** a sum runs over a set defined by a parity or divisibility test on $\lfloor nx \rfloor$ **DO** ask first how long a run of consecutive integers can fail the test **BECAUSE** for $x < 1$ the counter $\lfloor nx \rfloor$ moves by 0 or 1, so a long run of one parity is a long stall and forces x to be small, which is a usable inequality

- **WHEN** the weights are r^n with $r \leq 1/2$ and you must minimise a subset sum **DO** schedule the forced elements as late as the constraint permits, then sum the periodic worst case as a geometric series **BECAUSE** each term dominates its own tail, so the comparison is decided at the first differing position and greedy scheduling is optimal
- **WHEN** the question says "largest L with $f \geq L$ for all x" instead of "the minimum of f" **DO** prove a universal bound and separately construct a family approaching it, and expect no minimiser **BECAUSE** that phrasing is the definition of an infimum, and the two halves need different tools
- **WHEN** the answer is a repeating binary or base r fraction **DO** read the period as the pattern of admissible indices and reverse engineer the parameter that produces it **BECAUSE** a repeating expansion is the fingerprint of a periodic extremal configuration, which is a blueprint for the construction

Backtesting

1. First useful cue: $x \in [0, 1)$ with the index starting at $n = 1$, giving $\lfloor x \rfloor = 0$ and the unconditional term $1/2$; immediately after it, the step bound $\lfloor (n+1)x \rfloor - \lfloor nx \rfloor \in \{0, 1\}$.
2. Method the expression suggested: not evaluation of f, but a gap bound on S_x followed by a blockwise geometric sum, then an explicit family x_N realising the worst pattern.
3. Where you would hesitate, and the thought to reinforce: at the case split $x < 1/2$ against $x \geq 1/2$. The thought to reinforce is that a case split forced by the mechanism is not untidiness, it is the mechanism; the enemy cannot have both a long gap and starved early terms, and that is precisely why the answer exists.
4. The trap the question sets: hunting for a minimiser. The infimum is not attained, so any argument that assumes a minimising x exists will fail or, worse, will produce a plausible wrong constant such as $5/7$, the value at $x = 2/3$ itself.
5. What makes you recognise it faster next time: the pairing of base 2 weights with a floor parity test. Base 2 says place value logic; a floor of nx with $x < 1$ says slow counter; together they say gap analysis.
6. The reusable decision rule: when summing over an arithmetically defined set, bound the maximum gap, partition into blocks of that length, charge each block its worst element, then test achievability of the resulting periodic pattern.

Mirror drills

1. Replace 2^n by 3^n: find the largest L with $\sum_{n \in S_x} 3^{-n} \geq L$ for all $x \in [0, 1)$. *Answer:* $9/26$, by the same two lemmas and the general bound $r + r^{g+1}/(1 - r^g)$ with $r = 1/3$, $g = 3$; the same family $x = 2/3 - \delta$ shows it is sharp.
2. Let S be any set of positive integers meeting every block of three consecutive integers. Show $\sum_{n \in S} 2^{-n} \geq 1/7$ and that $1/7$ is best possible. *Answer:* charge each block $\{3k+1, 3k+2, 3k+3\}$ its largest element to get $\sum_{k \geq 0} 2^{-(3k+3)} = 1/7$; the set $\{3, 6, 9, \dots\}$ attains it.
3. Show that $f(x) > 4/7$ strictly for every $x \in [0, 1)$. *Answer:* equality forces $S_x = \{1, 4, 7, \dots\}$ by uniqueness of the binary expansion of $4/7$, which by the step lemma forces $\lfloor (3k+1)x \rfloor = 2k$ and $\lfloor (3k+3)x \rfloor = 2k+1$ for all k, giving $x \geq 2/3$ and $x < 2/3$ at once.

11.5 A5. Parity of the coprime count below a quarter

Let q be an odd positive integer, and let N_q denote the number of integers a such that $0 < a < q/4$ and $\gcd(a, q) = 1$. Show that N_q is odd if and only if q is of the form p^k with k a positive integer and p a prime congruent to 5 or 7 modulo 8.

Read it out loud

Say the conclusion first, because the conclusion is the cue. "A prime congruent to 5 or 7 modulo 8" is not a condition anybody invents for fun. Conditions on a prime modulo 8 come from exactly one place in mathematics: the supplementary laws of quadratic reciprocity. You know that $p \equiv \pm 1 \pmod 8$ is the criterion for 2 to be a quadratic residue and that $p \equiv 1 \pmod 4$ is the criterion for -1; multiply the two characters together and the classes $1, 3$ split from $5, 7$. So before touching the counting, you already know the target: N_q is odd exactly when -2 is a quadratic nonresidue modulo p. The problem is a quadratic residue problem wearing a counting costume.

Now say the hypothesis out loud. You are counting the integers coprime to q in the interval $(0, q/4)$, and you want only the parity of that count. Coprime residues, a fixed fraction of the modulus, and a parity: that is the signature of Gauss's lemma, which is the statement that multiplying every element of the first half of the residues by a fixed integer and counting how many wrap past the midpoint gives you a quadratic character as a sign. The link between the two costumes is one line of trivial algebra: $a < q/4$ if and only if $2a < q/2$. Multiplication by 2 carries the quarter interval onto the half interval, so N_q is precisely the count of a in the first half whose double does not wrap.

The fields wired in are quadratic residues (Legendre symbol, Euler's criterion, the two supplementary laws), elementary group theory modulo q (Euler's theorem, the Chinese remainder theorem, square roots of unity in $(\mathbb{Z}/p^k)^\times$), and the totient function. The counting flavour is a decoy; a Mobius inversion route exists and works, but it hides the reason the answer looks the way it does.

Within thirty seconds you should have written $a < q/4 \iff 2a < q/2$, the definition $A = \{a : 1 \le a \le (q-1)/2,\ \gcd(a, q) = 1\}$ with $|A| = \varphi(q)/2$, and the note that $q/4$ is never an integer because q is odd, so there are no boundary cases to worry about anywhere.

Approaches, ranked worst to best

1. **Tabulate and match.** *Cue:* a clean numerical claim about all odd q. *Circuit:* compute N_q for $q \le 100$ and check the classification. Useful for confidence and for catching a misread (for instance that $q = 1$ gives $N_q = 0$, which must land on the even side). It proves nothing, and a "verify the pattern" write-up scores zero.

2. **Direct evaluation of N_q by a closed formula.** *Cue:* an interval count with a coprimality condition. *Circuit:* inclusion and exclusion over the prime divisors of q. You can write $N_q = \sum_{d \mid q} \mu(d) \lfloor q/(4d) \rfloor$ and try to evaluate it. As an exact formula it is unusable; the floors do not simplify.

3. **Mobius inversion reduced modulo 2.** *Cue:* only the parity is wanted, so all coefficients live in $\mathbb{F}_2$. *Circuit:* squarefree divisors survive, $\mu(d)^2 = 1$, and $\lfloor (q/d)/4 \rfloor \bmod 2$ depends only on $q/d \bmod 8$. This is the published first solution and it is entirely correct: reduce the divisor sum modulo 2, note that $\lfloor m/4 \rfloor$ is odd exactly when $m \equiv 5, 7 \pmod 8$, then pair up the squarefree divisors to show the count of bad divisors is even unless q is a prime power. It ranks here because the pairing argument needs three separate cases and, more importantly, because it never explains why 5 and 7 modulo 8 appear. You finish with a proof and no understanding.

4. **Gauss's lemma with multiplier 2, generalised to composite modulus.** *Cue:* $q/4$ with coprimality and a parity, plus a conclusion stated modulo 8. *Circuit:* doubling the first half of the residues and counting wrap-arounds is a signed permutation, and the sign is a quadratic character. Prove the congruence $(-1)^{N_q} \equiv (-2)^{\varphi(q)/2} \pmod q$, then read off both directions: for q with two or more prime factors the right side collapses to 1 by the Chinese remainder theorem, and for $q = p^k$ it is ± 1 and Euler's criterion identifies it as the Legendre symbol of -2. Shortest, and it explains the answer: the condition $p \equiv 5, 7 \pmod 8$ is the condition that -2 is a nonresidue. This is the ideal route.

The ideal solution

Throughout q is odd. Handle $q = 1$ at once: there is no integer a with $0 < a < 1/4$, so $N_1 = 0$ is even, and 1 is not of the stated form. Assume $q \geq 3$ from here on, and set

$$A = \left\{ a \in \mathbb{Z} : 1 \leq a \leq \tfrac{q-1}{2},\ \gcd(a, q) = 1 \right\}.$$

Since the integers in $\{1, \dots, q-1\}$ coprime to q come in pairs $\{a, q-a\}$ with exactly one member in each half, $|A| = \varphi(q)/2$. Because q is odd, $q/4$ is not an integer, so

$$N_q = \#\{a \in A : a < q/4\} = \#\{a \in A : 2a < q/2\}.$$

Lemma 11.5.1. *Let $n = \#\{a \in A : 2a > q/2\}$. Then $2^{\varphi(q)/2} \equiv (-1)^n \pmod q$ and $N_q = \varphi(q)/2 - n$.*

Proof. For $a \in A$ we have $2 \leq 2a \leq q-1$, so $2a$ is congruent modulo q to exactly one element $r(a)$ of $\{-(q-1)/2, \dots, -1, 1, \dots, (q-1)/2\}$, namely $r(a) = 2a$ if $2a < q/2$ and $r(a) = 2a - q$ if $2a > q/2$ (equality is impossible since q is odd). Thus $r(a) < 0$ for exactly n values of a.

The map $a \mapsto |r(a)|$ sends A to A: the value lies in $\{1, \dots, (q-1)/2\}$ and is coprime to q because $2a$ and $q - 2a$ are. It is injective: if $|r(a)| = |r(b)|$ then $2a \equiv \pm 2b \pmod q$, so $a \equiv \pm b \pmod q$ since 2 is invertible; the sign $-$ is impossible because $0 < a + b < q$, so $a = b$. A injection from the finite set A to itself is a bijection.

Multiplying over A and using $2a \equiv r(a)$,

$$2^{|A|} \prod_{a \in A} a \ \equiv\ \prod_{a \in A} r(a) \ =\ (-1)^n \prod_{a \in A} |r(a)| \ =\ (-1)^n \prod_{a \in A} a \pmod q.$$

The product $\prod_{a \in A} a$ is coprime to q, hence invertible, and cancelling it gives $2^{\varphi(q)/2} \equiv (-1)^n$. Finally A is split by the condition $2a < q/2$ into the N_q elements counted by the problem and the n remaining ones, so $N_q = |A| - n = \varphi(q)/2 - n$. □

Corollary 11.5.2. *For every odd $q \geq 3$, $(-1)^{N_q} \equiv (-2)^{\varphi(q)/2} \pmod q$.*

Proof. $(-1)^{N_q} = (-1)^{\varphi(q)/2 - n} = (-1)^{\varphi(q)/2}(-1)^n \equiv (-1)^{\varphi(q)/2} 2^{\varphi(q)/2} = (-2)^{\varphi(q)/2} \pmod q$. □

Since $q \geq 3$, the residues $+1$ and -1 are distinct modulo q, so the corollary determines the parity of N_q completely: N_q is odd if and only if $(-2)^{\varphi(q)/2} \equiv -1 \pmod q$.

Proposition 11.5.3. *If q has at least two distinct prime factors, then N_q is even.*

Proof. Write $q = \prod_{i=1}^{r} p_i^{k_i}$ with $r \geq 2$. Fix i. Then

$$\frac{\varphi(q)}{\varphi(p_i^{k_i})} = \prod_{j \neq i} \varphi(p_j^{k_j})$$

is even, because each p_j is odd and so each $\varphi(p_j^{k_j}) = p_j^{k_j - 1}(p_j - 1)$ is even, and there is at least one factor in the product. Hence $\varphi(p_i^{k_i})$ divides $\varphi(q)/2$. Since $\gcd(-2, q) = 1$, Euler's theorem modulo $p_i^{k_i}$ gives $(-2)^{\varphi(q)/2} \equiv 1 \pmod{p_i^{k_i}}$. This holds for every i, so by the Chinese remainder theorem $(-2)^{\varphi(q)/2} \equiv 1 \pmod q$, and N_q is even. □

Proposition 11.5.4. *If $q = p^k$ with p an odd prime and $k \geq 1$, then N_q is odd if and only if $\left(\frac{-2}{p}\right) = -1$, that is, if and only if $p \equiv 5$ or $7 \pmod 8$.*

Proof. Put $u = (-2)^{\varphi(q)/2}$. By Euler's theorem $u^2 = (-2)^{\varphi(q)} \equiv 1 \pmod{p^k}$, so $p^k \mid (u-1)(u+1)$. As p is odd, p cannot divide both $u-1$ and $u+1$ (their difference is 2), so p^k divides one of them: $u \equiv \pm 1 \pmod{p^k}$.

To decide the sign, reduce modulo p. With $\varphi(q)/2 = p^{k-1}(p-1)/2$,

$$u \equiv \left((-2)^{(p-1)/2}\right)^{p^{k-1}} \equiv \left(\frac{-2}{p}\right)^{p^{k-1}} = \left(\frac{-2}{p}\right) \pmod p,$$

by Euler's criterion and because p^{k-1} is odd. Since $+1 \not\equiv -1 \pmod p$, the sign of u modulo p^k is exactly $\left(\frac{-2}{p}\right)$.

Finally evaluate the character by the two supplementary laws, $\left(\frac{-1}{p}\right) = (-1)^{(p-1)/2}$ and $\left(\frac{2}{p}\right) = (-1)^{(p^2-1)/8}$:

$p \bmod 8$	1	3	5	7
$\left(\frac{-1}{p}\right)$	$+1$	-1	$+1$	-1
$\left(\frac{2}{p}\right)$	$+1$	-1	-1	$+1$
$\left(\frac{-2}{p}\right)$	$+1$	$+1$	-1	-1

So $u \equiv -1$, and hence N_q is odd, exactly when $p \equiv 5$ or $7 \pmod 8$. □

Combining the three cases $q = 1$, q with at least two prime factors, and $q = p^k$: the count N_q is odd if and only if $q = p^k$ with $k \geq 1$ and p a prime congruent to 5 or 7 modulo 8. ■

Layer 1: WHAT

Block 1. Read $p \equiv 5, 7 \pmod 8$ as $\left(\frac{-2}{p}\right) = -1$. The target is a quadratic character, so the proof must end in Euler's criterion.

Block 2. Read $0 < a < q/4$ with $\gcd(a, q) = 1$ and a parity as Gauss's lemma with multiplier 2, using $a < q/4 \iff 2a < q/2$.

Block 3. Set up A, the half-system of coprime residues, note $|A| = \varphi(q)/2$, and let n count the doublings that wrap past the midpoint.

Block 4. Run the Gauss product: $a \mapsto |r(a)|$ permutes A, so multiplying $2a \equiv r(a)$ over A and cancelling gives $2^{\varphi(q)/2} \equiv (-1)^n$.

Block 5. Convert n into N_q using $N_q + n = |A|$, and absorb the resulting sign $(-1)^{\varphi(q)/2}$ into the base to get $(-1)^{N_q} \equiv (-2)^{\varphi(q)/2}$.

Block 6. Composite case: for each prime power factor the complementary totient factor is even, so $\varphi(p_i^{k_i}) \mid \varphi(q)/2$ and the power is 1 locally; the Chinese remainder theorem makes it 1 globally, so N_q is even.

Block 7. Prime power case: the power squares to 1, and modulo p^k the only square roots of 1 are ± 1; reduce modulo p and apply Euler's criterion; finish with the supplementary laws table.

Juxtaposition. The nominal route is to count directly: inclusion and exclusion over the primes dividing q, reduce the divisor sum modulo 2, evaluate $\lfloor m/4 \rfloor$ mod 2 by residue classes modulo 8, then pair up squarefree divisors. It works, it takes three cases of pairing, and it produces the classes $5, 7$ as an unexplained arithmetic coincidence. The bifurcation point is the line where you meet $q/4$: the nominal route counts inside the interval, the ideal route multiplies the interval by 2. Everything else follows from that single decision.

Layer 2: HOW

Block 1. **Cue.** A congruence condition modulo 8 imposed on a prime, in the conclusion of the problem. **Why natural.** Modulus 8 for primes has one standard source: $\left(\frac{2}{p}\right)$ depends on p modulo 8 and $\left(\frac{-1}{p}\right)$ on p modulo 4, so any character built from ± 1 and ± 2 is a function of p modulo 8. The classes $\{1,3\}$ against $\{5,7\}$ are the split of $\left(\frac{-2}{p}\right)$, which is a two second table lookup. Reading the answer first is legitimate and fast: the conclusion tells you which theorem the problem is about. **Link to the next block.** Knowing the destination is a quadratic character, look for the standard machine that turns a count into a character.

Block 2. **Cue.** Coprime residues, an interval that is a simple fraction of the modulus, and only the parity of the count. **Why natural.** That machine is Gauss's lemma. Its content is that doubling a half system of residues permutes the half system up to signs, and the sign of the permutation is $\left(\frac{2}{p}\right)$; the number of sign flips is a count of residues in an interval. Here the interval is $(0, q/4)$, and $a < q/4$ is the same statement as $2a < q/2$, so the problem is handing you the non-flipping count for the multiplier 2. The algebra is trivial; recognising which count it is, is the whole move. **Link to the next block.** You now need the lemma for general odd q, not just for primes, so check whether the proof used primality.

Block 3. **Cue.** The problem says only that q is odd, and Gauss's lemma is usually stated for a prime. **Why natural.** Read the standard proof and see what it actually uses: it needs 2 invertible, it needs the pairing $\{a, q-a\}$, and it needs to cancel a product of units. All three survive for any odd q, because the elements coprime to q form a group. This is a habit worth naming: when you want a theorem in more generality, do not look for a stronger theorem, re-read the proof of the one you have and see which hypotheses were load-bearing. **Link to the next block.** Write the lemma for general q and be careful about one thing only: the conclusion is a congruence, not an identity of symbols.

Block 4. **Cue.** A product over a set that a map permutes. **Why natural.** Cancelling a product of units is the standard engine of Euler's criterion, Wilson's theorem and Gauss's lemma alike. It works because a bijection of a finite set does not change a product, so any discrepancy between the two ways of writing the product is forced into the signs. **Link to the next block.** The lemma gives you n, but the problem asks about N_q, so convert.

Block 5. **Cue.** A is split into exactly two pieces by the condition $2a < q/2$. **Why natural.** $N_q + n = |A| = \varphi(q)/2$ is an identity, not an estimate, so the conversion is exact and costs one line. Absorbing $(-1)^{\varphi(q)/2}$ into the base to make $(-2)^{\varphi(q)/2}$ is not cosmetic: it turns two separate characters into one, and -2 is precisely the number whose residue class splits $\{1,3\}$ from $\{5,7\}$ modulo 8, which you predicted in Block 1. Seeing your prediction appear is the confirmation that the route is right. **Link to the next block.** Now the whole problem is the single question: when is $(-2)^{\varphi(q)/2} \equiv -1 \pmod{q}$?

Block 6. **Cue.** The conclusion singles out prime powers, so something must fail for composite q. **Why natural.** A power of a unit is -1 only if the unit has even order and the exponent is a suitable odd multiple of half that order. Modulo a product of coprime factors, the Chinese remainder theorem forces the answer to be -1 in every component simultaneously, which is hard; and with two or more factors the exponent $\varphi(q)/2$ is still a multiple of each local group order, because the complementary totient factor supplies the missing 2. That is the exact reason prime powers are special: only there is $\varphi(q)/2$ half of the local order rather than a whole multiple of it. **Link to the**

next block. The remaining case is $q = p^k$, where the group is cyclic and ± 1 are the only square roots of unity.

Block 7. Cue. $u^2 \equiv 1$ with $q = p^k$ odd. **Why natural.** You do not need cyclicity of $(\mathbb{Z}/p^k)^\times$ for this; $p^k \mid (u-1)(u+1)$ and $\gcd(u-1, u+1) \mid 2$ with p odd force p^k to divide one factor. Preferring the elementary argument is not pedantry, it is insurance: you can write it correctly under pressure. Then reducing modulo p identifies the sign, because p^{k-1} is odd so it does not disturb a ± 1. **Link to the next block.** Finish with Euler's criterion and the supplementary laws.

Layer 3: WHY

The mathematical truth underneath is that N_q was never a counting problem. Doubling is the map $a \mapsto 2a$ on the coprime residues; restricted to a half system it is a permutation up to signs, and the number of sign changes is what Gauss's lemma converts into a multiplicative character. The count of residues in $(0, q/4)$ is exactly the number of doublings that do not wrap around, so its parity is a signed determinant of a permutation in disguise. Once you see that, the appearance of 5 and 7 modulo 8 is not a coincidence to be verified but a consequence to be predicted.

The second truth is a warning worth installing permanently. The natural guess is that the parity of N_q is given by the Jacobi symbol $\left(\frac{-2}{q}\right)$. That is false. Gauss's lemma generalises to composite modulus as a congruence, $2^{\varphi(q)/2} \equiv (-1)^n \pmod q$, but not as an identity of symbols: for $q = 21$ the Jacobi symbol $\left(\frac{2}{21}\right) = \left(\frac{2}{3}\right)\left(\frac{2}{7}\right) = -1$, while the sign count gives $(-1)^n = +1$. The Jacobi symbol is multiplicative by construction and forgets whether a number is actually a square; the congruence remembers. Keeping the statement as a congruence is what makes the composite case collapse to $+1$, which is the entire content of the "only if" direction.

The third truth is the structural reason prime powers are the exceptional case. For q with at least two prime factors, $\varphi(q)/2$ is still divisible by every local group order $\varphi(p_i^{k_i})$, because each of those is even and there are at least two of them, so half of the product is still a multiple of each factor. In group language, the exponent of $(\mathbb{Z}/q)^\times$ divides $\varphi(q)/2$ as soon as the group is not cyclic. So the value $(-2)^{\varphi(q)/2}$ has no room to be anything but 1. The information you want lives in the two element group $\{\pm 1\}$ of square roots of unity, and only a cyclic group has exactly two of them.

The discernment mechanism to train: read the conclusion for the theorem, then read the hypothesis for the machine that produces that kind of conclusion. Modulus 8 on a prime means a quadratic character. A parity of a count of coprime residues in an interval means Gauss's lemma or its lattice point cousin (Eisenstein's proof of reciprocity counts lattice points under a line the same way). When both appear in one problem, the interval and the multiplier must match, and here they do after multiplying by 2. The examiners nudged you by choosing 4 as the denominator rather than 3 or 6, and by asking for parity rather than for the count.

Revolving concepts

Euler's totient and its arithmetic. φ is multiplicative, $\varphi(p^k) = p^{k-1}(p-1)$, and $\varphi(m)$ is even for every $m \geq 3$. The coprime residues split into pairs $\{a, q-a\}$, so exactly $\varphi(q)/2$ of them lie in the lower half. See Book 1, field: Number theory, topic: totient function.

Euler's theorem and Euler's criterion. For $\gcd(b, q) = 1$, $b^{\varphi(q)} \equiv 1 \pmod q$; for an odd prime p and $p \nmid b$, $b^{(p-1)/2} \equiv \left(\frac{b}{p}\right) \pmod p$. The second is the bridge from an exponentiation to a symbol, and it is the last step of the ideal solution. See Book 1, field: Number theory, topic: quadratic residues.

The supplementary laws. $\left(\frac{-1}{p}\right) = (-1)^{(p-1)/2}$ and $\left(\frac{2}{p}\right) = (-1)^{(p^2-1)/8}$, giving $\left(\frac{-2}{p}\right) = +1$ for $p \equiv 1, 3 \pmod 8$ and -1 for $p \equiv 5, 7 \pmod 8$. Memorise the combined table, not just the two separate laws, because -2 and ± 2 appear constantly. See Book 1, field: Number theory, topic: quadratic reciprocity.

Gauss's lemma, stated in the generality you need. Let q be odd, b coprime to q, and A a half system of coprime residues. If n is the number of $a \in A$ whose product ba reduces into the negative half, then $b^{\varphi(q)/2} \equiv (-1)^n \pmod q$. For q prime this is $\left(\frac{b}{p}\right) = (-1)^n$; for composite q it is a congruence only. See Book 1, field: Number theory, topic: Gauss's lemma.

Square roots of unity. In $\mathbb{Z}/p^k$ with p odd, $u^2 \equiv 1$ forces $u \equiv \pm 1$, because p divides at most one of $u-1$ and $u+1$. In $\mathbb{Z}/q$ with r distinct odd prime factors there are 2^r square roots of unity, which is exactly why the sign information is destroyed for composite q. See Book 1, field: Abstract algebra, topic: units modulo a prime power.

The counting recursion, generalised. The alternative route rests on a genuine recursion, Mobius inversion for counting coprime elements. For any $y > 0$,

$$\#\{a \le y : \gcd(a,q) = 1\} = \sum_{d \mid q} \mu(d) \left\lfloor \frac{y}{d} \right\rfloor,$$

which with $y = q/4$ and $\lfloor \lfloor q/4 \rfloor / d \rfloor = \lfloor q/(4d) \rfloor$ gives $N_q = \sum_{d \mid q} \mu(d) \lfloor (q/d)/4 \rfloor$. The general form is: any count with a coprimality side condition is the Mobius transform of the count without it, and the inversion is solved by summing over squarefree divisors only, since μ vanishes elsewhere. Reduced modulo 2 this becomes $N_q \equiv \sum_{d \mid q \text{ squarefree}} \lfloor (q/d)/4 \rfloor$, and $\lfloor m/4 \rfloor$ is odd exactly for $m \equiv 5, 7 \pmod 8$. The same inversion computes counts of coprime integers in any interval and is the standard first move for such problems. See Book 1, field: Number theory, topic: Mobius inversion.

Circuit Bank

- **WHEN** a condition on a prime is stated modulo 8 (or modulo 4) **DO** translate it into $\left(\frac{2}{p}\right)$, $\left(\frac{-1}{p}\right)$ or $\left(\frac{-2}{p}\right)$ before doing anything else **BECAUSE** the supplementary laws are essentially the only standard reason those moduli attach to primes, so the conclusion names the theorem you must reach
- **WHEN** you need the parity of the number of residues coprime to q inside an interval $(0, q/m)$ **DO** multiply the interval by m to land on a half system and run Gauss's lemma with multiplier m **BECAUSE** Gauss's lemma is exactly a wrap-around count, and the parity of a count is the sign of a signed permutation
- **WHEN** a standard theorem is stated for a prime but you need it for general q **DO** re-read its proof and check which steps used primality, rather than hunting for a stronger theorem **BECAUSE** most residue arguments use only that the coprime residues form a group, so they survive verbatim; but record whether the conclusion survives as a symbol identity or only as a congruence
- **WHEN** a statement singles out prime powers among all moduli **DO** look for the step where cyclicity or the count of square roots of unity is used **BECAUSE** only for $q = p^k$ does $(\mathbb{Z}/q)^\times$ have exactly two square roots of unity, so only there can a half power carry a sign
- **WHEN** you derive $(-1)^N \equiv u \pmod q$ with $u = \pm 1$ and $q \ge 3$ **DO** conclude the parity of N immediately **BECAUSE** $+1$ and -1 are distinct modulo any $q \ge 3$, so a congruence between signs is an equality of signs

Backtesting

1. First useful cue: the classes 5 and 7 modulo 8 in the conclusion, which are the nonresidue classes of -2; the confirming cue is $0 < a < q/4$ together with coprimality and parity, which is the wrap-around count in Gauss's lemma for multiplier 2.

2. Method the expression suggested: the congruence $(-1)^{N_q} \equiv (-2)^{\varphi(q)/2} \pmod{q}$, then a split into $q = 1$, q with two or more prime factors (Chinese remainder theorem), and $q = p^k$ (Euler's criterion).
3. Where you would hesitate, and the thought to reinforce: at the generalisation of Gauss's lemma to composite q. Reinforce that the proof only ever uses that the coprime residues form a group under multiplication and that 2 is invertible, and that the price of generality is that the conclusion is a congruence, not a Jacobi symbol identity.
4. The trap the question sets: writing $(-1)^{N_q} = \left(\frac{-2}{q}\right)$ as a Jacobi symbol. That is false for composite q ($q = 21$ is the smallest counterexample), and it would make the composite case appear to give odd values.
5. What makes you recognise it faster next time: the pairing of a quarter interval with the multiplier 2. Quarter and double, third and triple, sixth and sextuple: an interval of length q/m with a parity question is Gauss's lemma with multiplier m.
6. The reusable decision rule: to find the parity of a count of coprime residues in a sub-interval, scale the interval up to a half system, count the wrap-arounds, and read the sign as a power of the multiplier; then handle composite moduli by the Chinese remainder theorem and prime powers by Euler's criterion.

Mirror drills

1. Use Gauss's lemma with multiplier 2 for a prime p to prove the second supplementary law. *Answer:* the number of $a \leq (p-1)/2$ with $2a > p/2$ is $(p-1)/2 - \lfloor p/4 \rfloor$; evaluating this modulo 2 over the four classes of p modulo 8 gives $(-1)^{(p^2-1)/8}$.
2. Show $N_q = \sum_{d \mid q} \mu(d) \lfloor (q/d)/4 \rfloor$ and use it to compute N_{45}. *Answer:* the divisors with $\mu \neq 0$ are $1, 3, 5, 15$, giving $11 - 3 - 2 + 0 = 6$, which is even, consistent with 45 not being a prime power.
3. Compute N_{243} and N_{125} and check them against the classification. *Answer:* $243 = 3^5$ with $3 \equiv 3 \pmod 8$ gives $60 - 20 = 40$, even; $125 = 5^3$ with $5 \equiv 5 \pmod 8$ gives $31 - 6 = 25$, odd.

11.6 A6. A determinant identity for intertwined matrices

Let n be a positive integer. Suppose that A, B, and M are $n \times n$ matrices with real entries such that $AM = MB$, and such that A and B have the same characteristic polynomial. Prove that $\det(A - MX) = \det(B - XM)$ for every $n \times n$ matrix X with real entries.

Read it out loud

Three hypotheses, one conclusion, and every word is load-bearing. The relation $AM = MB$ says that M intertwines A and B: as a linear map, M carries the action of B on the source to the action of A on the target. If M were invertible this would say $A = MBM^{-1}$, that is, A and B are similar. The extra hypothesis, that A and B have the same characteristic polynomial, would then be free, since similar matrices have the same characteristic polynomial. So the hypothesis is telling you loudly that M is not assumed invertible, and that the characteristic polynomial hypothesis exists to repair exactly the damage that the kernel of M does.

Test the invertible case in ten seconds, because it is one line and it calibrates everything: if M is invertible then $A = MBM^{-1}$ and

$$A - MX = MBM^{-1} - MX = M(B - XM)M^{-1},$$

so the two matrices are similar and their determinants agree. That computation is the skeleton of the whole problem. It also tells you what the general statement should be: not merely that the determinants agree, but that $A - MX$ and $B - XM$ have the same characteristic polynomial, since in the invertible case they are similar.

Now read the shape of the conclusion. On the left the product is MX, on the right it is XM. The two sides differ by the order of a product and nothing else. There is exactly one elementary fact in linear algebra that trades the order of a product for nothing, and it is $\det(I - UV) = \det(I - VU)$, equivalently that UV and VU have the same characteristic polynomial. So the target is clear: manipulate $\det(A - MX)$ until it contains $\det(I - UV)$ for some U and V, swap, and manipulate back.

The fields wired in are determinants and characteristic polynomials, the rank normal form, and the polynomial identity principle (an identity of polynomials that holds off a finite set of parameter values holds everywhere). The last of these is the piece most likely to be missing from an A-level trained toolkit, and it is the piece that makes the proof short.

Approaches, ranked worst to best

1. **Expand both determinants.** *Cue:* a determinant identity to be proved. *Circuit:* Leibniz expansion or cofactor induction. Both sides are sums of $n!$ signed products in the entries of five matrices. There is no route from here to a proof, and no partial credit lives at the end of it.

2. **Perturb M to make it invertible.** *Cue:* the invertible case is one line, so approximate. *Circuit:* density of GL_n in M_n plus continuity of det. This is the first idea everyone has and it fails, for a reason worth learning permanently. The set of triples with $AM = MB$ is not open: replacing M by $M + \varepsilon I$ destroys the hypothesis unless $A = B$, and there may be no invertible M at all satisfying it (take A and B with disjoint spectra, where the only intertwiner is $M = 0$). Density arguments work on a variable that is free, never on one that is constrained. Ranks above brute force only because it identifies the correct skeleton.

3. **Compare traces of all powers.** *Cue:* characteristic polynomial equality is equivalent to equality of power sums in characteristic zero. *Circuit:* Newton's identities plus cyclic invariance of the trace. Show $\mathrm{Tr}((A - MX)^k) = \mathrm{Tr}((B - XM)^k)$ for all k by expanding, using $\mathrm{Tr}(PQ) = \mathrm{Tr}(QP)$ to rotate each word, and using $AM = MB$ to commute every M past every A. It is correct

and it is honest work: each word $A^{i_1}MXA^{i_2}\cdots$ turns into the matching word in B and XM. The costs are that it needs characteristic zero (Newton's identities divide by k) and that the bookkeeping of words is long under time pressure.

4. **Rank normal form of M.** *Cue:* M is the only degenerate object, so normalise it. *Circuit:* choose U, V invertible with $U^{-1}MV$ block diagonal $\operatorname{diag}(I_r, 0)$. Write everything in the two block decomposition. The relation $AM = MB$ forces $A_{11} = B_{11}$, $A_{21} = 0$ and $B_{12} = 0$, so both A' and B' are block triangular; the characteristic polynomial hypothesis then forces A_{22} and B_{22} to have the same characteristic polynomial, and both determinants evaluate to $\det(A_{11} - X_{11})\det(A_{22})$. This is the most transparent proof: it shows precisely where each hypothesis is spent, namely that the invertible part of M is handled by similarity and the kernel part is handled by the characteristic polynomial. It costs more ink than the winner and needs care with which matrix is conjugated by U and which by V.

5. **Sylvester's identity plus a shift by tI.** *Cue:* MX against XM, and a determinant that wants an inverse you are not given. *Circuit:* $\det(I - UV) = \det(I - VU)$, and equality of polynomials in t off a finite set. Replace A, B by $A - tI, B - tI$, which preserves both hypotheses, use $M = (A - tI)M(B - tI)^{-1}$ for the finitely many t where $B - tI$ is invertible, factor, swap with Sylvester, factor back, and finish by comparing two polynomials of degree n in t that agree at infinitely many points. Six lines, no case analysis, and it proves the stronger characteristic polynomial statement for free. This is the ideal route.

The ideal solution

Lemma 11.6.1. *For any $n \times n$ matrices U and V over a commutative ring,* $\det(I - UV) = \det(I - VU)$.

Proof. Work with the $2n \times 2n$ block matrix $N = \begin{pmatrix} I & U \\ V & I \end{pmatrix}$ and compute two unitriangular reductions:

$$\begin{pmatrix} I & 0 \\ -V & I \end{pmatrix} N = \begin{pmatrix} I & U \\ 0 & I - VU \end{pmatrix}, \qquad \begin{pmatrix} I & -U \\ 0 & I \end{pmatrix} N = \begin{pmatrix} I - UV & 0 \\ V & I \end{pmatrix}.$$

Both left factors have determinant 1, and the determinant of a block triangular matrix with square diagonal blocks is the product of the determinants of those blocks. Hence $\det N = \det(I - VU)$ and $\det N = \det(I - UV)$. □

Theorem 11.6.2. *Let A, B, M be real $n \times n$ matrices with $AM = MB$ and* $\det(A - tI) = \det(B - tI)$ *for all t. Then for every real $n \times n$ matrix X and every real t,*

$$\det\big((A - tI) - MX\big) = \det\big((B - tI) - XM\big).$$

In particular, taking $t = 0$, $\det(A - MX) = \det(B - XM)$, *and $A - MX$ and $B - XM$ have the same characteristic polynomial.*

Proof. Fix X and write $A_t = A - tI$ and $B_t = B - tI$. Define

$$p(t) = \det(A_t - MX), \qquad r(t) = \det(B_t - XM).$$

Each entry of $A_t - MX$ is a polynomial in t of degree at most 1, so p is a polynomial in t of degree at most n; likewise r. It therefore suffices to prove $p(t) = r(t)$ for all but finitely many real t, since then $p - r$ is a polynomial of degree at most n with infinitely many roots and so is identically zero.

The polynomial $\det(B - tI)$ is not identically zero (it has degree n with leading coefficient $(-1)^n$), so B_t is invertible for all but at most n values of t. Fix such a t.

From $AM = MB$ we get $A_tM = AM - tM = MB - tM = MB_t$, and therefore

$$M = A_tMB_t^{-1}.$$

Substituting this occurrence of M and factoring A_t out on the left,

$$\det(A_t - MX) = \det\left(A_t - A_t M B_t^{-1} X\right) = \det(A_t)\det\left(I - (MB_t^{-1})X\right).$$

Apply the lemma with $U = MB_t^{-1}$ and $V = X$, then factor B_t^{-1} out on the right:

$$\det\left(I - (MB_t^{-1})X\right) = \det\left(I - X(MB_t^{-1})\right) = \det\left((B_t - XM)B_t^{-1}\right) = \frac{\det(B_t - XM)}{\det(B_t)}.$$

The hypothesis that A and B have the same characteristic polynomial says exactly that $\det(A_t) = \det(B_t)$ for every t. Combining the two displays,

$$p(t) = \det(A_t - MX) = \frac{\det(A_t)}{\det(B_t)}\det(B_t - XM) = \det(B_t - XM) = r(t).$$

This holds for all but at most n values of t, hence for all t, and $t = 0$ gives the required identity. Since $p(t) = (-1)^n$ times the characteristic polynomial of $A - MX$ evaluated at t, and similarly for r, the two matrices have the same characteristic polynomial. □

Remark 11.6.3. Both hypotheses are sharp. Equality of determinants alone does not suffice: take

$$A = \begin{pmatrix} 0 & 0 \\ 0 & 1 \end{pmatrix}, \quad B = \begin{pmatrix} 0 & 0 \\ 0 & 2 \end{pmatrix}, \quad M = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad X = \begin{pmatrix} 3 & 0 \\ 0 & 1 \end{pmatrix}.$$

Then $AM = MB = 0$ and $\det A = \det B = 0$, but $\det(A - MX) = -3$ while $\det(B - XM) = -6$. The characteristic polynomials $t^2 - t$ and $t^2 - 2t$ differ, which is precisely what the proof needed. Note also that A and B need not be similar: with $A = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, $B = 0$ and $M = \left(\begin{smallmatrix} 1 & 0 \\ 0 & 0 \end{smallmatrix}\right)$ the hypotheses hold and A is not similar to B.

Layer 1: WHAT

Block 1. Test M invertible: $A - MX = M(B - XM)M^{-1}$, so the statement is a similarity statement in disguise, and the degenerate M is the entire difficulty.

Block 2. Name the shape: the two sides differ only in the order of a product, so the target identity is $\det(I - UV) = \det(I - VU)$.

Block 3. To factor A out of $A - MX$ you need M written with a factor of A on the left. The relation $AM = MB$ gives exactly that, $M = AMB^{-1}$, provided B is invertible.

Block 4. B need not be invertible, so shift: replace A, B by $A - tI, B - tI$. This is the unique perturbation that preserves both hypotheses, since $A_t M = MB_t$ follows from $AM = MB$ and $\det A_t = \det B_t$ is the characteristic polynomial hypothesis itself.

Block 5. Run the chain: factor A_t out, swap with Sylvester, factor B_t^{-1} out, cancel $\det A_t$ against $\det B_t$.

Block 6. Finish by degree counting: both sides are polynomials of degree at most n in t agreeing off a finite set, so they agree everywhere; set $t = 0$.

Juxtaposition. The nominal route is to try to prove that $A - MX$ and $B - XM$ are similar, which is what the invertible case suggests, and to reach for a density argument in M when that fails. It fails twice: the matrices need not be similar, and M cannot be perturbed without destroying $AM = MB$. The bifurcation point is the sentence "B_t must be invertible for this step". The zombie response is to abandon the step. The trained response is to notice that invertibility is a generic property, to add the parameter t that makes it generic, and to check that the added parameter respects the hypotheses. Choosing the perturbation that preserves the hypotheses, rather than the one that makes the matrix you want invertible, is the whole discipline of this problem.

Layer 2: HOW

Block 1. **Cue.** M appears on the left in one determinant and on the right in the other, with no invertibility assumed anywhere. **Why natural.** Testing the invertible case is not optimism, it is measurement: it tells you what the statement means, it gives you the algebraic identity $A - MX = M(B - XM)M^{-1}$ to imitate, and it tells you which hypothesis is doing which job. Since the characteristic polynomial hypothesis is redundant when M is invertible, it must be there for the singular case only. **Link to the next block.** Knowing the target is a similarity-like identity, look for the standard fact that swaps the order of a product.

Block 2. **Cue.** MX against XM. **Why natural.** $\det(I - UV) = \det(I - VU)$ is the unique elementary statement that exchanges the order of a matrix product without cost, and it holds even when neither factor is invertible. Recognising it is a pattern match on the shape of the expression, not an insight. Its cousins are worth holding in the same slot: UV and VU have the same characteristic polynomial, and for rectangular U, V the nonzero eigenvalues match. **Link to the next block.** To reach the form $I - UV$ you must factor something out of $A - MX$, so ask what can be factored.

Block 3. **Cue.** The only relation available is $AM = MB$, and it has M on both sides. **Why natural.** If you want A to come out as a left factor of MX, you must rewrite M so that it starts with A; the relation rearranges to $M = AMB^{-1}$ and that is the only rearrangement it offers. Notice how the hypothesis is used: not as a symmetry, but as a substitution rule that inserts A where you need it. That is the standard way an intertwining relation is spent. **Link to the next block.** The rearrangement requires B^{-1}, which you have not been given, so the next move must manufacture invertibility.

Block 4. **Cue.** A required inverse that the hypotheses do not supply. **Why natural.** Invertibility is generic: $\det(B - tI)$ is a nonzero polynomial in t, so all but finitely many shifts of B are invertible. The reason to shift by tI specifically, rather than by anything else, is that this is the one perturbation under which both hypotheses survive: $A_t M = MB_t$ is immediate, and $\det A_t = \det B_t$ is a restatement of the characteristic polynomial hypothesis. In other words, the hypothesis you were given is precisely a licence to shift. When a hypothesis is stated for all t in disguise, the intended proof uses all t. **Link to the next block.** With B_t invertible, run the algebra.

Block 5. **Cue.** A product of three determinants with $\det A_t$ and $\det B_t$ at the two ends. **Why natural.** The cancellation $\det A_t / \det B_t = 1$ is where the characteristic polynomial hypothesis is spent, and it is spent exactly once. That is a good sign in a proof: each hypothesis should have a visible point of consumption. If a hypothesis is never consumed, either the proof is wrong or the statement is stronger than claimed. **Link to the next block.** You now have the identity for all but finitely many t, so upgrade.

Block 6. **Cue.** An identity proved off a finite exceptional set, with both sides polynomial. **Why natural.** Two polynomials of degree at most n that agree at more than n points are equal. Degree counting is the cheapest form of the polynomial identity principle and it avoids any appeal to topology or to Zariski density, so it costs one sentence and cannot be challenged by a grader. State the degree bound explicitly, because a grader will look for it. **Link to the next block.** Set $t = 0$ and observe that you have in fact proved the stronger characteristic polynomial statement.

Layer 3: WHY

The mathematics underneath is the theory of UV against VU. Sylvester's identity $\det(I - UV) = \det(I - VU)$ says that the two products have the same nonzero spectrum, and the block proof

shows why: the matrix $\left(\begin{smallmatrix} I & U \\ V & I \end{smallmatrix}\right)$ can be cleared in two directions, and the leftover block is $I - UV$ or $I - VU$ depending on which direction you clear. Every identity in which the order of a product flips is ultimately this statement, and it is why the problem is solvable at all: the conclusion is not an accident of determinants, it is a shadow of the fact that MX and XM are spectrally almost the same.

The second truth is the polynomial identity principle. A polynomial identity in several variables that holds on a set which is large in each variable holds identically. The cheap version, used here, is that a polynomial in one variable of degree at most n with more than n roots is zero. The general version, that a polynomial vanishing on a nonempty open subset of $\mathbb{R}^N$ is zero, is what lets you assume genericity in any variable that is free. The discipline is to identify which variables are free. Here t and X are free and M is constrained, which is the exact reason the perturbation must be applied to the shift and not to M.

The third truth is structural and is best seen through the rank normal form route. Write $M = U \left(\begin{smallmatrix} I_r & 0 \\ 0 & 0 \end{smallmatrix}\right) V^{-1}$. In these coordinates the relation $AM = MB$ forces both A and B into block triangular form with a shared top left block, and the two determinants both collapse to $\det(A_{11} - X_{11}) \det(A_{22})$ against $\det(B_{11} - X_{11}) \det(B_{22})$. The first factors agree because M acts invertibly there, so similarity does the work; the second factors agree only because A_{22} and B_{22} inherit the same characteristic polynomial, and X has disappeared from them entirely. So the honest content of the theorem is: on the part of the space where M is invertible, the identity is similarity; on the part killed by M, the variable X cannot reach and only the spectral data matter. Knowing that, you understand why the hypothesis could not be weakened to $\det A = \det B$, and the two by two example in the remark is the smallest witness.

The discernment mechanism to train: when a proof stalls on a missing invertibility, do not abandon the step and do not perturb blindly. Ask which perturbations preserve every hypothesis. Usually there is exactly one natural family, and the hypotheses themselves tell you which, because a hypothesis stated as an equality of characteristic polynomials is an equality for all t, and an equality for all t is an invitation to use t. The examiners nudged you toward this by stating the hypothesis as "the same characteristic polynomial" rather than the more natural sounding "the same eigenvalues with multiplicity"; the polynomial phrasing is the hint.

Revolving concepts

Characteristic polynomial. $\chi_C(t) = \det(tI - C)$, monic of degree n; its coefficients are the elementary symmetric functions of the eigenvalues, with $\operatorname{Tr}(C)$ and $\det(C)$ as the two extremes. Similar matrices have the same characteristic polynomial; the converse is false (a nonzero nilpotent Jordan block against the zero matrix). Cayley and Hamilton gives $\chi_C(C) = 0$. See Book 1, field: Linear algebra, topic: characteristic polynomial and similarity.

Sylvester's determinant identity. $\det(I_m - UV) = \det(I_n - VU)$ for U of size $m \times n$ and V of size $n \times m$; for square U, V this says $\chi_{UV} = \chi_{VU}$. Proof by clearing the block matrix $\left(\begin{smallmatrix} I & U \\ V & I \end{smallmatrix}\right)$ in two directions. Consequences worth storing: $\det(I + uv^T) = 1 + v^T u$ for column vectors, and the matrix determinant lemma $\det(C + uv^T) = \det(C)(1 + v^T C^{-1} u)$. See Book 1, field: Matrices and determinants, topic: Sylvester identity.

Intertwiners. $AM = MB$ says M is a homomorphism of $\mathbb{R}[x]$ modules from $(\mathbb{R}^n, B)$ to $(\mathbb{R}^n, A)$, so M maps generalised eigenspaces of B into generalised eigenspaces of A for the same eigenvalue; in particular if A and B have disjoint spectra then $M = 0$. This is the reason a perturbation of M is illegal: the solution set of $AM = MB$ in M is a linear subspace determined by the spectra, and it may miss GL_n entirely. See Book 1, field: Linear algebra, topic: Sylvester equation and intertwiners.

Rank normal form and block determinants. Every M of rank r factors as $U \left(\begin{smallmatrix} I_r & 0 \\ 0 & 0 \end{smallmatrix}\right) V^{-1}$ with U, V invertible; and $\det \left(\begin{smallmatrix} P & Q \\ 0 & S \end{smallmatrix}\right) = \det P \det S$ for square diagonal blocks. These two facts together are the whole of the alternative proof. See Book 1, field: Matrices and determinants, topic: block determinants.

The polynomial identity principle. A polynomial of degree at most d in one variable with more than d roots is identically zero; consequently an identity between polynomial expressions that holds for all but finitely many values of a parameter holds for all values, and one that holds on a nonempty open subset of $\mathbb{R}^N$ holds identically. This is how "assume the generic case" is made rigorous. See Book 1, field: Polynomials, topic: identity principle and genericity.

The recursion behind the trace route, generalised. The third published route needs that the traces of all powers determine the characteristic polynomial, which is Newton's identities: with $p_k = \operatorname{Tr}(C^k)$ and e_i the elementary symmetric functions of the eigenvalues,

$$p_k - e_1 p_{k-1} + e_2 p_{k-2} - \cdots + (-1)^{k-1} e_{k-1} p_1 + (-1)^k k\, e_k = 0,$$

a recursion which determines $e_1, \ldots, e_n$ from $p_1, \ldots, p_n$ one at a time. Its closed form is the generating function identity

$$\det(I - tC) = \exp\left(- \sum_{k \geq 1} \frac{\operatorname{Tr}(C^k)}{k} t^k \right),$$

obtained by taking logarithms of $\prod_i (1 - \lambda_i t)$. The generalisation to record: any symmetric function of eigenvalues can be computed from power sums by this recursion, and the division by k is exactly why the trace route needs characteristic zero (or characteristic exceeding n). The same recursion is the standard tool whenever a problem gives you $\operatorname{Tr}(C^k)$ for all k and asks for eigenvalues. See Book 1, field: Polynomials, topic: Newton identities and power sums.

Circuit Bank

- **WHEN** one side of an identity has UV and the other has VU **DO** reach for $\det(I - UV) = \det(I - VU)$ and reshape both sides into that form **BECAUSE** the only structural difference is the order of a product, and that identity is the elementary fact which trades order for nothing, valid without any invertibility
- **WHEN** a step needs an inverse the hypotheses do not provide **DO** shift by tI, prove the identity for the all but finitely many t where the inverse exists, then compare polynomials of degree at most n in t **BECAUSE** invertibility is generic, and equality of two degree n polynomials at more than n points is equality
- **WHEN** you must choose which object to perturb **DO** perturb only a free variable, never one constrained by a hypothesis, and prefer the perturbation under which every hypothesis is stable **BECAUSE** the solution set of a constraint like $AM = MB$ need not meet the good set at all, so a perturbation that breaks a hypothesis proves nothing
- **WHEN** a hypothesis looks stronger than the conclusion seems to need **DO** find the single line where it is consumed, and build a small counterexample to the weaker version **BECAUSE** every hypothesis in a well posed problem has exactly one point of consumption; if you cannot find it, you have either misused it or proved something stronger
- **WHEN** a matrix in the problem is the only object not assumed invertible **DO** try its rank normal form $\operatorname{diag}(I_r, 0)$ and read what the other hypotheses force on the blocks **BECAUSE** the rank normal form separates the part where the map is an isomorphism, handled by similarity, from the kernel, where only spectral data survive

Backtesting

1. First useful cue: MX on one side against XM on the other, which names Sylvester's identity; the confirming cue is that M is never assumed invertible while the one line invertible case is trivial, which locates the difficulty and explains why the characteristic polynomial hypothesis exists.

2. Method the expression suggested: substitute $M = A_t M B_t^{-1}$ from the shifted intertwining relation, factor, swap orders by Sylvester, factor back, cancel $\det A_t$ against $\det B_t$, and conclude by degree counting in t.
3. Where you would hesitate, and the thought to reinforce: at the demand for B_t^{-1}. Reinforce that invertibility is generic and that the correct perturbation is the one preserving both hypotheses, which the phrase "the same characteristic polynomial" is quietly recommending, since it is an equality for every t.
4. The trap the question sets: perturbing M towards invertibility. It is the first idea, it feels like the standard density argument, and it is illegal because $AM = MB$ is a constraint on M; worse, there are instances (disjoint spectra) where $M = 0$ is the only intertwiner.
5. What makes you recognise it faster next time: the pairing of an intertwining relation with a determinant whose two sides differ by the order of a product. Intertwining relations are substitution rules; order flips are Sylvester.
6. The reusable decision rule: when a determinant identity needs an inverse you do not have, introduce a spectral parameter t that makes the inverse generic, check the hypotheses survive the shift, prove the identity generically, and finish by comparing polynomials in t of bounded degree.

Mirror drills

1. Prove $\det(I_m - UV) = \det(I_n - VU)$ for rectangular U ($m \times n$) and V ($n \times m$), and deduce that UV and VU have the same nonzero eigenvalues with multiplicity. *Answer:* clear $\left(\begin{smallmatrix} I_m & U \\ V & I_n \end{smallmatrix}\right)$ in both directions using unitriangular factors of determinant 1; the eigenvalue statement follows by applying the identity with U replaced by U/t.
2. With $AM = MB$ and M invertible, show $A - MX$ and $B - XM$ are similar for every X. *Answer:* $A = MBM^{-1}$, so $A - MX = M(B - XM)M^{-1}$.
3. Produce A, B, M, X with $AM = MB$ and $\det A = \det B$ but $\det(A - MX) \neq \det(B - XM)$, and say which step of the proof breaks. *Answer:* $A = \operatorname{diag}(0,1)$, $B = \operatorname{diag}(0,2)$, $M = \operatorname{diag}(1,0)$, $X = \operatorname{diag}(3,1)$ give -3 against -6; the cancellation $\det A_t = \det B_t$ fails for $t \neq 0$.

11.7 B1. a perfect cube hides an integrating factor

Let f be a three times differentiable function (defined on $\mathbb{R}$ and real valued) such that f has at least five distinct real zeros. Prove that $f + 6f' + 12f'' + 8f'''$ has at least two distinct real zeros.

Read it out loud

This is a zero counting question dressed as a differential expression. You are given a hypothesis about zeros of f, nothing about a formula for f, and asked for a conclusion about zeros of a fixed combination of f and its first three derivatives. The only elementary tool that turns "many zeros of h" into "one fewer zero of h'" is Rolle's theorem, so the moment "distinct real zeros" is attached to a statement about derivatives, the reaction circuit fires: count with Rolle, do not compute with formulas, and do not look for an explicit f.

The second thing to read out loud is the coefficient string $1, 6, 12, 8$ attached to f, f', f'', f''' in ascending order of derivative. Four numbers on a degree three combination, in the ratio $1 : 6 : 12 : 8$, which is Pascal's row for degree three scaled by powers of two ($1, 3, 3, 1$ against $1, 2, 4, 8$ gives $1, 6, 12, 8$). That is the guiding light: the combination is not an arbitrary cubic in the derivative, it is the binomial expansion of a single linear term cubed, and a perfect power sitting inside a linear combination of derivatives is exactly the signal to reach for an integrating factor, because an integrating factor is the device that turns $h' + ch$ into a bare derivative of a product.

Within thirty seconds you should have $8t^3 + 12t^2 + 6t + 1 = (2t + 1)^3$ written on your paper, because that identity is the whole problem.

Approaches, ranked worst to best

1. **direct estimation of the combination.** *Cue:* a five zero hypothesis and a two zero conclusion, no formula in sight. *Circuit:* none, this is zombie mathematics. Trying to bound $f + 6f' + 12f'' + 8f'''$ with generic inequalities relating a function to its derivatives, with no use of the specific coefficients. This is the wrong kind of tool for the claim: existence of a zero is a qualitative statement, and it stalls immediately because there is no bound to write down without more information about f.

2. **apply Rolle separately to f, f', f'', f'''.** *Cue:* four objects, four zero counts $5, 4, 3, 2$. *Circuit:* Rolle's theorem, used without a target function. Rolle alone gives f' at least four zeros, f'' at least three, f''' at least two. The zero counts match the problem's numbers, which is exactly why this route is tempting, but it is a dead end: the zeros of f, f', f'', f''' are four different sets of points, and knowing each function vanishes somewhere says nothing about where the specific combination $f + 6f' + 12f'' + 8f'''$ vanishes. This ranks above the first approach because it finds the right theorem, but it never bundles the four pieces into one function, which is the actual content of the problem.

3. **integrating factor $e^{x/2}$ collapsing the cubic operator.** *Cue:* $1, 6, 12, 8$ factor as $(2t + 1)^3$. *Circuit:* product rule plus Rolle, chained three times on one function. Since $8t^3 + 12t^2 + 6t + 1 = (2t + 1)^3$ is a perfect cube, differentiating $e^{x/2}f$ three times reproduces the target combination exactly, up to a nonvanishing factor, in one substitution with no bookkeeping of several exponentials. This is the ideal route: minimal ink, every step forced once the factorisation is seen.

The ideal solution

Since $8t^3 + 12t^2 + 6t + 1 = (2t + 1)^3$, the target expression is 8 times $(2D + 1)^3$ applied to f, where $D = d/dx$; equivalently it is 64 times $(D + \frac{1}{2})^3 f$.

Let $g(x) = e^{x/2}f(x)$. Since $e^{x/2} > 0$ for every real x, g vanishes exactly where f vanishes, so g has at least five distinct real zeros.

Because f is three times differentiable, so is g, being a product of f with a smooth exponential; hence g, g', g'' are each differentiable and therefore continuous on all of $\mathbb{R}$. This is exactly what Rolle's theorem needs, so it applies at every stage below without any extra hypothesis being smuggled in.

Differentiate once:

$$g'(x) = e^{x/2}\left(\tfrac{1}{2}f(x) + f'(x)\right).$$

Let $x_1 < x_2 < \cdots < x_5$ be five of the zeros of g. Rolle's theorem supplies a zero of g' in each open interval (x_i, x_{i+1}), $i = 1, \ldots, 4$. These four intervals are pairwise disjoint, so the four zeros of g' they produce are pairwise distinct. Hence g' has at least four distinct real zeros.

Differentiate again:

$$g''(x) = e^{x/2}\left(\tfrac{1}{4}f(x) + f'(x) + f''(x)\right).$$

By the same argument, applied to the at least four zeros of g', g'' has at least three distinct real zeros.

Differentiate a third time:

$$g'''(x) = e^{x/2}\left(\tfrac{1}{8}f(x) + \tfrac{3}{4}f'(x) + \tfrac{3}{2}f''(x) + f'''(x)\right) = \tfrac{1}{8}e^{x/2}\big(f(x) + 6f'(x) + 12f''(x) + 8f'''(x)\big).$$

By Rolle applied once more to the at least three zeros of g'', g''' has at least two distinct real zeros.

Since $e^{x/2}$ never vanishes, $g'''(x) = 0$ exactly when $f(x) + 6f'(x) + 12f''(x) + 8f'''(x) = 0$. Therefore $f + 6f' + 12f'' + 8f'''$ has at least two distinct real zeros, as required.

Layer 1: WHAT

Block 1: read the coefficients $1, 6, 12, 8$ as $(2t+1)^3$, converting an opaque fourth order combination into a labelled operator $(2D+1)^3$. Block 2: set $g = e^{x/2}f$, so that multiplying by a never vanishing exponential preserves the zero set of f exactly. Block 3: differentiate g once and see that this reproduces $(D + \frac{1}{2})$ acting on f inside the exponential wrapper. Block 4: apply Rolle three times in a row, from g to g' to g'' to g''', since differentiating three times is exactly cubing the operator. Block 5: read off that g''' is, up to the nonvanishing factor $\frac{1}{8}e^{x/2}$, precisely the target combination, so its zero count transfers unchanged.

Juxtaposition: the nominal route (approach two above) also reaches "two zeros" through three applications of Rolle, but on the wrong function; it differentiates f itself and hopes the four resulting zero sets recombine at the end. The bifurcation point is the choice of what to differentiate. The zombie route starts differentiating before deciding what object it is building. The ideal route multiplies by $e^{x/2}$ first and only then starts differentiating, so the recombination is automatic at every step instead of being wished for at the last one.

Layer 2: HOW

Block 1. **Cue.** The coefficients $1, 6, 12, 8$ attached to f, f', f'', f'''. **Why natural.** Four coefficients on a degree three combination is Pascal's row for $n = 3$ scaled geometrically; checking $8t^3 + 12t^2 + 6t + 1 = (2t+1)^3$ costs one line and shows the operator is a perfect cube of a first order operator, not an arbitrary cubic. **Link to the next block.** A perfect cube of $(D + \frac{1}{2})$ is exactly what three product rule differentiations naturally produce, so the next move is to build the function that differentiates into that operator.

Block 2. **Cue.** The identity $\frac{d}{dx}\big(e^{cx}h(x)\big) = e^{cx}\big(h'(x) + ch(x)\big)$. **Why natural.** This is the only way, up to a scalar, to turn $h' + ch$ into a bare derivative, because $u' = cu$ forces u to be exponential; choosing $c = \frac{1}{2}$ matches the root found in Block 1 exactly. **Link to the next block.** Setting $g = e^{x/2}f$ makes the next differentiation reproduce $(D + \frac{1}{2})f$ automatically, and since $e^{x/2} > 0$ everywhere, g starts with the same five zeros as f for free.

Block 3. **Cue.** g has at least five distinct zeros, and Rolle's theorem is available because g is differentiable everywhere. **Why natural.** Between any two consecutive zeros of a differentiable function the derivative must vanish at least once; four disjoint intervals from five zeros force four distinct zeros of g'. **Link to the next block.** The identical argument applies to g' (three more zeros for g'') and then to g'' (two more zeros for g'''), because nothing in the argument used any property of g beyond differentiability and having several zeros.

Block 4. **Cue.** g''' has been computed explicitly and factors as $\frac{1}{8}e^{x/2}$ times the target expression. **Why natural.** $e^{x/2}$ never vanishes, so multiplying or dividing by it rescales a function without moving or removing any zero. **Link to the next block.** The two zeros of g''' transfer without loss to $f + 6f' + 12f'' + 8f'''$, which is exactly the claim to be proved, and the chain closes.

Layer 3: WHY

The truth underneath this problem is that a linear differential operator with constant coefficients and only real roots can always be realised, up to multiplication by a never vanishing exponential, as a pure iterated derivative, and Rolle's theorem is blind to such multiplication because it only cares about the location of zeros, never their size. So any statement of the shape "f has m zeros, show $L(D)f$ has at least $m - n$ zeros" for a real rooted degree n polynomial L reduces to the same three ingredients: factor L, peel off one root at a time with an exponential, apply Rolle once per root.

The discernment mechanism is algebraic and happens before any calculus: look at the coefficients and ask whether they belong to a polynomial with only real roots, here visibly a perfect cube. If they do not, this route is unavailable, because Rolle's theorem is a statement about real valued functions of a real variable with no direct analogue for a genuinely complex exponential factor; that signals a different problem entirely, not a harder version of this one.

The examiners nudged you toward the route by choosing $1, 6, 12, 8$ rather than an arbitrary quadruple. A solver who expands $(2t + 1)^3$ on seeing four coefficients on a third order combination finds the route in under a minute; a solver who treats the four terms as unrelated spends the full time on the second approach above and never closes the argument. Training this discernment means building the habit of factoring every coefficient list met in a differential context before touching a derivative, the same reflex as factoring a quadratic before completing the square.

Revolving concepts

Rolle's theorem: if h is continuous on $[a, b]$, differentiable on (a, b), and $h(a) = h(b)$, some $c \in (a, b)$ has $h'(c) = 0$. The corollary used three times above, that at least k distinct zeros of h force at least $k - 1$ distinct zeros of h' strictly between consecutive pairs, is the workhorse of every zero counting problem on the Putnam.

Integrating factors: for constant c, e^{cx} is the unique function u (up to a nonzero scalar) with $u' = cu$, and $\frac{d}{dx}(uh) = u(D + c)h$ for such u. This is the same device used to solve first order linear differential equations $y' + cy = Q(x)$; here it certifies a zero count instead of solving an equation.

Generalised recursion (the repeated structure in this problem): let $L(t) = \sum_{k=0}^{n} c_k t^k$ have real roots $r_1, \ldots, r_n$ counted with multiplicity. Multiplying by $e^{-r_1 x}$ and differentiating once realises $(D - r_1)f$ up to that exponential; rescaling by $e^{(r_1 - r_2)x}$ and differentiating again realises $(D - r_2)(D - r_1)f$; continuing for all n roots realises $L(D)f$, up to a leading coefficient and a nonvanishing exponential, after exactly n differentiations. Each differentiation costs exactly one zero via Rolle, so if f has $m > n$ distinct real zeros, $L(D)f$ has at least $m - n$ distinct real zeros. When the roots coincide, as here with $r_1 = r_2 = r_3 = -\frac{1}{2}$, every rescaling factor between steps is trivially 1, and the whole chain collapses to differentiating a single $e^{-r_1 x} f$ repeatedly, which is why this problem admits the one substitution shortcut; distinct real roots still work, but need a genuinely different exponential bridging each pair of steps.

Binomial coefficients as a diagnostic: recognising $1, 6, 12, 8$ as the expansion of $(2t+1)^3$ is the same pattern matching used whenever a coefficient list should be tested against $(at+b)^n$ before any calculus is attempted on it.

Circuit Bank

- **WHEN** a fixed linear combination of f and several derivatives sits between a zero count hypothesis and a zero count conclusion **DO** write the coefficients as a polynomial $L(t)$ in the operator $D = d/dx$ and check whether L has only real roots **BECAUSE** Rolle's theorem is the only elementary tool linking zero counts across a derivative, and it only reaches the target expression once that expression is exhibited as a pure derivative
- **WHEN** $L(D)$ factors with a single repeated real root r **DO** multiply f by one exponential e^{-rx} and differentiate that one product the full number of times, instead of juggling several exponentials **BECAUSE** a repeated root makes every intermediate rescaling factor between steps trivial, so one substitution absorbs the whole chain
- **WHEN** $L(D)$ factors with distinct real roots $r_1, \ldots, r_n$ **DO** peel off one root at a time: multiply by $e^{-r_1 x}$, differentiate, rescale by $e^{(r_1-r_2)x}$, differentiate again, and so on **BECAUSE** each factor of L needs its own matching exponential, and only the most recent derivative, not the original function, may be rescaled between steps
- **WHEN** a hypothesis gives m distinct zeros and a conclusion asks for exactly $m-n$ zeros where n is the order of the combination **DO** expect the proof to consist of exactly n applications of Rolle's theorem, no more and no fewer **BECAUSE** each application of Rolle costs exactly one zero, so the arithmetic $m \to m-n$ is the proof's structure, not a coincidence to verify separately

Backtesting

1. First useful cue: the coefficient string $1, 6, 12, 8$ on a third order combination.
2. Method the expression suggested: factor as a perfect cube $(2t+1)^3$, then peel it off with a single exponential integrating factor.
3. Where you would hesitate, and the thought to reinforce: whether multiplying by $e^{x/2}$ changes the zero count of f; reinforce that any strictly positive factor preserves a zero set exactly, so this step needs only a one line justification.
4. The trap the question sets: applying Rolle three times to f itself, getting the numerically satisfying chain $5 \to 4 \to 3 \to 2$ on four different functions, and mistaking that coincidence of arithmetic for a proof.
5. What makes you recognise it faster next time: any problem of the shape "f has m zeros, prove a linear combination of f and its derivatives has $m-n$ zeros" should trigger a check of the coefficients against $(at+b)^n$ before a single derivative is written.
6. The reusable decision rule: factor the coefficient polynomial in D first, choose the integrating factor matching its root or roots second, differentiate and apply Rolle third.

Mirror drills

1. If f is differentiable with at least two distinct real zeros, show $f + f'$ has at least one real zero. *Answer:* let $g = e^x f$; then $g' = e^x(f + f')$, and g has two zeros so g' has one by Rolle.

2. If f is twice differentiable with at least four distinct real zeros, show $f + 4f' + 4f''$ has at least two real zeros. *Answer:* $4t^2 + 4t + 1 = (2t+1)^2$, so with $g = e^{x/2}f$, $g'' = \frac{1}{4}e^{x/2}(f + 4f' + 4f'')$, and g has four zeros so g'' has two.
3. If f is twice differentiable with at least three distinct real zeros, show $f'' - 3f' + 2f$ has at least one real zero. *Answer:* $t^2 - 3t + 2 = (t-1)(t-2)$ has distinct real roots; chain two exponentials (multiply by e^{-x}, differentiate, rescale by e^{-x} again, differentiate) applying Rolle after each derivative to go from three zeros of f down to one.

11.8 B2. crossing off triples and their sums

Given the list of positive integers $1, 2, 3, 4, \ldots$, *take the first three numbers* $1, 2, 3$ *and their sum* 6, *and cross all four numbers off the list. Repeat with the three smallest remaining numbers* $4, 5, 7$ *and their sum* 16, *crossing all four off the list. Continue in this way, at each stage crossing off the three smallest remaining numbers together with their sum, and consider the sequence of sums produced:* $6, 16, 27, 36, \ldots$. *Prove or disprove that some number in this sequence has base ten representation ending in* 2015.

Read it out loud

Strip away the story and this is an operationally defined sequence: you are not handed a formula for s_n, you are handed a rule (take three, append the sum, delete four) and asked a question about the residue of s_n modulo 10^4 for some n. That mismatch, a rule with no formula set against a question purely about digits, is itself the command term. Whenever a Putnam problem defines a sequence procedurally and then asks about its behaviour modulo a power of ten, it is never inviting you to run the procedure until you get lucky. It is inviting you to prove the procedure is eventually periodic, or periodic enough to steer, and then to solve a small congruence.

Read "prove or disprove" for what it is here: exam theatre. An ending-digit claim about an infinite sequence is almost always provable by exhibiting one term, and disproving it would need a genuine obstruction, some invariant forbidding the digits 2015 forever, which is a much harder thing to manufacture for a sieve this loosely constrained. Default to "prove", and go looking for a construction, not a proof of impossibility.

The first useful move costs thirty seconds and no insight at all: extend the given data by hand. Continuing past 36: $11, 12, 13 \to 36$ (given), then $14, 15, 17 \to 46$ (note 16 is already gone, it was a sum), then $18, 19, 20 \to 57$, then $21, 22, 23 \to 66$, then $24, 25, 26 \to 75$. Write the differences $s_{n+1} - s_n$: $10, 11, 9, 10, 11, 9, \ldots$. That is the guiding light. The sums are not exactly arithmetic, but they wobble around an arithmetic trend with period 3, and a wobble of period 3 riding on a step size close to 10 is exactly the fingerprint of a self-referential block structure. This sentence, three steps remove twelve numbers of data but a block of ten integers only has ten slots, has to be self-generated the moment you see it, in the arena, with nothing but this differenced data in front of you.

Fields embedded: Sequences and recurrences dressed in a combinatorics costume. The real content being tested is whether you can turn an operational rule into a provable block invariant by strong induction, and whether "ends in 2015" becomes, in your hands, a congruence you can actually solve.

Approaches, ranked worst to best

1. **brute simulation to the target.** *Cue:* the specific four digit tail in "prove or disprove". *Circuit:* none, this is zombie maths. it costs unbounded ink: the first term that ends in 2015 turns out to be $s_{1201} = 12015$, well over a thousand steps of hand simulation before any hit. Even granting a correct answer, this route earns no marks for insight and is not finishable in the arena.

2. **spot the pattern, state it, skip the proof.** *Cue:* the visibly near-arithmetic differences $10, 11, 9, 10, 11, 9, \ldots$. *Circuit:* pattern recognition with no induction behind it. this finds the right idea fast but leaves every claim about s_n for large n unjustified. A grader has no reason to believe a guessed formula holds at $n = 4201$ just because it held for $n \leq 8$; this loses most of the marks despite having the right instinct.

3. **strong induction on blocks of three steps proving a self-referential decade lemma, then a linear congruence.** *Cue:* the period-3 wobble riding on a step size of 10. *Circuit:* the block

invariant is exactly the periodicity the rule was begging you to find. the ideal solution. It proves the pattern in one clean induction and needs only ONE controllable sub-case of the wobble, not a full classification, because the problem only demands existence.

The ideal solution

Index the steps from 0: at step n let $a_n < b_n < c_n$ be the three numbers removed and $s_n = a_n + b_n + c_n$ the accompanying sum, so $s_0 = 6$, $s_1 = 16$, $s_2 = 27$, $s_3 = 36, \ldots$ exactly as given. Group steps into blocks of three indexed by $q = 0, 1, 2, \ldots$, block q consisting of steps $3q, 3q+1, 3q+2$.

Claim 11.8.1. *For every integer $q \geq 0$:*

(i) $a_{3q} = 10q + 1$, $b_{3q} = 10q + 2$, $c_{3q} = 10q + 3$, *so* $s_{3q} = 30q + 6$;

(ii) $s_q = 10q + h_q$ *for some* $h_q \in \{5, 6, 7\}$, *and* $\{a_{3q+1}, b_{3q+1}, c_{3q+1}\} = \{10q + 4, 10q + 5, 10q + 6, 10q + 7\} \setminus \{s_q\}$, *so* $s_{3q+1} = 40q + 22 - s_q$, *equal to* $30q + 17$, $30q + 16$, $30q + 15$ *according as* $h_q = 5, 6, 7$;

(iii) $a_{3q+2} = 10q + 8$, $b_{3q+2} = 10q + 9$, $c_{3q+2} = 10q + 10$, *so* $s_{3q+2} = 30q + 27$.

Proof. Strong induction on q. For $q = 0$, steps $0, 1, 2$ are given directly: $1, 2, 3 \to 6$ (so $h_0 = 6$), $4, 5, 7 \to 16$, $8, 9, 10 \to 27$, matching (i) to (iii).

Fix $q \geq 1$ and assume the Claim for every $q' < q$. Two consequences follow.

First, every number removed during steps $0, \ldots, 3q - 1$ lies in $\{1, \ldots, 10q\}$, and all $10q$ of them are removed. Each earlier block $q' < q$ removes nine numbers directly, $\{10q' + 1, \ldots, 10q' + 10\} \setminus \{s_{q'}\}$ by (i) to (iii), and by (ii) at index q' the missing tenth number $s_{q'}$ also lies inside that same decade. So decade q' is entirely crossed off by the end of block q', for every $q' < q$, giving $\{1, \ldots, 10q\}$ fully crossed off after step $3q - 1$. It is worth repeating why nothing else interferes: any sum $s_{q'}$ with $q \leq q' \leq 3q - 1$ already computed lies, by (ii), in the later decade $\{10q' + 1, \ldots, 10q' + 10\}$ with $q' \geq q$, hence strictly above $10q + 3$, and cannot masquerade as one of the next three integers. So the three smallest survivors at the start of step $3q$ are exactly $10q + 1, 10q + 2, 10q + 3$, proving (i) and giving $s_{3q} = 30q + 6$.

Second, since $q \geq 1$ we have $q < 3q$, so step q has already run by step $3q$; step q belongs to block $\lfloor q/3 \rfloor < q$, covered by the hypothesis, so s_q already has a fixed value $10q + h_q$ with $h_q \in \{5, 6, 7\}$. After step $3q$ removes $10q + 1, 10q + 2, 10q + 3$, the candidates $10q + 4, 10q + 5, 10q + 6, 10q + 7$ have exactly one entry, s_q, already crossed off (no other sum can land here, by the same decade-disjointness argument). So step $3q + 1$ removes the remaining three of these four, giving $s_{3q+1} = 40q + 22 - s_q$, proving (ii). In every case the untouched top of the decade, $10q + 8, 10q + 9, 10q + 10$, is now the full set of survivors, so step $3q + 2$ removes them, proving (iii). This closes the induction. □

The Claim answers everything. Parts (i) and (iii) fix the units digits of s_{3q} and s_{3q+2} at 6 and 7, so only s_{3q+1} can ever end in the digit 5 that 2015 needs, and it does so exactly when $h_q = 7$. Part (iii) supplies an unconditional witness for this: $s_{3q'+2} = 30q' + 27 = 10(3q' + 2) + 7$, so $h_{3q'+2} = 7$ automatically, for every q', with no dependence on deeper history. Choosing $q \equiv 2 \pmod 3$ therefore guarantees $h_q = 7$ and $s_{3q+1} = 30q + 15$, at no extra cost.

Solve the tail: $30q + 15 \equiv 2015 \pmod{10^4} \iff 3q \equiv 200 \pmod{1000}$. Since $3 \times 667 \equiv 1 \pmod{1000}$, this gives $q \equiv 200 \times 667 \equiv 400 \pmod{1000}$. Combined with $q \equiv 2 \pmod 3$ (moduli coprime, so the Chinese remainder theorem applies): write $q = 1000j + 400$; since $1000 \equiv 1$ and $400 \equiv 1 \pmod 3$, this needs $j + 1 \equiv 2 \pmod 3$, i.e. $j \equiv 1 \pmod 3$. Take $j = 1$, giving $q = 1400$, and indeed $1400 \equiv 2 \pmod 3$.

With $q = 1400$: $N = 3q + 1 = 4201$ and

$$s_{4201} = 30(1400) + 15 = 42015,$$

whose base ten representation ends in 2015. The statement is proved, not disproved: $\boxed{42015}$ is an explicit term of the sequence with this property.

Layer 1: WHAT

Block 1: define a_n, b_n, c_n, s_n and group steps into blocks of three indexed by q. Block 2: extend the given data by hand to expose the near-arithmetic wobble $10, 11, 9, 10, 11, 9, \ldots$. Block 3: recognise the wobble as a decade colliding with itself, one number in each block of ten was already used up earlier as a sum. Block 4: prove the decade lemma by strong induction, tracking carefully why no stray sum can interfere. Block 5: read off that only the middle sub-case of each block can end in the digit 5, and that one sub-case, $h_q = 7$, is available unconditionally. Block 6: turn "ends in 2015" into a pair of congruences and solve them by the Chinese remainder theorem to name an explicit witness, 42015.

Juxtapose against the nominal route: the zombie approach simulates the sieve term by term, hunting for a hit that does not arrive until $s_{1201} = 12015$, well past a thousand steps, with no argument for why it must arrive at all. The bifurcation point is Block 3: the instant you stop asking what the next term is and start asking why the differenced pattern has period exactly 3, you have left brute force behind for good.

Layer 2: HOW

Block 1. **Cue.** The problem hands you a rule, not a formula, and asks a question purely about digits. **Why natural.** An operational definition plus a digit-ending question is the standard signature of "find the eventual period, then solve a congruence"; no other tool turns a rule into digits. **Link to the next block.** Naming a_n, b_n, c_n, s_n turns a story into an object you can induct on.

Block 2. **Cue.** The given data already shows four terms, too few to trust, just enough to suggest. **Why natural.** Extending by hand costs a few lines and either kills a wrong guess immediately or hands you the exact period to prove. **Link to the next block.** The difference sequence $10, 11, 9, 10, 11, 9$ has period 3, and three steps removing four numbers of data against a block of ten is a ratio worth interrogating, not admiring.

Block 3. **Cue.** Each block of three steps produces twelve numbers of data, but a decade has only ten slots. **Why natural.** Two of those twelve are sums, and sums are large; if a sum from one decade happens to have a value landing inside a later decade, it pre-empties a slot there before that decade is ever reached directly. **Link to the next block.** The quantum leap is naming the self-reference precisely: block q's missing tenth number should be s_q itself, computed earlier, or in the same breath when $q = 0$.

Block 4. **Cue.** A claim depending on every smaller index and on the current index itself is a strong induction with a lookahead. **Why natural.** The base case is forced by direct computation; the step needs the decade fact for all $q' < q$ to clear the ground below $10q$, and the value fact for q itself, legitimate precisely because $q < 3q$ once $q \geq 1$. **Link to the next block.** Once proved, every s_n is known up to the single unresolved bit $h_q \in \{5, 6, 7\}$.

Block 5. **Cue.** The question only asks for existence, not a full description of the sequence. **Why natural.** Chasing the general rule for h_q when $q \equiv 1 \pmod 3$ would need one more layer of recursion, extra ink this problem never asked for; the $q \equiv 2 \pmod 3$ sub-case is unconditional and free, predicated on nothing but part (iii) of the Claim. **Link to the next block.** Fix attention on $s_{3q+1} = 30q + 15$, valid whenever $q \equiv 2 \pmod 3$.

Block 6. **Cue.** "Ends in 2015" is a statement about residues modulo 10^4. **Why natural.** A linear congruence in q modulo 1000, intersected via the Chinese remainder theorem with $q \equiv 2 \pmod 3$,

is guaranteed a solution because the moduli are coprime. **Link to the next block.** Solve, substitute, exhibit $s_{4201} = 42015$, done.

Layer 3: WHY

The truth underneath this problem is a resonance. The rule consumes three integers and manufactures one new one, their sum, every step, so across m steps roughly $3m$ integers are consumed directly while the m sums produced grow roughly linearly at a rate close to the consumption rate. When the output stream's growth rate is close to a whole multiple of the input rate, here a block of three steps growing sums by about 30 while consuming a decade of 10, the outputs and the frontier of unconsumed integers are forced to interleave in a fixed, periodic way. This is the same phenomenon behind Beatty sequences, three-distance theorems, and other greedy constructions on $\mathbb{N}$: whenever removal and reinsertion rates are commensurate, the long-run behaviour collapses onto a finite set of local pictures.

The discernment mechanism: when a sequence is defined by an operation rather than a formula, and the question concerns residues, compute enough terms to see whether the process is nearly arithmetic; if it is, suspect the near-arithmetic behaviour is exact modulo the deviation's own period, then hunt for the reason using the resonance above. Train this by deliberately extending any "sequence defined by an operation" problem by at least six terms before attempting a formula, purely to expose the period, the way $6, 16, 27, 36$ was extended to $46, 57, 66, 75$ here.

What the examiners were nudging toward: giving exactly four terms, no more, is a dare. Four terms is too few to trust a period-3 phenomenon by eye alone, on purpose, so a candidate who only pattern-matches without proving anything is exposed, while a candidate willing to extend the data and then prove the extension earns the full ten marks.

Revolving concepts

Strong induction with a lookahead: ordinary induction assumes the statement at $n-1$ to get n; here the step at block q needs the statement at every $q' < q$, to clear ground, and separately the value fact at index q itself, legitimate because $q < 3q$ for $q \geq 1$ means index q has already been resolved by an earlier block. See Book 1, field: Sequences and recurrences, topic: strong induction. Whenever a recursively defined object at position N needs information from a position that is a function of N itself, check that the referenced position is provably earlier in the construction order before invoking it; that check is the entire content of the induction.

Digit-ending conditions as congruences: the base ten representation of x ends in the digit string D exactly when $x \equiv D \pmod{10^{(\text{digits in } D)}}$, together with $x \geq D$. See Book 1, field: Number theory, topic: congruences and modular arithmetic. Solving $aq \equiv b \pmod{m}$ with $\gcd(a, m) = 1$ is done by multiplying through by the inverse of a modulo m, found by the extended Euclidean algorithm; combining two congruences with coprime moduli is the Chinese remainder theorem, which guarantees a solution unique modulo the product of the moduli.

Greedy self-consuming sequences: any process shaped like "remove the k smallest surviving elements of $\mathbb{N}$, adjoin a new element built from them, remove that too, repeat" is a recognisable Putnam genre. If the adjoined element's typical size grows linearly in the step count at a rate commensurate with $k+1$, the process is eventually periodic when viewed in blocks of the right length, and the correct method is always the same two moves used here: extend by hand to expose the period, then prove a block invariant by strong induction.

Circuit Bank

- **WHEN** a sequence is defined operationally and the question asks about residues or ending digits **DO** extend the data by hand for at least six terms and difference it, looking for a periodic

wobble around an arithmetic trend **BECAUSE** a periodic wobble is provable by a block-level strong induction, and only the induction, not the raw terms, can be trusted at the scale the question demands

- **WHEN** a "prove or disprove" question concerns whether some residue or digit pattern occurs, ever, in a loosely constrained integer process **DO** default to assuming it is true and hunt for one explicit witness **BECAUSE** disproving requires a genuine obstruction, rare and hard to manufacture for sieves this free, while proving needs only one construction
- **WHEN** an inductive step needs a fact about an index that looks earlier but is really tied to the current one, such as s_q needing to be known while proving something about block q **DO** check explicitly that the referenced index is resolved strictly before the point of use, here $q < 3q$ for $q \geq 1$, and treat equality as the base case **BECAUSE** skipping this check is the single most common way this style of induction silently becomes circular
- **WHEN** a problem only asks for existence of one witness, not a full classification **DO** solve only the branch of any case split that is easiest to pin down unconditionally, and ignore the rest **BECAUSE** full classification costs ink the problem is not paying for, and the discarded branches, here h_q for $q \equiv 0, 1 \pmod 3$, are never needed for an existence claim

Backtesting

1. First useful cue: an operationally defined sequence paired with a question purely about the last four digits.
2. Method the expression suggested: extend the given terms by hand and difference them to expose a periodic wobble.
3. Where you would hesitate, and the thought to reinforce: trusting the pattern without proof; the thought is that a strong induction with a lookahead is available precisely because the referenced index q always sits inside an earlier block once $q \geq 1$.
4. The trap the question sets: giving only four terms, too few to responsibly pattern match, to punish anyone who does not extend the data and prove the extension.
5. What makes you recognise it faster next time: noticing the consumption rate, three per step, and the decade size, ten, are close multiples, the numeric fingerprint of a resonant, eventually periodic sieve.
6. The reusable decision rule: when a rule-defined sequence is asked about modulo a power of ten, extend by hand, find the period by differencing, prove a block invariant by strong induction, then solve the resulting congruence for a single witness.

Mirror drills

1. Starting from $1, 2, 3, 4, \ldots$, repeatedly cross off the two smallest remaining numbers and their sum, three numbers per step. Show the resulting sequence of sums is eventually periodic modulo a suitable power of ten. *Answer:* each step removes three numbers and each sum lands inside a block of size close to a small multiple of the step count; the same block invariant technique applies, with the block length worked out by matching removal rate to sum size.
2. A sequence is defined by $t_0 = 1$ and $t_{n+1} = t_n + (\text{digit sum of } t_n)$. Does some t_n end in the digits 2026? *Answer:* this process is not obviously periodic, digit sum is not linear in t_n, so the resonance argument above does not transfer directly; flag this as a structurally different, harder genre rather than forcing the same technique.

3. Solve $17q \equiv 5 \pmod{2015}$ for all integer q, then combine with $q \equiv 1 \pmod{4}$ by the Chinese remainder theorem to give the general solution modulo 4×2015. *Answer:* $\gcd(17, 2015) = 1$ gives a unique inverse of 17 mod 2015, found by the extended Euclidean algorithm, then combine exactly as in the main solution.

11.9 B3. which arithmetic progression matrices have AP powers

Let S be the set of all 2×2 real matrices

$$M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$$

whose entries a, b, c, d, in that order, form an arithmetic progression. Find all matrices M in S for which there is some integer $k > 1$ such that M^k is also in S.

Read it out loud

Four unknowns, tied together by "in arithmetic progression, in that order." Structurally that is two linear equations in four unknowns, dressed up as a combinatorial condition so it does not look like one at first glance. The first useful move, available within thirty seconds, is to write a, b, c, d as $a, a+t, a+2t, a+3t$: the sentence handed you this substitution for free, and it collapses four unknowns to two.

The second cue sits in "there is some integer $k > 1$ such that M^k is also in S": you are being asked whether a HIGH power of a matrix lands back inside a thin, linearly defined slice of the space of 2×2 matrices. Computing M^k symbolically by repeated multiplication is exactly the situation the trace and determinant quadratic identity exists to shortcut: every 2×2 matrix satisfies $M^2 = \tau M - \delta I$ with $\tau = \operatorname{tr} M$, $\delta = \det M$, which turns an exponentially growing computation into a two-term linear recursion before a single eigenvector is ever written down.

Fields embedded: Matrices and determinants on the surface, Linear algebra and Sequences and recurrences underneath. What is really being tested is whether you reach for trace and determinant before reaching for diagonalisation, and whether you can test membership in a linearly defined set by evaluating its defining equations rather than rebuilding a basis for it.

Approaches, ranked worst to best

1. **symbolic brute force.** *Cue:* "is some power in S". *Circuit:* none, pure computation. compute $M^2, M^3, \ldots$ directly in a, t and impose the arithmetic-progression conditions on raw entries at each step. The algebra explodes well before $k = 5$ and gives no reason to stop looking, since a negative result at small k proves nothing about all k.

2. **explicit diagonalisation.** *Cue:* a 2×2 real matrix, powers wanted. *Circuit:* eigenvalues and a change of basis. find M's eigenvalues and eigenvectors, write $M = PDP^{-1}$, compute $M^k = PD^kP^{-1}$ explicitly, then impose the arithmetic-progression condition on the result. This is correct and is essentially the official first solution's route; it works, but it drags an explicit $\sqrt{\tau^2 + 8t^2}$ and a recursively defined auxiliary sequence through the entire rest of the argument, costing more ink than the problem needs.

3. **trace/determinant reduction plus a functional-evaluation test for membership.** *Cue:* $M^2 = \tau M - \delta I$ holding for every 2×2 matrix, and S being linearly defined. *Circuit:* no eigenvector is ever computed; only the sign and magnitude of the eigenvalues matter, never their exact values. the ideal solution: every power collapses to $t_k M + u_k I$, and S-membership collapses to the single condition $u_k = 0$ by evaluating S's defining equations at the identity matrix.

The ideal solution

Write the entries as $a, a+t, a+2t, a+3t$ for real a, t, forced directly by "in arithmetic progression, in that order." Then

$$\tau := \operatorname{tr} M = a + d = 2a + 3t, \qquad \delta := \det M = ad - bc = a(a+3t) - (a+t)(a+2t) = -2t^2.$$

For any 2×2 matrix $M = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$, direct expansion gives $M^2 = \begin{pmatrix} a^2+bc & b(a+d) \\ c(a+d) & d^2+bc \end{pmatrix}$, which equals $\tau M - \delta I = \begin{pmatrix} (a+d)a-(ad-bc) & (a+d)b \\ (a+d)c & (a+d)d-(ad-bc) \end{pmatrix}$ entrywise, since $(a+d)a-(ad-bc) = a^2+bc$ and likewise for the last entry. So

$$M^2 = \tau M - \delta I$$

for every 2×2 matrix, with no special property of M required.

Case $t = 0$**.** Then $a = b = c = d$ and $M = aJ$ where $J = \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$ satisfies $J^2 = 2J$, so $J^k = 2^{k-1}J$ for all $k \geq 1$ by induction, giving $M^k = a^k 2^{k-1} J$: a matrix with all four entries equal, trivially an arithmetic progression with common difference 0. So every $M = aJ$, $a \in \mathbb{R}$, works for every $k \geq 2$.

Case $t \neq 0$**.** Then $\delta = -2t^2 \neq 0$, so M is invertible, and $b = a + t \neq a$ means M is not a scalar matrix, so $\{I, M\}$ is linearly independent. Repeated use of $M^2 = \tau M - \delta I$ gives, for every $k \geq 1$,

$$M^k = t_k M + u_k I, \qquad t_1 = 1,\ u_1 = 0, \qquad t_{k+1} = \tau t_k + u_k, \quad u_{k+1} = -\delta t_k,$$

so $u_k = -\delta t_{k-1} = 2t^2 t_{k-1}$ for $k \geq 2$.

Characterise S by its defining equations directly. Writing a general matrix as $N = \begin{pmatrix} p & q \\ u & v \end{pmatrix}$, "$p, q, u, v$ in arithmetic progression" means $q - p = u - q = v - u$, i.e.

$$L_1(N) := p - 2q + u = 0, \qquad L_2(N) := q - 2u + v = 0.$$

These are linear functionals; $M \in S$ gives $L_1(M) = L_2(M) = 0$, and direct evaluation on the identity gives $L_1(I) = 1 - 0 + 0 = 1$ and $L_2(I) = 0 - 0 + 1 = 1$. So for any scalars x, y,

$$L_i(xM + yI) = xL_i(M) + yL_i(I) = y \qquad (i = 1, 2).$$

Both equations return the same number, 1, when evaluated at I: I's entries $1, 0, 0, 1$ read the same forwards and backwards, and that palindromic symmetry is exactly why both differences see the same value. Applying this with $x = t_k$, $y = u_k$: $M^k \in S \iff u_k = 0$. Membership collapses from two conditions to one, because they coincide on I.

So, for $k \geq 2$, $M^k \in S \iff t_{k-1} = 0$. The sequence t_j obeys $t_{j+1} = \tau t_j + 2t^2 t_{j-1}$, whose characteristic equation $x^2 - \tau x - 2t^2 = 0$ is exactly M's own characteristic equation, with roots

$$\lambda_{1,2} = \frac{\tau \pm \sqrt{\tau^2 + 8t^2}}{2},$$

real and distinct (discriminant $\tau^2 + 8t^2 > 0$ automatically, since $t \neq 0$), with $\lambda_1 \lambda_2 = -2t^2 < 0$: the eigenvalues of M always have opposite sign, whatever τ is. Since $t_1 = 1$, $t_2 = \tau = \lambda_1 + \lambda_2$ match the standard closed form for this recursion, $t_j = \dfrac{\lambda_1^j - \lambda_2^j}{\lambda_1 - \lambda_2}$ for every $j \geq 1$.

So $t_{k-1} = 0 \iff \lambda_1^{k-1} = \lambda_2^{k-1}$. Taking absolute values, $|\lambda_1|^{k-1} = |\lambda_2|^{k-1}$; since $x \mapsto x^{k-1}$ is injective on $[0, \infty)$ for $k - 1 \geq 1$, this forces $|\lambda_1| = |\lambda_2|$, which combined with $\lambda_1 \lambda_2 < 0$ forces $\lambda_2 = -\lambda_1$, i.e. $\tau = 0$. Conversely, if $\tau = 0$ then $\lambda_2 = -\lambda_1$ and $\lambda_1^{k-1} = \lambda_2^{k-1} = (-1)^{k-1} \lambda_1^{k-1}$ holds, as $\lambda_1 \neq 0$, exactly when $k - 1$ is even, i.e. k is odd.

So, for $t \neq 0$: some $k > 1$ works if and only if $\tau = 0$, and then every odd $k \geq 3$ works. $\tau = 2a + 3t = 0$ gives $a = -\frac{3t}{2}$, so the entries are $-\frac{3t}{2}, -\frac{t}{2}, \frac{t}{2}, \frac{3t}{2}$, i.e. $M = \frac{t}{2} \begin{pmatrix} -3 & -1 \\ 1 & 3 \end{pmatrix}$ for $t \neq 0$; taking $t = 0$ here only reproduces the zero matrix, already covered by the first family.

Combining both cases, with $B = \begin{pmatrix} -3 & -1 \\ 1 & 3 \end{pmatrix}$:

$$\boxed{M \in \left\{ \lambda \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} : \lambda \in \mathbb{R} \right\} \cup \{\mu B : \mu \in \mathbb{R}\}}$$

are exactly the matrices of S with $M^k \in S$ for some integer $k > 1$.

Layer 1: WHAT

Block 1: collapse a, b, c, d to (a, t) via the arithmetic-progression substitution. Block 2: derive τ, δ and the universal identity $M^2 = \tau M - \delta I$. Block 3: split into $t = 0$, trivial and always valid, and $t \neq 0$, the real content. Block 4: characterise S by its two linear equations and notice both agree on I, collapsing membership to $u_k = 0$. Block 5: solve $t_{k-1} = 0$ via the eigenvalues' sign and magnitude alone, never their exact values. Block 6: translate $\tau = 0$ back into matrix form and state the union of two families.

Juxtapose against the nominal route: diagonalise M explicitly and carry $\sqrt{\tau^2 + 8t^2}$ through every subsequent line, as the official first solution does. The bifurcation point is Block 4: the decision to test membership by evaluating S's own defining equations at I, rather than re-deriving an explicit basis for S and arguing linear independence against it.

Layer 2: HOW

Block 1. **Cue.** Four unknowns tied together by "in arithmetic progression, in that order." **Why natural.** That phrase is exactly two linear equations in four unknowns; the natural response is to substitute the general term of an arithmetic progression directly and cut the unknowns to two. **Link to the next block.** Every entry is now written in a and t alone, so trace and determinant are the next natural objects.

Block 2. **Cue.** A 2×2 matrix and a question about its powers. **Why natural.** Every 2×2 matrix obeys $M^2 = \tau M - \delta I$, a one-line consequence of expanding both sides; there is no reason to reach for eigenvectors before checking whether this alone is enough. **Link to the next block.** Repeated multiplication turns this into a two-term linear recursion for the coefficients of $M^k = t_k M + u_k I$.

Block 3. **Cue.** $t = 0$ makes M a scalar multiple of the all-ones matrix, a degenerate case worth clearing first. **Why natural.** Degenerate cases often satisfy a target condition trivially and for a different reason than the generic case; separating them stops one proof doing two jobs at once. **Link to the next block.** With $t = 0$ disposed of, the whole remaining question is about $t \neq 0$, where M is invertible and non-scalar.

Block 4. **Cue.** S itself is defined by two linear equations, and M^k always lives in $\operatorname{span}\{M, I\}$. **Why natural.** Testing whether a vector in a known two-dimensional space lies in a subspace defined by linear equations is a direct evaluation, not a search; evaluating both equations at I is the cheapest possible test. **Link to the next block.** Both equations return the same constant 1 on I, so membership collapses to the single condition $u_k = 0$.

Block 5. **Cue.** $u_k = 0$ for $k \geq 2$ is the same as $t_{k-1} = 0$, a statement about a linear recursion's closed form. **Why natural.** The recursion's characteristic roots are M's own eigenvalues, always real, discriminant automatically positive, and of opposite sign, product $= \delta < 0$; equal powers of opposite-signed reals is a magnitude-then-sign argument, not a computation. **Link to the next block.** $t_{k-1} = 0$ happens exactly when $\tau = 0$ and k is odd, translating a matrix-power question into one clean scalar condition.

Block 6. Cue. $\tau = 0$ is one linear equation in a, t. **Why natural.** Substituting back converts the abstract condition into an explicit one-parameter family of matrices, matching the concrete $t = 0$ family already found. **Link to the next block.** Union the two families for the final answer.

Layer 3: WHY

The truth underneath: any 2×2 real matrix obeys a two-term linear recursion for its powers, immediate from the trace/determinant quadratic and needing no eigenvector, so M^k always has the shape $t_k M + u_k I$. A codimension-two condition like "lies in S", restricted to this two-dimensional slice, can be at most a two-equation linear system in (t_k, u_k); here it degenerates to one equation because S's two defining functionals coincide on I. That coincidence is predicated on nothing more exotic than I's entries reading the same forwards and backwards, $1, 0, 0, 1$, which is exactly the symmetry "arithmetic progression" itself respects: reversing a, b, c, d to d, c, b, a negates the common difference but preserves the progression.

The discernment mechanism: when a problem restricts to a linear subspace of matrices and asks about powers staying inside it, look first for the smallest invariant subspace containing the powers you care about, here $\operatorname{span}\{M, I\}$, forced by the degree two trace/determinant identity, then test the original subspace's defining equations on a convenient spanning set of that invariant subspace rather than rebuilding a basis for it. Train this by reaching for trace and determinant, or the minimal polynomial in general, before reaching for eigenvectors, whenever "matrix power stays in a linearly defined set" appears.

What the examiners were nudging toward: phrasing the constraint as entries "in arithmetic progression, in that order" rather than stating the equations outright is a small act of misdirection, reading like a combinatorial condition and tempting a search for integer patterns, when the intended reduction is purely linear algebraic, two equations, four unknowns, straight down to two free parameters, available in the arena without ever factoring $\tau^2 + 8t^2$.

Revolving concepts

The 2×2 trace/determinant quadratic $M^2 = \tau M - \delta I$: every 2×2 real matrix satisfies its own characteristic equation, checkable by direct expansion in one line. See Book 1, field: Matrices and determinants, topic: characteristic polynomial. This is a genuine recursion once multiplied through by M^{k-1}, giving $M^{k+1} = \tau M^k - \delta M^{k-1}$, and it generalises to $n \times n$ matrices via the degree n characteristic polynomial, giving a length n linear recursion for the coefficients of M^k in the basis $I, M, \ldots, M^{n-1}$.

Linear functionals cutting out a subspace, evaluated at a fixed vector: if $V = \{N : L_1(N) = \cdots = L_r(N) = 0\}$ and $x \in V$, then for fixed y, $sx + ty \in V \iff t \cdot L_i(y) = 0$ for every i; when several $L_i(y)$ coincide, membership collapses to fewer independent conditions. See Book 1, field: Linear algebra, topic: subspaces and dual functionals.

Closed forms for order two linear recurrences via characteristic roots: $x_{j+1} = px_j + qx_{j-1}$ with distinct roots λ_1, λ_2 of $z^2 - pz - q = 0$ gives $x_j = A\lambda_1^j + B\lambda_2^j$, fit from the first two terms; here it produces $t_j = (\lambda_1^j - \lambda_2^j)/(\lambda_1 - \lambda_2)$. See Book 1, field: Sequences and recurrences, topic: auxiliary equation method. A repeated root λ gives $x_j = (A + Bj)\lambda^j$; complex conjugate roots give a modulus times a cosine or sine envelope; the same machinery computes ordinary Fibonacci-type sequences and any linear recursion met on the Putnam, this problem's t_k, u_k included.

Sign and reality of eigenvalues from the determinant alone: a real 2×2 matrix with negative determinant has two real eigenvalues of strictly opposite sign, since the discriminant $\tau^2 - 4\delta$ is then automatically positive and the product of the roots equals $\delta < 0$. See Book 1, field: Matrices and determinants, topic: eigenvalues of 2×2 matrices. This single sign fact, not the eigenvalues' exact values, is what closes the problem.

Circuit Bank

- **WHEN** a matrix's entries obey a single linear constraint, arithmetic progression, equal entries, trace zero, and a question asks about powers staying inside that constraint **DO** reduce to the trace/determinant recursion $M^k = t_k M + u_k I$ before touching the specific constraint **BECAUSE** this shrinks the entire power sequence to two scalar sequences that any linear condition can then be tested against directly
- **WHEN** a target subspace S is defined by linear equations L_i **DO** test membership of $tM + uI$ by evaluating each L_i at I directly **BECAUSE** $L_i(tM + uI) = tL_i(M) + uL_i(I) = uL_i(I)$ whenever $M \in S$, collapsing the whole question to the sign of u against the known constants $L_i(I)$
- **WHEN** two real numbers raised to the same positive integer power are claimed equal **DO** split the claim into equal magnitude, injective on $[0, \infty)$, and compatible sign, treated separately **BECAUSE** odd and even powers treat sign differently, and merging the two steps is where errors hide
- **WHEN** a 2×2 real matrix has negative determinant **DO** conclude immediately, with no computation, that its eigenvalues are real and of opposite sign **BECAUSE** the discriminant $\tau^2 - 4\delta$ is then automatically positive and the product of the roots is $\delta < 0$

Backtesting

1. First useful cue: "in arithmetic progression, in that order" applied to four matrix entries, a linear constraint dressed as a combinatorial one.
2. Method the expression suggested: substitute $b, c, d = a + t, a + 2t, a + 3t$, cutting to two free parameters, then reach for trace and determinant rather than eigenvectors.
3. Where you would hesitate, and the thought to reinforce: wanting to diagonalise M because that is how powers are usually computed; the thought is that the trace/determinant recursion already gives every power in $\text{span}\{M, I\}$ without naming an eigenvector.
4. The trap the question sets: "there is some integer $k > 1$" invites chasing specific small k by direct computation, which finds $k = 3$ for one family quickly but gives no reason no other k, or no other matrix, works; only the sign and magnitude argument on the eigenvalues closes the case fully.
5. What makes you recognise it faster next time: seeing a linear constraint on matrix entries, together with a "does some power stay in the constraint" question, as asking you to find the invariant two-dimensional slice $\{tM + uI\}$ and test the constraint's defining equations on I alone.
6. The reusable decision rule: reduce matrix powers via trace and determinant before anything else, characterise the target set by its defining linear equations, evaluate those equations at the identity, and resolve the resulting scalar recursion by the sign and magnitude of the characteristic roots rather than their exact values.

Mirror drills

1. Let $N = \begin{pmatrix} a & b \\ c & d \end{pmatrix}$ be real with $a + d = 0$. Show N^2 is always a scalar multiple of I. *Answer:* immediate from $N^2 = \tau N - \delta I$ with $\tau = 0$, so $N^2 = -\delta I$.
2. Let T be the set of 2×2 real matrices with $a = d$, equal diagonal entries. For which $M \in T$ does $M^k \in T$ for some $k > 1$? *Answer:* the defining equation is $L(N) = p - v = 0$ and

$L(I) = 1 - 1 = 0$, so every power of every $M \in T$ automatically stays in T; a degenerate case worth noting precisely because it contrasts with B3, where $L_i(I) \neq 0$ did the work.

3. For real $t \neq 0$, does $\begin{pmatrix} 0 & t \\ -t & 0 \end{pmatrix}$ have some power equal to a real scalar multiple of $\begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix}$? *Answer:* its eigenvalues are $\pm it$, not real, so the opposite-sign-real-eigenvalue argument breaks down entirely; check directly that every power is a scalar multiple of I or of the original matrix, never of the all-ones matrix.

11.10 B4. Ravi substitution decouples a triangle sum

Let T be the set of all triples (a, b, c) of positive integers for which there exist triangles with side lengths a, b, c. Express

$$\sum_{(a,b,c)\in T} \frac{2^a}{3^b 5^c}$$

as a rational number in lowest terms.

Read it out loud

This is a three variable geometric sum taken not over a box but over the integer points of an open cone: the summand $2^a 3^{-b} 5^{-c}$ is completely separable in a, b, c, and every ounce of difficulty has been pushed into the index set. Say that out loud, because it decides everything. If the index set were $a, b, c \geq 1$ with no further condition, you would write the answer in one line as a product of three geometric series. So the whole problem is: change coordinates until the index set becomes a product again.

The fields wired into the question are tiny in number and that is the tell. Geometry supplies exactly one fact, the triangle inequality. Analysis supplies exactly one fact, that a sum of nonnegative terms may be regrouped and reordered at will. Everything else is the geometric series. The examiners are not testing whether you can sum $\sum r^n$; they are testing whether, on being handed a constraint, you reroute into coordinates in which the constraint dissolves.

The first useful cue is the phrase "for which there exist triangles with side lengths a, b, c". That is the command term. It does not say $a < b + c$; it says triangle. A triangle is the object whose defining data is three positive side lengths satisfying three strict inequalities, and the canonical way to store those three inequalities without carrying them around is to name their slacks. Within thirty seconds you should have written

$$x = b + c - a, \qquad y = c + a - b, \qquad z = a + b - c,$$

and underneath it, "$x, y, z > 0$, integers, and $y + z = 2a$, $z + x = 2b$, $x + y = 2c$". That last line is where the whole solution lives, and it also contains the one trap: over the integers the correspondence is not onto, because x, y, z must all have the same parity.

The second cue is the specific bases $2, 3, 5$ with 2 upstairs and $3, 5$ downstairs. That asymmetry is deliberate. It is there so that the answer is not symmetric and so that you cannot fudge the bookkeeping; it also guarantees convergence, because the constraint $a < b + c$ is exactly what keeps 2^a from beating $3^b 5^c$.

Approaches, ranked worst to best

1. **Slice by perimeter.** *Cue:* "triples of positive integers", a Diophantine flavour. *Circuit:* count lattice points level by level, then sum the counting function. You set $a + b + c = N$, count the triangles of perimeter N (Alcuin's sequence, nearest integer to $N^2/48$ or $(N + 3)^2/48$), and try to reassemble. It fails instantly because the summand is not a function of N: $2^a 3^{-b} 5^{-c}$ separates the variables, so slicing by their sum is exactly the wrong slice. Maximum ink, zero progress.

2. **Fix a first.** *Cue:* a is the only exponent in the numerator. *Circuit:* innermost sum over the variable that looks most special. For fixed a the pairs (b, c) must satisfy $b + c > a$ and $|b - c| < a$, which is a wedge intersected with a half plane. You end up splitting the (b, c) quadrant into three regions and summing a two dimensional geometric series over each. It works, but you will write four times as much as you need and each region is an index error waiting to happen.

3. **Fix b and c, geometric in a.** *Cue:* for fixed b, c the legal a form a contiguous interval $|b-c| < a < b+c$. *Circuit:* when one variable ranges over an interval of consecutive integers, do that sum first. This is the honest workhorse and it is the official first solution. You get

$$S = \sum_{b,c \geq 1} \frac{2^{b+c} - 2^{|b-c|+1}}{3^b 5^c},$$

then you must split on the sign of $b-c$ to remove the modulus, giving $S_1 = 85/231$ over $b \leq c$ and $S_2 = 34/77$ over $b > c$, and $S_1 + S_2 = 17/21$. Correct, fully rigorous, and it costs four geometric series plus two reindexings. It ranks below the ideal because the modulus was never necessary: it is an artefact of refusing to change coordinates.

4. **Ravi substitution.** *Cue:* "there exist triangles with side lengths". *Circuit:* name the slacks of the constraining inequalities and use them as coordinates. The three inequalities become $x, y, z > 0$; the summand becomes a product of powers of $1/15$, $2/5$, $2/3$; the only residue of the integrality is a parity condition which splits the sum into two copies of the same product. Three geometric series, one factor of $17/2$, done.

5. **Cone decomposition.** *Cue:* the index set is the set of lattice points of a simplicial cone. *Circuit:* write the cone as a lattice translate of the free monoid on its extreme rays. This is the same computation seen from above, and it is what you should install as the general theory. It also produces the closed form

$$\sum_{(a,b,c) \in T} x^a y^b z^c = \frac{xyz(1+xyz)}{(1-xy)(1-yz)(1-zx)},$$

of which the problem is the case $(x, y, z) = (2, \frac{1}{3}, \frac{1}{5})$. It ranks alongside the ideal solution rather than above it, because on the day you want the shortest correct page, not the most general one.

The ideal solution

Every term of the sum is positive, so the sum is a well defined element of $[0, \infty]$ and may be reordered and regrouped freely; the computation below produces a finite value, which retrospectively justifies every rearrangement.

Lemma 11.10.1. *The map $\Phi(a,b,c) = (b+c-a,\ c+a-b,\ a+b-c)$ is a bijection from T onto*

$$L = \{(x,y,z) \in \mathbb{Z}^3_{>0} :\ x \equiv y \equiv z \pmod 2\},$$

with inverse $(x,y,z) \mapsto \left(\frac{y+z}{2}, \frac{z+x}{2}, \frac{x+y}{2}\right)$.

Proof. Positive integers a, b, c are the side lengths of an actual triangle exactly when the three strict inequalities $a < b+c$, $b < c+a$, $c < a+b$ hold, that is, exactly when $x, y, z > 0$. If $(a,b,c) \in T$ then x, y, z are integers with $y+z = 2a$, $z+x = 2b$, $x+y = 2c$, so each of the three pairwise sums is even, forcing x, y, z to share a parity; hence $\Phi(T) \subseteq L$. Conversely if $(x,y,z) \in L$ then $y+z$, $z+x$, $x+y$ are even, so the stated inverse produces positive integers a, b, c, and these satisfy the three strict inequalities because $x, y, z > 0$. The two maps are visibly mutually inverse. □

Now rewrite the summand. Under the correspondence,

$$\frac{2^a}{3^b 5^c} = 2^{(y+z)/2}\, 3^{-(z+x)/2}\, 5^{-(x+y)/2},$$

so the sum splits over the two parity classes of L.

All even. Put $x = 2u, y = 2v, z = 2w$ with $u, v, w \geq 1$. Then the exponents are $v + w$, $-(w + u)$, $-(u + v)$ and the term regroups as

$$2^{v+w}3^{-(w+u)}5^{-(u+v)} = \Big(\frac{1}{15}\Big)^u \Big(\frac{2}{5}\Big)^v \Big(\frac{2}{3}\Big)^w.$$

Each of u, v, w now appears in exactly one factor: the sum has become a product.

All odd. Put $x = 2u - 1$, $y = 2v - 1$, $z = 2w - 1$ with $u, v, w \geq 1$. Then the exponents are $v + w - 1$, $-(w + u - 1)$, $-(u + v - 1)$, so the term is $\frac{3\cdot 5}{2}$ times the even term with the same (u, v, w):

$$2^{v+w-1}3^{-(w+u-1)}5^{-(u+v-1)} = \frac{15}{2}\Big(\frac{1}{15}\Big)^u \Big(\frac{2}{5}\Big)^v \Big(\frac{2}{3}\Big)^w.$$

The factor $\frac{15}{2} = 3 \cdot 5 \cdot 2^{-1}$ is not an accident: passing from the even class to the odd class subtracts $(1, 1, 1)$ from (a, b, c), and the weight $2^a 3^{-b} 5^{-c}$ scales by $2^{-1}3^1 5^1$ under that shift.

Therefore, writing Σ for the common product,

$$\sum_{(a,b,c)\in T} \frac{2^a}{3^b 5^c} = \Big(1 + \frac{15}{2}\Big)\Sigma = \frac{17}{2}\,\Sigma,$$

$$\Sigma = \sum_{u\geq 1}\Big(\frac{1}{15}\Big)^u \sum_{v\geq 1}\Big(\frac{2}{5}\Big)^v \sum_{w\geq 1}\Big(\frac{2}{3}\Big)^w = \frac{1/15}{14/15}\cdot\frac{2/5}{3/5}\cdot\frac{2/3}{1/3} = \frac{1}{14}\cdot\frac{2}{3}\cdot 2 = \frac{2}{21}.$$

All three ratios lie in $(0, 1)$, so each series converges and the value is finite. Hence

$$\sum_{(a,b,c)\in T} \frac{2^a}{3^b 5^c} = \frac{17}{2}\cdot\frac{2}{21} = \boxed{\frac{17}{21}}.$$

Finally, 17 is prime and $17 \nmid 21$, so $\gcd(17, 21) = 1$ and $17/21$ is already in lowest terms, as the question demanded.

Layer 1: WHAT

Block 1. Read the summand: it is separable. Read the index set: it is not a product. Conclude that the entire task is a change of index.

Block 2. Note all terms are positive, so no convergence anxiety and no ordering discipline is required; the only risk left is bookkeeping.

Block 3. Translate "there exist triangles" into the three strict inequalities and immediately name their slacks x, y, z.

Block 4. Determine the exact image of the substitution over $\mathbb{Z}$: positive, and all of one parity. Prove it is a bijection, since the whole computation is a transport of the sum along that bijection.

Block 5. Substitute in each parity class, and check that each new variable appears in exactly one base. It does; the sum factorises.

Block 6. Sum three geometric series, multiply by the parity factor $1 + \frac{15}{2}$, and check lowest terms.

Juxtaposition. The nominal route agrees with the ideal route through Block 2 and then diverges. The zombie writes "for fixed b, c we need $|b - c| < a < b + c$" and starts summing 2^a over that interval. He is not wrong; he is committed. From that line onwards he owes a case split on the sign of $b - c$, two double sums, four geometric series and a reindex of the inner limit. The bifurcation point is the exact line where the triangle condition is written down. Write it as an interval for a and you have chosen the long road; write it as three positivity statements about x, y, z and you have chosen the short one. Same fact, two encodings, ten times the ink.

Layer 2: HOW

Block 1. **Cue.** The summand is $2^a 3^{-b} 5^{-c}$, a product of one factor per variable, while the index set couples all three. **Why natural.** A sum of a separable function over a product set equals the product of the one dimensional sums. That is the only closed form mechanism available here, so the strategy is forced: make the index set a product. You do not need inspiration to see this; you need to notice which half of the expression is hard. **Link to the next block.** Having decided to reindex, you next need permission to reindex.

Block 2. **Cue.** All terms are positive. **Why natural.** For nonnegative families, the sum over any index set is the supremum of finite partial sums, so it is independent of order and of grouping. This is the nonnegative case of Tonelli's theorem and it is the reason you may split by parity, substitute, and factor without a single epsilon. **Link to the next block.** With permission secured, choose the coordinates.

Block 3. **Cue.** The command term is "triangle", not "inequality". **Why natural.** A triangle inequality is a statement that a quantity is positive. The natural coordinates for a region cut out by linear inequalities are the linear forms themselves: set each inequality's slack to be a new variable. Geometrically, x, y, z are twice the tangent lengths from the vertices to the incircle, which is why they are positive precisely for genuine triangles and vanish exactly in the degenerate case. This is the Ravi substitution and it is a permanent circuit, not a one off. **Link to the next block.** The substitution is over the reals; you now have to ask what it does over the integers.

Block 4. **Cue.** The inverse formulas carry a 2 in the denominator. **Why natural.** Any time an inverse map divides, integrality becomes a congruence condition. Here $y + z = 2a$ forces $y \equiv z$, and similarly for the other pairs, so all three are congruent modulo 2. Missing this is the only way to get the wrong answer with the right method: you would sum over all positive x, y, z and pick up half integer exponents. **Link to the next block.** The image splits into two classes; handle them by the same substitution.

Block 5. **Cue.** Each parity class is a translate of $2\mathbb{Z}^3_{>0}$. **Why natural.** Reparametrising a lattice coset by $u, v, w \geq 1$ is the standard way to turn a congruence condition into a free index. After substitution, u appears only through $3^{-u}5^{-u}$, v only through $2^v 5^{-v}$, w only through $2^w 3^{-w}$: the exponent map, read in the new coordinates, is diagonal. That diagonality is precisely what the Ravi substitution was bought for. **Link to the next block.** Two classes, identical product, differing by one scalar.

Block 6. **Cue.** The odd class is the even class shifted by $-(1,1,1)$ in (a,b,c) coordinates. **Why natural.** The weight is multiplicative in (a,b,c), so a translation of the index set multiplies the whole sum by the weight of the translation vector. Here that weight is $2^{-1}3^1 5^1 = 15/2$, so the two classes contribute in the ratio $1 : \frac{15}{2}$, and the total factor is $\frac{17}{2}$. You can see the answer's numerator 17 being born here, before you have summed anything. **Link to the next block.** Three geometric series and a gcd sentence finish it.

Layer 3: WHY

The mathematical truth beneath this route is that T is the set of lattice points in the interior of a simplicial cone. In $\mathbb{R}^3$, the conditions $a < b + c$, $b < c + a$, $c < a + b$ cut out the cone spanned by the three vectors

$$g_1 = (0,1,1), \qquad g_2 = (1,0,1), \qquad g_3 = (1,1,0),$$

since $(a,b,c) = \frac{x}{2}g_1 + \frac{y}{2}g_2 + \frac{z}{2}g_3$. A cone spanned by three linearly independent vectors is called simplicial, and the fundamental structural fact is this: if g_1, g_2, g_3 generate a sublattice of finite index m in $\mathbb{Z}^3$, then the lattice points of the cone are the disjoint union of m translates of the free monoid generated by the g_i, one translate for each lattice point of the half open fundamental parallelepiped. Here

$$\det \begin{pmatrix} 0 & 1 & 1 \\ 1 & 0 & 1 \\ 1 & 1 & 0 \end{pmatrix} = 2,$$

so the index is 2, the parallelepiped contains exactly two lattice points, and those two translates are precisely the two parity classes. The representative of the nontrivial class is $\frac{1}{2}(g_1+g_2+g_3) = (1,1,1)$, the equilateral triple: the smallest triangle is literally the extra generator you need.

Because the weight $w(a,b,c) = 2^a 3^{-b} 5^{-c}$ is a monoid homomorphism from $(\mathbb{Z}^3, +)$ to $(\mathbb{R}_{>0}, \times)$, summing it over the cone is now automatic:

$$\sum_{(a,b,c)\in T} w = \Big(\sum_{\text{box points } p} w(p) \Big) \prod_{i=1}^{3} \frac{1}{1 - w(g_i)} \quad \text{after the appropriate shift,}$$

and with $w(g_1) = \frac{1}{15}$, $w(g_2) = \frac{2}{5}$, $w(g_3) = \frac{2}{3}$, $w(1,1,1) = \frac{2}{15}$ you get $\frac{2}{21} + \frac{5}{7} = \frac{17}{21}$ a second time. Two independent derivations agreeing is how you should sign off a numerical answer under exam conditions.

The discernment mechanism, stated so you can train it: when a sum's summand is separable and its index set is not, all the work is in the index set, and the only structures whose lattice sums factorise are boxes and simplicial cones. So ask, in this order: (i) is my region a box in disguise? (ii) is it a cone, and what are its extreme rays? (iii) what is the index of the lattice they generate, because that index is the number of extra terms I will have to carry. In this problem the answers are no, yes with rays $(0,1,1), (1,0,1), (1,1,0)$, and 2. The factor $\frac{17}{2}$ and the final 17 are consequences of that index 2.

What the examiners nudged with: they wrote "there exist triangles with side lengths" rather than an inequality, which is a direct instruction to use triangle structure; they chose three distinct prime bases so that no accidental symmetry could rescue a sloppy split; and they asked for lowest terms, which is a quiet reminder that the numerator 17 is a real prime and not a slip. Train the trigger by generalising it: whenever a question restricts to "sides of a triangle", substitute the slacks, even when the question is an inequality rather than a sum, because that is the coordinate system in which the constraint is invisible.

Revolving concepts

Triangle inequality and the Ravi substitution. For positive reals, a, b, c are the sides of a nondegenerate triangle iff $x = b + c - a$, $y = c + a - b$, $z = a + b - c$ are all positive; then x, y, z are twice the tangent lengths from the vertices to the incircle, $x + y + z = a + b + c$, and $s - a = x/2$ where s is the semiperimeter. Over $\mathbb{Z}$ the image is exactly the positive triples of equal parity. Heron's formula becomes $16K^2 = (x+y+z)xyz$, which is the reason Ravi trivialises most triangle inequalities. See Book 1, field: Geometry, topic: triangle inequality and Ravi substitution.

Simplicial cones and lattice points. A cone generated by linearly independent $g_1, \dots, g_d \in \mathbb{Z}^d$ meets $\mathbb{Z}^d$ in $|\det(g_i)|$ translates of $\mathbb{Z}_{\geq 0}g_1 + \cdots + \mathbb{Z}_{\geq 0}g_d$; the translates are indexed by the lattice points of the half open parallelepiped. This is the elementary core of Ehrhart theory and of Stanley's theory of cone generating functions. See Book 1, field: Generating functions, topic: cone decompositions.

Geometric series as the simplest recursion. $\sum_{n\geq 1} r^n = \frac{r}{1-r}$ for $|r| < 1$ is the solved case of the first order linear recursion $x_{n+1} = rx_n$. The generalised recursion you should carry is $x_{n+1} = rx_n + s$: its fixed point is $x^* = s/(1-r)$ when $r \neq 1$, the substitution $y_n = x_n - x^*$ makes it homogeneous,

and hence $x_n = (x_0 - x^*)r^n + x^*$. Every multivariable geometric sum in this problem is a product of solved first order recursions, one per extreme ray, and that is exactly why the cone had to be simplicial: one generator, one recursion. See Book 1, field: Sequences and recurrences, topic: first order linear recurrences.

Nonnegative rearrangement. For $f : I \to [0, \infty]$, $\sum_{i \in I} f(i) := \sup\{\sum_{i \in F} f(i) : F \subseteq I \text{ finite}\}$ is independent of any ordering, and iterated sums over a product index set agree with the double sum. This is Tonelli for counting measure and it is the licence used three times above. Contrast with conditionally convergent series, where reordering can produce any value. See Book 1, field: Real analysis, topic: Tonelli and unordered summation.

Lowest terms. A rational is in lowest terms when numerator and denominator are coprime; here 17 is prime and does not divide 21. Cheap to say, and the question explicitly asks for it. See Book 1, field: Number theory, topic: gcd and reduced fractions.

Circuit Bank

- **WHEN** a sum or an inequality whose hypothesis is "a, b, c are the sides of a triangle" **DO** substitute $x = b + c - a$, $y = c + a - b$, $z = a + b - c$ and record "positive, and over $\mathbb{Z}$ all of one parity" **BECAUSE** the triangle conditions are three linear inequalities and their slacks are coordinates in which the constraint set is an orthant, which is the only shape a separable sum can be summed over
- **WHEN** a separable summand over a coupled index set **DO** spend your effort exclusively on reindexing, not on the summand **BECAUSE** a separable function over a product set factorises instantly, so the entire difficulty is the geometry of the index set
- **WHEN** a modulus such as $|b - c|$ inside a constraint **DO** look for a change of variables that removes it rather than splitting into cases **BECAUSE** each case split doubles the ink and each half carries its own reindexing risk, and a modulus is usually the shadow of a badly chosen coordinate
- **WHEN** a substitution whose inverse involves division **DO** immediately compute the image lattice and state the congruence condition **BECAUSE** integrality is preserved forwards but not backwards, and the congruence classes are exactly the extra terms you must carry
- **WHEN** an index set that is a lattice cone **DO** find the extreme rays, compute the determinant of the generator matrix, and expect that many box terms **BECAUSE** the determinant is the index of the sublattice, and the number of box points is the number of geometric products you will have to add
- **WHEN** the phrase "as a rational number in lowest terms" **DO** finish with an explicit gcd sentence **BECAUSE** it costs one line and it is a marked instruction

Backtesting

1. First useful cue: the phrase "there exist triangles with side lengths a, b, c", paired with the observation that the summand is separable while the index set is not.
2. Method the expression suggested: transport the sum along the Ravi bijection into coordinates where the constraint becomes an orthant, then sum three independent geometric series.
3. Where you would hesitate, and the thought to reinforce: at the parity condition. The hesitation is "is the substitution onto?". Reinforce with the sentence "an inverse that divides by 2 always leaves a congruence behind", and check the image explicitly before summing anything.

4. The trap the question sets: two traps. The obvious one is the modulus $|b-c|$ that appears if you sum over a first, which costs a case split. The lethal one is forgetting parity, which gives a clean looking but wrong product with half integer exponents.
5. What makes you recognise it faster next time: the shape "separable summand, cone index set". Once seen, the reflex is to name the extreme rays and the sublattice index before writing a single series.
6. The reusable decision rule: when a constraint is a system of linear inequalities, make its slacks your coordinates; the sum then factorises up to a finite correction of size $|\det|$.

Mirror drills

1. For positive reals with $|xy|, |yz|, |zx| < 1$, prove $\sum_{(a,b,c)\in T} x^a y^b z^c = \dfrac{xyz(1+xyz)}{(1-xy)(1-yz)(1-zx)}$, and check that $(x,y,z) = (2, \frac{1}{3}, \frac{1}{5})$ returns $17/21$. *Answer:* the extreme rays $(0,1,1), (1,0,1), (1,1,0)$ carry weights yz, zx, xy and the single nontrivial box point $(1,1,1)$ carries weight xyz; sum the two translates and add.
2. Evaluate $\sum_{(a,b,c)\in T} 2^{-(a+b+c)}$. *Answer:* put $x = y = z = \frac{1}{2}$ in the drill above: $\dfrac{(1/8)(9/8)}{(3/4)^3} = \dfrac{9/64}{27/64} = \dfrac{1}{3}$.
3. Show that the number of ordered triples $(a,b,c) \in T$ with $a+b+c = 2n$ is $\binom{n-1}{2}$. *Answer:* $x+y+z = 2n$ is even, so the all odd class is impossible; write $x = 2u$, $y = 2v$, $z = 2w$ with $u, v, w \geq 1$ and $u+v+w = n$, giving $\binom{n-1}{2}$ compositions.

11.11 B5. Key permutations and a shift operator annihilator

Let P_n be the number of permutations π of $\{1, 2, \ldots, n\}$ such that

$$|i - j| = 1 \text{ implies } |\pi(i) - \pi(j)| \leq 2$$

for all i, j in $\{1, 2, \ldots, n\}$. Show that for $n \geq 2$, the quantity

$$P_{n+5} - P_{n+4} - P_{n+3} + P_n$$

does not depend on n, and find its value.

Read it out loud

Write the permutation as the word $\pi(1)\pi(2)\cdots\pi(n)$. Then $|i - j| = 1$ says "i and j are consecutive positions", so the condition says: consecutive letters of the word differ by at most 2. That is not a permutation condition at all; it is an edge condition. So P_n counts the words that list $1, \ldots, n$ walking along the graph G_n on the vertex set $\{1, \ldots, n\}$ in which i and j are joined exactly when $0 < |i - j| \leq 2$. In one sentence: P_n is the number of directed Hamiltonian paths in the square of a path. Say that out loud and the fields declare themselves: graph theory for the structure, combinatorics for the bijections, and the theory of linear recurrences with constant coefficients for the endgame.

Now read the thing you are asked to prove, because it is the single largest cue on the page. You are told that $P_{n+5} - P_{n+4} - P_{n+3} + P_n$ is constant. In shift operator language, with $(Ex)_n = x_{n+1}$, you are told that

$$(E^5 - E^4 - E^3 + 1)P = \text{constant},$$

and you should factorise that polynomial before you count a single permutation:

$$x^5 - x^4 - x^3 + 1 = (x - 1)(x + 1)(x^3 - x^2 - 1).$$

That is a map of the whole solution. It says: somewhere inside there is a sequence obeying $u_m = u_{m-1} + u_{m-3}$ up to an additive constant, the constant is handled by the factor $(x - 1)$, and there is a factor of 2 floating about (the root -1) that will turn out to be the reversal symmetry. You now know what you are hunting for. That is worth more than any amount of enumeration.

The first move to write down within thirty seconds is the degree observation: in G_n the vertex n has neighbours $n - 1$ and $n - 2$ only, and $n - 1$ has neighbours $n - 3, n - 2, n$ only. Extremes are rigid. In a graph whose edges are "close in value", the largest and smallest values have the fewest options, so they are where the structure is, and deleting the largest value is the natural induction.

Approaches, ranked worst to best

1. **Compute and guess.** *Cue:* "find its value". *Circuit:* tabulate, fit, assert. Enumerate $P_n = 1, 2, 6, 12, 20, 34, 56, 88, 136, 208$, notice the combination is 4, and write "hence the value is 4". On the Putnam that is one point at best. You have not shown independence of n; you have observed it. It ranks last because it is not an argument, and because it tempts you to stop.

2. **Subset transfer matrix.** *Cue:* a permutation built left to right. *Circuit:* dynamic programming over the set of used values. You track which values have been used and the last value. The state space has size $2^n \cdot n$; it computes P_n but proves nothing about a fifth order recurrence. It stalls because the state is unbounded, and an unbounded state can never produce a fixed order linear recurrence.

3. **Split by the position of the smallest values.** *Cue:* 1 has only two neighbours, 2 and 3. *Circuit:* condition on where 1 sits and peel. Count $P_n = 2P'_n$ where P'_n has 1 before 2, then split P'_n by whether the word starts with 1. This works, and yields $Q'_n = Q'_{n-1} + Q'_{n-3} + 1$ together with $R'_n = R'_{n-1} + Q'_{n-2}$. It ranks third because the second recursion needs a delicate case analysis of where the letter 4 can sit, and losing one case there loses the problem silently.
4. **Delete the largest letter and classify the top.** *Cue:* n has exactly two possible neighbours, so deletion cannot break the word. *Circuit:* peel the most constrained element and let the fibre of the peeling map define your state space. Deleting n always leaves a legal word, so the counting is governed by how many legal slots a word of length n offers the letter $n+1$; that number depends only on how n and $n-1$ sit relative to each other and to the two ends. Enumerating those configurations is forced, and there are exactly five. This is the ideal route: the state space is discovered rather than guessed.
5. **Generating function.** *Cue:* a linear recurrence with constant coefficients. *Circuit:* sum the state vector recursion into a rational function. Once you have the five state transfer, $\sum P_n x^n$ is rational with denominator $(1-x)(1-x^2)(1-x-x^3)$ and the assertion follows by clearing denominators. It is the same mathematics as the ideal route with more bookkeeping, so it is a good check and a poor first draft.

The ideal solution

Throughout, identify a permutation π with the word $\pi(1)\pi(2)\cdots\pi(n)$, and let $\mathcal{P}_n$ be the set of words listing $1,\dots,n$ in which consecutive letters differ by at most 2; so $P_n = |\mathcal{P}_n|$. Equivalently $\mathcal{P}_n$ is the set of directed Hamiltonian paths of the graph G_n on $\{1,\dots,n\}$ with $i \sim j$ iff $0 < |i-j| \le 2$. Reversal is an involution on $\mathcal{P}_n$ without fixed points for $n \ge 2$, so

$$P_n = 2\Sigma_n, \qquad \Sigma_n := \#\{\text{undirected Hamiltonian paths of } G_n\}.$$

Lemma 11.11.1 (Deletion)**.** *For $n \ge 3$, deleting the letter n from a word in $\mathcal{P}_n$ leaves a word in $\mathcal{P}_{n-1}$.*

Proof. The neighbours of n in G_n are $n-1$ and $n-2$ and nothing else. If n is at an end of the word, deleting it changes no consecutive pair. If n is interior, its two word neighbours lie in $\{n-1, n-2\}$, hence are $n-1$ and $n-2$ in some order; closing the gap creates the new consecutive pair $(n-1, n-2)$, whose difference is $1 \le 2$. All other pairs are untouched. □

Lemma 11.11.2 (Insertion slots)**.** *Let $n \ge 3$ and $\sigma \in \mathcal{P}_n$. Every word of $\mathcal{P}_{n+1}$ arises from exactly one $\sigma \in \mathcal{P}_n$ by inserting the letter $n+1$, and the legal slots in σ are: before the first letter if $\sigma(1) \in \{n-1, n\}$; after the last letter if $\sigma(n) \in \{n-1, n\}$; and between two adjacent letters that are n and $n-1$ in either order. Hence*

$$P_{n+1} = \sum_{\sigma \in \mathcal{P}_n} c(\sigma), \quad c(\sigma) = [\sigma(1) \in \{n-1, n\}] + [\sigma(n) \in \{n-1, n\}] + [\,n \text{ and } n-1 \text{ adjacent in } \sigma\,].$$

Proof. By the previous lemma deletion maps $\mathcal{P}_{n+1}$ to $\mathcal{P}_n$, and inserting $n+1$ back into the slot it came from inverts it, so the fibres are exactly the sets of legal slots. A slot is legal precisely when every new consecutive pair it creates is legal, and the only neighbours of $n+1$ are n and $n-1$; at an end that is one condition, in the interior it is two. □

So the entire problem reduces to a question with a finite answer: how can n and $n-1$ sit inside a word of $\mathcal{P}_n$? Since n has only two possible neighbours, the list is short.

Lemma 11.11.3 (Top configurations)**.** *Let $n \ge 3$ and let p be an undirected Hamiltonian path of G_n. Exactly one of the following holds, and in each case exactly one of the two words representing p has the stated shape.*

(U) *the word ends* $\dots, n-1, n$;

(V) *the word ends* $\dots, n-2, n, n-1$;

(W) *the word is* $n-1, \dots, n-2, n$, *that is, it ends* $\dots, n-2, n$ *and begins with* $n-1$;

(T) *the word ends* $\dots, n-3, n-1, n-2, n$;

(S) *both* n *and* $n-1$ *are interior and the word contains* $\dots, n-1, n, n-2, \dots$.

Writing U_n, V_n, W_n, T_n, S_n *for the number of paths of each type,* $\Sigma_n = U_n + V_n + W_n + T_n + S_n$.

Proof. The path neighbours of n lie in $\{n-1, n-2\}$.

If n is an endpoint of p it has one path neighbour. If that neighbour is $n-1$ we are in case (U). If it is $n-2$, consider $n-1$, whose path neighbours lie in $\{n-3, n-2, n\}$; the value n is unavailable, since n's only path neighbour is $n-2$. If $n-1$ is the other endpoint of p we are in case (W). Otherwise $n-1$ is interior with path neighbours exactly $n-3$ and $n-2$; then $n-2$ is interior with path neighbours $n-1$ and n, so the path terminates $\dots, n-3, n-1, n-2, n$, which is case (T).

If n is interior its two path neighbours are exactly $n-1$ and $n-2$. If $n-1$ is an endpoint of p, orient the word to finish at $n-1$ and it ends $\dots, n-2, n, n-1$, which is case (V). If $n-1$ is interior we are in case (S).

The five cases are mutually exclusive: (U),(W),(T) have n as an endpoint and (V),(S) have n interior; among the first three, (U) is separated by the neighbour of n and (W),(T) by whether $n-1$ is an endpoint; among the last two, by whether $n-1$ is an endpoint. Finally the orientation is forced in each case: in (U),(W),(T) the word must terminate at n; in (V) it must terminate at $n-1$; in (S) both $n, n-1$ are interior and they are adjacent, so exactly one of the two orientations reads $n-1$ then n. □

Lemma 11.11.4 (Transfer)**.** *For* $n \geq 3$,

$$U_{n+1} = U_n + W_n + T_n, \quad V_{n+1} = U_n, \quad W_{n+1} = W_n, \quad T_{n+1} = V_n, \quad S_{n+1} = S_n + V_n.$$

Proof. Take a word of $\mathcal{P}_{n+1}$ in each class and delete the letter $n+1$; the classification lemma applied at level n says what must remain, and the insertion lemma says the correspondence is a bijection.

(U) at level $n+1$: the word ends $\dots, n, n+1$. Deleting $n+1$ leaves a word of $\mathcal{P}_n$ ending in n, which by the classification is of type (U), (W) or (T); conversely appending $n+1$ to any word ending in n is legal and lands in (U). Hence $U_{n+1} = U_n + W_n + T_n$.

(V) at level $n+1$: the word ends $\dots, n-1, n+1, n$. Deleting leaves a word ending $\dots, n-1, n$, type (U); conversely insert $n+1$ between the final $n-1$ and n. Hence $V_{n+1} = U_n$.

(W) at level $n+1$: the word is $n, \dots, n-1, n+1$. Deleting leaves $\sigma \in \mathcal{P}_n$ beginning with n and ending with $n-1$. Since σ begins with n, its second letter is $n-1$ or $n-2$; it cannot be $n-1$, which is the last letter and $n \geq 3$. So σ begins $n, n-2$, hence its reverse ends $\dots, n-2, n$ and begins with $n-1$: type (W). Conversely reverse a (W) word and append $n+1$. Hence $W_{n+1} = W_n$.

(T) at level $n+1$: the word ends $\dots, n-2, n, n-1, n+1$. Deleting leaves a word ending $\dots, n-2, n, n-1$, type (V); conversely append $n+1$ to a (V) word, which is legal since $|(n-1) - (n+1)| = 2$. Hence $T_{n+1} = V_n$.

(S) at level $n+1$: both $n+1$ and n are interior and the word contains $\dots, n, n+1, n-1, \dots$. Deleting $n+1$ leaves a word containing $\dots, n, n-1, \dots$ with n interior. If $n-1$ is now an endpoint it must be the final letter, and the word is of type (V); if $n-1$ is interior, then reversing gives a word with $n-1, n$ adjacent in that order and both interior, that is type (S). The two sub cases are distinguished by whether $n-1$ is the final letter, so they are disjoint, and both insertions are legal and land in (S). Hence $S_{n+1} = S_n + V_n$. □

At $n = 3$ all six words are legal, and $U_3 = V_3 = W_3 = 1$, $T_3 = S_3 = 0$, so $\Sigma_3 = 3$ and $P_3 = 6$. In particular $W_3 = 1$, and $W_{n+1} = W_n$ gives

$$W_n = 1 \quad \text{for all } n \geq 3.$$

Adding the five transfer identities,

$$\Sigma_{n+1} = (U_n + W_n + T_n) + U_n + W_n + V_n + (S_n + V_n) = \Sigma_n + U_n + V_n + W_n,$$

so, writing $d_n := \Sigma_{n+1} - \Sigma_n$,

$$d_n = U_n + V_n + 1 \qquad (n \geq 3). \tag{1}$$

Next, collapse the five sequences to one. Set $U_1 := 0$ and $U_2 := 1$; then $V_n = U_{n-1}$ and $T_n = V_{n-1} = U_{n-2}$ hold for all $n \geq 3$ (for $n \geq 4$ they are the transfer identities, and at $n = 3$ they read $V_3 = U_2 = 1$ and $T_3 = U_1 = 0$, which is correct). Substituting into the first transfer identity,

$$U_{n+1} = U_n + W_n + T_n = U_n + U_{n-2} + 1 \qquad (n \geq 3),$$

and the case $n = 3$ reads $U_4 = U_3 + U_1 + 1 = 2$, which is correct. Thus

$$U_m = U_{m-1} + U_{m-3} + 1 \qquad (m \geq 4), \qquad U_1 = 0,\ U_2 = 1,\ U_3 = 1. \tag{2}$$

By (1) and $V_n = U_{n-1}$, $d_n = U_n + U_{n-1} + 1$ for $n \geq 3$. Now compute, for $n \geq 3$,

$$\begin{aligned} d_{n+4} - d_{n+2} - d_{n+1} - d_n &= U_{n+4} + U_{n+3} - U_{n+2} - 2U_{n+1} - 2U_n - U_{n-1} - 2 \\ &= 2U_{n+3} - U_{n+2} - U_{n+1} - 2U_n - U_{n-1} - 1 \\ &= U_{n+2} - U_{n+1} - U_{n-1} + 1 \\ &= 2, \end{aligned}$$

where the three simplifications used (2) in the forms $U_{n+4} = U_{n+3} + U_{n+1} + 1$, then $U_{n+3} = U_{n+2} + U_n + 1$, then $U_{n+2} = U_{n+1} + U_{n-1} + 1$. Since

$$\Sigma_{n+5} - \Sigma_{n+4} - \Sigma_{n+3} + \Sigma_n = d_{n+4} - (d_{n+2} + d_{n+1} + d_n),$$

we get $\Sigma_{n+5} - \Sigma_{n+4} - \Sigma_{n+3} + \Sigma_n = 2$, hence

$$P_{n+5} - P_{n+4} - P_{n+3} + P_n = 4 \qquad (n \geq 3).$$

It remains to treat $n = 2$ by hand. Both words of length 2 are legal, so $P_2 = 2$. From $\Sigma_3 = 3$ and the transfer identities one gets the table

n	3	4	5	6	7	8	9
U_n	1	2	4	6	9	14	21
V_n	1	1	2	4	6	9	14
W_n	1	1	1	1	1	1	1
T_n	0	1	1	2	4	6	9
S_n	0	1	2	4	8	14	23
Σ_n	3	6	10	17	28	44	68
P_n	6	12	20	34	56	88	136

so $P_7 - P_6 - P_5 + P_2 = 56 - 34 - 20 + 2 = 4$ as well. Therefore the quantity is independent of n for all $n \geq 2$ and its value is

$$\boxed{4}.$$

Layer 1: WHAT

Block 1. Rewrite the condition as "consecutive letters differ by at most 2", that is, as a Hamiltonian path count in G_n.

Block 2. Factorise the target operator $x^5 - x^4 - x^3 + 1 = (x-1)(x+1)(x^3 - x^2 - 1)$ and record the prediction: a hidden sequence with $u_m = u_{m-1} + u_{m-3} + \text{const}$, plus a factor of 2 from a symmetry.

Block 3. Use reversal to halve the count: $P_n = 2\Sigma_n$.

Block 4. Observe that n has only two neighbours, so deleting n never breaks a word. Deletion is therefore a map $\mathcal{P}_{n+1} \to \mathcal{P}_n$ whose fibres are the legal insertion slots.

Block 5. Read off which data the slot count depends on: the first letter, the last letter, and whether n and $n-1$ are adjacent. That is the state.

Block 6. Enumerate the possible top configurations. There are exactly five, and the enumeration is forced by the degree of n.

Block 7. Write the five transfer identities by deletion, note $W_n = 1$, and collapse to the single recursion $U_m = U_{m-1} + U_{m-3} + 1$.

Block 8. Do the annihilator algebra, get 2 for Σ and 4 for P, and dispose of $n = 2$ by direct computation.

Juxtaposition. The nominal route agrees with Block 1 and then goes straight to counting: tabulate P_n, spot 4, attempt to justify a fifth order recurrence by "the structure is clearly eventually periodic". The bifurcation point is the line immediately after the graph reformulation. The zombie asks "what is P_n?". The ideal asks "what does P_{n+1} know about P_n, and what is the smallest amount of information about a word that determines its contribution?". The second question has a bounded answer and therefore a proof; the first does not.

Layer 2: HOW

Block 1. **Cue.** $|i-j| = 1$ appears as a hypothesis on positions. **Why natural.** Positions differing by 1 are consecutive letters of the word, so the hypothesis is about adjacency, not about the permutation as a function. Once you translate "adjacent positions carry values within 2" into "the word is a walk on a graph whose edges are value pairs within 2", you gain endpoints, degrees, interiors and deletion, none of which exist in permutation language. **Link to the next block.** A graph reading immediately invites you to look at vertices of small degree.

Block 2. **Cue.** The problem hands you a specific linear combination and asserts it is constant. **Why natural.** A statement of the form "$L(P)$ is constant for a fixed linear combination L" is a statement that a particular polynomial in the shift operator annihilates P up to a constant. Factorising that polynomial costs one minute and tells you which recursions to hunt for. Here the cubic factor $x^3 - x^2 - 1$ is the fingerprint of $u_m = u_{m-1} + u_{m-3}$, the linear factor $x - 1$ absorbs an additive constant, and $x + 1$ is the parity artefact of a doubling. **Link to the next block.** Knowing you need a 2 and a cubic, look for the symmetry that supplies the 2.

Block 3. **Cue.** The defining condition is symmetric under reversing the word. **Why natural.** Reversal preserves the set of consecutive pairs, so it acts on $\mathcal{P}_n$; it has no fixed point for $n \geq 2$ because the first and last letters differ. A fixed point free involution partitions the set into pairs, so counting orbits halves the case analysis for the rest of the argument. This is where the factor $x + 1$ in the annihilator comes from. **Link to the next block.** With orbits as the objects, ask which orbit invariants control the induction.

Block 4. **Cue.** In G_n the vertex n has degree 2. **Why natural.** Deleting a vertex of a path normally destroys it, but if the deleted vertex's only possible neighbours are themselves adjacent, the path closes up. That is exactly the situation for n, whose neighbours $n-1$ and $n-2$ differ by 1. So deletion is total, which is the strongest possible form of an induction step: no exceptional cases. **Link to the next block.** If deletion is total, counting is fibre counting.

Block 5. **Cue.** The letter $n+1$ can only sit next to n and $n-1$. **Why natural.** The number of legal slots is therefore determined by three yes or no questions about σ: is the first letter in $\{n-1, n\}$, is

the last letter in $\{n-1, n\}$, are n and $n-1$ adjacent. A bounded question set is a bounded state space, and a bounded state space is exactly what produces a fixed order linear recurrence. This is the moment the fifth order recurrence in the question becomes believable. **Link to the next block.** Enumerate the states.

Block 6. **Cue.** n is either an endpoint or interior, and its neighbours are forced. **Why natural.** The case tree has depth two: first ask whether n is an endpoint, then ask where $n-1$ is. Because $n-1$ also has small degree, each branch collapses to a single shape. There is no creativity here at all; it is an exhaustive search over a two level tree, which is why five classes is the answer rather than some number you had to be clever to find. **Link to the next block.** Each class must now be pushed forwards.

Block 7. **Cue.** Every word at level $n+1$ has a unique parent at level n. **Why natural.** So prove the recurrences backwards, by deletion, not forwards by insertion. Deletion is deterministic and the class of the parent is forced by the shape of the child; insertion requires you to argue that you have found every child, which is where cases get lost. Writing the argument in the deletion direction is a proof craft decision, not a mathematical one, and it is worth several marks. **Link to the next block.** Five sequences, but four of them are shifts of one.

Block 8. **Cue.** $V_n = U_{n-1}$, $T_n = V_{n-1}$, $W_n = 1$. **Why natural.** Three of the five transfer identities are pure shifts, so the system has rank far below five. Substituting them into the identity for U produces $U_{n+1} = U_n + U_{n-2} + 1$, exactly the cubic predicted in Block 2. Then $\Sigma_{n+1} - \Sigma_n = U_n + U_{n-1} + 1$, and the requested combination is a telescoping of these first differences. The final algebra is three substitutions and it must land on a constant, because Block 2 said so. **Link to the next block.** Only the smallest index remains, and it is cheaper to check than to argue.

Layer 3: WHY

The mathematics beneath this route is the theory of linear recurrences read through the shift operator. Write E for the shift, $(Ex)_n = x_{n+1}$, and $\mathbf{1}$ for the constant sequence 1. Identity (2) says

$$(E^3 - E^2 - 1)U = \mathbf{1},$$

and identity (1) with $V_n = U_{n-1}$ says $E(E-1)\Sigma = (E+1)U + \mathbf{1}$. Apply $E^3 - E^2 - 1$ to the second and use the first, remembering that any polynomial $q(E)$ acts on $\mathbf{1}$ by multiplication by $q(1)$:

$$(E^3 - E^2 - 1)E(E-1)\Sigma = (E+1)\mathbf{1} + (1 - 1 - 1)\mathbf{1} = 2\mathbf{1} - \mathbf{1} = \mathbf{1}.$$

Cancelling the harmless E and multiplying by $(E+1)$,

$$(E-1)(E+1)(E^3 - E^2 - 1)\,\Sigma = (E+1)\mathbf{1} = 2\mathbf{1},$$

which is precisely $\Sigma_{n+5} - \Sigma_{n+4} - \Sigma_{n+3} + \Sigma_n = 2$, hence 4 for P. So the examiners' combination is not arbitrary at all: it is the product of the three natural annihilators, one for the intrinsic growth, one for the constant inhomogeneity, one for the reversal doubling.

Two consequences worth carrying away. First, the sharper statement is one order lower: cancelling the $(E+1)$ gives

$$P_{n+4} - 2P_{n+3} + P_{n+2} - P_{n+1} + P_n = 2,$$

which you can check on the table above at $n = 2$: $34 - 40 + 12 - 6 + 2 = 2$. The examiners preferred the fifth order form because all its coefficients are ± 1 and it looks like a competition statement; the fourth order form is the mathematically primitive one. Second, the growth rate. Substituting $U_m = v_m - 1$ into (2) gives the homogeneous $v_m = v_{m-1} + v_{m-3}$, so P_n grows like ψ^n where $\psi \approx 1.46557$

is the real root of $x^3 = x^2 + 1$, the supergolden ratio. The ratios 208/136, 314/208, 470/314, ... visibly settle towards it. That substitution is the general method: an inhomogeneous constant is removed by shifting to the constant particular solution, which exists precisely because 1 is not a root of $x^3 - x^2 - 1$.

The discernment mechanism has three parts, and each is trainable. (i) Read the target as an operator and factorise it first. A competition problem that hands you a specific linear combination is telling you the characteristic polynomial; refusing to look at it is refusing a free map. (ii) In any structure built on $\{1, \ldots, n\}$ with a locality constraint, peel the extreme element, because extremes have the smallest neighbourhoods and therefore the most rigid behaviour. Ask "does deleting it always leave a legal object?"; here the answer is yes precisely because $n-1$ and $n-2$ are themselves adjacent. (iii) Let the state space be discovered by the peeling rule rather than guessed. Write down the fibre size $c(\sigma)$ of the deletion map, see which features of σ it mentions, and define your classes to be exactly the equivalence classes of those features. You will never enumerate five classes by intuition; you will always enumerate them by asking what the insertion rule needs to know.

The nudges the examiners planted: the bound 2 rather than 1 (with 1 the answer is trivially 2 permutations for every n, so 2 is the smallest interesting threshold); the phrase "does not depend on n", which is a licence to prove a recurrence rather than a formula; and the range $n \geq 2$, which is a warning that the general argument will start at $n = 3$ and that the boundary case must be checked separately. These permutations are known in the literature as key permutations and the counting sequence is A003274 in the OEIS, but nothing in the solution needs that.

Revolving concepts

Hamiltonian paths in the square of a path. For a graph H, the square H^2 joins vertices at distance at most 2. Here $G_n = P_n^2$ where P_n is the path $1, 2, \ldots, n$. Degrees: the vertices 1 and n have degree 2, the vertices 2 and $n-1$ have degree 3, all others degree 4. The cut between $\{1, \ldots, k\}$ and $\{k+1, \ldots, n\}$ has only three edges, which is the structural reason a Hamiltonian path can cross it at most three times and hence why the count is so rigid. See Book 1, field: Discrete structures and graphs, topic: Hamiltonian paths and powers of paths.

Deletion and insertion bijections. The technique: find an element whose removal always yields a valid smaller object, then count the fibre of the removal map. The fibre size is the number of legal reinsertions, and the classes on which it is constant form the state space. This is the combinatorial version of a transfer matrix, and it is how almost every "prove the count satisfies a linear recurrence" problem is meant to be attacked. See Book 1, field: Combinatorics, topic: deletion and insertion arguments.

Fixed point free involutions. If a finite set carries an involution with no fixed points, its cardinality is twice the number of orbits. Reversal on $\mathcal{P}_n$ is such an involution for $n \geq 2$. Recognising the involution before the case analysis, not after, is what keeps the classification to five cases rather than ten. See Book 1, field: Combinatorics, topic: involutions and orbit counting.

Linear recurrences and the shift operator. With $(Ex)_n = x_{n+1}$, a linear recurrence with constant coefficients is a statement $q(E)x = 0$ or $q(E)x = f$. Key facts: $q(E)$ acts on the constant sequence $\mathbf{1}$ by multiplication by $q(1)$; polynomials in E commute, so annihilators may be applied in any order; the general homogeneous solution is spanned by $n^j\lambda^n$ over the roots λ of q with multiplicity; and an inhomogeneous constant $f = c\mathbf{1}$ has the constant particular solution $c/q(1)$ provided $q(1) \neq 0$.

Generalised recursion. The recursion appearing here is $u_m = u_{m-1} + u_{m-3} + 1$ with $u_1 = 0, u_2 = 1, u_3 = 1$. Substituting $u_m = v_m - 1$ gives $v_m = v_{m-1} + v_{m-3}$, the Narayana recursion, whose characteristic equation $x^3 = x^2 + 1$ has one real root $\psi \approx 1.46557$ (the supergolden ratio) and a conjugate pair of modulus less than 1; hence $u_m = C\psi^m + O(1)$. The same substitution solves every $u_m = \sum_i a_i u_{m-i} + c$ with $\sum a_i \neq 1$, and when $\sum a_i = 1$ the particular solution is linear in m instead, because 1 is then a root and the annihilator acquires a repeated factor $(E-1)^2$. See Book 1, field: Sequences and recurrences, topic: constant coefficient recurrences and annihilators.

Rational generating functions. A sequence satisfies a constant coefficient linear recurrence iff its generating function is rational with denominator the reversed characteristic polynomial. Here $\sum_{n\geq 2} P_n x^n$ has denominator $(1-x)(1-x^2)(1-x-x^3)$, and the three factors are the same three annihilators seen once more. See Book 1, field: Generating functions, topic: rational generating functions.

Circuit Bank

- **WHEN** a permutation condition that only constrains neighbouring positions **DO** redraw it as a Hamiltonian path count in a graph on the values **BECAUSE** "$|i-j|=1$" means "consecutive in the word", so the condition is an edge condition, and graph language hands you endpoints, degrees and deletion for free
- **WHEN** a problem that hands you a specific linear combination and asserts it is constant **DO** factorise the corresponding characteristic polynomial before counting anything **BECAUSE** the factors are a map of the solution: each one names a recursion, a symmetry or an inhomogeneity you must find
- **WHEN** a counting problem on $\{1,\ldots,n\}$ with a locality constraint **DO** peel the largest or smallest element and check whether deletion is total **BECAUSE** extremes have the smallest neighbourhoods, and deletion is total exactly when the possible neighbours of the peeled element are themselves mutually adjacent
- **WHEN** a deletion map whose fibre size you can write down **DO** define your state space to be the features of the object that the fibre size mentions, and nothing else **BECAUSE** a bounded feature list is a bounded state space, which is precisely what forces a fixed order linear recurrence
- **WHEN** a recurrence to prove between two levels **DO** argue in the deletion direction, from level $n+1$ down to level n **BECAUSE** deletion is deterministic and forces the parent's shape, whereas insertion obliges you to prove you have found every child, which is where cases go missing
- **WHEN** a linear recurrence with a constant inhomogeneous term and $q(1)\neq 0$ **DO** shift by the constant particular solution $c/q(1)$ **BECAUSE** the shifted sequence is homogeneous, so the characteristic polynomial applies directly and the growth rate becomes visible
- **WHEN** a general argument that only starts at $n=3$ but a claim asserted from $n=2$ **DO** check the boundary index by direct computation and say so **BECAUSE** a boundary case is cheaper to verify than to force the general machinery to cover, and leaving it unmentioned is a marked omission

Backtesting

1. First useful cue: "$|i-j|=1$ implies $|\pi(i)-\pi(j)|\leq 2$" read as an adjacency condition, together with the observation that the vertex n has only two neighbours.
2. Method the expression suggested: factorise $x^5-x^4-x^3+1=(x-1)(x+1)(x^3-x^2-1)$, then hunt for a state recursion of the shape $u_m = u_{m-1}+u_{m-3}+\text{const}$ and a symmetry supplying the factor 2.
3. Where you would hesitate, and the thought to reinforce: at the classification of top configurations, where it feels as if you might have missed a case. Reinforce with "the case tree is forced: is n an endpoint, then where is $n-1$", and note that n has degree 2 so there is nothing else to ask.

4. The trap the question sets: computing the numbers, seeing 4, and writing an answer that never proves independence of n. A second trap is starting the induction at $n = 2$ and quietly using identities that only hold from $n = 3$.
5. What makes you recognise it faster next time: the pairing of a locality constraint with a requested linear combination. That pairing always means transfer matrix, and the requested combination is the characteristic polynomial handed to you in advance.
6. The reusable decision rule: let the peeling map define the state space. Write the fibre size, list the features it depends on, enumerate those features exhaustively, and transfer.

Mirror drills

1. Let Q_n count the permutations with $|i - j| = 1 \Rightarrow |\pi(i) - \pi(j)| \leq 1$. Find Q_n. *Answer:* G_n becomes the path itself, whose only Hamiltonian paths are the two monotone orders, so $Q_n = 2$ for $n \geq 2$. This shows why the bound 2 is the smallest interesting threshold.
2. Let E_n be the number of words in $\mathcal{P}_n$ beginning with the letter 1. Show $E_n = E_{n-1} + E_{n-3} + 1$ for $n \geq 4$, with $E_1 = E_2 = 1$ and $E_3 = 2$. *Answer:* the second letter is 2 or 3. If 2, delete the leading 1 and shift down. If 3 and the third letter is 2, the fourth is forced to be 4, so delete $1, 3, 2$ and shift by 3. Otherwise 2 must be the final letter preceded by 4, and the remainder is the unique Hamiltonian path of G on $\{4, \ldots, n\}$ from 5 to 4, giving the single word $1, 3, 5, \ldots, 6, 4, 2$.
3. Prove that in G_m there is exactly one Hamiltonian path from 2 to 1, and deduce $\lim_{n\to\infty} P_{n+1}/P_n$. *Answer:* for $m \geq 3$ the vertex 1 must be entered from 3 and the vertex 2 must leave to 4; deleting 1 and 2 leaves the same problem on $\{3, \ldots, m\}$, so induct. The limit is the real root $\psi \approx 1.46557$ of $x^3 = x^2 + 1$.

11.12 B6. Alternating divisor sum hiding an arctangent square

For each positive integer k, let $A(k)$ be the number of odd divisors of k in the interval $[1, \sqrt{2k})$. Evaluate

$$\sum_{k=1}^{\infty} (-1)^{k-1} \frac{A(k)}{k}.$$

Read it out loud

An alternating series in k whose coefficient counts divisors of k, and the divisor condition is measured against $\sqrt{2k}$, that is, against k itself. Read that sentence out loud and the first instruction is already sitting there: the square root is not real. It is an artefact of describing a divisor pair by only one of its two members. Name the cofactor and the square root evaporates. So write $k = dm$ with d the odd divisor. Then $d < \sqrt{2k}$ is equivalent to $d^2 < 2k = 2dm$, that is, to $d < 2m$. Now substitute $d = 2\ell - 1$, which you do because d is odd and because odd numbers indexed by ℓ are how the arctangent series is indexed, and the condition becomes $2\ell - 1 < 2m$, that is, since both sides are integers and the left side is odd, simply $\ell \leq m$. One line of algebra has turned a divisor condition with an irrational bound into an ordering condition on a lattice.

So $A(k)$ is the number of pairs (ℓ, m) of positive integers with $\ell \leq m$ and $(2\ell - 1)m = k$. There is a pretty combinatorial reading of that which is worth carrying: a run of $2\ell - 1$ consecutive positive integers centred at m sums to $(2\ell - 1)m$, and it consists of positive integers exactly when its first term $m - \ell + 1$ is at least 1, that is, exactly when $\ell \leq m$. Thus $A(k)$ is the number of ways to write k as a sum of an odd number of consecutive positive integers. You do not need that reading to solve the problem, but it is the reason the count is natural rather than invented, and a count that has a reason is a count you will trust under exam pressure.

Now the fields. Number theory is present for exactly three lines, the three lines above. Everything after that is real analysis, and the real analysis is the entire question. Two warnings should fire immediately. First, $A(k) \geq 1$ for every k because $d = 1$ always qualifies, so $\sum_k A(k)/k$ diverges: the series is not absolutely convergent, and any rearrangement you perform is a theorem you owe the grader, not a notational convenience. Second, $A(k)$ is violently non monotone: it is 1 at $k = 2$ and 3 at $k = 45$, so $A(k)/k$ does not decrease to zero and the alternating series test does not apply to the series as written. Neither of the two convergence tools you own applies to the object you are handed. That is what B6 is testing. The number theory is the costume; the examination is on whether you notice that even the convergence must be established by hand, and then establish it.

Within thirty seconds you should have written three things: $k = (2\ell - 1)m$ with $\ell \leq m$; $(-1)^{k-1} = (-1)^{m-1}$ because $2\ell - 1$ is odd so $k \equiv m \pmod 2$; and the question of which of the two indices goes on the outside of the double sum. That third item is the whole strategic decision of the problem, and the sign line answers it for you: the sign depends only on m, so m must stay on the inside, where its alternation can be exploited.

Approaches, ranked worst to best

1. **Compute partial sums and recognise the constant.** *Cue:* the command term "Evaluate" promises a closed form. *Circuit:* numerical pattern matching. You will find the partial sums approach $0.6168\ldots$ and $\pi^2/16 = 0.61685\ldots$ is recognisable. It ranks last for two reasons. The convergence is of order $N^{-1/2}$, so four correct digits need about a million terms, which rules it out in the arena. More importantly it earns nothing: the entire content of B6 is the justification, and a number with no proof of convergence attached is worth zero.

2. **Dirichlet series and analytic continuation.** *Cue:* a divisor counting coefficient divided by k smells like $\sum a_k k^{-s}$ at $s = 1$. *Circuit:* write the generating Dirichlet series, factor it, evaluate. You would form $\sum_k A(k)k^{-s} = \sum_{\ell \leq m}((2\ell - 1)m)^{-s}$ and hope for a product of zeta and beta

functions. It does not factor, because the constraint $\ell \le m$ couples the two indices, and the alternating sign forces eta versions on top. It stalls precisely where you started: the value you want sits at the edge of the region of convergence, so you have imported the convergence problem into a heavier machine.

3. **Sum over m on the outside, producing odd harmonic numbers.** *Cue:* the inner sum over ℓ runs over the finite range $1 \le \ell \le m$, which looks safer than an infinite inner sum. *Circuit:* group the lattice by the larger coordinate. This gives $\sum_{m\ge1} \frac{(-1)^{m-1}}{m} O_m$ with $O_m = 1 + \frac{1}{3} + \cdots + \frac{1}{2m-1}$, and then $O_m = \int_0^1 \frac{1-u^{2m}}{1-u^2}\,du$ leads to $\int_0^1 \frac{\log(2/(1+u^2))}{1-u^2}\,du$. That integral is genuinely equal to $\pi^2/16$, so the route is not wrong. It ranks third because the integrand needs care at $u = 1$, the antiderivative is far harder than the one you get from the other ordering, and the regrouping needs exactly the same convergence argument as the winning route while paying more ink for the endgame. Choosing the finite inner sum feels safe and is the classic false economy.

4. **Sum over ℓ on the outside, so the inner sum is an alternating harmonic tail.** *Cue:* the sign $(-1)^{k-1}$ depends only on m, so m belongs where alternation can be used. *Circuit:* put the sign bearing index innermost, then bound its tail by its first term. The inner object becomes $T_\ell = \sum_{m\ge\ell} \frac{(-1)^{m-1}}{m}$, which satisfies $|T_\ell| \le 1/\ell$ by the alternating tail bound. That single inequality makes the outer sum absolutely convergent, converts the rearrangement from a leap of faith into a two line estimate, and hands you arctan at the end. This is the ideal route, and its selection is forced by one observation about where the minus sign lives.

The ideal solution

The answer is $\boxed{\frac{\pi^2}{16}}$.

Lemma 11.12.1. *For every positive integer k,*

$$A(k) = \#\{(\ell, m) \in \mathbb{Z}^2_{\ge1} : \ m \ge \ell,\ (2\ell-1)m = k\},$$

and for every such pair, $(-1)^{k-1} = (-1)^{m-1}$.

Proof. The odd divisors of k are exactly the integers $2\ell-1$ with $\ell \ge 1$ and $(2\ell-1) \mid k$; given such a divisor put $m = k/(2\ell-1)$. Since $2k = 2(2\ell-1)m$ is positive, the condition $2\ell - 1 < \sqrt{2k}$ is equivalent to $(2\ell-1)^2 < 2(2\ell-1)m$, hence to $2\ell - 1 < 2m$. As $2\ell-1$ and $2m$ are integers of opposite parity, this is equivalent to $2\ell \le 2m$, that is, to $\ell \le m$. The lower bound $2\ell - 1 \ge 1$ is automatic. Finally $2\ell - 1$ is odd, so $k = (2\ell-1)m \equiv m \pmod 2$. □

For $\ell \ge 1$ set

$$T_\ell = \sum_{m=\ell}^{\infty} \frac{(-1)^{m-1}}{m},$$

which converges by the alternating series test, and whose tail bound gives $|T_\ell| \le 1/\ell$. Define

$$S = \sum_{\ell=1}^{\infty} \frac{T_\ell}{2\ell-1}.$$

Since $2\ell - 1 \ge \ell$ we have $|T_\ell/(2\ell-1)| \le 1/\ell^2$, so S converges absolutely.

Claim 11.12.2. *The series $\sum_{k\ge1}(-1)^{k-1}A(k)/k$ converges, and its sum is S.*

Proof. Write $S_N = \sum_{k=1}^{N} (-1)^{k-1} A(k)/k$. By the Lemma, S_N is a finite sum over the lattice points (ℓ, m) with $\ell \le m$ and $(2\ell-1)m \le N$, the contribution of (ℓ, m) being $\frac{(-1)^{m-1}}{(2\ell-1)m}$. Such a point requires $\ell(2\ell-1) \le (2\ell-1)m \le N$, so only finitely many ℓ occur; let L be the largest ℓ with $\ell(2\ell-1) \le N$. If $\ell > \sqrt{N}$ then $\ell(2\ell-1) > \sqrt{N}(2\sqrt{N}-1) = 2N - \sqrt{N} \ge N$ for $N \ge 1$, so $L \le \sqrt{N}$; and $\ell(2\ell-1) \le 2\ell^2$ shows $L \ge \lfloor \sqrt{N/2} \rfloor$, so $L \to \infty$. Grouping the finite sum by ℓ,

$$S_N = \sum_{\ell=1}^{L} \frac{1}{2\ell-1} \sum_{m=\ell}^{M_\ell} \frac{(-1)^{m-1}}{m}, \qquad M_\ell = \left\lfloor \frac{N}{2\ell-1} \right\rfloor.$$

Therefore

$$S - S_N = \sum_{\ell > L} \frac{T_\ell}{2\ell-1} + \sum_{\ell=1}^{L} \frac{1}{2\ell-1} \sum_{m > M_\ell} \frac{(-1)^{m-1}}{m}.$$

For the first sum, $|T_\ell/(2\ell-1)| \le 1/\ell^2$ gives $\left|\sum_{\ell>L} T_\ell/(2\ell-1)\right| \le \sum_{\ell>L} \ell^{-2} \le 1/L$. For the second, each inner sum is an alternating tail, so it is bounded in absolute value by its first term $\frac{1}{M_\ell+1}$, and $M_\ell + 1 > N/(2\ell-1)$ gives $\frac{1}{M_\ell+1} < \frac{2\ell-1}{N}$. Hence the second sum is bounded in absolute value by

$$\sum_{\ell=1}^{L} \frac{1}{2\ell-1} \cdot \frac{2\ell-1}{N} = \frac{L}{N} \le \frac{1}{\sqrt{N}}.$$

Thus $|S - S_N| \le \frac{1}{L} + \frac{1}{\sqrt{N}} \to 0$, which proves both the convergence and the value. □

Next, an exact formula for T_ℓ. For $N \ge \ell$, summing the finite geometric progression $\sum_{m=\ell}^{N} (-1)^{m-1} t^{m-1} = \frac{(-1)^{\ell-1} t^{\ell-1} - (-1)^N t^N}{1+t}$ and integrating over $[0,1]$ termwise (a finite sum, so no justification is needed) gives

$$\sum_{m=\ell}^{N} \frac{(-1)^{m-1}}{m} = (-1)^{\ell-1} \int_0^1 \frac{t^{\ell-1}}{1+t}\, dt - (-1)^N \int_0^1 \frac{t^N}{1+t}\, dt.$$

The last integral is at most $\int_0^1 t^N\, dt = \frac{1}{N+1} \to 0$, so letting $N \to \infty$,

$$T_\ell = (-1)^{\ell-1} \int_0^1 \frac{t^{\ell-1}}{1+t}\, dt.$$

Now put $f_\ell(t) = \dfrac{(-1)^{\ell-1} t^{\ell-1}}{(2\ell-1)(1+t)}$, continuous on $[0,1]$, so that $S = \sum_{\ell \ge 1} \int_0^1 f_\ell(t)\, dt$. For each fixed $t \in [0,1]$ the moduli $\frac{t^{\ell-1}}{(2\ell-1)(1+t)}$ decrease to 0 as ℓ grows, so the alternating tail bound applies pointwise and uniformly:

$$\left| \sum_{\ell=L+1}^{\infty} f_\ell(t) \right| \le \frac{t^L}{(2L+1)(1+t)} \le \frac{1}{2L+1} \qquad \text{for all } t \in [0,1].$$

Hence $\sum_\ell f_\ell$ converges uniformly on $[0,1]$ to a continuous function F, and a uniformly convergent series of continuous functions on a compact interval may be integrated term by term. Therefore

$$S = \int_0^1 F(t)\, dt, \qquad F(t) = \frac{1}{1+t} \sum_{\ell=1}^{\infty} \frac{(-1)^{\ell-1} t^{\ell-1}}{2\ell-1}.$$

For $0 < t < 1$, writing $u = \sqrt{t} \in (0,1)$ and using the Gregory series $\arctan u = \sum_{\ell \ge 1} \frac{(-1)^{\ell-1} u^{2\ell-1}}{2\ell-1}$, valid for $|u| < 1$,

$$\sum_{\ell=1}^{\infty} \frac{(-1)^{\ell-1} t^{\ell-1}}{2\ell-1} = \frac{1}{u} \sum_{\ell=1}^{\infty} \frac{(-1)^{\ell-1} u^{2\ell-1}}{2\ell-1} = \frac{\arctan \sqrt{t}}{\sqrt{t}}.$$

So $F(t) = \frac{\arctan\sqrt{t}}{\sqrt{t}\,(1+t)}$ on $(0,1)$. Both sides are continuous on $[0,1]$, the right side extending by the value 1 at $t = 0$; since F is continuous as a uniform limit of continuous functions, the identity holds on all of $[0,1]$, and no appeal to Abel's theorem at the endpoint is required. Substituting $t = u^2$, $dt = 2u\,du$,

$$S = \int_0^1 \frac{\arctan\sqrt{t}}{\sqrt{t}\,(1+t)}\,dt = \int_0^1 \frac{\arctan u}{u(1+u^2)} \cdot 2u\,du = \int_0^1 \frac{2\arctan u}{1+u^2}\,du.$$

The integrand is $2gg'$ with $g = \arctan$, so it is the derivative of g^2, and

$$S = \Big[\arctan^2 u\Big]_0^1 = \Big(\frac{\pi}{4}\Big)^2 = \frac{\pi^2}{16}.$$

Layer 1: WHAT

Block 1. Name the cofactor: $k = dm$, so the condition $d < \sqrt{2k}$ becomes $d < 2m$; with $d = 2\ell - 1$ it becomes $\ell \le m$.

Block 2. Read the parity: $2\ell - 1$ is odd, so $(-1)^{k-1} = (-1)^{m-1}$. The sign lives on m alone.

Block 3. Choose the order of summation from Block 2: ℓ outside, m inside, so the inner sum is a tail of the alternating harmonic series.

Block 4. Bound that tail: $|T_\ell| \le 1/\ell$, whence the outer sum $S = \sum_\ell T_\ell/(2\ell - 1)$ converges absolutely.

Block 5. Prove that the original partial sums converge to S, by bounding $S - S_N$ in two pieces: the tail in ℓ and the truncation error in m.

Block 6. Replace $1/m$ by $\int_0^1 t^{m-1}\,dt$ so that the alternating tail becomes a geometric series inside an integral, giving $T_\ell = (-1)^{\ell-1}\int_0^1 t^{\ell-1}(1+t)^{-1}dt$.

Block 7. Exchange the outer sum with the integral, justified by uniform convergence from the same alternating tail bound used in Block 4.

Block 8. Recognise the resulting power series as $\arctan\sqrt{t}/\sqrt{t}$, substitute $u = \sqrt{t}$ so the Jacobian cancels the $1/u$, and read off $2\arctan u/(1+u^2)$ as the derivative of $\arctan^2 u$.

Juxtapose that with the nominal route. The zombie writes $A(k)$, tries to compute it for small k, notices $1, 1, 1, 2, 1, 2, 1, 2, 2, 2, \ldots$, searches for a multiplicative structure that is not there, gives up on structure, expands the sum as $\sum_{d \text{ odd}} \frac{1}{d}\sum_m \frac{(-1)^{dm-1}}{m}$ with no attention to whether that rearrangement is legal or to the range of m, and either gets a divergent object or asserts the interchange and loses most of the marks. The bifurcation point is one line long, and it is this: after writing $k = dm$, do you square the inequality $d < \sqrt{2k}$? Squaring is what removes k from the condition and replaces it by a comparison between d and m. Every good thing that follows, the ordering $\ell \le m$, the choice of summation order, the tail bound, is downstream of that square. The zombie keeps $\sqrt{2k}$ because it is what the question wrote, and therefore never gets a lattice.

Layer 2: HOW

Block 1. Cue. The divisor bound $\sqrt{2k}$ mentions k, and k is the thing being divided. A condition that mentions both a divisor and its dividend is a condition on the pair, not on the divisor. **Why natural.** A divisor is never alone: $d \mid k$ means $k = dm$, and the honest coordinates on the set you are counting are (d, m), not d and k. In those coordinates the bound $d < \sqrt{2dm}$ is homogeneous of degree one in the pair after squaring, so squaring cancels one factor of d and leaves $d < 2m$. Squaring is legitimate because both sides are positive. This is the same reflex that turns "divisors below $\sqrt{n}$" into "the smaller member of each divisor pair". **Link to the next block.** The condition $d < 2m$ with d odd is a statement about integers of opposite parity, so it tightens; parametrising the odd d as $2\ell - 1$ makes the tightening visible and hands you the index ℓ that the arctangent series wants.

Block 2. **Cue.** There is a factor $(-1)^{k-1}$ and you have just written k as a product. **Why natural.** The sign of a product is decided by the parity of its factors, and one of the two factors is odd by construction. So the sign is carried entirely by m. Always compute what a given sign depends on before deciding anything else about a double sum; it is one line and it determines the architecture. **Link to the next block.** Knowing that the sign depends on m alone tells you which index must remain inside the summation.

Block 3. **Cue.** An index whose sign alternates is an index you want adjacent to a summation sign, because alternation is the only free convergence you get in a non absolutely convergent problem. **Why natural.** If m is on the outside, the inner sum over ℓ is a partial odd harmonic sum, which grows like $\frac{1}{2}\log m$ and carries no cancellation; you have thrown away the alternation and bought a logarithm. If m is on the inside, the inner sum is $\sum_{m\geq\ell}(-1)^{m-1}/m$, an honest alternating tail with a bound you already own. The choice is not aesthetic; it decides whether the outer sum is absolutely convergent. **Link to the next block.** With m inside you now need one quantitative fact about alternating tails.

Block 4. **Cue.** An alternating series whose terms decrease in modulus to zero. **Why natural.** The alternating series estimate says the tail is bounded in modulus by its first omitted term, here $1/\ell$. Combined with $2\ell-1\geq\ell$ this gives a summand of modulus at most $1/\ell^2$, so the rearranged series is absolutely convergent. This is the pivot of the whole solution: it converts a conditionally convergent mess into an absolutely convergent series of well understood objects. **Link to the next block.** Absolute convergence of the rearranged sum is not the same as the original series converging to it. That gap must now be closed, explicitly.

Block 5. **Cue.** The original series is not absolutely convergent, so Riemann's rearrangement theorem is in force and the rearrangement is a claim, not a step. **Why natural.** The correct discipline is to compare partial sums, never to invoke a theorem you cannot name. The partial sum S_N is a finite sum over lattice points under the hyperbola $(2\ell-1)m\leq N$; the difference $S-S_N$ splits into the part of S with ℓ too large and the alternating tails in m that S_N has truncated. Both are bounded by the two facts already in hand: $|T_\ell|\leq 1/\ell$ for the first, and the alternating tail estimate for the second, where the truncation point M_ℓ is at least $N/(2\ell-1)$, so the $2\ell-1$ in the denominator of the outer sum is exactly cancelled. That cancellation is why the estimate is clean: the outer weight $1/(2\ell-1)$ and the inner error $(2\ell-1)/N$ multiply to $1/N$, and there are at most $\sqrt{N}$ terms. **Link to the next block.** With convergence secured, the problem is now the purely analytic evaluation of an absolutely convergent series.

Block 6. **Cue.** You need a closed form for $\sum_{m\geq\ell}(-1)^{m-1}/m$, and $1/m$ is the obstruction to summing a geometric series. **Why natural.** $1/m=\int_0^1 t^{m-1}\,dt$ is the standard device for removing a linear denominator: it turns $\sum a^m/m$ into $\int\sum a^m t^{m-1}$, which is geometric. Doing it on the finite partial sum first, and only then letting $N\to\infty$, means the only limit you must justify is $\int_0^1 t^N(1+t)^{-1}dt\to 0$, which is immediate. That ordering is deliberate: it avoids needing the monotone or dominated convergence theorem at this step. **Link to the next block.** You now have S as a sum of integrals and want it as an integral of a sum.

Block 7. **Cue.** A series of continuous functions on a compact interval, alternating in the summation index with moduli decreasing in that index uniformly in t. **Why natural.** The alternating tail bound is uniform in t because $t^{\ell-1}\leq 1$ and $(1+t)^{-1}\leq 1$ on $[0,1]$: the tail after L terms is at most $1/(2L+1)$ for every t. Uniform convergence of continuous functions on $[0,1]$ licenses termwise integration. The Weierstrass M test would fail here since $\sum 1/(2\ell-1)$ diverges, so it is the alternation, not the size, that buys uniformity. Notice you used the very same inequality twice, at Block 4 and here; that

is the sign of a well chosen decomposition. **Link to the next block.** The bracket you are left with is a power series in t with odd denominators, which is one substitution away from a function you know.

Block 8. Cue. Odd denominators $2\ell - 1$ with alternating signs and a variable that wants to be a square. **Why natural.** $\sum_{\ell \geq 1} (-1)^{\ell-1} x^{2\ell-1}/(2\ell-1) = \arctan x$ is the defining series of the arctangent, and the powers of t appearing are $t^{\ell-1} = x^{2\ell-2}$, one power short, which is exactly why the answer carries a $1/\sqrt{t}$. Substituting $t = u^2$ produces $dt = 2u\,du$, which cancels that $1/u$ and simultaneously turns $1/(1+t)$ into $1/(1+u^2) = \arctan'(u)$. Both odd denominator structures in the problem, the one from the odd divisor and the one from the alternating harmonic tail, are the same function; that is why the answer is a perfect square. **Link to the next block.** Nothing is left: $2\arctan u \cdot \arctan'(u)$ integrates to $\arctan^2 u$, so the value is $\arctan(1)^2 = \pi^2/16$.

Layer 3: WHY

Three mathematical truths hold this route up. The first is that a sum weighted by a divisor counting function is always a sum over a lattice region, and you should convert it before doing anything analytic. The identity $\sum_k \big(\sum_{d|k} c_{d,k/d}\big) x_k = \sum_{d,m} c_{d,m} x_{dm}$ is trivial as bookkeeping but decisive as strategy, because every constraint written in terms of k becomes a constraint relating d and m, and constraints relating two free variables define regions you can draw. Here the region is $\{\ell \leq m\}$, a half plane, which is the simplest region there is. If the constraint had been $d < \sqrt{3k}$ you would have got $2\ell - 1 < 3m$ and the same method with a slightly uglier region; the technique does not depend on the luck of the constant.

The second is that conditional convergence is a property of an ordering, not of a set of numbers. The family $\{(-1)^{m-1}/((2\ell-1)m)\}$ over $\ell \leq m$ is not summable in the unordered sense, so "the sum" means nothing until you fix an order, and the problem fixes it: increasing k. Every other order you write down is a different object until proved otherwise. The discipline that follows is mechanical and should be automatic: when you rearrange a non absolutely convergent series, write $S - S_N$ explicitly and bound it. Do not gesture at Fubini. In this problem the bound falls out because the outer weight $1/(2\ell-1)$ and the inner truncation error $(2\ell-1)/N$ are reciprocal, and the number of surviving ℓ is $O(\sqrt{N})$, giving an error of order $N^{-1/2}$. That is also the true rate, which is why the numerical approach was hopeless.

The third is that $\pi^2/16 = (\pi/4)^2$ is the square of the Leibniz value, and squares of series values arise from integrals of the form $\int 2gg'$. Once you see $\int_0^1 2\arctan u\, d(\arctan u)$ you understand the design of the problem: the examiners built a double sum whose two index families both encode odd reciprocals, one as $\sum(-1)^{\ell-1}t^{\ell-1}/(2\ell-1)$ and one as $1/(1+t)$, and the change of variable $t = u^2$ is what reveals them to be the same. The nudge they planted is the word "odd" in the definition of $A(k)$ together with the factor 2 inside $\sqrt{2k}$. That 2 has no arithmetic purpose; its only function is to make the inequality collapse to $\ell \leq m$ with no floor functions surviving. When a constant in a problem statement exists purely to make one inequality clean, it is telling you where the intended first move is.

How to train the discernment. Drill the question "which index carries the sign" on every double sum you meet, before choosing an order, and drill the question "is this family absolutely summable" before writing any interchange. Both are ten second checks. Then drill the alternating tail bound until it is reflex, because in this problem it does three separate jobs: it makes the outer series absolutely convergent, it controls the truncation error, and it supplies uniform convergence for the integral swap. A single inequality doing three jobs is the mark of the correct decomposition, and noticing that is how you tell in the arena that you are on the right road rather than a road.

Revolving concepts

Odd divisors and the divisor pair. A divisor d of k comes with a cofactor $m = k/d$; the map $d \mapsto (d, m)$ turns divisor sums into lattice sums. The hyperbola method (splitting $\sum_{dm \le x}$ at $d \le \sqrt{x}$) is the standard exploitation, and Dirichlet convolution $(f * g)(k) = \sum_{d|k} f(d)g(k/d)$ is its algebraic form. Note that A is not multiplicative, because the constraint mixes d with m; do not waste time hunting for an Euler product. See Book 1, field: Number theory, topic: divisor sums and the hyperbola method.

Sums of consecutive integers. A run of j consecutive positive integers starting at a sums to $j(2a + j - 1)/2$, and the two factors j and $2a + j - 1$ have opposite parity with product $2k$. Hence the number of representations of k as a sum of one or more consecutive positive integers equals the number of odd divisors of k, and restricting to odd j gives exactly $A(k)$. See Book 1, field: Number theory, topic: representations by consecutive integers.

Alternating series and their tails. If $b_1 \ge b_2 \ge \cdots \to 0$ then $\sum(-1)^{m-1}b_m$ converges and every tail is bounded in modulus by its first term. Two generalisations you must own: Dirichlet's test (bounded partial sums of one factor times a monotone null sequence) and its uniform version, which is what licenses termwise integration when the Weierstrass M test fails, as it does here.

Conditional convergence and rearrangement. Riemann's rearrangement theorem: a real series that converges but not absolutely can be reordered to converge to any prescribed value or to diverge. Consequence for exam writing: an interchange of summation order in a non absolutely convergent double family is a claim requiring proof, and the proof is a partial sum comparison.

Recursion and telescoping, generalised. The tails obey the first order recursion $T_\ell - T_{\ell+1} = \frac{(-1)^{\ell-1}}{\ell}$, and any sum weighted by tails can be resummed by Abel summation by parts: for $C_m = \sum_{\ell \le m} c_\ell$ and $T_\ell = \sum_{m \ge \ell} a_m$ one has $\sum_{\ell \ge 1} c_\ell T_\ell = \sum_{m \ge 1} a_m C_m$ whenever either side is legitimate. That identity is precisely the passage between the third ranked approach and the ideal one, with $c_\ell = 1/(2\ell - 1)$, $C_m = O_m$ and $a_m = (-1)^{m-1}/m$. The general solution method for a whole family of these problems is: express the inner tail as $T_\ell = (-1)^{\ell-1} \int_0^1 t^{\ell-1}(1 + t)^{-1} dt$, so that for any weights with $\sum |c_\ell|/\ell < \infty$,

$$\sum_{\ell \ge 1} c_\ell T_\ell = \int_0^1 \frac{C(-t)}{1+t}\, dt, \qquad C(x) = \sum_{\ell \ge 1} c_\ell x^{\ell-1}.$$

Here $C(x) = \sum x^{\ell-1}/(2\ell - 1)$, so $C(-t) = \arctan(\sqrt{t})/\sqrt{t}$. Change the weights and you change only the generating function.

Power series you must have memorised. $\arctan x = \sum_{\ell \ge 1}(-1)^{\ell-1}x^{2\ell-1}/(2\ell - 1)$ and $\log(1 + x) = \sum_{m \ge 1}(-1)^{m-1}x^m/m$ for $|x| < 1$, together with Abel's theorem for the endpoint $x = 1$. Also $1/m = \int_0^1 t^{m-1} dt$ and $1/(2\ell - 1) = \int_0^1 u^{2\ell-2} du$ as denominator removal devices. See Book 1, field: Real analysis, topic: power series and Abel summation.

Knowledge gaps, stated plainly. You cannot approach this problem at all without four items that are outside Edexcel Further Mathematics. One: the alternating series test with its tail estimate, in the quantitative form "the tail is at most the first omitted term". Two: the notion of absolute versus conditional convergence and Riemann's rearrangement theorem, without which you will not know that an interchange needs proof, and you will lose the problem while believing you have solved it. Three: uniform convergence of a series of functions and the theorem that it permits termwise integration on a compact interval; the Weierstrass M test alone is not enough here, so you need the alternating or Dirichlet form. Four: the divisor pair parametrisation as a habit rather than a fact. A fifth, optional but comfortable, is the Lebesgue monotone convergence theorem, which shortens the interchange if you have measure theory; the solution above avoids it deliberately so that the write up stands on first year analysis only.

Circuit Bank

- **WHEN** a sum over k whose coefficient counts divisors of k subject to an inequality that mentions k itself **DO** substitute $k = dm$ and square the inequality **BECAUSE** the square root is an artefact of naming only one member of the divisor pair, and squaring converts the condition into a comparison between d and m, which is a region in the lattice
- **WHEN** a double sum in which the sign of the term depends on only one of the two indices **DO** put that index on the inside and the other on the outside **BECAUSE** the inner sum is then alternating, its tail is bounded by its first term, and that bound is usually exactly what makes the outer sum absolutely convergent
- **WHEN** yourself about to interchange two summations, or a sum and an integral, in a family that is not absolutely summable **DO** write the difference of the two partial sums and bound it explicitly **BECAUSE** Riemann's rearrangement theorem says the value can be changed by reordering, so the interchange is a theorem you owe, and the bound is usually two lines once you have an alternating tail estimate
- **WHEN** a denominator m or $2\ell - 1$ blocking a geometric summation **DO** write $1/m = \int_0^1 t^{m-1} dt$ or $1/(2\ell - 1) = \int_0^1 u^{2\ell-2} du$ and take the sum inside **BECAUSE** integration against a monomial is the inverse of differentiating a power series, so it exchanges a linear denominator for one extra integration
- **WHEN** an integrand of the form $2g(u)g'(u)$, or an answer that looks like the square of a known series value **DO** write it as the derivative of g^2 and evaluate at the endpoints **BECAUSE** a squared constant such as $(\pi/4)^2$ almost never arises from a single series; it arises from an integral of a function against its own differential

Backtesting

1. First useful cue: the divisor bound $\sqrt{2k}$ mentions k, so the condition is really a condition on the divisor pair (d, m); write $k = dm$ and square.
2. Method the expression suggested: convert to a lattice sum over $\{(\ell, m) : \ell \leq m\}$, put the sign bearing index m innermost so the inner sum is an alternating harmonic tail, then use the integral representation of that tail and the arctangent series.
3. Where you would hesitate, and the thought to reinforce: at the rearrangement. The thought to reinforce is that the hesitation is correct and the cure is arithmetic, not philosophy: bound $S - S_N$ in two pieces, the discarded ℓ and the truncated m, using the same tail inequality twice.
4. The trap the question sets: the coefficient $A(k)$ is neither bounded nor monotone and $\sum A(k)/k$ diverges, so a candidate who treats the series as an ordinary alternating series, or who interchanges without comment, writes a solution that reaches the right number and scores very little. The number is not the problem; the ordering is.
5. What makes you recognise it faster next time: the pairing of an arithmetic coefficient with the factor $(-1)^{k-1}$ and a $1/k$. That combination is always a lattice sum with the sign attached to one coordinate, and the first two questions are always "which coordinate carries the sign" and "is this family absolutely summable".
6. The reusable decision rule: in a divisor weighted series, parametrise by the pair, order the double sum so that the alternating index is innermost, and pay for the reordering immediately with an explicit partial sum estimate.

Mirror drills

1. Show that the number of ways to write a positive integer k as a sum of one or more consecutive positive integers equals the number of odd divisors of k. *Answer:* a run of j terms starting at a sums to $j(2a+j-1)/2$, so $2k = j \cdot (2a+j-1)$ with the two factors of opposite parity and the second larger than the first; sending the representation to whichever factor is odd is a bijection onto the odd divisors of k.
2. Prove that $\sum_{m=\ell}^{\infty} \frac{(-1)^{m-1}}{m} = (-1)^{\ell-1} \int_0^1 \frac{t^{\ell-1}}{1+t}\,dt$ and deduce that its modulus is at most $1/\ell$. *Answer:* sum the finite geometric progression $\sum_{m=\ell}^{N}(-1)^{m-1}t^{m-1}$ inside the integral, note the error term is $\int_0^1 t^N(1+t)^{-1}dt \le \frac{1}{N+1}$, and bound $\int_0^1 t^{\ell-1}(1+t)^{-1}dt \le \int_0^1 t^{\ell-1}dt = 1/\ell$.
3. Evaluate $\sum_{m=1}^{\infty} \frac{(-1)^{m-1}}{m}\left(1+\frac{1}{3}+\frac{1}{5}+\cdots+\frac{1}{2m-1}\right)$. *Answer:* $\pi^2/16$; it is the Abel summation by parts partner of the sum in this problem, since $\sum_{\ell\ge1} c_\ell T_\ell = \sum_{m\ge1} a_m C_m$ with $c_\ell = 1/(2\ell-1)$ and $a_m = (-1)^{m-1}/m$, and the same convergence estimate applies.

Chapter 12

Putnam 2014

12.1 A1. Prime numerators in an exponential Taylor series

Prove that every nonzero coefficient of the Taylor series of

$$(1 - x + x^2)e^x$$

about $x = 0$ is a rational number whose numerator (in lowest terms) is either 1 or a prime number.

Read it out loud

A fixed quadratic multiplies e^x, and you are asked about the numerator of a coefficient once it is written in lowest terms. The instant "in lowest terms" appears you are not being asked for calculus, you are being asked for number theory wearing a Taylor series as a costume. The series is only the delivery mechanism; the content lives entirely in what a fraction loses when you cancel it down. That is the first cue, and it should fire within the first sentence: whenever a coefficient's reduced numerator or denominator is the object of the question, the circuit is "combine into one fraction over a factorial, then ask which primes the factorial can certify are gone", because a factorial is the one denominator whose prime content you know completely.

The fields wired into the question: e^x brings exponential generating functions and the fact $[x^n]e^x = 1/n!$; "numerator in lowest terms" brings elementary number theory, specifically the divisibility of factorials and the three ways a positive integer can be built (it is 1, it is prime, or it is composite). What is really being tested is whether you resist differentiating $(1 - x + x^2)e^x$ several times in search of a pattern, and instead read the product as three shifted copies of a series you already know. Within thirty seconds you should have $[x^n](1 - x + x^2)e^x = \frac{1}{n!} - \frac{1}{(n-1)!} + \frac{1}{(n-2)!}$, with the convention that a factorial of a negative integer is 0. Everything after that line is arithmetic.

Approaches, ranked worst to best

1. **repeated differentiation.** *Cue:* "Taylor series...about $x = 0$". *Circuit:* $f^{(n)}(0)/n!$ by the product rule. (costs enormous ink: the product rule regenerates all three terms at every order and gives no structural reason a numerator would be prime; you compute blind, hoping a pattern appears.)
2. **numerical pattern spotting.** *Cue:* "either 1 or a prime". *Circuit:* tabulate the first ten coefficients, factor the numerators, guess a rule. (a legitimate way to discover the closed form if you have not met this cue before, but on its own it is a conjecture, not a proof, since it never explains why composite $n - 1$ forces the numerator down.)
3. **shift the series then classify one integer.** *Cue:* $(1 - x + x^2)e^x$ read as three copies of e^x shifted by $0, 1, 2$. *Circuit:* write the coefficient as one fraction over $n!$, cancel the shared factor, classify $n - 1$ as prime, composite with unequal factors, or a prime square. (the ideal route: least ink, and every case is forced by the definitions of prime and composite, never by inspection.)

The ideal solution

Adopt the convention $1/k! = 0$ for negative integers k. Then for every $n \geq 0$,

$$c_n := [x^n](1 - x + x^2)e^x = \frac{1}{n!} - \frac{1}{(n-1)!} + \frac{1}{(n-2)!}.$$

Multiply by $n!$: the integer $n!\,c_n$ equals $1 - n + n(n-1) = (n-1)^2$, a check that holds at $n = 0$ ($1 = (-1)^2$) and $n = 1$ ($0 = 0^2$) exactly as it holds for larger n, so with no separate base case,

$$c_n = \frac{(n-1)^2}{n!} \qquad \text{for every } n \geq 0.$$

This gives $c_0 = 1$ (numerator 1) and $c_1 = 0$ (excluded, it is not a nonzero coefficient). Fix $n \geq 2$ and set $m = n - 1 \geq 1$. Since $n! = n \cdot m \cdot (m-1)!$, one copy of m cancels against $(n-1)^2 = m^2$:

$$c_n = \frac{m}{n\,(m-1)!}.$$

Consecutive integers m and n share no common factor, so the only further cancellation is between m and $(m-1)!$. Let $d = \gcd(m, (m-1)!)$; the reduced numerator is m/d, and the claim is $m/d \in \{1\} \cup \{\text{primes}\}$.

If $m = 1$, then $(m-1)! = 1$, $d = 1$, numerator 1. If m is prime, every factor of $(m-1)! = 1 \cdot 2 \cdots (m-1)$ is smaller than m, so none carries the prime m; hence $d = 1$ and the numerator is the prime m itself.

If m is composite, let p be its smallest prime factor. Were m/p to have a prime factor below p, that factor would divide m, contradicting the minimality of p; so every prime factor of m/p is at least p, forcing $m/p \geq p$. Two shapes follow. If $m/p > p$, then p and m/p are distinct integers, both at most $m-1$, so both occur as separate factors inside $(m-1)!$; hence $m = p \cdot (m/p)$ divides $(m-1)!$, giving $d = m$ and numerator 1. If $m/p = p$, then $m = p^2$. For $p \geq 3$, $2p < p^2$ (since $p(p-2) > 0$), so p and $2p$ are distinct integers at most $m-1$, and again m divides $(m-1)!$, giving numerator 1. For $p = 2$, $m = 4$ and $(m-1)! = 3! = 6$, so $d = \gcd(4, 6) = 2$ and the numerator is $4/2 = 2$, a prime. In every case the reduced numerator of c_n is 1 or a prime. ■

Layer 1: WHAT

1. **Block 1.** Read $(1 - x + x^2)e^x$ as three shifted copies of the exponential series and write $c_n = 1/n! - 1/(n-1)! + 1/(n-2)!$.
2. **Block 2.** Combine over $n!$ and simplify the integer numerator $1 - n + n(n-1)$ to the perfect square $(n-1)^2$, unifying every $n \geq 0$ into one formula.
3. **Block 3.** Cancel one copy of $m = n - 1$ against $n! = n \cdot m \cdot (m-1)!$, reducing the whole problem to one question about one integer m against $(m-1)!$.
4. **Block 4.** Classify m by its smallest prime factor into three shapes (prime; composite with unequal factors; a prime square) and read off the numerator in each.

The nominal route differentiates $(1 - x + x^2)e^x$ repeatedly, hunting a pattern in raw numbers, and never rises above a conjecture. The ideal route never differentiates: it reads the product as a shifted sum on sight and moves straight into divisibility bookkeeping. The bifurcation point is the opening reflex itself, the choice between "differentiate" and "shift and combine", made before a single number is computed.

Layer 2: HOW

Block 1. **Cue.** The factor $(1 - x + x^2)$ is fixed and finite, while $e^x = \sum x^k/k!$ has coefficients you already know. **Why natural.** Multiplying a finite polynomial into a power series only shifts and rescales that series; there is no genuine interaction for the product rule to rediscover term by term. **Link to the next block.** Three known terms sitting at three different factorial denominators are unreadable until forced onto one denominator.

Block 2. **Cue.** The denominators $n!$, $(n-1)!$, $(n-2)!$ differ only by the small multipliers n and $n(n-1)$. **Why natural.** Clearing to $n!$ turns the coefficient into (integer)$/n!$, and that integer is a quadratic in n worth simplifying on sight. **Link to the next block.** Seeing $1 - n + n(n-1) = (n-1)^2$ guarantees a whole factor of $(n-1)$ is available for cancellation against the visible $(n-1)$ inside $n!$.

Block 3. **Cue.** $n! = n(n-1)(n-2)!$ displays $(n-1)$ as an explicit factor. **Why natural.** Consecutive integers are always coprime, so n can never absorb any part of $(n-1)^2$; the only real contest is between $(n-1)$ and $(n-2)!$. **Link to the next block.** This isolates one question about one integer $m = n-1$: how much of m does $(m-1)!$ already contain?

Block 4. **Cue.** The three-way shape of a positive integer under its smallest prime factor. **Why natural.** Prime, composite with unequal factors, and prime square are the only shapes that classification produces, and each gives a fully determined answer. **Link to the next block.** No further block is needed; each case terminates in an explicit numerator.

Layer 3: WHY

The truth underneath this problem is a companion to Wilson's theorem: a composite integer m divides $(m-1)!$ for every m except $m = 4$. You do not need Wilson's theorem's statement about primes, only this shadow of it on the composite side, and you should be able to reprove that shadow from the definition of smallest prime factor in under a minute rather than quote it and hope. The discernment mechanism that finds this route is a standing question: the instant a numerator in lowest terms appears, ask what the denominator is built from and which primes it can certify absent. Factorials answer this completely, since every prime up to n divides $n!$ and none above it does, which is exactly why the examiners chose e^x rather than, say, a geometric series, whose coefficients carry no factorial structure to interrogate. Train this by refusing, on any future factorial-numerator problem, to compute a single term before writing the coefficient as (small polynomial in n)/(a factorial). The examiners nudged you here by asking for the numerator specifically: a question about the value of a coefficient invites calculus; a question about its reduced numerator is a standing invitation to prime factorisation.

Revolving concepts

Exponential generating functions. For $P(x) = \sum_{i=0}^{d} c_i x^i$, $[x^n]\big(P(x)e^x\big) = \sum_{i=0}^{d} c_i/(n-i)!$, with negative factorials contributing 0; this problem is the case $d = 2$, $P(x) = 1 - x + x^2$. Meet any fixed polynomial against e^x, $\sin x$, $\cos x$, or another series with known coefficients, and write this shifted sum first, before differentiating. See Book 1, field: Calculus and integrals, topic: Taylor and Maclaurin series.

Factorials and their primes. $n!$ is divisible by exactly the primes up to n, to the power given by Legendre's formula $\sum_{i \geq 1} \lfloor n/p^i \rfloor$; here you only needed that every integer from 1 to $n-1$ is an explicit factor of $(n-1)!$. See Book 1, field: Number theory, topic: factorials and prime valuations.

The composite-factorial lemma. Composite m divides $(m-1)!$ unless $m = 4$. This is the engine of the whole proof and is worth installing by name, though its proof needs nothing beyond the classification used above. See Book 1, field: Number theory, topic: Wilson's theorem and its refinements.

Coprimality of consecutive integers. $\gcd(k, k+1) = 1$ for every integer k, since any common divisor divides their difference 1; this is what guarantees n never interferes with the cancellation of $n-1$. See Book 1, field: Number theory, topic: greatest common divisor.

There is no recursion or telescoping in this problem: c_n is handed to you in closed form by the product rule for generating functions, not by an earlier coefficient of itself, so there is nothing to generalise along that axis.

Circuit Bank

- **WHEN** a fixed polynomial multiplies a series with known coefficients **DO** write the coefficient of x^n as the finite shifted sum on sight **BECAUSE** differentiating regenerates the same

information at far higher cost and hides the factorial structure you need

- **WHEN** a question names the numerator or denominator of a coefficient in lowest terms **DO** express the coefficient as a small polynomial in n over a factorial, then hunt the shared factor **BECAUSE** a factorial's prime content is completely known, so it is the one denominator that lets you certify which primes survive
- **WHEN** you must decide whether an integer m divides $(m-1)!$ **DO** classify m as prime, composite with unequal factors, or a prime square, and check the single exception $m = 4$ **BECAUSE** these three shapes exhaust every composite integer's smallest-prime-factor decomposition
- **WHEN** two quantities differ by consecutive integers **DO** treat them as coprime without further checking **BECAUSE** $\gcd(k, k+1) = 1$ always

Backtesting

1. First useful cue: "numerator...in lowest terms" attached to a Taylor coefficient, converting calculus into number theory before you differentiate anything.
2. Method the expression suggested: read $(1 - x + x^2)e^x$ as three shifted copies of e^x and combine directly.
3. Where you would hesitate, and the thought to reinforce: after reaching $(n-1)^2/n!$ the three-way case split looks like three separate proofs; it is one question, does m divide $(m-1)!$, asked three times.
4. The trap the question sets: assuming every composite $n-1$ gives numerator 1 and missing the single exception $m = 4$, which survives as the prime 2.
5. What makes you recognise it faster next time: any coefficient built from factorials should trigger "which primes can the denominator certify absent" before any computation.
6. The reusable decision rule: a numerator-in-lowest-terms question over a factorial denominator reduces to classifying one integer against $(m-1)!$, in at most three cases.

Mirror drills

1. (drill) Show that the coefficient of x^n in $(1+x)e^x$ has reduced numerator 1 or prime for every $n \geq 0$. *Answer:* the coefficient is $(n+1)/n!$; classify $m = n+1$ by its smallest prime factor exactly as in the ideal solution, with $m = 4$ (that is $n = 3$) the lone exception, giving numerator 2.
2. (drill) Prove directly that every composite $m > 4$ divides $(m-1)!$. *Answer:* let p be the smallest prime factor of m; if $m/p > p$ use the distinct factors p, $m/p \leq m-1$, and if $m = p^2$ with $p \geq 3$ use the distinct factors p, $2p \leq m-1$.
3. (drill) Write the coefficient of x^n in $(1 + x + x^2 + x^3)e^x$ as a single fraction over $n!$ without differentiating. *Answer:* $\dfrac{1 + n + n(n-1) + n(n-1)(n-2)}{n!}$, read off term by term exactly as in Block 1.

12.2 A2. Determinant of the reciprocal minimum matrix

Let A be the $n \times n$ matrix whose entry in the i-th row and j-th column is

$$\frac{1}{\min(i,j)}$$

for $1 \leq i, j \leq n$. Compute $\det(A)$.

Read it out loud

A matrix built entirely from $\min(i,j)$, and the command term is "compute", not "prove": the ideal route has to terminate in a formula clean enough to write down, not an existence statement. The moment $\min(i,j)$ appears inside a matrix entry a memory should surface: the plain matrix $M_{ij} = \min(i,j)$ is a well known object with determinant 1 for every n, and the reflex to treat this problem as "that matrix, just reciprocated" is exactly the trap on offer. Reciprocating entries does not commute with inverting or with taking determinants, so whatever you remember about M tells you about the shape of A and nothing about its determinant.

What the structure actually offers is this: row i of A reads $1, \frac{1}{2}, \ldots, \frac{1}{i-1}$ and then a run of the constant $\frac{1}{i}$ out to column n, because $\min(i,j) = j$ while $j < i$ and $\min(i,j) = i$ once $j \geq i$. Two rows built this way agree on a long initial stretch and disagree only in a constant tail, and that is the guiding light: any operation that subtracts one row from an adjacent one annihilates the entire agreeing stretch at a stroke. You should have this reflex within thirty seconds of seeing $\min(i,j)$ inside a matrix, before touching a single cofactor: try adjacent row differencing first, because a matrix whose entries are a function of min or max of the indices is built for exactly this.

The fields embedded are linear algebra (row operations, and the determinant of a triangular matrix) and a small amount of series manipulation (the leftover product of consecutive differences telescopes into a ratio of factorials). Nothing here needs eigenvalues or a general cofactor expansion; you need only that left multiplication by a unit triangular matrix leaves the determinant unchanged, and that a triangular matrix's determinant is the product of its diagonal.

Approaches, ranked worst to best

1. **Leibniz or raw cofactor expansion.** *Cue:* "compute $\det(A)$" with no other hint. *Circuit:* sum over all $n!$ permutations, or expand by minors along a row. (stalls at once: $n!$ dense terms with no visible cancellation, and repeated cofactor expansion regenerates a dense $(n-1) \times (n-1)$ minor at every step.)

2. **recall the classical $\min(i,j)$ determinant and try to adapt it.** *Cue:* the visibly famous kernel $\min(i,j)$. *Circuit:* quote $\det(M) = 1$ for $M_{ij} = \min(i,j)$ and try to relate A to M's inverse or adjugate. (a trap: entrywise reciprocal is not matrix inverse, so the classical fact about M says nothing directly about A; chasing the relationship costs ink and never closes.)

3. **border and recurse along the last row.** *Cue:* the matrix grows by one row and column as n increases. *Circuit:* expand along row n to find D_n in terms of D_{n-1}. (works in principle, but row n is not sparse, every entry is $1/i$ or $1/n$, so extracting the recursion needs several lines of algebra and more chances for a sign slip than the route below.)

4. **adjacent row differencing via a single elimination matrix.** *Cue:* row i agrees with row $i-1$ on a long initial stretch, by the definition of min. *Circuit:* left multiply by a fixed unit lower bidiagonal matrix to expose an upper triangular matrix in one identity. (the ideal route: one matrix identity, the determinant read straight off a diagonal, and the leftover product telescopes into factorials.)

The ideal solution

Write the rows of A as $v_1, \dots, v_n$. Since $\min(i,j) = j$ for $j < i$ and $\min(i,j) = i$ for $j \geq i$,

$$v_i = \Big(1, \tfrac{1}{2}, \dots, \tfrac{1}{i-1}, \underbrace{\tfrac{1}{i}, \tfrac{1}{i}, \dots, \tfrac{1}{i}}_{\text{columns } i \text{ to } n}\Big).$$

Let L be the $n \times n$ unit lower bidiagonal matrix with $L_{11} = 1$, $L_{kk} = 1$ and $L_{k,k-1} = -1$ for $k = 2, \dots, n$, and every other entry 0. Then L is lower triangular with every diagonal entry 1, so $\det L = 1$, and the rows of LA are

$$w_1 = v_1, \qquad w_k = v_k - v_{k-1} \ (k = 2, \dots, n),$$

computed directly from the original rows of A, with no question of which row was modified first. It remains to compute $\det(LA) = \det(L)\det(A) = \det(A)$.

Fix $k \geq 2$ and compare v_k with v_{k-1} column by column. For $j \leq k-2$, both $\min(k,j)$ and $\min(k-1,j)$ equal j, so the entries agree. For $j = k-1$, $\min(k, k-1) = k-1 = \min(k-1, k-1)$, so they agree there too. For $j \geq k$, $\min(k,j) = k$ while $\min(k-1,j) = k-1$ (since $j \geq k > k-1$), so the entries differ by the constant $\frac{1}{k} - \frac{1}{k-1}$ across every such column. Hence

$$w_k = \Big(\underbrace{0, \dots, 0}_{k-1}, \underbrace{\tfrac{1}{k} - \tfrac{1}{k-1}, \dots, \tfrac{1}{k} - \tfrac{1}{k-1}}_{\text{columns } k \text{ to } n}\Big), \qquad k = 2, \dots, n,$$

and $w_1 = (1, 1, \dots, 1)$. Every w_k vanishes strictly before column k, so LA is upper triangular, with diagonal entries 1, $\frac{1}{2} - 1$, $\frac{1}{3} - \frac{1}{2}$, $\dots$, $\frac{1}{n} - \frac{1}{n-1}$. Therefore

$$\det(A) = \det(LA) = 1 \cdot \prod_{k=2}^{n} \Big(\frac{1}{k} - \frac{1}{k-1}\Big) = \prod_{k=2}^{n} \frac{-1}{k(k-1)} = (-1)^{n-1} \cdot \frac{1}{\prod_{k=2}^{n} k} \cdot \frac{1}{\prod_{k=2}^{n} (k-1)}.$$

The two products telescope on sight, $\prod_{k=2}^{n} k = n!$ and $\prod_{k=2}^{n} (k-1) = (n-1)!$, so

$$\boxed{\det(A) = \frac{(-1)^{n-1}}{n!\,(n-1)!}}.$$

(This also covers $n = 1$, where the product over $k = 2, \dots, n$ is empty and $\det A = 1$, matching $A = (1)$.)

Layer 1: WHAT

1. **Block 1.** Read off the row structure of A directly from the definition of min: an increasing prefix, then a constant tail.
2. **Block 2.** Notice adjacent rows agree on the entire prefix and differ only in the tail, so a row difference should annihilate the agreement.
3. **Block 3.** Package every row difference into one matrix identity, left multiplication by a unit lower bidiagonal L with $\det L = 1$, so no order of operations needs justifying.
4. **Block 4.** Verify LA is upper triangular and read the determinant off the diagonal.
5. **Block 5.** Telescope the resulting product of consecutive reciprocal differences into a ratio of two factorials.

The nominal route reaches for cofactor expansion or the Leibniz sum, treating A as an opaque array of numbers with no exploitable shape; it never terminates in a clean formula without an unmotivated leap. The ideal route never expands anything: it reads the min structure as an instruction to difference rows, and packages that instruction as a single matrix so no step is ambiguous. The bifurcation point is the opening reflex, whether $\min(i,j)$ inside a matrix triggers "expand" or "difference adjacent rows", fired before a single entry is computed.

Layer 2: HOW

Block 1. **Cue.** Row i is constant from column i onward, because $\min(i, j)$ stops growing once j reaches i. **Why natural.** A matrix entry defined by a comparison of indices always has this "capped" shape: once one index dominates, further increase of the other index changes nothing. **Link to the next block.** Two rows differing only in where the cap kicks in must be nearly identical, which begs for a comparison between adjacent rows.

Block 2. **Cue.** Rows i and $i-1$ share the exact same values on columns 1 through $i-1$. **Why natural.** Both are governed by the smaller index on that range, since $j < i-1 < i$ forces $\min(i, j) = \min(i-1, j) = j$ for both rows simultaneously. **Link to the next block.** Subtracting row $i-1$ from row i should zero out that entire shared range in one operation, which is the definition of progress toward a triangular matrix.

Block 3. **Cue.** You need every row difference to use the original, unmodified row above it, or the zero pattern breaks. **Why natural.** If you instead updated rows sequentially in increasing order, row k would difference against an already changed row $k-1$, corrupting the result; encoding all the differences as one multiplication by L applies them simultaneously to the original A and removes the question entirely. **Link to the next block.** Once LA is written down as a genuine matrix, checking its shape is now pure verification, not a procedure that can go wrong in execution.

Block 4. **Cue.** Every row of LA vanishes strictly before its own diagonal position, by direct computation. **Why natural.** That is the definition of upper triangular, and for such a matrix the determinant is exactly the product of the diagonal, with no cofactor work left to do. **Link to the next block.** The only content still needed is to name the diagonal entries themselves.

Block 5. **Cue.** The diagonal entries are the consecutive differences $\frac{1}{k} - \frac{1}{k-1}$. **Why natural.** Combining over a common denominator turns each term into $-1/(k(k-1))$, and a product of such terms splits into two separate products, one running over k and one over $k-1$, each of which is visibly $n!$ or $(n-1)!$. **Link to the next block.** Nothing further is needed; the two telescoping products are the final answer.

Layer 3: WHY

The truth underneath this problem is that any matrix entry depending only on $\min(i, j)$ (or, symmetrically, only on $\max(i, j)$) encodes a "prefix agreement" between rows that is invisible to cofactor expansion but is exactly what row differencing is built to exploit. The discernment mechanism is to ask, on sight of min or max inside a matrix entry, whether two adjacent rows or columns must agree on a long run purely from the definition of the comparison, before attempting any other manipulation; if they do, row or column differencing is close to free and cofactor expansion is close to the most expensive option available. Packaging the differencing as multiplication by an explicit unit triangular L is the discipline that removes the one way this route can go wrong, namely applying the differences in an order that reuses an already modified row; a single matrix identity has no order to get wrong. The examiners nudged you toward this by choosing $\min(i, j)$ specifically rather than, say, $i + j$ or ij: those entries do not create the prefix agreement between rows that makes differencing free, so the choice of min is itself the guiding light, not an incidental feature of the problem.

Revolving concepts

Matrices built from min **of the indices.** If $c_1, c_2, \ldots, c_n$ is any sequence and $A_{ij} = c_{\min(i,j)}$, the same argument gives $\det(A) = c_1 \prod_{k=2}^{n}(c_k - c_{k-1})$: row 1 is constant at c_1, and row k minus row $k-1$ vanishes before column k and equals $c_k - c_{k-1}$ from column k onward. This problem is the case $c_k = 1/k$; the classical fact $\det(\min(i,j)) = 1$ is the case $c_k = k$, where every difference $c_k - c_{k-1}$ equals 1. See Book 1, field: Matrices and determinants, topic: structured determinants.

Determinant of a triangular matrix. For an upper or lower triangular matrix, det is the product of the diagonal entries; this is the single fact that turns "matrix reduced to triangular form" into "answer written down". See Book 1, field: Matrices and determinants, topic: triangular matrices.

Row operations as left multiplication. Replacing the rows of A according to a fixed linear rule is the same as forming LA for a matrix L built from that rule; if L is unit triangular, $\det L = 1$ and $\det(LA) = \det(A)$. Framing row reduction this way, rather than as a sequence of instructions, is what lets you certify a computation without worrying about the order in which steps were carried out. See Book 1, field: Linear algebra, topic: elementary matrices.

Nested determinants and a one-step recursion. The top left $(n-1) \times (n-1)$ block of A is exactly the $(n-1) \times (n-1)$ version of the same matrix, and the triangularisation is compatible across n, so $D_n = (\frac{1}{n} - \frac{1}{n-1})\, D_{n-1}$: each new size contributes exactly one new diagonal entry. Generalised, whenever a family of matrices is nested this way and reduces to triangular form compatibly, its determinants satisfy $D_n = (\text{new diagonal entry}) \times D_{n-1}$, a first order recursion solved immediately by telescoping the ratios, exactly as the product $\prod(c_k - c_{k-1})$ above does for any such family. See Book 1, field: Sequences and recurrences, topic: telescoping products.

A remark worth knowing, not needed for the proof. The official solution notes that the inverse of A is tridiagonal. This is not a coincidence: the continuous kernel $\min(x,y)$ on $[0,1]^2$ is the Green's function of $-u'' = f$ with $u(0) = 0$, $u'(1) = 0$, and matrices built from $\min(i,j)$ are the discrete analogue of that kernel, so their inverses inherit the tridiagonal shape of the discrete second derivative operator. See Book 1, field: Real analysis, topic: Green's functions.

Circuit Bank

- **WHEN** a matrix entry depends only on $\min(i,j)$ or $\max(i,j)$ **DO** try adjacent row or column differencing before any cofactor expansion **BECAUSE** two adjacent rows of such a matrix agree on the entire capped prefix and differ only on the remainder, so differencing produces triangularity almost for free

- **WHEN** you want to justify a chain of row operations without worrying about their order **DO** exhibit the single matrix L with known determinant such that LA is the matrix you want **BECAUSE** $\det(LA) = \det(L)\det(A)$ settles the computation in one line and removes any question of which row was updated first

- **WHEN** a matrix is triangular, by construction or after reduction **DO** read the determinant straight off the diagonal **BECAUSE** that is the entire content of the triangular determinant formula, with no cofactor work left

- **WHEN** a product of consecutive-integer reciprocal differences appears **DO** combine each term into a single fraction and split the product into two separately telescoping factorial products **BECAUSE** $\prod_{k=a}^{b} k$ collapses immediately to a ratio of factorials

- **WHEN** a family of matrices is nested, with the $(n-1)$-th sitting as the top left block of the n-th, and triangularises compatibly **DO** look for the one-step recursion $D_n = (\text{new diagonal entry}) \times D_{n-1}$ **BECAUSE** the new diagonal entry is the only genuinely new information the larger matrix contributes once the reduction pattern is fixed

Backtesting

1. First useful cue: $\min(i, j)$ sitting inside a matrix entry, which should trigger adjacent row differencing before any other tool.
2. Method the expression suggested: compare row i with row $i-1$ directly from the definition of min, rather than expanding by cofactors.
3. Where you would hesitate, and the thought to reinforce: whether the row differences can be done in any order; reinforce that they cannot in general, and that packaging them as one matrix L removes the question.
4. The trap the question sets: recalling that $\det(\min(i, j)) = 1$ classically and assuming the reciprocated matrix inherits some simple relation to that fact, when reciprocation does not commute with either inversion or determinants.
5. What makes you recognise it faster next time: seeing that a matrix entry compares two indices at all should trigger "check what adjacent rows or columns share" before any other move.
6. The reusable decision rule: for $A_{ij} = c_{\min(i,j)}$, the determinant is always c_1 times the telescoping product of $c_k - c_{k-1}$, so identify the sequence c_k and substitute.

Mirror drills

1. (drill) Show that $\det(M) = 1$ for every n, where $M_{ij} = \min(i, j)$. *Answer:* apply the general lemma with $c_k = k$: $c_1 = 1$ and $c_k - c_{k-1} = 1$ for every k, so $\det(M) = 1 \cdot 1^{n-1} = 1$.
2. (drill) Compute $\det(A)$ where $A_{ij} = \min(i, j)^2$. *Answer:* take $c_k = k^2$ in the general lemma: $c_1 = 1$ and $c_k - c_{k-1} = 2k - 1$, so $\det(A) = \prod_{k=2}^{n}(2k-1)$, the product of the odd numbers from 3 to $2n - 1$.
3. (drill) In the 3×3 case, compute what happens if you update rows in increasing order instead of using the original rows throughout: replace row 2 by row 2 minus row 1, then replace row 3 by (new) row 2 minus row 1. *Answer:* you get row $3 \to (1, 1, \frac{5}{6})$, not triangular at all, which is exactly why the ideal solution defines every w_k from the original rows rather than updating sequentially.

12.3 A3. Telescoping infinite product from a doubling recursion

Let $a_0 = 5/2$ and $a_k = a_{k-1}^2 - 2$ for $k \geq 1$. Compute

$$\prod_{k=0}^{\infty} \left(1 - \frac{1}{a_k}\right)$$

in closed form.

Read it out loud

A recursion of the shape $a_k = a_{k-1}^2 - 2$, and the instant you see a squaring combined with subtracting exactly 2, one specific memory should fire: this is the algebraic shadow of doubling an exponent. If $t = y + \frac{1}{y}$ for some y, then $t^2 - 2 = y^2 + \frac{1}{y^2}$, the very same shape with y replaced by y^2. The recursion is not really about squaring numbers, it is about squaring an auxiliary variable y_k whose exponent doubles at every step, and the first useful move is to find y_0 with $y_0 + \frac{1}{y_0} = a_0 = \frac{5}{2}$, which is $y_0 = 2$ (solve $2y_0^2 - 5y_0 + 2 = 0$).

The command term "compute...in closed form" attached to an infinite product tells you the second half of the strategy before you have used the recursion at all: an infinite product only closes cleanly when it telescopes, meaning each factor can be written as (something depending on the next index) divided by (that same something one index earlier), so every interior factor cancels and only the two ends survive. So the moment $\prod_{k=0}^{\infty}(\cdots)$ appears around a recursively defined sequence, the guiding lights are, in order: find the closed form for a_k first, then hunt for a telescoping certificate for the product, an auxiliary expression that turns the k-th factor into a ratio linking index k to index $k+1$.

The fields embedded are sequences and recurrences (the doubling substitution), infinite products and their convergence (you must justify that the product tends to a nonzero limit, not merely compute as if it does), and a short piece of algebraic identity hunting to build the telescoping certificate. Nothing here needs more than quadratic identities in y and $1/y$; the entire difficulty is structural recognition, not computation, which is exactly why a_0 was handed to you as the suggestive fraction $5/2$ rather than a decimal.

Approaches, ranked worst to best

1. **numerical estimation.** *Cue:* "compute...in closed form" tempts a decimal check. *Circuit:* compute $a_0, a_1, a_2, \ldots$ and the partial products numerically, notice convergence to a suspicious decimal, guess a fraction. (useful only as a sanity check on a candidate answer: it cannot certify the limit is exactly $3/7$ rather than a nearby value indistinguishable at finite precision, and it explains nothing.)

2. **partial products by induction on the recursion alone.** *Cue:* $a_{k+1} + 1 = (a_k - 1)(a_k + 1)$ falls out of the recursion by one line of algebra. *Circuit:* guess, from Π_0, Π_1 computed as exact fractions, a formula for $\prod_{k=0}^{n}(1 - 1/a_k)$ in terms of a_{n+1}, then prove it by induction, without ever finding a_k itself in closed form. (rigorous, and it avoids guessing the closed form for a_k, but the auxiliary identities needed are not suggested by anything visible in the recursion; finding them costs more search time and more luck than the route below.)

3. **doubling substitution, then a matching telescoping certificate.** *Cue:* $a_k = a_{k-1}^2 - 2$ is the algebraic shadow of $y \mapsto y^2$ under $y + 1/y$. *Circuit:* solve for y_0, get $a_k = y_0^{2^k} + y_0^{-2^k}$ in closed form, then reuse the same square-minus-something identity to build an explicit telescoping ratio for the product. (the ideal route: the closed form is forced by the shape of the recursion itself, and the certificate is then a two-line identity, not a guess.)

The ideal solution

Step 1, linearise the recursion. Solve $y_0 + \frac{1}{y_0} = a_0 = \frac{5}{2}$, i.e. $2y_0^2 - 5y_0 + 2 = 0 = (2y_0 - 1)(y_0 - 2)$, giving $y_0 = 2$ (the root $y_0 = \frac{1}{2}$ gives the same sequence by the symmetry $y \leftrightarrow 1/y$). Claim: $a_k = y_0^{2^k} + y_0^{-2^k}$ for every $k \geq 0$. This holds at $k = 0$ by construction. If it holds at $k - 1$, write $z = y_0^{2^{k-1}}$, so $a_{k-1} = z + \frac{1}{z}$, and

$$a_k = a_{k-1}^2 - 2 = \left(z + \tfrac{1}{z}\right)^2 - 2 = z^2 + \tfrac{1}{z^2} = y_0^{2^k} + y_0^{-2^k},$$

which closes the induction. With $y_0 = 2$, $a_k = 2^{2^k} + 2^{-2^k}$.

Step 2, certify convergence. Since $2^{2^k} \geq 2$, every $a_k \geq 2 > 1$, so each factor $1 - \frac{1}{a_k}$ lies in $(0, 1)$. For an infinite product of terms $1 - \epsilon_k$ with $\epsilon_k \in [0, 1)$, the product converges to a strictly positive limit exactly when $\sum_k \epsilon_k$ converges. Here $\epsilon_k = 1/a_k \leq 2^{-2^k}$, and $\sum_k 2^{-2^k}$ converges (it is dominated by a geometric series), so the infinite product in the problem is a genuine, nonzero, well defined limit, not a formal expression.

Step 3, build the telescoping certificate. Let $y_k = y_0^{2^k}$, so $a_k = y_k + y_k^{-1}$ and $y_{k+1} = y_k^2$. For any $t \neq 0$,

$$(t + t^{-1})^2 - 1 = t^2 + 1 + t^{-2},$$

which rearranges, since $t + t^{-1} - 1$ is one factor of a difference of squares against $t + t^{-1} + 1$, into

$$t - 1 + t^{-1} = \frac{t^2 + 1 + t^{-2}}{t + 1 + t^{-1}} = \frac{P(t^2)}{P(t)}, \qquad P(t) := t + 1 + t^{-1}.$$

Separately, since $t^2 - t^{-2} = (t - t^{-1})(t + t^{-1})$,

$$t + t^{-1} = \frac{t^2 - t^{-2}}{t - t^{-1}} = \frac{Q(t^2)}{Q(t)}, \qquad Q(t) := t - t^{-1},$$

valid because $t = y_k \geq 2$ so $Q(t) \neq 0$ throughout. Since $1 - \frac{1}{a_k} = \dfrac{y_k - 1 + y_k^{-1}}{y_k + y_k^{-1}}$, substituting $t = y_k$ into both identities gives

$$1 - \frac{1}{a_k} = \left.\frac{P(t^2)/P(t)}{Q(t^2)/Q(t)}\right|_{t=y_k} = \frac{P(y_{k+1})}{P(y_k)} \cdot \frac{Q(y_k)}{Q(y_{k+1})}.$$

Step 4, telescope and take the limit. The partial product collapses immediately:

$$\prod_{k=0}^{n}\left(1 - \frac{1}{a_k}\right) = \prod_{k=0}^{n} \frac{P(y_{k+1})}{P(y_k)} \cdot \frac{Q(y_k)}{Q(y_{k+1})} = \frac{P(y_{n+1})}{P(y_0)} \cdot \frac{Q(y_0)}{Q(y_{n+1})}.$$

As $n \to \infty$, $y_{n+1} = 2^{2^{n+1}} \to \infty$, and both $P(y_{n+1}) = y_{n+1} + 1 + y_{n+1}^{-1}$ and $Q(y_{n+1}) = y_{n+1} - y_{n+1}^{-1}$ are dominated by y_{n+1} itself, so their ratio tends to 1. Hence

$$\prod_{k=0}^{\infty}\left(1 - \frac{1}{a_k}\right) = \frac{Q(y_0)}{P(y_0)} = \frac{y_0 - y_0^{-1}}{y_0 + 1 + y_0^{-1}} = \frac{2 - \frac{1}{2}}{2 + 1 + \frac{1}{2}} = \frac{3/2}{7/2} = \boxed{\frac{3}{7}}.$$

Layer 1: WHAT

1. **Block 1.** Recognise $a_k = a_{k-1}^2 - 2$ as squaring in disguise under $t = y + 1/y$, and solve for y_0 from a_0.
2. **Block 2.** Prove $a_k = y_0^{2^k} + y_0^{-2^k}$ by one line of induction, converting a numeric recursion into a doubling exponent.

3. **Block 3.** Confirm each factor lies in $(0,1)$ and that $\sum 1/a_k$ converges, certifying the infinite product is a genuine nonzero limit before computing it.
4. **Block 4.** Reuse the identity that produced the closed form to build a telescoping certificate P, Q linking consecutive doublings.
5. **Block 5.** Telescope the partial product to two boundary terms and take the limit as the far boundary's ratio tends to 1.

The nominal route computes decimals, notices $0.428571\ldots$, and stops, or else hunts directly for a partial-product formula in terms of a_{n+1} by unmotivated trial and error on the raw recursion. The ideal route finds the closed form for a_k first, a forced two-line induction, and then builds the telescoping certificate from the exact same algebraic identity used to prove that closed form, rather than searching for a second, unrelated idea. The bifurcation point is the opening reflex: whether "square, then subtract 2" triggers the substitution $y + 1/y$ before a single decimal is computed.

Layer 2: HOW

Block 1. **Cue.** The right-hand side of the recursion is exactly $(\cdot)^2 - 2$. **Why natural.** This is the one quadratic shift for which $t = y + 1/y$ linearises the map completely: $(y + 1/y)^2 - 2 = y^2 + 1/y^2$ has no leftover cross term, so squaring t corresponds exactly to squaring y. **Link to the next block.** Once $y_0 + 1/y_0 = a_0$ is solved, the only remaining question is whether the pattern persists, which is what induction is for.

Block 2. **Cue.** A claimed closed form matches the recursion at the base case $k = 0$. **Why natural.** Induction is the certificate that a closed form matches a recursion for every k, and here the inductive step is the same one-line algebra as Block 1, applied to $y_0^{2^{k-1}}$ instead of y_0. **Link to the next block.** With a_k explicit, the k-th factor of the product is now a fixed rational function of $y_0^{2^k}$, ready for algebraic manipulation rather than numerical estimation.

Block 3. **Cue.** The product is asked for over infinitely many factors. **Why natural.** An infinite product of terms in $(0,1)$ can converge to 0 if the terms shrink too slowly to 1; the correct test is whether $\sum(1 - \text{factor})$ converges, and here it does easily because a_k grows doubly exponentially. **Link to the next block.** With convergence certified, computing "the closed form" is now computing an honest limit, and any valid method, including a telescoping partial product, is permitted.

Block 4. **Cue.** The factor $1 - 1/a_k$, written in terms of y_k, is $\dfrac{y_k - 1 + y_k^{-1}}{y_k + y_k^{-1}}$. **Why natural.** An infinite product only closes cleanly when consecutive factors interlock, so you look for two auxiliary functions, one whose value at y_{k+1} appears in this factor's numerator and one whose value at y_{k+1} appears in its denominator, both built from the identity already used in Block 1. **Link to the next block.** Finding P and Q this way turns the entire infinite product into two boundary evaluations instead of infinitely many factors.

Block 5. **Cue.** The partial product collapses to $P(y_{n+1})/P(y_0)$ times $Q(y_0)/Q(y_{n+1})$. **Why natural.** As $y_{n+1} \to \infty$, the $+1$ inside P and the sign flip inside Q become negligible against the dominant term y_{n+1}, so their ratio tends to 1. **Link to the next block.** Only the two starting values $P(y_0)$ and $Q(y_0)$ survive, giving the closed form directly with no further work.

Layer 3: WHY

The truth underneath this problem is that $u \mapsto u^2 - 2$ is, up to the substitution $u = y + 1/y$ (equivalently $u = 2\cosh\theta$ for $u \geq 2$, or $u = 2\cos\theta$ for $|u| \leq 2$), literally squaring, and squaring is the one operation whose k-fold iterate you can write down immediately: iterate it on an exponent and the exponent doubles, with no approximation and no accumulating error. The discernment mechanism is to treat any recursion of the shape $t_k = t_{k-1}^2 - c$ as a search for the substitution $t = y + c'/y$ that turns it into $y \mapsto y^2$; the constant $c = 2$ is the signature to watch for, because $(y + 1/y)^2 - 2 = y^2 + 1/y^2$ matches with coefficient exactly 1 on both sides, whereas $c = 1$ or $c = 3$ would not linearise this cleanly. Once the closed form exists, an infinite product built from a doubling-exponent sequence should immediately suggest a telescoping certificate, because such a sequence's k-th term is always a fixed rational function of $y_0^{2^k}$, and the certificate is typically the same identity, reused, that produced the closed form in the first place. The examiners nudged you toward all of this by choosing $a_0 = \frac{5}{2}$, a fraction with a visibly clean partner $y_0 = 2$ under $y + 1/y$, rather than a value that would hide the substitution.

Revolving concepts

The recursion $t_k = t_{k-1}^2 - 2$. For $t_0 = y_0 + 1/y_0$ (real $y_0 \geq 1$, so $t_0 \geq 2$) the closed form is $t_k = y_0^{2^k} + y_0^{-2^k}$; for $t_0 = 2\cos\theta_0$ ($|t_0| \leq 2$) it is $t_k = 2\cos(2^k\theta_0)$. Both are the statement that $u \mapsto u^2 - 2$ is conjugate to doubling, and both are the k-fold iterate of the degree 2 Chebyshev polynomial: $t_k = 2T_{2^k}(t_0/2)$, using the composition law $T_m \circ T_n = T_{mn}$. This applies to any iterated quadratic recursion of exactly this shape, including computing cos of dyadically many multiples of an angle and Lucas-Lehmer-flavoured recursions. See Book 1, field: Polynomials, topic: Chebyshev polynomials.

Telescoping certificates for a product over an iterated map. To evaluate $\prod_k f(y_0^{2^k})$ in closed form, look for auxiliary functions P, Q with $f(t) = P(t^2)/P(t) \cdot Q(t)/Q(t^2)$; the partial product then collapses to $P(y_{n+1})/P(y_0) \cdot Q(y_0)/Q(y_{n+1})$ regardless of how complicated f looks termwise. See Book 1, field: Sequences and recurrences, topic: telescoping products.

Convergence of an infinite product. For $\epsilon_k \in [0, 1)$, $\prod_k (1 - \epsilon_k)$ converges to a strictly positive limit if and only if $\sum_k \epsilon_k$ converges; this is checked by comparing $\log(1 - \epsilon_k) \sim -\epsilon_k$ for small ϵ_k against the convergent series $\sum \epsilon_k$. See Book 1, field: Real analysis, topic: infinite products.

Sum of a number and its reciprocal. $y + 1/y = 2\cosh(\ln y)$ for $y > 0$, and for $y = e^{i\theta}$ this is $2\cos\theta$; the identity $(y + 1/y)^2 - 2 = y^2 + 1/y^2$ is exactly the double angle formula for cosh or cos written multiplicatively. Recognising this form on sight is what lets you treat an algebraic recursion and a trigonometric double angle identity as the same fact seen twice. See Book 1, field: Complex numbers, topic: hyperbolic and circular functions.

Circuit Bank

- **WHEN** a recursion has the shape $t_k = t_{k-1}^2 - 2$ **DO** substitute $t_0 = y_0 + 1/y_0$ and conclude $t_k = y_0^{2^k} + y_0^{-2^k}$ **BECAUSE** $(y + 1/y)^2 - 2 = y^2 + 1/y^2$ turns squaring the recursion into squaring the auxiliary variable, doubling its exponent
- **WHEN** an infinite product is built from a sequence whose exponent doubles at every step **DO** hunt for auxiliary functions P, Q with (term)$= P(\text{next})/P(\text{current}) \cdot Q(\text{current})/Q(\text{next})$ **BECAUSE** such a certificate collapses the whole product to the two boundary values at the first and, in the limit, the last index
- **WHEN** you are asked for the value of an infinite product **DO** check first that $\sum(1 - \text{term}_k)$ converges **BECAUSE** this is exactly the condition for the product to converge to a nonzero limit, and skipping it leaves the closed form unjustified even if the arithmetic is right

- **WHEN** an expression of the form $y + 1/y$ appears **DO** recognise it as $2\cosh(\ln y)$, or $2\cos\theta$ when it lies in $[-2, 2]$ **BECAUSE** this connects any algebraic identity you find to the matching angle doubling identity, often easier to recall under pressure

Backtesting

1. First useful cue: the exact combination "square, then subtract 2", the algebraic signature of doubling under $y + 1/y$.
2. Method the expression suggested: solve $y_0 + 1/y_0 = a_0$ for y_0, then prove $a_k = y_0^{2^k} + y_0^{-2^k}$ by one line of induction.
3. Where you would hesitate, and the thought to reinforce: after the closed form, the jump to a telescoping certificate looks unmotivated; reinforce that the certificate reuses the exact identity $(t + 1/t)^2 - 1 = t^2 + 1 + 1/t^2$ that proved the recursion.
4. The trap the question sets: computing decimals, spotting $0.428571\ldots$, and stopping, without proving the limit is exactly $3/7$ rather than a value indistinguishable at finite precision.
5. What makes you recognise it faster next time: any recursion of the shape variable squared minus a constant should immediately suggest hunting for the matching sum-with-reciprocal substitution before anything numeric.
6. The reusable decision rule: an infinite product over a doubling-exponent sequence telescopes via a certificate built from the same identity that produced the closed form; find that identity once and reuse it twice.

Mirror drills

1. (drill) Let $b_0 = 3$ and $b_k = b_{k-1}^2 - 2$. Find b_k in closed form. *Answer:* solve $y_0 + 1/y_0 = 3$, giving $y_0 = \frac{3+\sqrt{5}}{2}$, so $b_k = y_0^{2^k} + y_0^{-2^k}$.
2. (drill) With $a_0 = 5/2$ as in the problem, compute $\prod_{k=1}^{\infty}\left(1 - \frac{1}{a_k}\right)$, starting the product one index later. *Answer:* the same telescoping now starts at $y_1 = y_0^2 = 4$, giving $Q(4)/P(4) = (4 - \frac{1}{4})/(4 + 1 + \frac{1}{4}) = \frac{15}{21} = \frac{5}{7}$; check it against the full product by dividing: $(3/7) \div (1 - 1/a_0) = (3/7) \div (3/5) = 5/7$.
3. (drill) Show $(2\cos\theta)^2 - 2 = 2\cos(2\theta)$, and hence that iterating $t \mapsto t^2 - 2$ on $t_0 = 2\cos\theta_0$ gives $2\cos(2^k\theta_0)$. *Answer:* $4\cos^2\theta - 2 = 2(2\cos^2\theta - 1) = 2\cos 2\theta$ by the double angle formula, then induct on k exactly as in Block 2 of Layer 2.

12.4 A4. Minimising the mass at zero from three moments

Suppose X is a random variable that takes on only nonnegative integer values, with $E[X] = 1$, $E[X^2] = 2$, and $E[X^3] = 5$. (Here $E[Y]$ denotes the expectation of the random variable Y.) Determine the smallest possible value of the probability of the event $X = 0$.

Read it out loud

Say the structure aloud: a distribution on the nonnegative integers, three prescribed power moments, minimise the mass sitting at the origin. Once you say that, the probabilistic costume falls off. The unknown is not a number and it is not a random variable in any analytic sense; the unknown is the sequence $p_0, p_1, p_2, \ldots$ of nonnegative reals, $p_n = P(X = n)$. Every piece of data is linear in that sequence, and so is the thing you are asked to minimise. This is a linear programme, and it is an infinite dimensional one only in appearance.

Within thirty seconds the page should carry this and nothing else:

$$\sum_{n\geq 0} p_n = 1, \quad \sum_{n\geq 0} np_n = 1, \quad \sum_{n\geq 0} n^2 p_n = 2, \quad \sum_{n\geq 0} n^3 p_n = 5, \quad p_n \geq 0, \quad \text{minimise } p_0.$$

The fields wired into the question are linear programming duality, the classical moment problem, and, underneath both, polynomial positivity on a discrete set. The word "integer" is not decoration. It is the entire content of the problem: you will see below that if X were allowed arbitrary nonnegative real values with the same three moments, the mass at 0 could be pushed to zero exactly. The integers are what stop it at $\frac{1}{3}$.

Now the cue that decides the whole solution. You are given the moments of degree 1, 2 and 3 (and degree 0, which is the total mass). So the set of test functions whose expectation you can compute exactly, without knowing anything else about X, is precisely the set of polynomials of degree at most 3. That is your entire toolkit, stated exactly. The quantity you want is $E[\mathbf{1}[X = 0]]$, the expectation of an indicator. Thus the problem, fully undisguised, is: find the cubic that hugs the indicator of $\{0\}$ from below as tightly as possible on $\mathbb{Z}_{\geq 0}$. Everything after this sentence is bookkeeping.

Approaches, ranked worst to best

1. **Guess distributions and compare.** *Cue:* "determine the smallest possible value" reads like a search. *Circuit:* no circuit at all, this is consumption. You can stumble onto $\left(\frac{1}{3}, \frac{1}{2}, 0, \frac{1}{6}\right)$ in a few minutes of fiddling, and then you have zero of the ten marks, because a minimum claimed without a proof of optimality is a guess. Ranked last because it feels like progress.

2. **Cauchy and Schwarz on the positive part.** *Cue:* a first moment, a second moment and a tail probability in the same problem. *Circuit:* the second moment tail bound. From $1 = E[X\mathbf{1}[X \geq 1]]^2 \leq E[X^2]P(X \geq 1)$ you get $p_0 \leq \frac{1}{2}$. Correct, useless: it bounds p_0 from above when you were asked for a bound from below, and it never touches the third moment, which is the only reason the answer is not 0.

3. **Lagrange multipliers on a truncated support.** *Cue:* constrained optimisation of a smooth objective. *Circuit:* the multiplier circuit. You must first prove the optimum is finitely supported, which is most of the work, and then you are doing by calculus what the next two methods do by inspection. It costs ink and it hides the inequality that actually carries the proof.

4. **Probability generating function with Taylor remainder.** *Cue:* an integer valued random variable, so encode it as a power series. *Circuit:* the generating function circuit: $G(z) = E[z^X]$ turns moments into derivatives at $z = 1$ and turns p_0 into $G(0)$. Here $G(1) = G'(1) = G''(1) = G'''(1) = 1$, since the factorial moments $E[X]$, $E[X(X-1)]$, $E[X(X-1)(X-2)]$ are 1, 1, 1.

Taylor to third order about $z = 1$ evaluated at $z = 0$ gives $G(0) = 1 - 1 + \frac{1}{2} - \frac{1}{6} + \frac{1}{24}G''''(c) = \frac{1}{3} + \frac{1}{24}G''''(c) \geq \frac{1}{3}$, because a power series with nonnegative coefficients has nonnegative derivatives on $[0, 1)$. This is a genuinely good route and it generalises instantly to more moments. It ranks second only because the one sided derivatives at $z = 1$ need an Abel type justification before you may write them, and a grader is entitled to want that line.

5. **Dual certificate: the Lagrange basis cubic.** *Cue:* the data are exactly the moments of degree 0 to 3, and the target is the expectation of an indicator. *Circuit:* linear programming duality made concrete: to bound a linear functional of an unknown measure, test the measure against a polynomial you are allowed to evaluate. Best. Four lines, no analysis, and the certificate writes itself from the interpolation nodes, as shown below.

The ideal solution

The answer is $\boxed{\frac{1}{3}}$.

Write $p_n = P(X = n)$. Since $E[X^3] = 5$ is finite, $\sum n^3 p_n$ converges, hence so do $\sum n^2 p_n$, $\sum n p_n$ and $\sum p_n$, and every expectation below is a finite linear combination of these four convergent sums.

Let

$$q(x) = -\frac{(x-1)(x-2)(x-3)}{6} = -\frac{x^3 - 6x^2 + 11x - 6}{6},$$

the unique cubic taking the value 1 at $x = 0$ and 0 at $x = 1, 2, 3$.

Lemma 12.4.1. *For every integer $n \geq 0$ we have $q(n) \leq \mathbf{1}[n = 0]$.*

Proof. At $n = 0$, $q(0) = -\frac{(-1)(-2)(-3)}{6} = 1$, so the inequality holds with equality. At $n = 1, 2, 3$ we have $q(n) = 0 \leq 0$. For an integer $n \geq 4$ the three factors $n - 1, n - 2, n - 3$ are all positive, so $(n-1)(n-2)(n-3) > 0$ and $q(n) < 0$. There are no other nonnegative integers. □

Taking expectations in the pointwise inequality of the lemma (both sides have finite expectation, and the inequality holds at every point of the support of X),

$$E\left[q(X)\right] \leq E\left[\mathbf{1}[X = 0]\right] = p_0.$$

The left hand side is computable from the data alone:

$$E\left[q(X)\right] = -\frac{E[X^3] - 6E[X^2] + 11E[X] - 6}{6} = -\frac{5 - 12 + 11 - 6}{6} = -\frac{-2}{6} = \frac{1}{3}.$$

Therefore $p_0 \geq \frac{1}{3}$ for every admissible X.

The bound is attained. Let X take the values $0, 1, 3$ with probabilities $\frac{1}{3}, \frac{1}{2}, \frac{1}{6}$ respectively. Then

$$\sum p_n = \tfrac{1}{3} + \tfrac{1}{2} + \tfrac{1}{6} = 1, \quad E[X] = \tfrac{1}{2} + \tfrac{3}{6} = 1, \quad E[X^2] = \tfrac{1}{2} + \tfrac{9}{6} = 2, \quad E[X^3] = \tfrac{1}{2} + \tfrac{27}{6} = 5,$$

so this X satisfies every hypothesis and has $P(X = 0) = \frac{1}{3}$. Hence the smallest possible value of $P(X = 0)$ is $\frac{1}{3}$, and it is a minimum, not merely an infimum.

Layer 1: WHAT

Block 1. Rename the unknown: the object being optimised is the sequence (p_n), and every constraint is linear in it.

Block 2. Identify the admissible test functions: the given data determine $E[f(X)]$ for exactly one family, the polynomials of degree at most 3.

Block 3. Restate the target: $p_0 = E\left[\mathbf{1}[X = 0]\right]$, so you want a cubic lying below $\mathbf{1}[n = 0]$ on $\mathbb{Z}_{\geq 0}$ with the largest possible expectation.

Block 4. Build that cubic by interpolation: force it to equal the indicator at $n = 0, 1, 2, 3$. That is the Lagrange basis polynomial q for the node 0 on the node set $\{0, 1, 2, 3\}$.

Block 5. Check the only thing interpolation does not give you for free: that $q(n) \leq 0$ persists for every integer $n \geq 4$. It does, because past the last node all three factors are positive and the sign in front is negative.

Block 6. Take expectations, read off $p_0 \geq \frac{1}{3}$.

Block 7. Produce a distribution attaining it, which the nodes have already told you where to look: the mass must sit where q is tight.

The nominal route runs: parametrise $p_0, p_1, p_2, \ldots$, truncate the support at some guessed level, solve four linear equations in several unknowns, vary the free parameters, and try to argue that the minimum over the truncation is the minimum overall. It produces pages, it produces the right number, and it never produces a proof, because the truncation was never justified.

The bifurcation point is the line where you decide what your unknown is. If you write "let the support be $\{0, a, b, c\}$", you are on the zombie route. If you write "the data give me $E[f(X)]$ for every cubic f, so let me choose f well", you are on the ideal route. The zombie route treats the distribution as the unknown to be pinned down; the ideal route treats the distribution as an adversary and hands it an inequality it cannot escape.

Layer 2: HOW

Block 1. Cue. The phrase "takes on only nonnegative integer values" plus three numerical expectations plus "smallest possible value". **Why natural.** The data are values of linear functionals on an unknown measure, and the target is another linear functional. Optimisation of a linear functional subject to linear constraints and a positivity constraint is a linear programme, and the reflex for a linear programme is always the same: do not chase the primal optimum, construct a dual certificate. **Link to the next block.** Once you accept you are constructing a certificate, the next question is forced: a certificate made of what?

Block 2. Cue. You are given moments of degree $0, 1, 2, 3$ and nothing else. **Why natural.** A certificate must be a function whose expectation you can evaluate without knowing X. That set is exactly $\{$polynomials of degree at most $3\}$, no larger and no smaller. Naming the admissible set precisely is the move; most candidates never write this sentence and so never know what they are allowed to use. **Link to the next block.** You now need to compare a cubic with the quantity you want, so write that quantity as an expectation too.

Block 3. Cue. p_0 is a probability, and every probability is the expectation of an indicator. **Why natural.** $p_0 = E[\mathbf{1}[X = 0]]$ puts the target and the toolkit in the same currency. Any cubic q with $q \leq \mathbf{1}[\,\cdot = 0\,]$ pointwise on $\mathbb{Z}_{\geq 0}$ gives $E[q(X)] \leq p_0$, and $E[q(X)]$ is a known number. The problem has become: maximise a known number over such cubics. **Link to the next block.** So construct the best such cubic.

Block 4. Cue. A cubic has four degrees of freedom, and you want it to agree with a $0/1$ pattern. **Why natural.** Four coefficients, four interpolation conditions. Prescribe $q(0) = 1$ and $q(1) = q(2) = q(3) = 0$: the cubic that matches the indicator on the first four nonnegative integers exactly. This is the natural greedy attempt, and it is the only one worth trying first, because a certificate can only be optimal if it is tight wherever the extremal distribution has mass, and the extremal distribution of a four constraint linear programme has at most four atoms. The interpolation nodes and the support of the optimum are the same object seen from the two sides of duality. **Link to the next block.** Interpolation guarantees the inequality at the nodes and says nothing beyond them, so check beyond them.

Block 5. **Cue.** The nodes stop at 3, the integers do not. **Why natural.** $q(n) = -\frac{(n-1)(n-2)(n-3)}{6}$ is written in factored form precisely so this check is one sentence: for $n \geq 4$ every factor is positive. Had the sign gone the other way, the certificate would have failed and you would have known immediately that the answer involved larger atoms. **Link to the next block.** The certificate is verified, so cash it in.

Block 6. **Cue.** A verified pointwise inequality between two functions of X. **Why natural.** Expectation is monotone. The arithmetic $5 - 12 + 11 - 6 = -2$ is the only computation in the problem. **Link to the next block.** A lower bound is half an answer. The question said "smallest possible value", which asserts attainment; you must supply it.

Block 7. **Cue.** Equality in the certificate requires the mass to sit where $q(n) = \mathbf{1}[n = 0]$, that is, on $\{0, 1, 2, 3\}$. **Why natural.** Solve the three moment equations on that support. The solution is forced: $p_2 = 0$, $p_1 = \frac{1}{2}$, $p_3 = \frac{1}{6}$, $p_0 = \frac{1}{3}$. The vanishing of p_2 is not a coincidence to be explained; duality only demands that the support be contained in the tight set, not equal to it. **Link to the next block.** Bound plus construction equals a complete solution.

Layer 3: WHY

The mathematical truth underneath is the duality between measures and functions. A measure you cannot see is pinned down only by the functionals you are allowed to integrate against it. Give me the moments up to degree d and you have given me the linear map $f \mapsto E[f(X)]$ on polynomials of degree at most d, and nothing else. Any bound on any other functional must therefore be obtained by sandwiching that functional between polynomials of degree at most d. This is the entire theory of the truncated moment problem in one sentence, and it is why the certificate is a polynomial and why its degree is 3.

The discernment mechanism, the thing to install so that you choose this route in the arena and not the truncation route: count the constraints. Four linear constraints, including the normalisation, mean the extreme points of the feasible set have at most four atoms; a linear objective is minimised at an extreme point; so the optimum is a four atom distribution, and the natural guess for where those atoms sit is the smallest available integers, because moments this small ($E[X] = 1$) forbid mass far out. That reasoning gives you the primal guess. Complementary slackness then tells you that the dual certificate must be tight at every atom of the optimum, and a cubic tight at four prescribed points is unique. In other words, the interpolation nodes were never a guess: they were dictated by the number of constraints.

Now the sentence the model answers never write. Why is the answer $\frac{1}{3}$ and not 0? Because of integrality, and only because of integrality. The certificate q satisfies $q(n) \leq 0$ for integers $n \geq 1$, but it fails badly on the reals: $q\left(\frac{1}{2}\right) = -\frac{(-\frac{1}{2})(-\frac{3}{2})(-\frac{5}{2})}{6} = \frac{5}{16} > 0$, and indeed q is positive across the whole interval $(0, 1)$, where all three factors are negative. So the same argument collapses for real valued X, and it should, because for real valued X the mass at 0 can be zero: let X take the two values $\frac{3 \pm \sqrt{5}}{2}$ with probabilities $\frac{5 \mp \sqrt{5}}{10}$. Those two values are the roots of $t^2 - 3t + 1$, and the moment recursion $m_{k+1} = 3m_k - m_{k-1}$ with $m_0 = 1$, $m_1 = 1$ gives $m_2 = 2$, $m_3 = 5$ exactly. All three moments match, and $P(X = 0) = 0$. The examiners chose the numbers $1, 2, 5$ so that a two point real distribution exists with no mass at the origin, and then restricted to the integers so that the nearest admissible atoms, 1 and 3, force one third of the mass down to zero. That is the trap and the design.

To train the discernment: every time you meet "given these moments, bound this probability", write down the two lists side by side, the functionals you know and the functional you want, and then interpolate. The training reps are the mirror drills below.

Revolving concepts

Moment. The number $E[X^k]$. The map $k \mapsto E[X^k]$ determines the distribution when the moments do not grow too fast (Carleman's condition), and for a bounded or finitely supported variable it always does. What orbits this term for the Putnam is the truncated version: finitely many moments never determine the distribution, they only determine a convex set of distributions, and every question of the form "how large or small can this be" is an optimisation over that convex set. See Book 1, field: Probability, topic: moment problems.

Factorial moment. $E[X(X-1)\cdots(X-k+1)]$, the natural moment for integer valued variables, because it is the kth derivative of the probability generating function at 1. Here $E[X] = 1$, $E[X(X-1)] = 2-1 = 1$, $E[X(X-1)(X-2)] = 5-6+2 = 1$: the data are much prettier in this basis than in the power basis, which is the signal that generating functions are close to the surface. See Book 1, field: Probability, topic: generating functions.

Probability generating function. $G(z) = E[z^X] = \sum_{n\geq 0} p_n z^n$, a power series with nonnegative coefficients, radius of convergence at least 1, $G(1) = 1$, $G(0) = p_0$, and $G^{(k)}(1)$ the kth factorial moment. Nonnegativity of coefficients gives nonnegativity of every derivative on $[0, 1)$, which is exactly the ingredient the Taylor remainder route needs. See Book 1, field: Generating functions, topic: probability generating functions.

Lagrange basis polynomial. For nodes $x_0, \ldots, x_d$ the polynomial $\ell_i(x) = \prod_{m\neq i} \frac{x-x_m}{x_i-x_m}$ is the unique polynomial of degree at most d with $\ell_i(x_j) = \delta_{ij}$. Our q is ℓ_0 for nodes $0, 1, 2, 3$. The orbiting fact worth carrying: interpolation is how you manufacture a polynomial with a prescribed $0/1$ pattern, and prescribed sign patterns beyond the nodes are then read off from the factored form. See Book 1, field: Polynomials, topic: interpolation.

Linear programming duality and complementary slackness. Minimising $c^{\mathsf{T}}p$ over $p \geq 0$ with $Ap = b$ has the dual: maximise $b^{\mathsf{T}}y$ over y with $A^{\mathsf{T}}y \leq c$. Weak duality is the single line we actually used; complementary slackness (the dual constraint is tight at every index where the primal has mass) is the navigational tool that told us where the nodes go. See Book 1, field: Inequalities, topic: linear programming duality.

The recursion in the background, generalised. A measure supported on the two roots of $t^2 - st + p = 0$ has moments satisfying the linear recursion $m_{k+1} = s\,m_k - p\,m_{k-1}$, because the roots satisfy $t^{k+1} = st^k - pt^{k-1}$ and expectation is linear. Generally, a measure supported on the r roots of a monic $f(t) = t^r - c_1 t^{r-1} - \cdots - c_r$ has $m_{k+r} = c_1 m_{k+r-1} + \cdots + c_r m_k$ for all $k \geq 0$, and conversely a moment sequence satisfying a linear recursion of order r with the right positivity is realised by an r atom measure, the atoms being the roots of the characteristic polynomial and the weights being solved from the first r moments. That is the general machine behind "find the extremal finitely supported distribution": build the characteristic polynomial from the moment recursion, take its roots as atoms. It is also the Hankel determinant story: a moment sequence comes from an r atom measure exactly when the Hankel matrices of size $r+1$ become singular. See Book 1, field: Sequences and recurrences, topic: linear recurrences and characteristic polynomials.

Circuit Bank

- **WHEN** you see a quantity bounded using only finitely many moments of an unknown measure **DO** write down the exact family of test functions the data allow, then hunt for the best member of that family **BECAUSE** finitely many moments determine a linear functional on exactly that family, so every valid bound factors through it
- **WHEN** you see a probability you must bound **DO** rewrite it as the expectation of an indicator **BECAUSE** it puts the target in the same currency as the data, which is the only way an inequality between them can be written
- **WHEN** you need a polynomial with a prescribed sign or value pattern on a discrete set **DO** interpolate at the pattern and keep the answer in factored form **BECAUSE** the factored form

makes the sign beyond the last node a one line check, and interpolation is the only construction with the right number of degrees of freedom

- **WHEN** you see a linear objective, linear equality constraints and a nonnegativity constraint **DO** count the equalities, conclude the optimum has at most that many atoms, and make the dual certificate tight exactly there **BECAUSE** extreme points of such a feasible set have support size at most the number of equalities, and complementary slackness ties the certificate to the support
- **WHEN** you prove a lower bound for an extremal problem **DO** immediately construct the object attaining it before moving on **BECAUSE** "smallest possible value" is two claims, and a bound alone is at best half the marks

Backtesting

1. First useful cue: three moments and nonnegative integer values, so the data are a linear functional on cubics and the answer is a linear programme.
2. Method the expression suggested: test the unknown distribution against the cubic that interpolates the indicator of $\{0\}$ at the nodes $0, 1, 2, 3$.
3. Where you would hesitate, and the thought to reinforce: at the moment of choosing the nodes. Reinforce this: four constraints means four atoms means four nodes; there is nothing to choose.
4. The trap the question sets: the numbers $1, 2, 5$ are exactly the moments of a two point real distribution with no mass at the origin, so any argument that does not use integrality must fail, and several natural inequalities (Cauchy and Schwarz among them) do not use it and give bounds in the wrong direction.
5. What makes you recognise it faster next time: the shape "$E[X^k]$ given for $k \le d$, bound a probability" is always the truncated moment problem, and the certificate is always a polynomial of degree d.
6. The reusable decision rule: to bound an unknown measure, do not describe the measure; hand it a test function it cannot beat.

Mirror drills

1. Let X be a nonnegative integer valued random variable with $E[X] = 1$ and $E[X^2] = 2$ only. Show that $P(X = 0)$ can be made arbitrarily small, so that the third moment is doing all the work in A4. *Answer:* for $N \ge 3$ take $p_1 = 1 - \frac{1}{N-1}$, $p_N = \frac{1}{N(N-1)}$, $p_0 = \frac{1}{N}$; both moments check, and $p_0 \to 0$.
2. Let X be nonnegative integer valued with $E[X] = 1$, $E[X^2] = 2$. Prove $P(X = 0) \le \frac{1}{2}$ and find the equality case. *Answer:* $1 = E[X\mathbf{1}[X \ge 1]]^2 \le E[X^2]P(X \ge 1)$ gives $P(X \ge 1) \ge \frac{1}{2}$; equality forces X constant on $\{X \ge 1\}$, hence $X \in \{0, 2\}$ with probabilities $\frac{1}{2}, \frac{1}{2}$.
3. Same three moments as A4, but now X may take arbitrary nonnegative real values. Show $P(X = 0) = 0$ is achievable. *Answer:* put the mass on the roots of $t^2 - 3t + 1$, that is $\frac{3 \pm \sqrt{5}}{2}$, with weights $\frac{5 \mp \sqrt{5}}{10}$; the recursion $m_{k+1} = 3m_k - m_{k-1}$ from $m_0 = m_1 = 1$ gives $m_2 = 2$ and $m_3 = 5$.

12.5 A5. Coprimality forced by disjoint root annuli

Let

$$P_n(x) = 1 + 2x + 3x^2 + \cdots + nx^{n-1}.$$

Prove that the polynomials $P_j(x)$ and $P_k(x)$ are relatively prime for all positive integers j and k with $j \neq k$.

Read it out loud

"A family of polynomials whose coefficients are $1, 2, 3, \ldots, n$, and the claim is that no two members share a factor." Two structural facts are audible in that sentence before you write anything. First, the coefficients are positive and strictly increasing. Second, the family is nested: P_j starts with the whole of P_i and then continues. Both facts will be load bearing, and the second one is the reason the problem is about a family rather than about one polynomial.

Coprimality is not really an algebraic condition here, it is a geometric one. Two polynomials with rational coefficients are relatively prime exactly when they have no common root in $\mathbb{C}$: a nonconstant common divisor has a complex root by the fundamental theorem of algebra, and conversely a common root z makes the minimal polynomial of z over $\mathbb{Q}$ a common divisor. So the task is: prove that the root sets of P_j and P_k are disjoint. That is a strong demand. You are not being asked to compute a single root; you are being asked to confine the roots of each member of an infinite family to a region so precisely that different members get disjoint regions.

The fields wired in are polynomial algebra (gcd, content, Gauss), complex root location, and elementary modulus inequalities. What is really being tested is whether you own a tool that converts coefficient data directly into root location. There is exactly one cheap such tool, and it is the theorem of Enestrom and Kakeya. If it is not installed, install it now; it is four lines to prove and it recurs.

Within thirty seconds you should have written: the roots of P_i satisfy $|z| \leq \frac{i-1}{i}$, and $P_j - P_i = x^i\left((i+1) + (i+2)x + \cdots\right)$.

Approaches, ranked worst to best

1. **Factorise the P_n or compute resultants.** *Cue:* "relatively prime" is a divisibility statement, so attack divisibility. *Circuit:* the school circuit: to show no common factor, find the factors. Worst by a distance. Nobody knows an easy proof that P_n is irreducible (Kedlaya records it as expected but apparently hard), and the resultant of P_i and P_j is a determinant of size $i + j - 2$ with no visible structure. This route consumes the entire session.

2. **Compute roots numerically and observe separation.** *Cue:* a claim about roots. *Circuit:* reconnaissance. Legitimate as reconnaissance, worthless as a proof: there are infinitely many pairs. It does tell you the truth quickly, namely that the roots of P_i sit well inside the unit disc and drift outwards as i grows, which is the shape of the real argument.

3. **Telescope to a closed form and separate curves.** *Cue:* the coefficients $1, 2, 3, \ldots$ are the derivative of a geometric series. *Circuit:* the generating function circuit. From $(1-x)^2 P_n(x) = 1 - (n+1)x^n + nx^{n+1}$, a root $z \neq 1$ of P_n lies on the curve $|z|^n\,|n + 1 - nz| = 1$. Showing these curves are pairwise disjoint means studying $E_z(t) = t\log|z| + \log|t + 1 - tz|$ as a function of a real variable t and proving it has at most one positive zero, via a concavity computation. Fully rigorous, roughly a page, and every line of it is calculus you must not slip in.

4. **One auxiliary real function and Rolle.** *Cue:* a common root would satisfy two equations of the same shape with different exponents. *Circuit:* the circuit: when several exponents satisfy one identity, interpolate the exponent and count zeros. Set $w = z^{-1}$; then $P_n(z) = 0$ forces $w^N = Nw - N + 1$ for $N = n + 1$, and this identity also holds trivially for $N = 0$ and

$N = 1$. Multiplying the identity by its conjugate and replacing the integer N by a real variable t produces

$$f(t) = |w|^{2t} - t^2|w|^2 + 2t(t-1)\mathrm{Re}(w) - (t-1)^2$$

which vanishes at the four values $t \in \{0, 1, i+1, j+1\}$, while $f'''(t) = (2\log|w|)^3|w|^{2t}$ never vanishes provided $|w| \neq 1$. Rolle then caps the number of zeros at three. Beautiful and it generalises (it proves the rank 3 statement about the matrix with rows 1, exponents, powers). It ranks second only because you must invent f under time pressure and you must dispose of $|w| = 1$ separately, a case the published solution passes over.

5. **Two disjoint annuli from Enestrom and Kakeya.** *Cue:* coefficients positive and increasing, and the family is closed under differences. *Circuit:* the annulus circuit: consecutive coefficient ratios bracket the modulus of every root. Best. After a four line lemma the problem dies in three lines, with no calculus and no case analysis beyond one degenerate index.

The ideal solution

We show that P_j and P_k have no common root in $\mathbb{C}$ when $j \neq k$; that is equivalent to the assertion. Indeed, if $d = \gcd(P_j, P_k)$ in $\mathbb{Q}[x]$ were nonconstant it would have a complex root, which would be a common root of P_j and P_k. (In $\mathbb{Z}[x]$ nothing changes: both polynomials have constant term 1, hence are primitive, so by Gauss's lemma their coprimality in $\mathbb{Q}[x]$ gives coprimality in $\mathbb{Z}[x]$.)

Lemma 12.5.1 (Enestrom and Kakeya)**.** *Let $f(x) = a_0 + a_1x + \cdots + a_dx^d$ with $d \geq 1$ and $0 < a_0 \leq a_1 \leq \cdots \leq a_d$. Then every complex root z of f satisfies $|z| \leq 1$.*

Proof. Expanding,

$$(x-1)f(x) = a_dx^{d+1} - \sum_{t=1}^{d}(a_t - a_{t-1})x^t - a_0.$$

Suppose $f(z) = 0$ with $|z| > 1$. Then $a_dz^{d+1} = \sum_{t=1}^{d}(a_t - a_{t-1})z^t + a_0$, and since $|z|^t \leq |z|^d$ for $1 \leq t \leq d$ and $1 \leq |z|^d$, and since every $a_t - a_{t-1} \geq 0$,

$$a_d|z|^{d+1} \leq \sum_{t=1}^{d}(a_t - a_{t-1})|z|^t + a_0 \leq |z|^d\left(\sum_{t=1}^{d}(a_t - a_{t-1}) + a_0\right) = a_d|z|^d.$$

Hence $|z| \leq 1$, a contradiction. □

Corollary 12.5.2. *Let $f(x) = a_0 + a_1x + \cdots + a_dx^d$ with $d \geq 1$ and every $a_t > 0$. Put*

$$r = \min_{0\leq t\leq d-1}\frac{a_t}{a_{t+1}}, \qquad R = \max_{0\leq t\leq d-1}\frac{a_t}{a_{t+1}}.$$

Then every complex root z of f satisfies $r \leq |z| \leq R$.

Proof. The polynomial $g(x) = f(Rx)$ has coefficients $b_t = a_tR^t$, which are positive and satisfy $b_t/b_{t+1} = \frac{a_t}{a_{t+1}R} \leq 1$, so $b_0 \leq b_1 \leq \cdots \leq b_d$. By the lemma every root w of g has $|w| \leq 1$, and the roots of f are the numbers Rw, giving $|z| \leq R$. For the other bound, let $f^*(x) = x^df(1/x) = a_d + a_{d-1}x + \cdots + a_0x^d$, of degree d since $a_0 \neq 0$. Its roots are exactly the reciprocals of the roots of f (no root of f is 0, again since $a_0 \neq 0$). Its consecutive coefficient ratios are the numbers a_{s+1}/a_s, whose maximum is $1/r$, so by the case already proved every root w of f^* has $|w| \leq 1/r$, that is, every root z of f has $|z| \geq r$. □

Now let $1 \leq i < j$ and suppose $z \in \mathbb{C}$ satisfies $P_i(z) = P_j(z) = 0$.
If $i = 1$ then $P_1(x) = 1$ has no root at all, so no such z exists. Assume from here that $i \geq 2$.

The polynomial P_i has $\deg P_i = i-1 \geq 1$ and coefficients $a_t = t+1 > 0$ for $0 \leq t \leq i-1$, with consecutive ratios $a_t/a_{t+1} = \frac{t+1}{t+2}$, increasing in t. The corollary therefore gives

$$|z| \leq \max_{0\leq t\leq i-2} \frac{t+1}{t+2} = \frac{i-1}{i}. \tag{1}$$

Next subtract. We have

$$P_j(x) - P_i(x) = \sum_{m=i+1}^{j} m x^{m-1} = x^i Q(x), \qquad Q(x) = \sum_{t=0}^{j-i-1} (i+1+t)\, x^t,$$

and z is a root of $P_j - P_i$. Since $P_i(0) = 1 \neq 0$ we have $z \neq 0$, so $Q(z) = 0$.

If $j = i+1$ then Q is the nonzero constant $i+1$, which has no roots: contradiction.

If $j \geq i+2$ then $\deg Q = j-i-1 \geq 1$ and Q has positive coefficients $b_t = i+1+t$ with consecutive ratios $b_t/b_{t+1} = \frac{i+1+t}{i+2+t}$, again increasing in t. The corollary gives

$$|z| \geq \min_{0\leq t\leq j-i-2} \frac{i+1+t}{i+2+t} = \frac{i+1}{i+2}. \tag{2}$$

But $i(i+1) = i^2+i > i^2+i-2 = (i-1)(i+2)$, so $\frac{i+1}{i+2} > \frac{i-1}{i}$, and (1) and (2) are incompatible. Contradiction.

Hence P_i and P_j have no common complex root, so $\gcd(P_j, P_k) = 1$ for all positive integers $j \neq k$.

Layer 1: WHAT

Block 1. Translate "relatively prime" into "no common complex root", with the fundamental theorem of algebra as the bridge and Gauss's lemma as the footnote for $\mathbb{Z}[x]$.

Block 2. Read the coefficient pattern of P_i: positive, increasing, ratios $\frac{t+1}{t+2}$. Conclude every root lies in the closed disc of radius $1 - \frac{1}{i}$.

Block 3. Use the nesting: $P_j - P_i = x^i Q(x)$, and Q is again a polynomial with positive increasing coefficients, only with the index shifted up by i.

Block 4. A common root is nonzero (constant term 1), so it is a root of Q, whose coefficient ratios all exceed $\frac{i+1}{i+2}$; hence its modulus is at least $1 - \frac{1}{i+2}$.

Block 5. Compare the two windows. They are disjoint because the ratio sequence $\frac{t+1}{t+2}$ is strictly increasing. Contradiction.

Block 6. Sweep the degenerate indices: $i = 1$ (no roots at all) and $j = i+1$ (the difference is a monomial).

The nominal route runs: expand both polynomials, run the Euclidean algorithm on P_i and P_j, watch the remainders lose all structure after one step, then attempt irreducibility of P_n and stall permanently. Notice that the nominal route and the ideal route share their first move, subtraction: $\gcd(P_i, P_j) = \gcd(P_i, P_j - P_i)$. The bifurcation point is what you do with $P_j - P_i$. The zombie divides again. The ideal solution stops dividing and asks where the roots of the two objects live, because the difference has landed back in the same class of polynomials with a shifted index, and a class with a shifted index means a shifted root window. Subtraction here is not a step of the Euclidean algorithm; it is a way of manufacturing a second polynomial whose roots are pinned to a different annulus.

Layer 2: HOW

Block 1. Cue. The words "relatively prime" applied to two explicit polynomials with integer coefficients. **Why natural.** Divisibility in $\mathbb{Q}[x]$ is invisible; roots are visible. The gcd of two polynomials is unchanged under field extension (the Euclidean algorithm does not know which field it is in), so a nonconstant gcd over $\mathbb{Q}$ survives over $\mathbb{C}$, where it must have a root. Coprimality

and disjoint root sets are therefore the same statement, and only one of the two is something you can measure. **Link to the next block.** Now you need to measure root sets, so look at what the question gives you about the roots. It gives you the coefficients and nothing else.

Block 2. Cue. Coefficients $1, 2, 3, \ldots, n$: all positive, monotonically increasing. **Why natural.** Positivity plus monotonicity is the exact hypothesis of the one classical theorem that turns coefficients into root location for free. The intuition it formalises is the balance principle: a root of $\sum a_t x^t$ occurs where two terms $a_t z^t$ and $a_{t+1} z^{t+1}$ are comparable in size, that is, where $|z| \approx a_t / a_{t+1}$; so all roots must live between the smallest and largest such ratio. Enestrom and Kakeya makes the approximate statement exact. **Link to the next block.** You now have a disc for P_i. One disc is not a contradiction, so you need a second, incompatible statement about the same z.

Block 3. Cue. Two members of a nested family and a hypothesis about both. **Why natural.** If z kills P_i and P_j it kills every linear combination, and in a nested family the cheapest combination is the difference, because it cancels the shared head and leaves the tail. The tail $\sum_{m=i+1}^{j} m x^{m-1}$ still has positive increasing coefficients, so the tool from Block 2 applies to it as well. This is the whole design of the problem: the family is chosen so that differences stay inside the class. **Link to the next block.** A common root of P_i and P_j is a root of the tail, but only after you strip the factor x^i.

Block 4. Cue. The tail is visibly divisible by x^i, and the constant term of P_i is 1. **Why natural.** Stripping x^i is legitimate only because $z \neq 0$, and $z \neq 0$ because $P_i(0) = 1$. Write that sentence; it costs six words and it is the sort of step whose absence a grader notices. Once stripped, Q has ratios starting at $\frac{i+1}{i+2}$, and the corollary now supplies a lower bound for $|z|$, which is the opposite kind of bound to Block 2. **Link to the next block.** Two bounds on the same modulus, in opposite directions, must be compared.

Block 5. Cue. An upper bound $\frac{i-1}{i}$ and a lower bound $\frac{i+1}{i+2}$. **Why natural.** The ratios $\frac{t+1}{t+2}$ increase strictly towards 1, so the head's largest ratio is strictly smaller than the tail's smallest ratio. In scale language: the roots of the head sit at a strictly smaller scale than the roots of the tail, because the head's coefficients grow proportionally faster than the tail's. The gap between $\frac{i-1}{i}$ and $\frac{i+1}{i+2}$ never closes, and does not need to be large. **Link to the next block.** Only the boundary indices remain.

Block 6. Cue. Every bound above needed $\deg \geq 1$. **Why natural.** $i = 1$ makes P_i constant and $j = i + 1$ makes Q constant. In both cases the conclusion is immediate and in both cases an argument that ignores them is incomplete. Sweeping degenerate indices is not pedantry on the Putnam; it is where marks go.

Layer 3: WHY

The truth under this problem is that the modulus of a root of a polynomial with positive coefficients is controlled by the consecutive coefficient ratios, and by nothing else. Look at $f(z) = \sum a_t z^t$ with all $a_t > 0$ and ask when cancellation is even possible. If $|z|$ is much smaller than every ratio a_t / a_{t+1}, the constant term dominates every other term put together and $f(z)$ cannot vanish; if $|z|$ is much larger than every ratio, the top term dominates. So roots can only live in the window bracketed by the smallest and largest ratio. This is the archimedean shadow of the Newton polygon, the same object that in the p-adic world tells you the valuations of the roots from the valuations of the coefficients; Kedlaya's remark that this problem also yields to a p-adic valuation argument is exactly that duality showing itself.

The discernment mechanism, the thing that makes you pick this route in the arena: ask what the question has given you that is unusual. It has given you a coefficient sequence that is monotone.

Monotone coefficients are worth nothing for algebra and everything for root location, so the question is a root location question. Then ask what varies as the index moves. The ratio sequence $\frac{t+1}{t+2}$ increases strictly, so different truncations of the same coefficient sequence have root windows that slide outwards. Two objects with disjoint windows cannot share a root. That is the whole proof, and it is visible before any lemma is written.

Say the generalisation out loud, because it is the real theorem and it costs nothing. Let $a_0, a_1, a_2, \ldots$ be positive reals whose consecutive ratios a_t/a_{t+1} are strictly increasing (equivalently, $a_t a_{t+2} < a_{t+1}^2$ for all t, strict log concavity). Then the truncations $F_n(x) = a_0 + a_1 x + \cdots + a_{n-1}x^{n-1}$ are pairwise relatively prime. The proof is the one above verbatim: the head's window is $\left[\frac{a_0}{a_1}, \frac{a_{i-2}}{a_{i-1}}\right]$, the tail's window starts at $\frac{a_i}{a_{i+1}}$, and strict monotonicity of the ratios makes $\frac{a_i}{a_{i+1}} > \frac{a_{i-2}}{a_{i-1}}$. Our case is $a_t = t+1$, where $(t+1)(t+3) = t^2+4t+3 < (t+2)^2$. Knowing this is what converts a one off Putnam problem into a circuit: you now recognise the entire class of families to which the argument applies, and you will recognise it the next time the coefficients are $1, 2, 4, 8, \ldots$ or $1!, 2!, 3!, \ldots$ or any strictly log concave sequence.

To train it: whenever a problem hands you an explicit coefficient sequence, compute the consecutive ratios before doing anything else. That single reflex converts most root location problems into arithmetic.

Revolving concepts

Relatively prime polynomials. $\gcd(f, g)$ is a nonzero constant. Orbiting facts: the gcd is computed by the Euclidean algorithm, so it does not change when the field is enlarged; coprimality is equivalent to the Bezout identity $uf + vg = 1$ with u, v in the same polynomial ring; and by Gauss's lemma, primitive polynomials in $\mathbb{Z}[x]$ that are coprime in $\mathbb{Q}[x]$ are coprime in $\mathbb{Z}[x]$. See Book 1, field: Polynomials, topic: gcd and Bezout.

Enestrom and Kakeya theorem. As proved above, together with its scaling corollary. The proof technique deserves its own slot: multiply by $(x-1)$ to telescope the coefficients into differences, then use the triangle inequality on the differences. That "multiply by $(x-1)$ to convert coefficients into consecutive differences" move recurs across the whole subject. See Book 1, field: Polynomials, topic: root location.

Reverse polynomial. $f^*(x) = x^{\deg f} f(1/x)$ reverses the coefficient list and inverts the roots. It is the standard device for converting an upper bound theorem into a lower bound theorem, and it costs one line. See Book 1, field: Polynomials, topic: transformations of roots.

Other root location tools. Cauchy's bound: every root of $x^d + c_{d-1}x^{d-1} + \cdots + c_0$ satisfies $|z| \leq 1 + \max_t |c_t|$. Rouché's theorem: if $|f - g| < |g|$ on a circle then f and g have the same number of zeros inside. Both are worth having; Rouché in particular proves the same disc statement here with more machinery. See Book 1, field: Complex numbers, topic: argument principle.

The closed form and the recursion, generalised. The family satisfies $P_{n+1}(x) = P_n(x) + (n+1)x^n$, a first order recursion in the index whose increment is a monomial; summing the recursion is exactly the telescoping that gives $P_j - P_i = \sum_{m=i+1}^{j} m x^{m-1}$. The closed form comes from differentiating a geometric sum: $P_n(x) = \frac{d}{dx}\left(1 + x + \cdots + x^n\right) = \frac{1-(n+1)x^n+nx^{n+1}}{(1-x)^2}$. Generalise the pattern: for any sequence (a_t), the truncation family $F_n = \sum_{t<n} a_t x^t$ satisfies $F_{n+1} - F_n = a_n x^n$, so a common root of two members is a root of a contiguous block $\sum_{i \leq t < j} a_t x^t$, and if the block is again a member of a class with known root windows you get a contradiction for free. When (a_t) itself satisfies a linear recursion with characteristic polynomial χ, the generating function is $\frac{\text{numerator}}{\chi^*(x)}$ and every truncation is a numerator of a partial sum, which gives closed forms of the same shape as the one above. See Book 1, field: Sequences and recurrences, topic: generating functions of truncations.

Interpolating an exponent. The route ranked fourth turns the integer exponent n into a real variable t and applies Rolle to bound the number of solutions. The orbiting statement, worth memorising: an equation $e^x = P(x)$ with P real of degree d has at most $d+1$ real solutions, since differentiating

$d+1$ times kills P; more generally, counting zeros by repeated differentiation is the elementary end of the theory of o-minimality. See Book 1, field: Real analysis, topic: zero counting by Rolle.

Circuit Bank

- **WHEN** you are asked to prove two polynomials are relatively prime **DO** convert immediately to "no common complex root" and start locating roots **BECAUSE** gcd is invisible but roots are measurable, and the Euclidean algorithm does not care which field you work over, so the two statements are identical
- **WHEN** you see a polynomial with all coefficients positive **DO** compute the consecutive ratios and quote the annulus theorem **BECAUSE** a root can only occur where two consecutive terms balance, so every root modulus lies between the smallest and the largest consecutive ratio
- **WHEN** you have two members of a nested family sharing a hypothesis **DO** subtract them and identify the class the difference belongs to **BECAUSE** a common root kills every linear combination, and in a nested family the difference cancels the shared head and hands you a second object with a shifted index
- **WHEN** you strip a factor of x^m from an equation **DO** first say why the root is nonzero, citing the constant term **BECAUSE** stripping x^m is division by zero unless $z \neq 0$, and the constant term of the other polynomial is usually the cheapest reason
- **WHEN** you have an upper bound and a lower bound for the same quantity from two different sources **DO** write both, then compare with one cross multiplication **BECAUSE** contradiction arguments in root location almost always take this shape, and the gap needed is never large

Backtesting

1. First useful cue: the coefficients are positive and increasing, so the problem is about where the roots sit, not about factorisation.
2. Method the expression suggested: bound the roots of P_i above by the largest consecutive coefficient ratio, then bound the roots of the tail $P_j - P_i$ below by its smallest consecutive ratio.
3. Where you would hesitate, and the thought to reinforce: after subtracting, when the difference is not another member of the family. Reinforce this: it does not need to be a member, it only needs to be in the same class, and the class is "positive increasing coefficients".
4. The trap the question sets: it invites you to prove something about P_n alone (irreducibility), which is hard and unnecessary. The statement is about pairs, so the argument must use both indices at once.
5. What makes you recognise it faster next time: coefficient sequences with strictly increasing consecutive ratios have truncations with disjoint root annuli, hence pairwise coprime truncations. This is the general theorem; $1, 2, 3, \ldots$ is one instance.
6. The reusable decision rule: coefficients to ratios, ratios to annuli, annuli to disjointness.

Mirror drills

1. Prove the Enestrom and Kakeya lemma from scratch in four lines. *Answer:* multiply by $(x-1)$ to get $a_d x^{d+1} - \sum_{t\geq 1}(a_t - a_{t-1})x^t - a_0$; if $|z| > 1$ the triangle inequality bounds the tail by $a_d|z|^d$, forcing $a_d|z|^{d+1} \leq a_d|z|^d$.
2. Locate all complex roots of $1 + x + 2x^2 + 3x^3 + 5x^4 + 8x^5$ inside an annulus. *Answer:* the consecutive ratios are $1, \frac{1}{2}, \frac{2}{3}, \frac{3}{5}, \frac{5}{8}$, so every root satisfies $\frac{1}{2} \leq |z| \leq 1$.

3. Show without any root location that P_k and P_{k+1} are relatively prime. *Answer:* $P_{k+1} - P_k = (k+1)x^k$, so a common root would have to be 0, but $P_k(0) = 1$.

12.6 A6. Diagonal nonvanishing bounded by a tensor pairing

Let n be a positive integer. What is the largest k for which there exist $n \times n$ matrices $M_1, \dots, M_k$ and $N_1, \dots, N_k$ with real entries such that for all i and j, the matrix product $M_i N_j$ has a zero entry somewhere on its diagonal if and only if $i \neq j$?

Read it out loud

"Two families of square matrices, and the only thing anyone ever looks at is the diagonal of a product; the diagonal must be nowhere zero on the matched pairs and must have a zero somewhere on every mismatched pair; how long can the families be?" Say that and then say the one identity that governs everything: the tth diagonal entry of $M_i N_j$ is

$$(M_i N_j)_{tt} = (\text{row } t \text{ of } M_i) \cdot (\text{column } t \text{ of } N_j).$$

Nothing else about the matrices is ever used. So M_i is not a matrix in this problem, it is an ordered list of n vectors in $\mathbb{R}^n$, and N_j is another ordered list of n vectors in $\mathbb{R}^n$, and the matrix product is a red herring dressed as structure. Invertibility is never required, determinants never appear, and you should register both facts as permissions rather than as omissions.

Restated in the honest language: for each $t \in \{1, \dots, n\}$ you have a pairing, call it channel t, between the vectors $r_{i,t}$ (rows) and $c_{j,t}$ (columns). The condition is that all n channels agree simultaneously exactly on the diagonal $i = j$. Each channel is a bilinear form on $\mathbb{R}^n \times \mathbb{R}^n$, and there are n of them, and the conditions are combined by an AND. An AND of nonvanishing conditions is a single nonvanishing condition on a product. That sentence is the problem.

The fields wired in are linear algebra (dimension and independence), multilinear algebra (tensor powers), and the linear algebra method of extremal combinatorics. What is being tested is whether you know that the maximum size of a family carrying a dual family is a dimension, and whether you can find the right space.

Within thirty seconds you should have written $(M_i N_j)_{tt} = r_{i,t} \cdot c_{j,t}$, the reformulation $\prod_{t=1}^{n} r_{i,t} \cdot c_{j,t} \neq 0 \iff i = j$, and the small cases $n = 1 \Rightarrow k = 1$, $n = 2 \Rightarrow k = 4$.

Approaches, ranked worst to best

1. **Rank and determinant bookkeeping on the M_i alone.** *Cue:* the objects are matrices, so use matrix invariants. *Circuit:* the linear algebra reflex aimed at the wrong space. There are only n^2 entries in each matrix, so this route generates bounds that are polynomial in n, or no bound at all. It cannot reach an exponential answer, and everything it computes is invariant under changes of basis that the problem is not invariant under (the product of diagonal entries is not a similarity invariant). Worst.

2. **Colour each mismatched pair by a witnessing diagonal position and run pigeonhole or Ramsey.** *Cue:* "has a zero entry somewhere" is a disjunction over n positions. *Circuit:* the extremal combinatorics reflex. For each $i \neq j$ choose $t(i,j)$ with $(M_i N_j)_{t(i,j)\,t(i,j)} = 0$; you now have an n colouring of ordered pairs. This is a real idea, and inside one colour class the vectors $r_{i,t}$ and $c_{j,t}$ do form a useful configuration. But the colour classes interact, and reassembling them multiplicatively is exactly what the tensor product does properly; done by hand it stalls or produces a far weaker bound.

3. **Small cases and pattern matching.** *Cue:* "what is the largest k" with n free. *Circuit:* the reconnaissance circuit. $n = 1$ forces $k = 1$ (if $m_1 n_1 \neq 0$ and $m_2 n_2 \neq 0$ then $m_1 n_2 \neq 0$), and building $n = 2$ by hand gives 4. Then $1, 4, \dots$ against n suggests n^n rather than n^2 only if you also test $n = 2$ against the alternatives, so this is not decisive alone. It is essential, though: it tells you the answer is exponential, and it tells you which space to look for.

4. **Product of the diagonal entries, then a dual family in the tensor power.** *Cue:* an AND of n nonvanishing bilinear conditions, and an answer that must be exponential. *Circuit:* the linear algebra method: bound a family by embedding it in a vector space and exhibiting a dual family that pairs nonzero exactly on the diagonal. Best, and it is the only route here that is short. It also hands you the construction, because the proof shows exactly where the bound can be tight.

The ideal solution

The answer is $\boxed{n^n}$.

Throughout, $e_1, \dots, e_n$ is the standard basis of $\mathbb{R}^n$, and $\mathcal{A} = \{1, \dots, n\}^n$ is the set of n tuples $A = (a_1, \dots, a_n)$ with entries in $\{1, \dots, n\}$, so $|\mathcal{A}| = n^n$.

Step 1: the condition, restated. For any $n \times n$ matrices M and N, the tth diagonal entry of MN is $\sum_{a=1}^{n} M_{ta} N_{at}$. Since a product of real numbers vanishes exactly when one factor vanishes, the requirement of the problem is equivalent to

$$D(i,j) := \prod_{t=1}^{n} (M_i N_j)_{tt} \neq 0 \iff i = j. \tag{$*$}$$

Step 2: the construction, giving $k = n^n$. For $A \in \mathcal{A}$ let M_A be the matrix whose tth row is e_{a_t}, and let $N_A = M_A^{\mathsf{T}}$, so that the tth column of N_A is e_{a_t}. Then for $A, B \in \mathcal{A}$,

$$(M_A N_B)_{tt} = e_{a_t} \cdot e_{b_t} = \begin{cases} 1 & a_t = b_t \\ 0 & a_t \neq b_t. \end{cases}$$

If $A = B$ every diagonal entry equals 1, so there is no zero on the diagonal. If $A \neq B$ then $a_t \neq b_t$ for some t, and that diagonal entry is 0. So the n^n pairs (M_A, N_A) satisfy the requirement, indexed by $\mathcal{A}$.

Step 3: the upper bound. Suppose $M_1, \dots, M_k$ and $N_1, \dots, N_k$ satisfy $(*)$. Work in the space $V = \mathbb{R}^{\mathcal{A}}$ of real vectors indexed by $\mathcal{A}$, with the standard inner product; $\dim V = n^n$. For each i define $u_i, v_i \in V$ coordinatewise by

$$u_i(A) = \prod_{t=1}^{n} (M_i)_{t,a_t}, \qquad v_i(A) = \prod_{t=1}^{n} (N_i)_{a_t,t} \qquad (A = (a_1, \dots, a_n) \in \mathcal{A}).$$

Then for all i, j,

$$u_i \cdot v_j = \sum_{A \in \mathcal{A}} \prod_{t=1}^{n} (M_i)_{t,a_t} (N_j)_{a_t,t} = \prod_{t=1}^{n} \left(\sum_{a=1}^{n} (M_i)_{t,a} (N_j)_{a,t} \right) = \prod_{t=1}^{n} (M_i N_j)_{tt} = D(i,j).$$

The middle equality is the distributive law: expanding a product of n sums of n terms each produces exactly one term for each way of choosing an index a_t from the tth sum, that is, one term for each $A \in \mathcal{A}$.

By $(*)$ we have $u_i \cdot v_j = 0$ for $i \neq j$ and $u_i \cdot v_i \neq 0$. Suppose $\sum_{i=1}^{k} c_i u_i = 0$ with $c_i \in \mathbb{R}$. Taking the inner product with v_j gives

$$0 = \sum_{i=1}^{k} c_i (u_i \cdot v_j) = c_j (u_j \cdot v_j),$$

and since $u_j \cdot v_j \neq 0$ we get $c_j = 0$, for every j. Hence $u_1, \dots, u_k$ are linearly independent in V, so

$$k \leq \dim V = n^n.$$

Combining Steps 2 and 3, the largest such k is n^n.

Layer 1: WHAT

Block 1. Strip the matrices: only $(M_iN_j)_{tt} = r_{i,t} \cdot c_{j,t}$ matters, so the data are n independent channels of vector pairings.

Block 2. Replace the disjunction "some diagonal entry is zero" by the single scalar equation "the product of the diagonal entries is zero". One number $D(i,j)$ now carries the whole hypothesis.

Block 3. Recognise the shape of $(*)$: a family and a dual family pairing to nonzero exactly on the diagonal. That shape always caps the family size by the dimension of the ambient space.

Block 4. Notice the obstruction: D is not bilinear in (M_i, N_j), it is a product of n bilinear forms, so there is no ambient space yet. Manufacture one by expanding the product: the monomials are indexed by tuples $A \in \{1,\dots,n\}^n$, so the ambient space is $\mathbb{R}^{n^n}$.

Block 5. Run the standard independence argument: $k \le n^n$.

Block 6. Achieve the bound by making each index a codeword A and each channel a test for agreement in one coordinate, using $e_a \cdot e_b = \delta_{ab}$.

The nominal route: try to build large families by hand, discover n^2 or 2^n examples, then try to bound k with rank arguments about the M_i, since the M_i live in an n^2 dimensional space and that is the only dimension count in sight. It fails permanently, because the truth is exponential and the space of matrices is not the space the problem lives in.

The bifurcation point is Block 2, and it is one keystroke wide. Faced with "has a zero entry somewhere on its diagonal", the zombie keeps the disjunction and starts case work over which position vanishes. The ideal solution multiplies the diagonal entries together, because over a field the statement "one of these is zero" and the statement "their product is zero" are the same statement, and the second one is a single algebraic equation while the first is a case analysis of size n. Every good thing afterwards is downstream of that one move.

Layer 2: HOW

Block 1. **Cue.** The hypothesis mentions only the diagonal of the product. **Why natural.** The diagonal entry $(M_iN_j)_{tt}$ depends only on row t of M_i and column t of N_j. Different t involve disjoint data. So the problem decomposes into n channels which share nothing except the index set $\{1,\dots,k\}$. Recognising a decomposition into independent factors is what makes an exponential answer plausible: independent factors multiply. **Link to the next block.** If the channels are independent, the way to combine them is a product, not a case split.

Block 2. **Cue.** "Has a zero entry somewhere" is a disjunction. **Why natural.** Over a field, a product vanishes precisely when a factor does. Turning n conditions into one is not cosmetic: it converts the hypothesis into an algebraic identity that can be manipulated, and it makes $(*)$ symmetric and finite. Any time a hypothesis is an OR of equalities, multiply; any time it is an AND of equalities, add squares or stack into a vector. **Link to the next block.** The resulting statement, $D(i,j) \neq 0 \iff i = j$, is a shape you must recognise instantly.

Block 3. **Cue.** A number attached to every ordered pair of indices, nonzero exactly on the diagonal. **Why natural.** This is the signature of a dual pair of families. The reusable fact: if $u_i \cdot v_j = 0$ for $i \neq j$ and $u_i \cdot v_i \neq 0$, then the u_i are linearly independent, because pairing a vanishing linear combination against v_j isolates c_j. So the family size is at most the dimension of whatever space the u_i live in. This is the linear algebra method, and it is the standard weapon whenever a combinatorial family is constrained by an "exactly on the diagonal" condition. **Link to the next block.** The lemma needs vectors. You have matrices and a degree n expression. So build the vectors.

Block 4. **Cue.** $D(i,j)$ is a product of n bilinear forms, hence multilinear of degree n in each argument. **Why natural.** A product of linear things is linear on a bigger space: this is the entire

purpose of the tensor product. Concretely you do not need the abstract construction, you need its coordinates. Expand $\prod_t \left(\sum_a (M_i)_{t,a}(N_j)_{a,t}\right)$ by the distributive law and look at what indexes the resulting monomials: one choice of a_t per factor, that is, a tuple $A \in \{1,\dots,n\}^n$. Each monomial factors as (something depending only on M_i) times (something depending only on N_j). That factorisation is the definition of u_i and v_j, and it turns the multilinear D into an honest inner product on $\mathbb{R}^{n^n}$. **Link to the next block.** With the vectors in hand, Block 3 applies verbatim.

Block 5. Cue. Linear independence in a space of known dimension. **Why natural.** $k \le \dim V = n^n$. Note that the argument never used any property of the matrices beyond $(*)$, which is the sign that the bound is structural rather than accidental. **Link to the next block.** A bound is half an answer; now saturate it.

Block 6. Cue. The bound is $n^n = |\{1,\dots,n\}^n|$, and the upper bound proof identified the coordinates with tuples. **Why natural.** The proof itself tells you where equality can hold: the u_i must be n^n independent vectors, and the cheapest independent set is a basis, so aim for u_i to be the standard basis vectors of $\mathbb{R}^{\mathcal{A}}$. That happens exactly when each M_i has rows drawn from $\{e_1,\dots,e_n\}$, so that the product $\prod_t (M_i)_{t,a_t}$ is 1 for one tuple and 0 for the rest. Read backwards: index the pairs by tuples, let channel t test agreement in coordinate t, and use orthonormality $e_a \cdot e_b = \delta_{ab}$ to make agreement detectable. So the construction is not a separate act of invention; it is the equality case of the bound.

Layer 3: WHY

Two mathematical truths carry this problem, and they are worth more than the problem.

The first is the linear algebra method: to bound the size of a family of combinatorial objects, map the objects to vectors in a space of computable dimension and prove the vectors are independent; the bound is the dimension. The most common way to prove independence is to exhibit a dual family with $u_i \cdot v_j$ nonzero exactly on the diagonal, and that is precisely what the hypothesis of this problem hands you for free. The archetype in the literature is the set pair inequality of Bollobas, which follows the same skeleton. Once you have met three examples of this method you stop seeing "largest k" questions as combinatorics and start seeing them as dimension counts.

The second is that products of linear forms become linear on tensor powers. A bilinear form on $\mathbb{R}^n \times \mathbb{R}^n$ is an inner product after a change of coordinates; a product of n such forms, applied to n independent pairs of vectors, is an inner product on the n fold tensor power, whose dimension is n^n. The coordinates of a pure tensor $w_1 \otimes \cdots \otimes w_n$ are $\prod_t (w_t)_{a_t}$, which is exactly the $u_i(A)$ we wrote down, and the pairing satisfies $\langle w_1 \otimes \cdots \otimes w_n,\ y_1 \otimes \cdots \otimes y_n\rangle = \prod_t \langle w_t, y_t\rangle$. If you prefer to avoid the word tensor entirely, the distributive law does the same job in one line, which is what the ideal solution does. Naming the concept is still worth it, because the name is what lets you predict the answer before proving it.

Now the discernment mechanism, which is the part you must be able to run in the arena. You compute $n = 1$ and get 1; you build $n = 2$ by hand and get 4. Two data points do not determine the pattern, but they do determine the search: the answer is at least 4 for $n = 2$, so the ambient space you are looking for has dimension at least 4, and the natural spaces functorially built out of $\mathbb{R}^n$ are few. The nth tensor power has dimension n^n; the nth symmetric power has dimension $\binom{2n-1}{n}$; the nth exterior power has dimension 1; the space of matrices has dimension n^2. For $n = 2$ these are 4, 3, 1, 4. The construction rules out the exterior power and the symmetric power, and the difference between n^2 and n^n shows at $n = 3$, where the tuple construction gives 27 and the matrix space would allow only 9. So the answer names its own space, and the space tells you the proof. That is the discernment: for an extremal problem with linear algebraic data and an exponential answer, identify the space whose dimension matches the small cases, then find the map into it.

The knowledge gaps, stated without cosmetics, because this is an A6. If tensor products and their dimension are not installed, you can still write the construction and state the answer, and you will have a partial solution and nothing more; the upper bound is the entire difficulty. If the dual family lemma is not installed, you will not know that the hypothesis $D(i,j) \neq 0 \iff i = j$ is worth anything. Install, in this order: dual families and independence; the distributive expansion of a product of sums; tensor powers and their dimension; then the linear algebra method as a named strategy. None of it is deep; all of it is beyond A level and beyond NOIR PREMIERE.

Revolving concepts

Diagonal of a product. $(MN)_{tt} = \sum_a M_{ta} N_{at}$, the pairing of row t with column t. Its sum over t is the trace, which is a similarity invariant; its product over t is not an invariant of anything, which is the structural signal that this problem is combinatorial and not representation theoretic, and that basis changing manoeuvres will not help. See Book 1, field: Matrices and determinants, topic: products and traces.

Dual (biorthogonal) families. If $u_1, \dots, u_k$ and $v_1, \dots, v_k$ lie in an inner product space with $u_i \cdot v_j = 0$ for $i \neq j$ and $u_i \cdot v_i \neq 0$, then each family is linearly independent and $k \leq \dim$. Orbiting facts: a basis and its dual basis are the equality case; the same argument in a bilinear (not necessarily symmetric or definite) pairing works verbatim; the rank of the matrix $\left[u_i \cdot v_j\right]$ is k because it is diagonal with nonzero diagonal, and rank is at most the dimension of either space. See Book 1, field: Linear algebra, topic: independence via a dual family.

Tensor power. $V^{\otimes n}$ has basis $e_{a_1} \otimes \cdots \otimes e_{a_n}$ indexed by $A \in \{1, \dots, n\}^n$, so $\dim(\mathbb{R}^n)^{\otimes n} = n^n$; the pure tensor $w_1 \otimes \cdots \otimes w_n$ has A coordinate $\prod_t (w_t)_{a_t}$; and the induced pairing multiplies the factorwise pairings. The comparison points worth holding are $\dim \operatorname{Sym}^n(\mathbb{R}^n) = \binom{2n-1}{n}$ and $\dim \Lambda^n(\mathbb{R}^n) = 1$. See Book 1, field: Linear algebra, topic: multilinear algebra.

Distributive expansion. $\prod_{t=1}^{n} \left(\sum_{a=1}^{n} x_{t,a}\right) = \sum_{A \in \{1,\dots,n\}^n} \prod_{t=1}^{n} x_{t,a_t}$. This is the elementary content of the tensor identity and the reason the exponent is n^n: one free choice per factor. See Book 1, field: Combinatorics, topic: product principle.

The multiplicativity recursion, generalised. The structure here is a product of independent channels, and independent channels multiply the count. State it as a recursion: let $K(m)$ be the largest family of pairs of m tuples of vectors in $\mathbb{R}^n$ with $\prod_{t \leq m} r_{i,t} \cdot c_{j,t} \neq 0 \iff i = j$. Adjoining one more channel multiplies the ambient dimension by n and, by the same tensor argument, gives $K(m) \leq n\,K(m-1)$, with $K(0) = 1$ and equality achievable at every step by appending one more coordinate to the codeword. Hence $K(m) = n^m$, and the problem is $m = n$. The same recursion proves the rectangular generalisation Elkies notes: if the M_i are $m \times n$ and the N_j are $n \times m$, the diagonal has m entries, the ambient space is $(\mathbb{R}^n)^{\otimes m}$, and the answer is n^m. Generalise once more: if channel t pairs $\mathbb{R}^{d_t}$ with $\mathbb{R}^{d_t}$, the answer is $\prod_t d_t$, because the tensor product of the channels has that dimension and codewords over alphabets of sizes d_t realise it. See Book 1, field: Combinatorics, topic: product constructions and dimension bounds.

The linear algebra method. Named strategy: bound a combinatorial family by exhibiting an injective map to independent vectors. Its most quoted instance is the set pair inequality: if (A_i, B_i) are pairs of sets with $A_i \cap B_i = \emptyset$ and $A_i \cap B_j \neq \emptyset$ for $i \neq j$, with $|A_i| = a$ and $|B_i| = b$, then the number of pairs is at most $\binom{a+b}{a}$. Note the identical skeleton: a condition that holds exactly off the diagonal, converted into independence. See Book 1, field: Combinatorics, topic: the linear algebra method.

Circuit Bank

- **WHEN** a hypothesis says "some one of these quantities is zero" **DO** multiply them and say "the product is zero" **BECAUSE** over a field the two statements are identical, and one algebraic equation can be manipulated while a disjunction of n cases can only be split

- **WHEN** you meet a quantity attached to ordered pairs of indices that is nonzero exactly when the indices agree **DO** find vectors realising it as an inner product and conclude the family is linearly independent **BECAUSE** pairing a vanishing linear combination with the dual vector isolates one coefficient, so the family size is at most the dimension of the space
- **WHEN** the quantity you want to realise as an inner product is a product of several bilinear forms **DO** expand the product by the distributive law and index the resulting monomials by tuples **BECAUSE** a product of linear forms is linear on the tensor power, and the monomial index set is exactly a basis of that power
- **WHEN** an extremal answer looks exponential in n and the data are linear algebraic **DO** compute the small cases and match the numbers against n^n, n^2, $\binom{2n-1}{n}$, 2^n **BECAUSE** the bound will be the dimension of a space built from $\mathbb{R}^n$, and matching the small case identifies which space before you prove anything
- **WHEN** a question asks "what is the largest k" **DO** write the answer on line one, then a construction, then a matching bound **BECAUSE** these are two independent proofs and a grader marks them independently, so mixing them costs marks even when the mathematics is right

Backtesting

1. First useful cue: only diagonal entries appear, and $(M_i N_j)_{tt}$ couples row t with column t alone, so the problem is n independent channels rather than a matrix problem.
2. Method the expression suggested: multiply the n diagonal entries into one scalar, realise that scalar as an inner product of two vectors in a space of dimension n^n, and read off independence.
3. Where you would hesitate, and the thought to reinforce: at the moment the product of bilinear forms refuses to be bilinear. Reinforce this: products linearise on tensor powers, and in coordinates that is nothing but expanding brackets.
4. The trap the question sets: the word matrix. It invites determinant, rank and invertibility reflexes, all of which are useless because the product of diagonal entries is not invariant under any change of basis.
5. What makes you recognise it faster next time: the pattern "nonzero exactly on the diagonal" is a dual family, and dual families are always bounded by a dimension; the only work is identifying the space.
6. The reusable decision rule: an AND of nonvanishing conditions is one nonvanishing product, and a product of pairings is a pairing on the tensor product.

Mirror drills

1. Write out four explicit pairs of 2×2 matrices achieving $k = 4$ for $n = 2$. *Answer:* index by $(a_1, a_2) \in \{1,2\}^2$, let M have rows e_{a_1}, e_{a_2} and $N = M^{\mathsf{T}}$; the four matrices are I, $\left(\begin{smallmatrix} 1 & 0 \\ 1 & 0 \end{smallmatrix}\right)$, $\left(\begin{smallmatrix} 0 & 1 \\ 0 & 1 \end{smallmatrix}\right)$ and $\left(\begin{smallmatrix} 0 & 1 \\ 1 & 0 \end{smallmatrix}\right)$.
2. Let $u_1, \dots, u_k$ and $v_1, \dots, v_k$ be vectors in $\mathbb{R}^d$ with $u_i \cdot v_j = 0$ for $i \neq j$ and $u_i \cdot v_i \neq 0$. Prove $k \leq d$. *Answer:* if $\sum c_i u_i = 0$, pair with v_j to get $c_j(u_j \cdot v_j) = 0$, so all $c_j = 0$; independent vectors in $\mathbb{R}^d$ number at most d.
3. Solve the same problem when the M_i are $m \times n$ and the N_j are $n \times m$, so that $M_i N_j$ is $m \times m$. *Answer:* n^m, by codewords in $\{1, \dots, n\}^m$ and the same pairing in $(\mathbb{R}^n)^{\otimes m}$.

12.7 B1. Unique over-expansions and the missing zero digit

A base 10 *over-expansion of a positive integer* N *is an expression of the form*

$$N = d_k 10^k + d_{k-1} 10^{k-1} + \cdots + d_0 10^0$$

with $d_k \neq 0$ *and* $d_i \in \{0, 1, 2, \ldots, 10\}$ *for all* i*. For instance, the integer* $N = 10$ *has two base* 10 *over-expansions:* $10 = 10 \cdot 10^0$ *and the usual base* 10 *expansion* $10 = 1 \cdot 10^1 + 0 \cdot 10^0$*. Which positive integers have a unique base* 10 *over-expansion?*

Read it out loud

This is a positional representation question wearing a uniqueness classification mask. The field underneath is elementary number theory, the arithmetic of base ten place value, but the command term is not compute or find, it is which N, so the real task is a full classification proved as an if and only if. The structural feature the examiners planted is not subtle: they enlarged the legal digit alphabet from the usual $\{0, 1, \ldots, 9\}$ to $\{0, 1, \ldots, 10\}$, a single extra value at the top, then handed you a worked example of exactly what that extra value buys you. Read that example as a gift, not as decoration.

Within the first thirty seconds you should have written down the exchange this extra digit enables: $10 \cdot 10^i = 1 \cdot 10^{i+1}$, one unit at position $i + 1$ traded for ten units at position i. That trade produces a legal alternative expansion only when position i currently reads 0, so there is room to receive ten units without breaking the digit bound, and position $i + 1$ currently reads at least 1, so there is a unit there to give up. Look again at the worked example: $N = 10$ has usual digits $(d_1, d_0) = (1, 0)$, and the trade converts it to $(0, 10)$, exactly the pattern just described. The examiners are not hiding the mechanism, they are demonstrating it once so that you generalise it.

What is really being tested is proof discipline on a two directional claim, not cleverness. Once the extra freedom is seen to live entirely in that one exchange identity, the classification writes itself: a unique over-expansion should correspond exactly to having no zero digit in the usual expansion, since a zero digit is the only place the exchange can land. What remains is to prove both directions properly, existence of a second expansion whenever a zero digit is present and uniqueness whenever it is not, without conflating them into one argument that serves neither direction well.

Approaches, ranked worst to best

1. **Enumerate small cases by hand.** *Cue:* the worked example for $N = 10$ invites you to try $N = 1$ up to about 100 and tally over-expansions. *Circuit:* brute force digit string search. (this finds the pattern quickly but proves neither direction; every hour spent enumerating is thrown away the moment the real argument is written, and a coincidence across the first twenty cases is easy to mistake for a theorem).

2. **Count the total number of over-expansions as a function of the zero digits.** *Cue:* noticing that $N = 1000$, with three zero digits, plainly has more than two over-expansions. *Circuit:* a run length or generating function count over blocks of consecutive zero digits. (this answers a harder question than the one asked; the problem wants unique versus not unique, not the exact count, and the exact count is not even a simple function of the number of zero digits, so this route spends ink solving a different problem).

3. **Split into existence and uniqueness, each with its native tool.** *Cue:* the phrase which positive integers is a classification, and classifications proved as an if and only if always split cleanly. *Circuit:* explicit local exchange for existence, residue driven induction on digit count for uniqueness. (this is the direct route: one line constructs the second expansion whenever a zero digit is present, and one induction, peeling the last digit by its forced residue modulo 10, rules out any second expansion whenever no zero digit is present).

The ideal solution

Claim. A positive integer N has a unique base 10 over-expansion if and only if its usual base 10 expansion contains no digit equal to 0.

Write the usual base 10 expansion of N as $N = d_k 10^k + \cdots + d_0 10^0$ with $d_k \neq 0$ and $d_i \in \{0, \ldots, 9\}$ for all i; every over-expansion of N is compared against this fixed reference.

Existence of a second expansion when some digit is 0. Suppose $d_i = 0$ for some $i < k$, and let i be the largest such index. Because i is the largest index below k at which the digit vanishes, the digit immediately above it cannot also vanish without contradicting that maximality, so $d_{i+1} > 0$. Replace the two terms $d_{i+1}10^{i+1} + d_i 10^i = d_{i+1}10^{i+1}$ by $(d_{i+1} - 1)10^{i+1} + 10 \cdot 10^i$; both expressions equal the same value since $10 \cdot 10^i = 10^{i+1}$, and the new digit at position i, namely 10, is legal. If $i + 1 < k$, the digit $d_{i+1} - 1$ sitting at position $i + 1$ is a perfectly legal interior digit, possibly 0, since only the top digit of an expansion is required to be nonzero. If $i + 1 = k$ and $d_k = 1$, the coefficient at the old top position becomes 0; simply read the new expansion with top index i instead of k, since the coefficient 10 at position i is nonzero. Either way this is a second, genuinely different over-expansion of N, so N does not have a unique one.

Uniqueness when no digit is 0. Suppose $d_i \geq 1$ for every $i = 0, \ldots, k$. Induct on k, the number of digits of N minus one. If $k = 0$ then $N = d_0 \in \{1, \ldots, 9\}$, and any over-expansion of N is a single digit $e_0 \in \{0, \ldots, 10\}$ with $e_0 = N$; this is forced, so the expansion is unique.

Let $k \geq 1$ and suppose the claim holds for every positive integer whose usual expansion has fewer than $k + 1$ digits. Let $N = e_l 10^l + \cdots + e_0 10^0$ be any over-expansion of N, with $e_l \neq 0$ and every $e_i \in \{0, \ldots, 10\}$. Reducing modulo 10, $e_0 \equiv N \equiv d_0 \pmod{10}$. Since $d_0 \in \{1, \ldots, 9\}$, the only value of e_0 in $\{0, \ldots, 10\}$ congruent to d_0 modulo 10 is $e_0 = d_0$ itself, since the other candidate $d_0 + 10$ already exceeds 10. So $e_0 = d_0$ is forced, and

$$\frac{N - d_0}{10} = e_l 10^{l-1} + \cdots + e_1 10^0$$

exhibits $e_l, \ldots, e_1$ as an over-expansion of $(N - d_0)/10$. But $(N - d_0)/10 = d_k 10^{k-1} + \cdots + d_1 10^0$ is precisely the usual expansion of this smaller integer, which has k digits, all nonzero. By the induction hypothesis this smaller integer has a unique over-expansion, its own usual one, so $l - 1 = k - 1$ and $e_i = d_i$ for $i = 1, \ldots, k$. Together with $e_0 = d_0$, the over-expansion of N is exactly the usual one, and since it was arbitrary, N's over-expansion is unique.

Both directions close the claim: N has a unique base 10 over-expansion exactly when its usual expansion has no zero digit.

Layer 1: WHAT

Block 1: fix the usual expansion as the reference and state the claim as an if and only if. Block 2: for existence, locate the largest zero digit and perform one local exchange via $10 \cdot 10^i = 10^{i+1}$. Block 3: handle the bookkeeping when that exchange empties the old top coefficient. Block 4: for uniqueness, set up induction on the number of digits. Block 5: force the last digit of any rival expansion by a residue argument. Block 6: reduce to the induction hypothesis on the smaller integer and conclude both the digits and the top index agree.

The nominal route hunts for one master argument that proves both directions at once, or tries to count every over-expansion of N in general, which answers more than was asked. The bifurcation point is the decision to split the if and only if into two independently toolled arguments before writing a single line of algebra, construction for existence, induction for uniqueness. A second, smaller bifurcation sits inside the uniqueness half: peeling the *last* digit by its forced residue modulo 10, rather than reasoning from an unknown top index downward, since the last digit is the only one pinned by the value of N alone.

Layer 2: HOW

Block 1. **Cue.** The command term "which positive integers" demands a classification, and the problem already supplies both a positive instance (the definition) and a negative one (the worked example) of that classification. **Why natural.** Fixing the usual expansion as a reference object is natural because it is the one expansion every positive integer already has for free; every other over-expansion is naturally described as a deviation from it. **Link to the next block.** With a fixed reference in hand, the two directions of the iff become two separate, comparable questions about that reference.

Block 2. **Cue.** The digit alphabet is the usual one plus exactly one extra value, 10, sitting one unit above the usual ceiling of 9. **Why natural.** A single extra legal value can only ever be reached through a single identity relating adjacent positions, $10 \cdot 10^i = 10^{i+1}$, so the entire search for a second expansion reduces to asking where that identity's two preconditions, a receiving zero and a paying nonzero neighbour, can both be met. **Link to the next block.** Choosing the *largest* zero index below the top, rather than any zero index, is what guarantees the paying neighbour is available, which sets up the bookkeeping in the next block.

Block 3. **Cue.** The exchange can land exactly at the old top position. **Why natural.** An over-expansion's defining rule is only that its own top digit is nonzero, not that the top index matches the usual expansion's, so a coefficient collapsing to 0 at the old top simply means the new top index is one smaller. **Link to the next block.** This closes the existence direction cleanly and hands over to the uniqueness direction, where the top index of a rival expansion is genuinely unknown in advance.

Block 4. **Cue.** No digit is 0, so the exchange identity from Block 2 has nowhere to land. **Why natural.** Induction on the number of digits is the natural tool whenever a claim about arbitrary length digit strings needs to be reduced to a claim about a strictly shorter one. **Link to the next block.** The base case, a single nonzero digit, is immediate, so the entire burden falls on relating a $(k+1)$-digit rival expansion to a k-digit one.

Block 5. **Cue.** Any rival expansion's constant term is only visible modulo 10 from the value of N itself. **Why natural.** The digit bound $\{0, \ldots, 10\}$ contains exactly one representative of each residue class modulo 10 except the class of 0, which contains two, 0 and 10; since N's own last digit is nonzero, only the smaller representative survives the bound. **Link to the next block.** Once the last digit is forced to agree with the usual expansion, everything above it is literally an over-expansion of the strictly smaller integer $(N - d_0)/10$.

Block 6. **Cue.** The smaller integer inherited from Block 5 has one fewer digit and, since it is just N with its nonzero last digit stripped, still has no zero digits. **Why natural.** The induction hypothesis applies verbatim to this smaller integer, and uniqueness there means uniqueness of its entire digit string, including its length. **Link to the next block.** Matching lengths forces the rival expansion's top index to equal k, and matching digits forces every coefficient to equal the usual one, closing the induction.

Layer 3: WHY

The digit bound lemma, $\sum_{i<k} 9 \cdot 10^i = 10^k - 1 < 10^k$, is what makes ordinary positional representations rigid: no combination of digits below a position can ever match or exceed one unit at that position. Over-expansion adds exactly one unit of slack per position, and that lemma shows the slack can only ever be absorbed where a position is already empty; anywhere a usual digit is

nonzero, the position is already too full to receive a compensating unit without a further, cascading adjustment that the induction shows never actually materialises once the bottom digit is pinned. The discernment mechanism to train is this: whenever a digit alphabet is widened by exactly one extra value at one end, name the single carry identity that value encodes, then ask precisely where in a legal digit string that identity's preconditions can both hold. That localises the entire phenomenon before any casework begins. Rehearsing the same question with the alphabet widened at the opposite end, as in the first mirror drill below, is what turns "spot the zero" into a reflex that also recognises which digit plays the analogous role when the direction of the carry is reversed. The examiners' worked instance for $N = 10$ is not decoration; it is a fully explicit demonstration of the exchange identity, planted so that a solver who tests it against the digit bound lemma reaches the classification without any outside machinery.

Revolving concepts

Positional base b representation. Every positive integer has a unique expansion with digits in $\{0, \ldots, b-1\}$; the proof is exactly the digit bound lemma above with 10 replaced by b. See Book 1, field: Number theory, topic: base representations.

The carrying identity and enlarged alphabets. Widening the alphabet to $\{0, \ldots, b\}$ introduces exactly the identity $b \cdot b^i = b^{i+1}$ as the sole source of extra representations.

Proposition 12.7.1. *Fix an integer $b \geq 2$ and enlarge the digit alphabet to $\{0, 1, \ldots, b\}$. A positive integer N has a unique base b over-expansion if and only if the usual base b expansion of N has no digit equal to 0.*

The proof is the argument above verbatim with b in place of 10 throughout; nothing about the base 10 was ever load bearing.

Residue induction, peeling the last digit. The general reduction $N \mapsto (N - d_0)/b$, with d_0 the residue of N modulo b, turns any claim about arbitrary length base b digit strings into an induction on the number of digits, with a one digit base case. This is the same self-similar step that underlies converting between number bases by repeated division, and mixed radix systems more generally.

Other uniqueness of representation theorems, a genuine knowledge gap. Not every uniqueness of representation result comes from a value bound like the one here. Zeckendorf's theorem states that every positive integer has a unique representation as a sum of non-consecutive Fibonacci numbers; the constraint enforcing uniqueness there is an adjacency rule on which terms may appear together, not a bound on any single digit. Knowing that a second family of such theorems exists, governed by a structurally different kind of constraint, is worth installing before meeting it cold on a later problem.

Circuit Bank

- **WHEN** a positional digit alphabet is enlarged by exactly one extra value at one end **DO** write down the single carrying identity that extra value encodes and find its two preconditions on adjacent digits **BECAUSE** non-uniqueness of representation can only ever be generated by that one identity, so naming it first turns the whole problem into locating where its preconditions hold
- **WHEN** a claim has the shape "for all N, property $P(N)$ holds if and only if digit condition Q" **DO** split immediately into the two directions and choose an independent tool for each, typically an explicit construction for the existence direction and an induction for the uniqueness direction **BECAUSE** conflating the two directions into one combined argument obscures which lemma is doing the load bearing work and tends to smuggle in an unjustified assumption on one side
- **WHEN** a uniqueness claim about base b digit strings of unknown length must be proved by induction **DO** induct on the number of digits and peel off the last digit via its forced residue

modulo b, not the leading digit **BECAUSE** the residue of N modulo b is the one digit value pinned by N alone, independent of any assumption about how many digits a rival expansion uses

- **WHEN** a bound on individual digits needs to be turned into a statement about the whole number **DO** compare the sum of all lower digits to the geometric total $\sum_{i<k}(b-1)b^i = b^k - 1$ rather than bounding term by term **BECAUSE** this exact identity is what proves a single higher digit dominates every combination of digits below it, the standard lemma behind every base representation uniqueness argument

Backtesting

1. First useful cue: the problem hands you a worked example of non-uniqueness for $N = 10$, manufactured exactly at a zero digit, and that placement is the cue, not the arithmetic.
2. Method the expression suggested: the enlarged digit alphabet suggested naming the one carrying identity $10 \cdot 10^i = 10^{i+1}$ it encodes, then localising rather than globally searching.
3. Where you would hesitate, and the thought to reinforce: worrying whether a rival expansion's top index could differ from the usual one; the fix is to notice the induction hypothesis's uniqueness already includes uniqueness of length, so the top index is pinned automatically once every digit below it is pinned.
4. The trap the question sets: assuming the total number of over-expansions is some simple function of how many zero digits are present, when the question never asked for a count at all, only for the binary classification unique versus not.
5. What makes you recognise it faster next time: seeing any digit alphabet fattened by one extra value at one end as an instant instruction to find the single identity that extra value encodes.
6. The reusable decision rule: when a positional system's legal digit set is $\{0, \ldots, b\}$ instead of $\{0, \ldots, b-1\}$, check the usual expansion for zero digits, because a zero digit is precisely the only slot the extra value's carrying identity can occupy.

Mirror drills

1. (drill) Which positive integers have a unique base 2 over-expansion, with digit alphabet $\{0, 1, 2\}$ in place of the usual $\{0, 1\}$? *Answer:* exactly those whose usual binary expansion has no zero digit, that is $N = 2^{k+1} - 1$ for some $k \geq 0$; the identity $2 \cdot 2^i = 2^{i+1}$ plays the role of $10 \cdot 10^i = 10^{i+1}$ verbatim.
2. (drill) How many distinct base 10 over-expansions does $N = 100$ have in total? *Answer:* exactly 3: the usual $(1, 0, 0)$, and two more with top index one smaller, $(d_1, d_0) = (10, 0)$ and $(9, 10)$; the count is not simply twice the number of zero digits, so always enumerate directly rather than pattern matching a formula for the count.
3. (drill) Exhibit a base 10 over-expansion of $N = 1000$ other than the usual one, and explain why the exchange has to be applied at the digit in the hundreds place rather than the tens or units place. *Answer:* $(d_2, d_1, d_0) = (10, 0, 0)$, since $1000 = 10 \cdot 10^2$; trying the exchange at the tens or units place fails because the digit immediately above it is also 0, violating the precondition that the paying neighbour be nonzero.

12.8 B2. Maximising a weighted integral under an area constraint

Suppose that f is a function on the interval $[1,3]$ such that $-1 \le f(x) \le 1$ for all x and $\int_1^3 f(x)\,dx = 0$. How large can $\int_1^3 \frac{f(x)}{x}\,dx$ be?

Read it out loud

Two hypotheses are given about f, a pointwise box bound and a single integral condition, and then a different integral is asked about. That pairing, a box bound plus one linear integral constraint, together with a question that is itself linear in f, is the signature of an extremal problem whose best function is not smooth at all but a step function sitting on the boundary of the box. The field underneath is real analysis and inequalities, and the caveat worth clarifying immediately is that $\int f(x)/x\,dx$ presupposes f is integrable in the first place, since the problem never says so explicitly; that assumption is used from the first line onward.

The first useful cue is the kernel $1/x$ itself: it is strictly decreasing on $[1,3]$. A linear objective wants f to be as large as the box allows wherever the kernel is large and as small as the box allows wherever the kernel is small, since that is the only way to correlate f with $1/x$ as strongly as the box permits. The single constraint $\int f = 0$ over an interval of length 2 then forces the region where $f = 1$ and the region where $f = -1$ to have equal measure, one unit each, once you accept that the extremal f should only ever take the two boundary values. Combine the two observations, decreasing kernel and an equal length split, and the candidate writes itself within the first thirty seconds: f should be 1 on the left half $[1,2]$ and -1 on the right half $(2,3]$, since a decreasing kernel is always largest on an initial segment.

What is really being tested is whether you trust that heuristic enough to state the candidate immediately and then spend all your remaining ink proving it, rather than trying to discover it through algebra. The examiners are not asking for calculus of variations; they are asking whether you recognise a bang bang extremal problem on sight and can verify a single guessed function with an elementary comparison.

Approaches, ranked worst to best

1. **Guess and check with smooth test functions.** *Cue:* no structural cue is used at all, only the temptation to try something explicit like a cosine or a linear ramp. *Circuit:* plugging in candidates and comparing numbers. (this never proves an upper bound, only ever produces lower bounds on the answer, and after several candidates you still do not know whether a better one exists).

2. **Calculus of variations with a Lagrange multiplier.** *Cue:* a linear objective under one linear equality constraint is the classic setting for a multiplier. *Circuit:* Pontryagin style reasoning, $f = \operatorname{sign}(1/x - \lambda)$ pointwise, then solve $\int f = 0$ for λ. (this correctly discovers the candidate switching function and even finds $\lambda = 1/2$, but turning "the unconstrained pointwise optimum is a sign function" into a rigorous statement over an infinite dimensional constraint set costs real machinery that a Putnam write-up should not need to import).

3. **Integration by parts against the primitive of f.** *Cue:* the constraint $\int_1^3 f = 0$ makes $F(x) = \int_1^x f(t)\,dt$ vanish at both endpoints. *Circuit:* trade a rough pointwise bound on f for a tight Lipschitz bound on F, then bound the transformed integral directly. (this is genuinely elegant and fully competition legal, but needs one honest remark about approximating a merely integrable f by continuous functions before the boundary terms in the parts formula are justified without qualification).

4. **Guess the extremal from the bathtub heuristic, then verify by one direct comparison.** *Cue:* the same decreasing kernel and equal split cue as above. *Circuit:* subtract the guessed extremal

from f and find one pointwise inequality valid on the whole domain. (this is the direct route: the guess costs no ink to state, and the verification is a single inequality integrated once, with no machinery beyond monotonicity of $1/x$).

The ideal solution

Define

$$g(x) = \begin{cases} 1 & 1 \le x \le 2 \\ -1 & 2 < x \le 3. \end{cases}$$

First check g is admissible: $-1 \le g \le 1$ trivially, and $\int_1^3 g\,dx = \int_1^2 1\,dx + \int_2^3 (-1)\,dx = 1 - 1 = 0$, so g satisfies exactly the hypotheses imposed on f.

Let $h = g - f$. Since f satisfies $\int_1^3 f = 0$ by hypothesis and $\int_1^3 g = 0$ as just checked, $\int_1^3 h\,dx = 0$. The key claim is the single pointwise inequality

$$h(x)\left(\frac{1}{x} - \frac{1}{2}\right) \ge 0 \qquad \text{for every } x \in [1,3].$$

On $[1,2]$, $x \le 2$ gives $1/x - 1/2 \ge 0$, and $g(x) = 1 \ge f(x)$ gives $h(x) = 1 - f(x) \ge 0$; a product of two nonnegative numbers is nonnegative. On $(2,3]$, $x > 2$ gives $1/x - 1/2 < 0$, and $g(x) = -1 \le f(x)$ gives $h(x) = -1 - f(x) \le 0$; a product of two nonpositive numbers is again nonnegative. This proves the claim on all of $[1,3]$, with the two cases exactly covering the domain and agreeing (trivially) at $x = 2$.

Integrate the claim over $[1,3]$:

$$0 \le \int_1^3 h(x)\left(\frac{1}{x} - \frac{1}{2}\right) dx = \int_1^3 \frac{h(x)}{x}\,dx - \frac{1}{2}\int_1^3 h(x)\,dx = \int_1^3 \frac{h(x)}{x}\,dx,$$

using $\int_1^3 h = 0$ from above. So $\int_1^3 h(x)/x\,dx \ge 0$, that is,

$$\int_1^3 \frac{f(x)}{x}\,dx \le \int_1^3 \frac{g(x)}{x}\,dx.$$

Compute the right side directly:

$$\int_1^3 \frac{g(x)}{x}\,dx = \int_1^2 \frac{dx}{x} - \int_2^3 \frac{dx}{x} = \ln 2 - (\ln 3 - \ln 2) = 2\ln 2 - \ln 3 = \ln\frac{4}{3}.$$

So $\int_1^3 f(x)/x\,dx \le \ln(4/3)$ for every admissible f, and this bound is attained by $f = g$ itself, which is admissible as checked above. Hence the largest possible value is

$$\boxed{\ln\frac{4}{3}}.$$

Layer 1: WHAT

Block 1: read the two hypotheses as a box bound plus one linear constraint, and the kernel $1/x$ as monotonic, which together flag a bang bang extremal problem. Block 2: use the equal length split forced by $\int f = 0$ together with the decreasing kernel to name the candidate g before doing any verification. Block 3: check g is admissible. Block 4: form $h = g - f$ and note $\int h = 0$ for free. Block 5: find the single pointwise inequality $h(x)(1/x - 1/2) \ge 0$ that holds on the whole domain at once. Block 6: integrate that one inequality and compute the resulting bound explicitly, then exhibit equality.

The nominal route tries to discover g by algebra, either by direct optimisation over an ansatz or by grinding through a Lagrange multiplier calculation until a switching function falls out, and

then re-derives the bound by splitting the domain into $[1,2]$ and $(2,3]$ at the integration stage as two separate estimates. The bifurcation point is the decision to hunt for one inequality valid on the entire domain, built by comparing the kernel $1/x$ against its own value $1/2$ at the switching point, rather than comparing f and g separately on each half. That one shift by $1/2$ is what collapses two case dependent estimates into a single integrated line.

Layer 2: HOW

Block 1. **Cue.** A pointwise box bound $-1 \leq f \leq 1$ appears together with exactly one linear integral constraint, and the target is itself linear in f. **Why natural.** Linear objectives over a set defined by a box bound and finitely many linear constraints are always optimised at an extreme point of that set, and here the extreme points are precisely the functions taking only the values ± 1 almost everywhere. **Link to the next block.** Accepting that the optimal f is boundary valued turns the search for it into a search for one set, the region where $f = 1$.

Block 2. **Cue.** The kernel $1/x$ is strictly decreasing on $[1,3]$, and the constraint $\int f = 0$ over a length 2 interval is a statement purely about measure. **Why natural.** A decreasing kernel is maximised on any fixed measure set by choosing the leftmost such set, and the constraint fixes that measure at exactly one unit on each side by matching the plus and minus regions. **Link to the next block.** These two facts together pin down the candidate completely, an initial segment of length one carrying the value 1 and the rest carrying -1, with the actual switching point falling exactly at $x = 2$.

Block 3. **Cue.** Any candidate produced by a heuristic must be checked against the problem's own hypotheses before it is used in a comparison. **Why natural.** g was built to have measure one on each side by construction, so its own integral vanishing is really the same fact as the construction, restated. **Link to the next block.** Knowing g is admissible means f and g can now be compared on equal footing.

Block 4. **Cue.** Two functions both satisfying the same linear constraint invite subtraction. **Why natural.** Subtracting cancels the constraint itself, $\int h = 0$ costs nothing to state once both integrals are known to vanish, and this is exactly the piece of information the final integration step needs. **Link to the next block.** With $\int h = 0$ available for free, only a pointwise fact about h is still needed to finish.

Block 5. **Cue.** g was built by comparing the kernel $1/x$ to the single value $1/2$, so that same value is the natural pivot to test h against. **Why natural.** g and f differ only where the box bound is active, and the box bound is active in a direction that always matches the sign of $1/x - 1/2$ by the very construction of g; this is not a coincidence, it is what defines g. **Link to the next block.** Once the sign of h is known to match the sign of $1/x - 1/2$ everywhere, their product is nonnegative everywhere, a single fact valid on the whole domain rather than two facts valid on two halves.

Block 6. **Cue.** A pointwise inequality that holds everywhere integrates to a global inequality in one step. **Why natural.** Expanding the product recovers exactly $\int h/x$ minus a multiple of $\int h$, and the second term is already known to vanish, so the entire inequality collapses onto the one integral that was actually wanted. **Link to the next block.** This closes the upper bound, and exhibiting $f = g$ as an admissible function achieving it upgrades the bound from a supremum to an attained maximum.

Layer 3: WHY

Two facts from outside the elementary toolkit are doing the discovery work here even though neither is needed in the final write-up. The first is the bang bang principle: a linear functional maximised over a convex set defined by a pointwise bound and linear constraints attains its maximum at an extreme point of that set, and for the set $\{f : |f| \leq 1\}$ the extreme points are exactly the ± 1-valued functions, because any f taking a value strictly between -1 and 1 on a set of positive measure could be perturbed within that set, keeping the constraint fixed while strictly increasing the objective unless the kernel happens to be constant there. The second is the bathtub principle: among all subsets of a fixed measure, the integral of a decreasing function is maximised by the leftmost such subset, since moving any mass rightward within a fixed measure set only ever trades a larger kernel value for a smaller one. Neither fact needs to be cited by name in a write-up, because once the candidate they jointly produce is in hand, an elementary direct comparison finishes the job with no further machinery; this is exactly why the ideal solution above never mentions either principle.

The discernment mechanism worth training is this: whenever a problem hands you a pointwise box bound together with one linear integral constraint and then asks you to bound a second linear integral, immediately identify the kernel of that second integral and ask whether it is monotonic on the domain. If it is, the extremal is forced to be a step function switching at whatever point makes the measure split match the constraint, and the entire remaining task is an elementary verification, not a search. What the examiners nudged toward is precisely this reflex: the numbers in the problem, the interval $[1, 3]$ of length 2 and the constraint value 0, are chosen so that the split is the cleanest possible one, an exact half and half, so that a solver who has the reflex reaches the candidate in seconds and can spend the remaining time on the one inequality that actually needs proving.

Revolving concepts

Bathtub principle for monotonic kernels. For a fixed measure m and a function ϕ decreasing on the domain, $\int_S \phi$ over measurable S with $|S| = m$ is maximised by taking S to be an initial segment of the domain, and minimised by taking a terminal segment. See Book 1, field: Inequalities, topic: rearrangement and the bathtub principle.

Bang bang principle for linear functionals. A linear functional of f, maximised subject to $|f| \leq 1$ pointwise and one linear integral constraint, attains its maximum at a ± 1-valued function; this is the same idea that appears as Pontryagin's maximum principle in optimal control, where the optimal control is always found on the boundary of its admissible range. Genuine knowledge gap: this is a real theorem from optimisation and control theory, not elementary calculus, and is worth installing precisely because it explains discovery, not verification.

Kernel independence of the optimal set. If the linear constraint on f forces a fixed measure split of the domain, maximising against any strictly decreasing kernel is achieved by the same leftmost set, regardless of which decreasing kernel is used; only the resulting numeric value depends on the kernel. Flipping the kernel to strictly increasing flips the optimal set to the rightmost one instead. This is why the mirror drills below can reuse the same g for a different kernel but must flip g when the kernel's monotonicity flips.

Trading a rough bound for a smooth one via a primitive. Setting $F(x) = \int_1^x f$, the hypotheses give $F(1) = F(3) = 0$ and, since $|f| \leq 1$, F is 1-Lipschitz, so $|F(x)| \leq \min(x - 1, 3 - x)$. Integrating $f(x)/x$ by parts converts the integral into one involving $F(x)/x^2$, where this sharper bound on F can be used directly; this is the official second solution's route and is worth knowing as a general reflex whenever a rough pointwise bound on f is not tight enough to finish an estimate involving a smooth kernel.

Circuit Bank

- **WHEN** a linear integral objective is paired with a pointwise box bound on f and one linear integral constraint **DO** expect the extremal to be a ± 1-valued step function, and look for the single switching point that the constraint forces **BECAUSE** a linear objective over a box shaped constraint set is always optimised at an extreme point, and the extreme points of this particular set are exactly the boundary valued functions
- **WHEN** the objective's kernel is monotonic on the domain **DO** place the value $+1$ on the piece of the domain where the kernel is largest, sized exactly to the measure the constraint demands **BECAUSE** a monotonic kernel is always maximised on a fixed measure set by an initial or terminal segment, never a scattered set, so only the split point remains unknown
- **WHEN** a guessed extremal g needs to be verified against every competitor f subject to the same two hypotheses **DO** subtract, $h = g - f$, and search for one pointwise inequality against the kernel shifted by its value at the switching point, valid on the whole domain at once **BECAUSE** a single global inequality integrates in one line, while splitting the domain into pieces at the integration stage repeats the same sign argument twice for no benefit
- **WHEN** a rough pointwise bound on f is not tight enough once multiplied by a kernel and integrated **DO** consider the primitive $F(x) = \int f$, exploit any vanishing boundary values it inherits from a given integral constraint, and integrate by parts **BECAUSE** a primitive of a bounded function is Lipschitz, converting a rough pointwise bound into a genuinely tight one near the boundary

Backtesting

1. First useful cue: a pointwise box bound on f together with a single integral condition equal to zero is exactly the signature of a bang bang extremal problem, before any calculus is attempted.
2. Method the expression suggested: dividing by x marks the kernel as strictly decreasing, and a decreasing kernel paired with an equal measure constraint always points to a single switching point at the midpoint of the domain's length, not the midpoint of the kernel's range.
3. Where you would hesitate, and the thought to reinforce: wondering whether a smooth candidate might beat the discontinuous step function g; reinforce by recalling that the objective and constraint are both linear in f, and linear problems over a box are never improved by moving away from the boundary.
4. The trap the question sets: assuming the answer wants a formula in terms of a general f, when actually a single explicit numeric bound is being asked for, achieved by one specific, easily written down function.
5. What makes you recognise it faster next time: seeing $-1 \le f \le 1$ paired with $\int f =$ a fixed constant as an instant cue for a bang bang extremal, independent of which kernel the actual objective uses.
6. The reusable decision rule: when bounding $\int f\phi$ under $|f| \le 1$ and a fixed value of $\int f$, guess a ± 1-valued step function switching at the point that matches the required measure split, then prove optimality with one pointwise comparison against the kernel shifted by its value at that point.

Mirror drills

1. (drill) Under the same two hypotheses on f, how large can $\int_1^3 f(x)/x^2\,dx$ be? *Answer:* the same switching function g is still optimal, since $1/x^2$ is still strictly decreasing on $[1, 3]$ and the

measure split forced by the constraint has not changed; the bound is $\int_1^2 dx/x^2 - \int_2^3 dx/x^2 = 1/2 - 1/6 = 1/3$.

2. (drill) If the constraint is changed to $\int_1^3 f(x)\,dx = 1$ instead of 0, where does the switching point move, and what is the new bound on $\int_1^3 f(x)/x\,dx$? *Answer:* writing the plus region's length as m_+, $m_+ - (2 - m_+) = 1$ gives $m_+ = 3/2$, so the switch moves to $x = 2.5$; the new bound is $2\ln(2.5) - \ln 3 = \ln(25/12)$.

3. (drill) True or false: if the kernel in the objective were x itself instead of $1/x$, the same g, equal to 1 on $[1,2]$, would still maximise $\int_1^3 f(x)\,x\,dx$. *Answer:* false; x is increasing, not decreasing, on $[1,3]$, so the optimal set flips to the rightmost half, giving the flipped function, -1 on $[1,2]$ and 1 on $(2,3]$, as the true maximiser.

12.9 B3. distinct primes among entries force rank two

Let A be an $m \times n$ matrix with rational entries. Suppose that there are at least $m + n$ distinct prime numbers among the absolute values of the entries of A. Show that the rank of A is at least 2.

Read it out loud

Say the hypothesis back to yourself in two separate pieces, because the problem is built from two pieces of mathematics that have nothing to do with each other until the very last line. The first piece is arithmetic: at least $m + n$ of the entries, read as absolute values, are distinct primes. The second piece is a bound on an integer invariant, the rank, phrased as a lower bound rather than an exact value: you are not asked to compute the rank, only to rule out the two smallest possibilities for it.

Look at the number $m + n$ itself before you look at anything else. It is not a fixed constant like 5 or 100, and it is not a product like mn. It is rows plus columns, and a matrix problem that states its hypothesis as rows plus columns is all but naming its own method: build a graph whose two vertex classes are the rows and the columns, of size m and n, so that $m + n$ is the vertex count of something you can draw. The moment you notice that coincidence of counts is the moment you should have written down, within the first thirty seconds, the sentence "one vertex per row, one vertex per column, an edge for every prime."

The fields wired into this problem are three, and they are deliberately kept apart until the final step. Linear algebra supplies the classification of rank at most 1 as an explicit two-factor formula for every entry. Graph theory supplies the exact threshold at which a graph stops being able to avoid a cycle, which is precisely when its edge count reaches its vertex count. Number theory, specifically unique factorisation, supplies the contradiction, because a product of pairwise distinct primes and a perfect square of a rational number are mutually exclusive descriptions of the same number. None of the three fields alone proves anything; the problem is testing whether you can chain them without dropping a piece.

What is really being tested is a discernment skill, not a technique: recognising which branch of mathematics a given numerical threshold belongs to. A weaker candidate reads "$m + n$" as an arbitrary bound and tries to attack the problem with linear algebra alone, hunting for two rows or two columns that are visibly independent. That hunt has nowhere to stand, because the primes need not sit inside any single 2×2 block of the matrix.

Approaches, ranked worst to best

1. **Hunt directly for a nonvanishing 2×2 minor.** *Cue:* "rank at least 2" restated as "some 2×2 minor is nonzero". *Circuit:* brute inspection of whichever two prime entries you notice first. (why it ranks here: with m and n arbitrary and the $m + n$ primes scattered anywhere in the grid, nothing tells you which two rows and which two columns to compare, and in general no two of the primes need share both a row and a column with two others, so this approach has no way to even choose its first move.)

2. **Embed the entries in the valuation vector space over all primes and run a dimension count.** *Cue:* "distinct primes" read as coordinates of a vector. *Circuit:* the valuation homomorphism from the positive rationals under multiplication into the group of integer sequences indexed by primes. (why it ranks here: this is correct and it generalises further than the problem asks, but you pay for that generality with a heavier setup, an infinite-dimensional ambient space, a hyperplane cut out by a sum-zero condition on auxiliary coefficients, and a dimension bound you must justify with care before it reproduces the same $m + n$ threshold that the graph route hands you immediately.)

3. **Model rows and columns as a bipartite graph and invoke the forest bound.** *Cue:* the count "$m+n$" is literally rows plus columns. *Circuit:* a simple graph with as many edges as vertices cannot be a forest, so it contains a cycle. (best: the least machinery of the three, every step forced by the last, and the threshold $m+n$ falls out of a single elementary lemma instead of being reverse-engineered from a dimension count.)

The ideal solution

Suppose, for contradiction, that the rank of A is at most 1.

If A is the zero matrix, none of its entries has prime absolute value, since 0 is not prime, so zero distinct primes occur among the entries of A. This already contradicts the hypothesis, because $m, n \geq 1$ forces $m+n \geq 2$ primes to be required. So from now on, $\operatorname{rank}(A) = 1$ exactly.

A rational matrix of rank exactly 1 is an outer product with rational factors. To see this, pick a row index i_0 for which row i_0 is not the zero row, and set $b_j = A_{i_0,j}$ for $j = 1, \ldots, n$, so that b is a nonzero vector in $\mathbb{Q}^n$. Because $\operatorname{rank}(A) = 1$, every row of A lies in the one dimensional row space spanned by row i_0, so for each row index i there is a scalar a_i with row i equal to a_i times row i_0. Row i_0 has some nonzero entry, in column j_0 say, and reading off that single column forces $a_i = A_{i,j_0} / A_{i_0,j_0}$, a ratio of two rational numbers, hence itself rational. By construction $A_{ij} = a_i b_j$ for every i and j, with $a \in \mathbb{Q}^m$ and $b \in \mathbb{Q}^n$ both nonzero.

Now build a bipartite graph G: one vertex for each row index $1, \ldots, m$, one vertex for each column index $1, \ldots, n$, giving $m+n$ vertices in total. For each of the at least $m+n$ distinct primes occurring among the $|A_{ij}|$, fix one entry realising that prime and join its row vertex to its column vertex. An entry occupies a single position in the grid and carries a single absolute value, so two different primes are always realised at two different positions, hence give two different edges. So G has at least $m+n$ edges on exactly $m+n$ vertices.

Lemma 12.9.1. *A finite simple graph with V vertices and no cycle has at most $V-1$ edges.*

Proof. Induct on V. For $V = 1$ there is no edge to draw, and $0 = V - 1$. For $V > 1$, if the graph has no edges the bound is immediate. Otherwise take a path of maximum length in the graph and let v be one of its endpoints. Every neighbour of v already lies on the path: a neighbour outside the path would extend it, contradicting maximality, and a neighbour further along the path than the vertex adjacent to v would close a cycle with the path itself, contradicting acyclicity. So v's only neighbour is the next vertex on the path, meaning v has degree 1. Delete v: the remaining graph still has no cycle, has $V-1$ vertices, so by induction has at most $V-2$ edges, and restoring v restores at most one edge, giving at most $V-1$ in total. □

Since G has $m+n$ vertices and at least $m+n$ edges, the lemma rules out G being a forest, so G contains a cycle. Every edge of G joins a row to a column, so G is bipartite, and a bipartite graph has no odd cycle; write the cycle as row vertices $\rho_1, \ldots, \rho_k$ and column vertices $\gamma_1, \ldots, \gamma_k$, indices read modulo k, visited in the order $\rho_1, \gamma_1, \rho_2, \gamma_2, \ldots, \rho_k, \gamma_k, \rho_1$, with $k \geq 2$ since a simple bipartite graph has no cycle shorter than 4.

The $2k$ edges of this cycle were chosen for $2k$ distinct primes, say $|A_{\rho_t,\gamma_t}| = p_{2t-1}$ and $|A_{\rho_{t+1},\gamma_t}| = p_{2t}$ for $t = 1, \ldots, k$, all $2k$ of the p's pairwise distinct. Write $r_t = |a_{\rho_t}|$ and $c_t = |b_{\gamma_t}|$, so that $p_{2t-1} = r_t c_t$ and $p_{2t} = r_{t+1} c_t$. Multiply all $2k$ primes together:

$$p_1 p_2 \cdots p_{2k} = \prod_{t=1}^{k} (r_t c_t)(r_{t+1} c_t) = \left(\prod_{t=1}^{k} r_t \right) \left(\prod_{t=1}^{k} r_{t+1} \right) \left(\prod_{t=1}^{k} c_t \right)^2 .$$

The map $t \mapsto t+1 \pmod{k}$ permutes $\{1, \ldots, k\}$, so $\prod_t r_{t+1} = \prod_t r_t$, and the right side collapses to $\left(\prod_t r_t \prod_t c_t\right)^2$, the square of a positive rational number q.

Write $q = s/t$ in lowest terms with s, t positive integers. Then $p_1 \cdots p_{2k}\, t^2 = s^2$, so $t^2 \mid s^2$; since $\gcd(s,t) = 1$ gives $\gcd(s^2, t^2) = 1$, this forces $t = 1$, so $q = s$ is a positive integer and $p_1 \cdots p_{2k} = s^2$

is a perfect square of an integer. Every prime in the factorisation of a perfect square occurs to an even exponent. But $p_1, \ldots, p_{2k}$ are pairwise distinct primes, so their product has each of them occurring to the exponent exactly 1, which is odd. This contradicts the previous sentence.

The contradiction shows that $\operatorname{rank}(A) \leq 1$ is impossible, so $\operatorname{rank}(A) \geq 2$, as required.

Layer 1: WHAT

1. Negate the goal: assume $\operatorname{rank}(A) \leq 1$ and aim for a contradiction; dispose of rank 0 immediately, since the zero matrix has no prime entries at all.
2. Convert rank 1 into the rigid formula $A_{ij} = a_i b_j$ for explicit rational vectors a, b.
3. Convert the counting hypothesis into a graph: rows and columns as vertices, one edge per distinct prime, exactly $m + n$ vertices and at least $m + n$ edges.
4. Invoke the forest bound to extract a cycle, necessarily of even length at least 4 because the graph is bipartite.
5. Substitute the outer product formula around the cycle and telescope the product of the $2k$ distinct primes, using a cyclic shift of index, into an explicit perfect square.
6. Contradict unique factorisation: a squarefree product of distinct primes cannot equal a rational square.

The nominal route stops before block 3. It tries to answer "rank at least 2" by pointing at two rows and two columns directly and computing a 2×2 determinant, which works instantly on hand-picked examples where the primes happen to cluster in a small block, and fails to generalise the moment m and n grow, because nothing in the hypothesis tells you which rows and columns to compare. The bifurcation point is block 3: the instant you read "$m + n$" as rows plus columns rather than as an arbitrary number, and reach for a graph on exactly those two vertex classes instead of continuing to search the grid by eye.

Layer 2: HOW

Block 1. **Cue.** An integer invariant, the rank, is bounded below by a small integer, 2. **Why natural.** Integer invariants bounded below are most cheaply attacked by ruling out every smaller value, because each smaller value usually has a closed, rigid description, whereas "rank at least 2" itself has no single formula to manipulate. **Link to the next block.** Ruling out rank 0 costs one sentence; what remains is the single case rank 1, which is exactly the case with a known rigid form.

Block 2. **Cue.** The phrase "rank at most 1", the negation you just introduced. **Why natural.** Rank at most 1 is definitionally the statement that the row space has dimension at most 1, which means every row is a scalar multiple of one fixed row; writing that out coordinatewise gives $A_{ij} = a_i b_j$, and rationality of the scalars follows for free by dividing two of A's own rational entries. **Link to the next block.** Every entry's absolute value is now a product of exactly two numbers, $|a_i|$ and $|b_j|$, which is the only algebraic handle the rest of the proof needs.

Block 3. **Cue.** The number "$m + n$" appearing as a hypothesis, exactly rows plus columns. **Why natural.** A bound phrased as rows plus columns, rather than a fixed constant or a product, is the fingerprint of a bipartite graph argument, because rows and columns are the two natural vertex classes of a matrix and $m + n$ is their combined size; no other natural object attached to an $m \times n$ grid has that count. **Link to the next block.** The arithmetic hypothesis, a count of primes, becomes a purely combinatorial one, an edge count against a vertex count, and the primes' actual values are set aside until the last block.

Block 4. **Cue.** Edge count at least equal to vertex count in a simple graph. **Why natural.** A cycle is the only obstruction to a graph being a forest, and forests are exactly the graphs that achieve the extremal bound of $V - 1$ edges on V vertices; reaching V edges leaves no forest able to hold them, so a cycle is forced, and because every edge in this particular graph crosses from a row to a column, that cycle cannot be odd. **Link to the next block.** The cycle supplies a closed, alternating list of row and column indices, exactly the shape needed to set up one product taken all the way round.

Block 5. **Cue.** A product taken cyclically around a closed alternating sequence of row and column values. **Why natural.** Each row value r_t is used in exactly two of the $2k$ factors around the cycle, once as the "current" row and once as the "next" row, so multiplying every factor counts each r_t and each c_t exactly twice; reindexing a product over all k indices by a cyclic shift does not change its value, and that single fact is what lets the two copies of the r's collapse into one squared term. **Link to the next block.** The product now has an explicit closed form, the square of a rational number, ready to be set against the other thing you separately know about that same product.

Block 6. **Cue.** A product of pairwise distinct primes on one side, a perfect square on the other side, of the same equation. **Why natural.** Unique factorisation assigns every rational number one well defined exponent at each prime; a square forces that exponent to be even everywhere, while a squarefree product of distinct primes forces it to be exactly 1, odd, at each of those primes, and no number carries two different exponents at the same prime. **Link to the next block.** There is no further block: this is the contradiction, and it closes the proof.

Layer 3: WHY

The route rests on a general pattern for proving an integer invariant is at least k: assume at most $k - 1$, find the exact classification of that boundary case, and break the classification with a piece of structure the boundary case cannot also satisfy. Here the boundary case, rank at most 1, has an unusually explicit classification, an outer product, and the structure that breaks it is arithmetic rather than further linear algebra, which is precisely why hunting for a second independent row inside the matrix itself never gets off the ground: the contradiction was never going to live inside linear algebra alone.

The discernment mechanism worth training is recognising which branch of mathematics owns a given numerical threshold before you commit to a method. A bound stated as a product, mn, would smell of counting entries or a double sum. A bound stated as a fixed constant would smell of an extremal example to pin down exactly. A bound stated as rows plus columns smells of a bipartite graph, because $m + n$ is the one count that a graph on rows and columns produces for free as its vertex total. Train this by asking, every time a matrix problem states a bound, "what natural object attached to this grid has exactly this many elements", before reaching for any technique.

The examiners' nudge sits in exactly one place: the linear algebra fact, rank at most 1 is an outer product, is the single piece of machinery you are assumed to know cold, and it is disposed of in one paragraph. The actual difficulty, the graph theory, is dressed in ordinary matrix language so that a candidate who only reaches for linear algebra stalls at "which two rows and columns do I compare", never suspecting that the real threshold belongs to a different field entirely.

Revolving concepts

Rank one matrices as outer products. A matrix A over a field has rank at most 1 exactly when $A = ab^T$ for some vectors a, b, possibly zero; when the entries of A lie in $\mathbb{Q}$, the vectors a, b can be chosen in $\mathbb{Q}$ as well, by solving for the scaling constants against two of A's own nonzero rational entries. The same idea extends one level up: rank at most r means A is a sum of r such outer products, and rank is unchanged by extending the field of scalars, which is why working over $\mathbb{Q}$

throughout never loses generality here. See Book 1, field: Linear algebra, topic: rank and matrix factorisations.

Unique factorisation and rational squares. Every nonzero rational number factors uniquely as a sign times a product of primes to integer exponents, and this assignment of an exponent to each prime, at each rational number, is additive under multiplication: the exponent of a product is the sum of the exponents of its factors. A positive rational number is the square of a rational number exactly when every one of its exponents is even, and that criterion, applied to a number that is separately known to be a product of pairwise distinct primes, each to the first power, is the exact contradiction this problem manufactures. See Book 1, field: Number theory, topic: unique factorisation and prime exponents.

Forests, trees and the edge against vertex threshold. A tree is a connected graph with no cycle; a forest is a disjoint union of trees. A tree on k vertices has exactly $k-1$ edges, provable by repeatedly deleting a vertex of degree 1, which always exists in a nonempty acyclic graph by taking a longest path and looking at its endpoint; a forest with several components inherits the bound $V-1$ on V vertices, with equality only for a single tree. The immediate corollary used here, a simple graph with at least as many edges as vertices contains a cycle, and in a bipartite graph that cycle has even length at least 4, since row to row or column to column edges do not exist. See Book 1, field: Discrete structures and graphs, topic: forests and the edge against vertex threshold. The leaf removal argument that proves the tree bound is itself a small induction, remove one vertex, apply the same statement to what remains, add back what you removed, and the identical shape of argument reappears whenever a finite structure is built by adding one relation at a time to a fixed set of objects: more relations than objects forces a redundancy, a cycle in a graph, a linear dependency among vectors once you have more of them than the ambient dimension. Recognising that a graph theory threshold and a linear algebra threshold are the same statement in different language is worth carrying forward as its own revolving concept. See Book 1, field: Linear algebra, topic: dependency thresholds and their graph theory analogue.

Circuit Bank

- **WHEN** a problem asks you to prove an integer invariant, rank, dimension, a count, is at least k **DO** assume it is at most $k-1$, find the exact classification of that boundary case, and break the classification with an independent piece of structure **BECAUSE** the boundary case is almost always far more rigid and describable in closed form than the open condition itself, so a contradiction against the rigid form costs less ink than attacking "at least k" directly
- **WHEN** a hypothesis in a matrix problem is stated as a count equal to the number of rows plus the number of columns **DO** build a bipartite graph with rows and columns as the two vertex classes and turn the count into an edge count **BECAUSE** $m+n$ is exactly the vertex count of that graph, so the hypothesis becomes a forest against cycle threshold, the one setting where "rows plus columns" is a forced number rather than an arbitrary one
- **WHEN** you need to know whether a simple graph on V vertices must contain a cycle **DO** check whether it has at least V edges **BECAUSE** a forest on V vertices has at most $V-1$ edges by the leaf removal induction, so V or more edges rules out being a forest
- **WHEN** a quantity is simultaneously known to be a product of pairwise distinct primes and forced by other algebra to equal a perfect square of a rational number **DO** declare the contradiction immediately through unique factorisation, comparing the exponent at any one of those primes **BECAUSE** a square has every exponent even, a squarefree product of distinct primes has every occurring exponent equal to 1, and no rational number carries two exponents at the same prime
- **WHEN** the rank of a matrix is at most 1 **DO** rewrite every entry as $A_{ij} = a_i b_j$ for explicit

vectors a, b over the same field as the entries **BECAUSE** this is exactly what "row space of dimension at most 1" means, and it turns every entry from a matrix condition you can only manipulate linearly into a two-factor product you can manipulate arithmetically

Backtesting

1. First useful cue: the number $m + n$ sitting in the hypothesis, instantly readable as rows plus columns rather than an arbitrary constant.
2. Method the expression suggested: a bipartite graph with rows and columns as vertex classes, one edge per distinct prime, compared against the forest bound.
3. Where you would hesitate, and the thought to reinforce: after finding the cycle, wondering whether the $2k$ chosen primes are still genuinely distinct at that point; reinforce that one grid position carries exactly one absolute value, so distinct primes could never have shared a position to begin with.
4. The trap the question sets: assuming two rows and two columns exist whose four intersection entries are all prime, so that a 2×2 minor can be read off directly; in general the $m + n$ primes are scattered with no such block guaranteed, and only the graph argument survives that scattering.
5. What makes you recognise it faster next time: seeing "the number of rows plus the number of columns" anywhere in a matrix hypothesis should fire the bipartite graph circuit before any other reading of the sentence is attempted.
6. The reusable decision rule: to prove a rank lower bound, negate to the largest forbidden rank, write out that rank's rigid classification explicitly, then hunt for a contradiction from an independent field, arithmetic here, but counting or algebra elsewhere.

Mirror drills

1. (drill) An $m \times n$ rational matrix has at least $m + n$ entries equal to 1 in absolute value, not necessarily at distinct values. Must its rank be at least 2? *Answer:* No. Take every entry equal to 1: the matrix has rank 1 and every one of its $mn \geq m + n$ entries (for $m, n \geq 2$) has absolute value 1. This isolates exactly why the primes in the original problem had to be distinct, not merely numerous.
2. (drill) Explain, using only the graph, why a 3×3 rational matrix can carry as many as 5 distinct primes among its entries while still having rank 1, but never 6. *Answer:* Five edges on six vertices can still form a forest, for instance a single path touching every row and column vertex once, so no cycle and no contradiction; six edges on six vertices cannot avoid a cycle.
3. (drill) A connected graph on 10 vertices has exactly 9 edges. Does it contain a cycle, and why does this feel like the boundary case of the Putnam problem above? *Answer:* No, a connected graph with $V - 1$ edges on V vertices is by definition a tree; this is the extremal forest case, the reason the Putnam hypothesis needed "at least $m + n$" primes and not "at least $m + n - 1$".

12.10 B4. All roots real for a log concave polynomial

Show that for each positive integer n, all the roots of the polynomial

$$\sum_{k=0}^{n} 2^{k(n-k)} x^k$$

are real numbers.

Read it out loud

Say it as it is built: this is a polynomial of degree n whose coefficients are all positive, all powers of two, and the exponent of the coefficient is $k(n-k)$, a concave quadratic in the index k. That last clause is the entire problem. You are not looking at a polynomial, you are looking at a parabola sitting in the exponent, and the parabola is symmetric about $k = n/2$.

Two things must be on your page within thirty seconds. First, $c_k = c_{n-k}$, so the coefficient list is palindromic and $x^n f_n(1/x) = f_n(x)$. Second, and far more important, compute the log concavity ratio, because positive coefficients plus a real rootedness claim always means log concavity is the hidden hypothesis:

$$\frac{c_k^2}{c_{k-1}c_{k+1}} = 2^{2k(n-k)-(k-1)(n-k+1)-(k+1)(n-k-1)} = 2^2 = 4.$$

Not "at least 4", not "roughly 4". Exactly 4, for every k and every n. That is a signature, and you should recognise it as one: the quantity $c_k^2 - 4c_{k-1}c_{k+1}$ is the discriminant of the local quadratic $c_{k-1} + c_k y + c_{k+1} y^2$, so this problem sits precisely on the boundary where every consecutive triple of terms is a perfect square. A problem placed exactly on a threshold is a problem where the naive estimate will give you zero and you will have to look at the next order term. Expect that, and it will not surprise you when it happens.

The fields wired in are: root location for real polynomials, the intermediate value theorem used as a counting device, log concavity of coefficient sequences, and finite alternating sums with rapidly decaying terms. What is really being tested is whether you know that "prove all roots are real" for a polynomial you cannot factor has exactly one elementary certificate, and whether you can choose the test points that make that certificate computable rather than merely conceivable.

Approaches, ranked worst to best

1. **Find the roots.** *Cue:* "all the roots" reads like an instruction to produce them. *Circuit:* solve the equation. Ranked last because it is the reflex of a candidate who has not read the coefficient. For $n = 3$ the palindrome gives $f_3(x) = (1+x)(x^2+3x+1)$ and you feel encouraged; for $n = 4$ the substitution $y = x + 1/x$ gives $y^2 + 8y + 14$ and you feel very encouraged; by $n = 6$ you need a cubic in y all of whose roots are real *and* of modulus at least 2, and there is no pattern to ride. It costs pages and it does not generalise, which is the definition of zombie maths.

2. **Newton's inequalities.** *Cue:* positive coefficients and a real rootedness claim. *Circuit:* check the standard necessary conditions. Newton's inequalities say that if $\sum a_k x^k$ is real rooted then $p_k^2 \geq p_{k-1}p_{k+1}$ for $p_k = a_k / \binom{n}{k}$. They are necessary and never sufficient, so verifying them proves nothing at all. Worth thirty seconds as a sanity check that the claim is not false, worth zero marks as a proof.

3. **Induct with the coefficient recursion.** *Cue:* c_k for n is 2^k times c_k for $n-1$. *Circuit:* turn a coefficient identity into a functional identity. Here $f_n(x) = f_{n-1}(2x) + x^n$, which is a clean identity and the fastest way to generate the family numerically. As a proof engine it stalls:

adding x^n to a real rooted polynomial destroys real rootedness in general, and there is no interlacing statement available to control it. Rank it above the previous two because it is structural rather than blind.

4. **Quote a log concavity theorem.** *Cue:* the ratio $c_k^2/(c_{k-1}c_{k+1})$ is a clean constant. *Circuit:* invoke Hutchinson or Kurtz. Hutchinson's theorem says that positive coefficients with $a_k^2 \geq 4a_{k-1}a_{k+1}$ for all k force all zeros to be real and negative; Kurtz's theorem gives simplicity from the strict inequality. This is exactly our situation, on the nose. It ranks high because it is the right idea, and it ranks below the winner because you may not assume a grader accepts an uncommon named theorem, and proving Hutchinson in general is harder than proving this one case. Know the theorem so that you know what to expect; write the case.

5. **Sign alternation at powers of two.** *Cue:* degree n, cannot factor, must show n real roots. *Circuit:* find $n+1$ points where the sign alternates and let the intermediate value theorem count. The winner. The second cue tells you which points: since $c_k x^k = 2^{k(n-k)+kt}$ up to sign when $x = -2^t$, and $k(n-k)+kt$ is the same parabola with its vertex slid to $k = (n+t)/2$, choosing $t = 2j-n$ puts the vertex exactly on the integer $k = j$ and turns every coefficient into a power of two indexed by $(k-j)^2$. The algebra collapses to a single universal sum with no n in it.

The ideal solution

Write $f_n(x) = \sum_{k=0}^{n} 2^{k(n-k)}x^k$, a polynomial of degree n with leading coefficient $2^{n(n-n)} = 1$.

For $n = 1$, $f_1(x) = 1 + x$. For $n = 2$, $f_2(x) = 1 + 2x + x^2 = (1+x)^2$. Both have all roots real, so assume $n \geq 3$ from now on.

Throughout, put

$$g(m) = 2^{-m^2} \qquad (m \in \mathbb{Z}),$$

so g is positive, even, and strictly decreasing in $|m|$, with $g(0) = 1$ and $g(1) = \frac{1}{2}$.

Lemma 12.10.1. *For $j = 0, 1, \ldots, n$ set $x_j = -2^{2j-n}$. Then*

$$f_n(x_j) = (-1)^j 2^{j^2}\, T_j, \qquad \textit{where} \qquad T_j = \sum_{m=-j}^{n-j} (-1)^m g(m).$$

Proof. For each k,

$$2^{k(n-k)}x_j^k = (-1)^k 2^{k(n-k)+k(2j-n)} = (-1)^k 2^{2jk-k^2} = (-1)^k 2^{j^2-(k-j)^2},$$

since $2jk - k^2 = j^2 - (k-j)^2$. Summing over k and substituting $m = k - j$ gives

$$f_n(x_j) = (-1)^j 2^{j^2} \sum_{m=-j}^{n-j} (-1)^m 2^{-m^2},$$

which is the claim. □

So every question about the sign of f_n at these $n+1$ points is a question about the sign of T_j, and T_j is a truncation of one fixed doubly infinite alternating sum. Note also the reflection: replacing m by $-m$ maps the index range $[-j,\ n-j]$ onto $[-(n-j),\ j]$, and g is even, so

$$T_j = T_{n-j} \qquad (0 \leq j \leq n). \tag{R}$$

This is the palindromic symmetry $x^n f_n(1/x) = f_n(x)$ in disguise, since $1/x_j = x_{n-j}$.

Claim 12.10.2. *If $n \geq 3$ then $T_j > 0$ for every $j \in \{0, 1, \ldots, n\}$.*

Proof. By (R) we may assume $2j \le n$. Split the sum at $|m| = j$:

$$T_j = u_j + R_j, \qquad u_j = \sum_{m=-j}^{j} (-1)^m g(m), \qquad R_j = \sum_{m=j+1}^{n-j} (-1)^m g(m),$$

where R_j is the empty sum exactly when $n = 2j$.

First, u_j. Using evenness of g and the exact identity $2g(1) = 1$,

$$u_j = g(0) + 2\sum_{m=1}^{j} (-1)^m g(m) = \big(1 - 2g(1)\big) + 2\sum_{m=2}^{j} (-1)^m g(m) = 2\sum_{m=2}^{j} (-1)^m g(m).$$

So $u_0 = 1$, $u_1 = 0$, and for $j \ge 2$ the sum $\sum_{m=2}^{j}(-1)^m g(m)$ is alternating with strictly decreasing positive terms and first term $g(2)$, hence lies between $g(2) - g(3)$ and $g(2)$. Therefore

$$j \ge 2 \implies u_j \ge 2\big(g(2) - g(3)\big) = 2\left(\tfrac{1}{16} - \tfrac{1}{512}\right) = \tfrac{31}{256} > \tfrac{1}{16}.$$

Second, R_j. When it is nonempty it is alternating with strictly decreasing positive terms and first term $g(j+1)$, so it has the sign of $(-1)^{j+1}$ and $|R_j| \le g(j+1)$.

Now three cases.

Case $j = 0$. $T_0 \ge u_0 - g(1) = 1 - \frac{1}{2} = \frac{1}{2} > 0$.

Case $j = 1$. Here $u_1 = 0$, so everything rests on $R_1 = \sum_{m=2}^{n-1}(-1)^m g(m)$. Since $n \ge 3$ this sum is nonempty, its first term is $+g(2)$, and by the alternating bound $R_1 \ge g(2) - g(3) > 0$. Hence $T_1 > 0$.

Case $j \ge 2$. $T_j \ge u_j - g(j+1) > \frac{1}{16} - 2^{-9} > 0$.

In all cases $T_j > 0$. □

Finish. The points $x_0 > x_1 > \cdots > x_n$ are $n+1$ distinct negative reals, and by the Lemma and the Claim,

$$\operatorname{sgn} f_n(x_j) = (-1)^j,$$

so the sign alternates along the list. By the intermediate value theorem, each of the n pairwise disjoint open intervals (x_j, x_{j-1}), $j = 1, \dots, n$, contains a root of f_n. That is n distinct real roots of a polynomial of degree n, so by the factor theorem these are all of them, each simple. Hence every root of f_n is real, and in fact

all n roots are real, simple and negative, one in each interval $(-2^{2j-n}, -2^{2j-2-n})$.

For $n = 1, 2$ the claim was checked directly. □

Layer 1: WHAT

Block 1. Read the exponent rather than the coefficient: $c_k = 2^{Q(k)}$ with $Q(k) = k(n-k)$ concave quadratic of second difference -2, equivalently $c_k^2 = 4c_{k-1}c_{k+1}$ exactly.

Block 2. Decide the certificate before doing any algebra: for a degree n polynomial that will not factor, "all roots real" is proved by exhibiting n sign changes and counting.

Block 3. Choose the test points so that the algebra collapses. Testing at $x = -2^t$ adds kt to the exponent, which only translates the vertex of the parabola; choose $t = 2j - n$ to park the vertex on the integer $k = j$.

Block 4. Evaluate. Every term becomes $\pm 2^{j^2-(k-j)^2}$, so after dividing by the largest term the whole problem becomes one universal sum $T_j = \sum(-1)^m 2^{-m^2}$ truncated at both ends.

Block 5. Estimate, and meet the threshold face on: the peak term 1 is exactly cancelled by its two neighbours, because $2 \cdot 2^{-1} = 1$. The sign is therefore decided by the $m = \pm 2$ terms, which is why $n \ge 3$ is needed and why $n = 2$ really does have a double root.

Block 6. Halve the work with the reflection $T_j = T_{n-j}$, so only $2j \le n$ needs estimating.

Block 7. Intermediate value theorem plus the factor theorem: n sign changes give n distinct real roots, and a degree n polynomial has no others.

Juxtaposition. The nominal route is: expand, spot the palindrome, substitute $y = x + 1/x$, get a polynomial of degree $\lfloor n/2 \rfloor$ in y, and now try to prove that all of *its* roots are real and of modulus at least 2. You have replaced one real rootedness problem by a harder one and lost the clean coefficient structure. The bifurcation point is the sentence "I will not look for the roots, I will look for the signs", taken the moment you notice you cannot factor. Everything after that is bookkeeping.

Layer 2: HOW

Block 1. **Cue.** Coefficients that are all powers of a single base, with an exponent that is visibly a quadratic in the index. **Why natural.** When coefficients are $b^{Q(k)}$, sums and ratios of coefficients are governed by differences of Q, not by Q itself. Computing the second difference of Q is therefore the cheapest possible probe, and it lands exactly on the log concavity ratio that controls real rootedness. **Link to the next block.** Knowing the ratio is exactly 4 tells you two things at once: real rootedness is plausible (this is the Hutchinson threshold), and any estimate you attempt will be tight rather than comfortable.

Block 2. **Cue.** "Show that all the roots are real", with no value asked for. **Why natural.** There are only three elementary certificates for real rootedness: exhibit the factorisation, realise the polynomial as a characteristic polynomial of a real symmetric matrix or a member of an orthogonal polynomial family, or exhibit enough sign changes. The first two need extra structure that is not on offer here, so the third is forced. This is elimination, not inspiration. **Link to the next block.** The certificate demands $n + 1$ test points. The rest of the problem is the choice of those points.

Block 3. **Cue.** A free parameter (which points?) meeting a coefficient family that is a power of two. **Why natural.** Choose the unknown to make the algebra trivial, not to make the geometry pretty. Substituting $x = -2^t$ replaces the exponent $k(n-k)$ by $k(n-k) + kt$, which is the same downward parabola in k with vertex at $(n+t)/2$. Since the terms are largest near the vertex and you want a clean dominant term, you want the vertex on an integer: $t = 2j - n$. Nothing here is inspired; it is completing the square in the exponent, which you would do to any quadratic. **Link to the next block.** Once the vertex is at $k = j$, the exponent is $j^2 - (k-j)^2$ and the whole family of $n+1$ evaluations is one sum in disguise.

Block 4. **Cue.** A sum whose terms all share the factor 2^{j^2} and whose signs alternate with k. **Why natural.** Normalise by the largest term before estimating anything. Dividing by $(-1)^j 2^{j^2}$ removes both j and n from the summand and leaves $\sum_m (-1)^m 2^{-m^2}$, with n and j surviving only in the truncation limits. A quantity that depends on the data only through its limits of summation is a quantity you can estimate once and reuse. **Link to the next block.** The estimate now needs only the shape of $g(m) = 2^{-m^2}$: positive, even, strictly decreasing away from 0, and dropping fast.

Block 5. **Cue.** $g(0) = 1$ and $2g(1) = 1$: the two neighbours cancel the peak exactly. **Why natural.** This is the promised threshold showing up in person. It is not bad luck; it is the identity $c_k^2 = 4c_{k-1}c_{k+1}$ read at the test point, and it means that the three central terms are the local quadratic $c_{k-1} + c_k y + c_{k+1} y^2$ evaluated at its own double root. When the leading estimate vanishes identically, the standard move is to go one term further out, and $2g(2) = 2^{-3}$ is the first surviving contribution. **Link to the next block.** Because the rescue comes from $m = \pm 2$, you need those indices to exist, which is exactly the condition $n \ge 3$; and $n \le 2$ must be checked by hand, where indeed f_2 has a repeated root.

Block 6. **Cue.** The truncation limits $-j$ and $n-j$ are symmetric under $j \mapsto n-j$, and the summand is even. **Why natural.** Symmetry of the data must be used before, not after, the estimating. Here it is free: one line of reindexing removes half the cases and, as a bonus, it re-derives the palindromic symmetry of f_n without any polynomial manipulation. **Link to the next block.** With $2j \le n$ assumed, the tail R_j starts at $m = j+1$ and is small compared with u_j except in the single case $j = 1$, which you then handle on its own.

Block 7. **Cue.** n sign changes in hand for a polynomial of degree n. **Why natural.** The intermediate value theorem produces roots; the factor theorem caps their number. Together they upgrade "at least n real roots" to "exactly n real roots, all simple, and no complex ones", which is strictly stronger than what was asked and costs no extra ink. **Link to the next block.** Nothing follows; state the interval for each root, since you have it for free and it shows the grader the roots are distinct.

Layer 3: WHY

The mathematical truth underneath is that real rootedness of a polynomial with positive coefficients is a statement about how strongly log concave the coefficient sequence is, and the threshold is the number 4 because $c_k^2 - 4c_{k-1}c_{k+1}$ is the discriminant of the local quadratic $c_{k-1} + c_k y + c_{k+1}y^2$. If the coefficients decay fast enough on both sides of any given index, then near the point where the kth term dominates, the polynomial behaves like that quadratic, and a quadratic with nonnegative discriminant has a real root. Hutchinson's theorem is the honest global version of this local picture. This problem is the extremal case: the discriminant is zero for every k, every local quadratic is a perfect square, and so the polynomial is as close to failing as a real rooted polynomial can be. That is precisely why the two neighbours cancel the peak exactly at each test point, and why the whole proof hangs on a term of size 2^{-4} rather than a term of size 1.

The discernment mechanism is a two step filter. Step one: when asked to prove a qualitative property of roots, list the certificates for that property and cross off the ones the given data cannot supply. For real rootedness the list is short, so this step is fast and decisive. Step two: when a certificate needs free parameters (here, the test points), choose the parameters to make the given structure degenerate, not to make the numbers small. The structure here is a parabola in the exponent, and a parabola degenerates when you slide its vertex onto a lattice point; that is what $x = -2^{2j-n}$ does. Train the second step by taking any coefficient family of the form $b^{Q(k)}$ and asking, before anything else, what substitution makes Q symmetric about an integer.

The examiners nudged in three ways. They wrote the coefficient as $2^{k(n-k)}$ rather than as a number, which is an instruction to read the exponent. They chose base 2, so that $2 \cdot 2^{-1} = 1$ produces exact cancellation rather than approximate cancellation, forcing you to be honest about the next order term instead of hand waving an inequality. And they chose $k(n-k)$ symmetric in $k \leftrightarrow n-k$, which hands you the reflection that halves the case analysis. Any base $b > 2$ in place of 2 makes the problem strictly easier, because then $2b^{-1} < 1$ and the peak term survives on its own.

Revolving concepts

Log concavity and the Hutchinson threshold. A positive sequence (c_k) is log concave if $c_k^2 \ge c_{k-1}c_{k+1}$. The relevant strengthening here is $c_k^2 \ge 4c_{k-1}c_{k+1}$: Hutchinson (1923) proved that this forces all zeros of $\sum c_k x^k$ to be real and negative, and Kurtz proved that strict inequality forces them also to be simple. Our sequence achieves equality everywhere, which is the extremal case. See Book 1, field: Polynomials, topic: real rootedness criteria.

Newton's inequalities. If $\sum_k a_k x^k$ of degree n has all real roots then $p_k^2 \ge p_{k-1}p_{k+1}$ where $p_k = a_k / \binom{n}{k}$. Necessary only. Useful for disproving real rootedness, never for proving it. See Book 1, field: Inequalities, topic: symmetric functions.

Eneström and Kakeya. If $0 < a_0 \le a_1 \le \cdots \le a_n$ then every complex root of $\sum a_k x^k$ has modulus at most 1; rescaling gives the annulus $r \le |z| \le R$ with $r = \min_k a_{k-1}/a_k$ and $R = \max_k a_{k-1}/a_k$ for any positive coefficients. Here $a_{k-1}/a_k = 2^{2k-n-1}$, so every root lies in $2^{1-n} \le |z| \le 2^{n-1}$, consistent with the intervals we produced. This is the same lemma used in the Elkies solution to 2014 A5, so install it once and bill it twice. See Book 1, field: Polynomials, topic: root location.

Palindromic polynomials. $c_k = c_{n-k}$ is equivalent to $x^n f(1/x) = f(x)$, which pairs each root z with $1/z$ and lets you substitute $y = x + 1/x$ when n is even. In our proof this symmetry appears as the identity $T_j = T_{n-j}$ and as the pairing $x_j \leftrightarrow x_{n-j}$ of test points. See Book 1, field: Polynomials, topic: reciprocal equations.

Alternating sums with rapidly decreasing terms. If $t_1 > t_2 > \cdots > t_N > 0$ then $\sum_{i=1}^{N} (-1)^{i-1} t_i$ lies strictly between $t_1 - t_2$ and t_1. That two sided bound, not merely the one sided Leibniz bound, is what makes both u_j and R_j controllable. See Book 1, field: Sequences and recurrences, topic: alternating series estimates.

Gaussian binomials and partial theta. The Gaussian binomial $\binom{n}{k}_q$ has degree exactly $k(n-k)$ in q, so $\sum_k 2^{k(n-k)} x^k$ is the leading exponent shadow of $\sum_k \binom{n}{k}_2 x^k$. Equally, the substitution $x = 2^{-n} y$ turns our polynomial into $\sum_{k=0}^{n} 2^{-k^2} y^k$, a truncation of the partial theta series $\sum_{k \ge 0} q^{k^2} y^k$ at $q = 1/2$. Recognising the normalised form is what tells you the answer cannot depend on n except through truncation. See Book 1, field: Generating functions, topic: q series.

The recursion, generalised. Since $2^{k(n-k)} = 2^{k(n-1-k)} \cdot 2^k$, we get $f_n(x) = f_{n-1}(2x) + x^n$, and by palindromy also $f_n(x) = 1 + x^n f_{n-1}(2/x)$. The general statement worth carrying: if $c_k^{(n)} = b^{\,Q(k,n)}$ with $Q(k,n) - Q(k,n-1) = \lambda k$ for a constant λ, then $f_n(x) = f_{n-1}(b^\lambda x) + c_n^{(n)} x^n$. The solution method for such a family is never to iterate the recursion but to use it as a consistency check on evaluations: it says $f_n(x_j) = f_{n-1}(2x_j) + x_j^n$ where $2x_j = -2^{2j-(n-1)}$ is again a test point for $n-1$, so the sign patterns for consecutive n are compatible. The same pattern (coefficient recursion becomes a dilation of the argument) recurs in every problem whose coefficients are q powers, including the Rogers and Szegő polynomials and every q binomial identity.

Circuit Bank

- **WHEN** you must prove every root of a degree n polynomial is real and you cannot factor it **DO** hunt for $n+1$ points at which the sign alternates, then invoke the intermediate value theorem and the factor theorem **BECAUSE** sign changes are the only elementary certificate that is both available without extra structure and strong enough to force all roots real

- **WHEN** coefficients have the form $b^{Q(k)}$ with Q a quadratic in the index **DO** test at $x = -b^t$ and complete the square in the exponent, choosing t so that the vertex sits on an integer **BECAUSE** multiplying by x^k adds a linear function of k to the exponent, which translates the parabola without changing its shape, so you may place its peak wherever you like

- **WHEN** a sum has one visibly dominant term **DO** divide through by that term before estimating anything **BECAUSE** normalising removes the parameters from the summand and leaves them only in the limits of summation, so one estimate serves every case

- **WHEN** the coefficients satisfy $c_k^2 = 4c_{k-1}c_{k+1}$ exactly **DO** expect the leading estimate to vanish identically and plan to use the next term out **BECAUSE** four is the discriminant threshold of the local quadratic $c_{k-1} + c_k y + c_{k+1} y^2$, so equality means that quadratic is a perfect square and its value at the double root is zero

- **WHEN** the data are symmetric under $k \mapsto n - k$ **DO** reindex to reduce to $2k \le n$ before estimating **BECAUSE** a symmetry used at the start halves the case analysis, whereas a symmetry noticed at the end saves nothing

Backtesting

1. First useful cue: the coefficient is written as a power of two with a quadratic exponent, so the object to analyse is the exponent $k(n-k)$, not the coefficient.
2. Method the expression suggested: exhibit n sign changes at test points chosen to complete the square in the exponent, namely $x = -2^{2j-n}$.
3. Where you would hesitate, and the thought to reinforce: at the moment the three central terms cancel to zero and it looks as though the method has failed. Reinforce this: an exact cancellation at the leading order is information, not failure, and it is exactly what the identity $c_k^2 = 4c_{k-1}c_{k+1}$ predicted. Go one term further out and continue.
4. The trap the question sets: $n = 2$ gives a genuine repeated root, so any argument claiming strict sign alternation for all n is false. The small cases are not decoration; they are where the threshold bites.
5. What makes you recognise it faster next time: the reflex of computing $c_k^2/(c_{k-1}c_{k+1})$ the instant you see positive coefficients plus a real rootedness claim. If the answer is a clean constant, the problem is a log concavity problem and the test points will be geometric.
6. The reusable decision rule: to prove real rootedness, choose test points that make one term dominate; to choose them, complete the square in the exponent of the coefficient.

Mirror drills

1. Show that all roots of $\sum_{k=0}^{n} 3^{k(n-k)}x^k$ are real, for every $n \geq 1$. *Answer:* the same test points $x_j = -3^{2j-n}$ give $\sum_m (-1)^m 3^{-m^2}$; now $2 \cdot 3^{-1} < 1$, so the peak term already dominates and $T_j \geq 1 - 2/3 - 2/3^4 - \cdots > 0$ with no threshold case and no need for $n \geq 3$. Base 2 was the only hard base.
2. Prove that $q_n(y) = \sum_{k=0}^{n} 2^{-k^2}y^k$ has all roots real and negative. *Answer:* $q_n(y) = f_n(2^{-n}y)$, so this is the same statement after a dilation; the test points become $y = -2^{2j}$.
3. Locate all roots of $\sum_{k=0}^{n} 2^{k(n-k)}x^k$ inside an annulus using only the coefficients. *Answer:* $a_{k-1}/a_k = 2^{2k-n-1}$ for $k = 1, \ldots, n$, so the Eneström and Kakeya corollary gives $2^{1-n} \leq |z| \leq 2^{n-1}$; compare with the intervals $(-2^{2j-n}, -2^{2j-2-n})$ produced by the sign argument.

12.11 B5. Parity of centralisers decides a matrix game

In the 75th annual Putnam Games, participants compete at mathematical games. Patniss and Keeta play a game in which they take turns choosing an element from the group of invertible $n \times n$ matrices with entries in the field $\mathbb{Z}/p\mathbb{Z}$ of integers modulo p, where n is a fixed positive integer and p is a fixed prime number. The rules of the game are:

(1) A player cannot choose an element that has been chosen by either player on any previous turn.

(2) A player can only choose an element that commutes with all previously chosen elements.

(3) A player who cannot choose an element on his or her turn loses the game.

Patniss takes the first turn. Which player has a winning strategy? (Your answer may depend on n and p.)

Read it out loud

Read the rules as algebra, not as a game. Rule (2) says that the set S of elements chosen so far is a commuting set, and that the legal moves are exactly the elements commuting with everything in S. That set of elements has a name: the centraliser $Z(S)$, and it is a subgroup. Rule (1) says you may not repeat. So the position is a set S, the move set is $Z(S) \setminus S$, and since every element of S commutes with every element of S we always have $S \subseteq Z(S)$. The game therefore ends exactly when $S = Z(S)$.

That last sentence is the whole problem, and you should have it inside thirty seconds. It says the total number of moves in a completed game is $|S| = |Z(S)|$, the order of a subgroup. This is a normal play game, so the loser is whoever faces an empty move set, which means the winner is decided by the parity of that group order and by nothing else. You do not need to know good moves from bad moves. You need to know which parities are reachable.

The fields wired in are finite group theory (centralisers, Lagrange, Cauchy), linear algebra over a finite field, and the small piece of combinatorial game theory that says a normal play game is decided by the parity of its length when that length is forced. What is really being tested is whether you will reach for Sprague and Grundy out of habit, or whether you will notice that this game has no independent components to add up and that its length is an invariant.

The first useful thing to write down after that is the observation that decides everything: $Z(S)$ shrinks as S grows, and each new centraliser is a subgroup of the previous one, so by Lagrange the orders divide each other. Parity is therefore a one way street. Once even, always even.

Approaches, ranked worst to best

1. **Play the small cases by hand.** *Cue:* "your answer may depend on n and p" invites experiments. *Circuit:* compute, then guess. $|GL_1(\mathbb{Z}/3\mathbb{Z})| = 2$ and $|GL_2(\mathbb{Z}/3\mathbb{Z})| = 48$; the second is already unplayable by hand and the third is out of the question. Ranked last as a proof method. Ranked first as a five minute sanity check: $n = 1$, $p = 2$ gives the trivial group, Patniss takes the identity and wins, which already tells you the answer is not "Keeta always".

2. **Sprague and Grundy.** *Cue:* an impartial normal play game. *Circuit:* compute Nim values of positions. The standard circuit, and here it misfires. Grundy theory earns its keep when a position splits into independent components whose values are added; this position is a single commuting set with no decomposition, and the position space is the set of all commuting subsets of a group of order roughly p^{n^2}. Recognising that a familiar circuit does not apply is itself a skill, and the tell is the absence of any sum structure.

3. **Pairing strategy for odd p.** *Cue:* the scalar matrix $-I$ is central and has order two. *Circuit:* answer every move g with a fixed involution applied to it, here $-g$. Excellent and almost free: it wins the odd case outright with no theory beyond the observation that $-g$ commutes with whatever g commutes with. It ranks below the winner only because it gives no purchase at all on $p = 2$, where no such universal involution exists.

4. **Parity of the centraliser chain.** *Cue:* the move set is a subgroup and the game ends when $S = Z(S)$. *Circuit:* count the length of the game instead of searching it. The winner. The length of a completed game is the order of the terminal subgroup; centralisers only shrink; subgroup orders only divide; therefore the parity gets locked as soon as either a chosen element has even order or some centraliser has odd order. This reduces the whole problem to a single question about the group: does some element have a centraliser of odd order?

The ideal solution

Let $G = GL_n(\mathbb{Z}/p\mathbb{Z})$, a finite group, and for a subset $S \subseteq G$ write

$$Z(S) = \{h \in G : hg = gh \text{ for all } g \in S\},$$

the centraliser of S, which is a subgroup of G (it is the intersection over $g \in S$ of the subgroups $Z(\{g\})$). Write S_t for the set of the first t chosen elements, so $S_0 = \emptyset$ and $Z(S_0) = G$.

Lemma 12.11.1. *The game terminates, and if it lasts N moves then $S_N = Z(S_N)$; in particular $N = |Z(S_N)|$ is the order of a subgroup of G.*

Proof. Each move removes one element from the finite set G, so the game terminates. By rule (2), any two chosen elements commute, so every element of S_t commutes with all of S_t, that is, $S_t \subseteq Z(S_t)$. By rules (1) and (2) the legal moves at position S_t are precisely the elements of $Z(S_t) \setminus S_t$. Hence the game ends at S_N if and only if $Z(S_N) \subseteq S_N$, which together with $S_N \subseteq Z(S_N)$ gives $S_N = Z(S_N)$. Finally $N = |S_N|$. □

Lemma 12.11.2. *If $S \subseteq S'$ then $Z(S') \leq Z(S)$, so $|Z(S')|$ divides $|Z(S)|$. Consequently, in any play of the game:*

(a) if $|Z(S_t)|$ is odd for some t, then N is odd;

(b) if some chosen element has even order, then N is even.

Proof. The inclusion is immediate from the definition, and the divisibility is Lagrange's theorem. For (a), $Z(S_N) \leq Z(S_t)$, so $N = |Z(S_N)|$ divides an odd number. For (b), if $g \in S_t$ has even order then $g \in S_N \subseteq Z(S_N)$, so $\operatorname{ord}(g)$ divides $N = |Z(S_N)|$ by Lagrange, and N is even. □

Proposition 12.11.3. *Patniss has a winning strategy if and only if there exists $g \in G$ with $|Z(\{g\})|$ odd.*

Proof. Since the players alternate with Patniss first, the Nth and final move is made by Patniss exactly when N is odd, and then Keeta is the player who cannot move. So Patniss wins if and only if N is odd.

Suppose such a g exists. Patniss opens with g. By Lemma 2(a) applied at $t = 1$, N is odd whatever either player does afterwards, so Patniss wins.

Suppose instead that $|Z(\{g\})|$ is even for every $g \in G$, and let g be Patniss's opening move. By Cauchy's theorem, the group $Z(\{g\})$ contains an element h of order 2. If $h \neq g$ then h is a legal reply for Keeta, and after it Lemma 2(b) forces N to be even. If $h = g$ then g itself has order 2 and Lemma 2(b) already forces N to be even; note that then $N = |Z(S_N)|$ is even and hence at least 2, so Keeta does have a legal reply and no special care is needed. Either way N is even and Keeta wins. □

Case p odd. The scalar matrix $-I$ lies in the centre of G, and $-I \neq I$ because $p > 2$, so $-I$ has order 2 and lies in $Z(\{g\})$ for every g. By Lagrange, $|Z(\{g\})|$ is even for every g, so by the Proposition Keeta wins. Explicitly, Keeta may answer every move g with $-g$: this commutes with everything that g commutes with because $-I$ is central, it is distinct from g because $2g = 0$ would force $g = 0$, and it is unchosen because the chosen set is stable under $x \mapsto -x$ after each of Keeta's replies.

Case $p = 2$. We exhibit g with $|Z(\{g\})|$ odd. Let $\mathbb{F}_{2^n}$ be the field with 2^n elements. As an $\mathbb{F}_2$ vector space it has dimension n, so choosing an $\mathbb{F}_2$ basis identifies it with the space of column vectors $\mathbb{F}_2^n$ and identifies $\mathbb{F}_2$ linear bijections of $\mathbb{F}_{2^n}$ with elements of G. The multiplicative group $\mathbb{F}_{2^n}^\times$ is cyclic; let α be a generator, and let $g \in G$ be the map $v \mapsto \alpha v$, which is $\mathbb{F}_2$ linear and bijective.

Suppose $h \in G$ commutes with g. Then $h(\alpha v) = \alpha h(v)$ for all v, and iterating gives $h(\alpha^i v) = \alpha^i h(v)$ for all $i \geq 0$; also $h(0 \cdot v) = 0 = 0 \cdot h(v)$. Every element of $\mathbb{F}_{2^n}$ is either 0 or a power of α, so $h(cv) = c\,h(v)$ for all $c \in \mathbb{F}_{2^n}$ and all v. Taking $v = 1$ gives $h(c) = c\,h(1)$ for all c, that is, h is multiplication by $\beta := h(1)$, and $\beta \neq 0$ since h is invertible. Conversely every such multiplication commutes with g. Hence

$$Z(\{g\}) \cong \mathbb{F}_{2^n}^\times, \qquad |Z(\{g\})| = 2^n - 1,$$

which is odd. By the Proposition, Patniss wins.

Patniss wins if $p = 2$; Keeta wins if $p > 2$; in both cases for every n.

Remark 12.11.4. Once $Z(S)$ is abelian the rest of the game is forced: every remaining element of $Z(S)$ commutes with every other, so play continues until all of them are used. With the g above, $Z(\{g\})$ is abelian and the game lasts exactly $2^n - 1$ moves. For $n = 1$, $p = 2$ the group is trivial, $2^1 - 1 = 1$, and Patniss wins by taking the identity.

Remark 12.11.5. The construction is not what separates $p = 2$ from odd p. The same g exists over any $\mathbb{F}_p$ and gives $|Z(\{g\})| = p^n - 1$. What separates the cases is only that $p^n - 1$ is odd precisely when $p = 2$. It is recorded in the official solution, and we do not prove it here, that for $p = 2$ the elements with centraliser of odd order are exactly those with n distinct eigenvalues in an algebraic closure; anything else fixes a nontrivial unipotent or a copy of $GL_k(\mathbb{F}_{2^d})$ with $k > 1$ inside its centraliser, and both have even order.

Layer 1: WHAT

Block 1. Translate the rules: position is the chosen set S, move set is $Z(S) \setminus S$.

Block 2. Rule (2) forces $S \subseteq Z(S)$; the chosen set is a commuting set sitting inside its own centraliser.

Block 3. Therefore the game ends exactly at $S = Z(S)$, so the length of a completed game is the order of a subgroup.

Block 4. Centralisers shrink as S grows and orders divide by Lagrange, so parity is a one way street: even stays even.

Block 5. Convert to a group question: Patniss wins if and only if some element has a centraliser of odd order, with Cauchy's theorem supplying the involution in the converse direction.

Block 6. Odd p: the centre contains the involution $-I$, so every centraliser has even order; Keeta wins.

Block 7. $p = 2$: realise $\mathbb{F}_2^n$ as the field $\mathbb{F}_{2^n}$ and take g to be multiplication by a generator of $\mathbb{F}_{2^n}^\times$; anything commuting with g is $\mathbb{F}_{2^n}$ linear, hence multiplication by a scalar, so $|Z(\{g\})| = 2^n - 1$ is odd and Patniss wins.

Juxtaposition. The nominal route treats this as a game: enumerate positions, look for a pairing or a mirroring strategy, try Grundy values, try to find "good" first moves by hand in $GL_2(\mathbb{Z}/2\mathbb{Z})$. It is not wrong so much as unbounded; there are far too many commuting subsets. The bifurcation point is one sentence: *stop asking which moves are good and ask how long the game is*. Once you ask

that, the terminal condition $S = Z(S)$ appears immediately, and with it a group order, and with it Lagrange.

Layer 2: HOW

Block 1. **Cue.** A rule of the form "you may only choose something compatible with everything already chosen". **Why natural.** Any such rule defines the move set as a closure operation applied to the current position. Naming that operation is the first move in every problem of this type, because a named operation has theorems attached, whereas an unnamed rule has none. Here the operation is the centraliser, and its first theorem is that it is a subgroup. **Link to the next block.** A move set that is a subgroup has an order, and orders have parity, so the game is now measurable.

Block 2. **Cue.** Every pair of chosen elements commutes, by rule (2) applied to whichever was chosen later. **Why natural.** This is the step people skip. The rule as written is asymmetric (the new element commutes with the old ones), and you must notice that it therefore makes the whole chosen set mutually commuting, so that $S \subseteq Z(S)$. Without this containment the terminal condition is not an equation and the count does not close. **Link to the next block.** With $S \subseteq Z(S)$ always true, the only way the move set can be empty is $S = Z(S)$.

Block 3. **Cue.** The game ends when the move set is empty. **Why natural.** Emptiness of $Z(S) \setminus S$ combined with $S \subseteq Z(S)$ is an equality of sets, not merely an inequality of sizes, and an equality of sets converts a dynamic quantity (the number of moves played) into a static one (the order of a subgroup). Converting a process into an invariant is the standard way to defeat a game whose state space is too large to search. **Link to the next block.** The number N of moves is now $|Z(S_N)|$, and normal play means Patniss wins if and only if N is odd.

Block 4. **Cue.** A nested chain $G = Z(S_0) \geq Z(S_1) \geq \cdots$ of subgroups. **Why natural.** Lagrange is the only tool needed: each order divides the previous, so an odd order can never become even and an even order can never become odd. Parity, unlike size, is monotone along the chain. This is why you never have to analyse the middle of the game. **Link to the next block.** So the parity of N is decided at the moment when the chain first becomes committed, and the first move is already such a moment if it makes the centraliser odd.

Block 5. **Cue.** "Even order" needs to be converted into a concrete element that a player can actually choose. **Why natural.** Cauchy's theorem is the standard bridge from a divisibility statement about $|H|$ to the existence of an element of order 2 in H. Keeta does not need a strategy for the whole game, she needs one element whose presence in S locks the parity for good, and an involution is exactly that. **Link to the next block.** The criterion is now purely about G: does some element have a centraliser of odd order? Answer it separately for odd p and for $p = 2$.

Block 6. **Cue.** p odd means $-1 \neq 1$ in $\mathbb{Z}/p\mathbb{Z}$, and scalar matrices are central. **Why natural.** The centre lies inside every centraliser, so a single central element of order 2 settles every case at once. When you need "this happens for all g", look for the reason that does not mention g; the centre is that reason. **Link to the next block.** For $p = 2$ the scalars are only $\{1\}$, the centre is trivial in the relevant sense, and this argument dies, which is precisely the signal that $p = 2$ is the other case rather than a special case of the same.

Block 7. **Cue.** You need an element of $GL_n(\mathbb{F}_2)$ whose centraliser has odd order, so you want the algebra of matrices commuting with it to be as far from a full matrix algebra as possible. **Why natural.** Centralisers measure how much a matrix fails to be generic, and the cleanest small case

is a matrix that generates a field. If $\mathbb{F}_2[g]$ is all of $\mathbb{F}_{2^n}$, then commuting with g upgrades to being linear over the big field, and the space of $\mathbb{F}_{2^n}$ linear endomorphisms of a one dimensional $\mathbb{F}_{2^n}$ space is just $\mathbb{F}_{2^n}$ itself. The identification $\mathbb{F}_2^n \cong \mathbb{F}_{2^n}$ is not a change of viewpoint for its own sake; it is the only way to make "commutes with g" say something about scalars. **Link to the next block.** $|Z(\{g\})| = 2^n - 1$ is odd, the criterion is met, and the case analysis is complete.

Layer 3: WHY

The mathematical truth beneath this problem is that a normal play game whose length is an invariant of the terminal position is not a game at all; it is a parity computation. Grundy theory exists to handle games that decompose into independent sums, and its machinery is wasted here because the position never decomposes. What replaces it is Lagrange's theorem, which controls how the move set can change: the move set is always a subgroup minus a subset, and subgroups of subgroups have orders that divide. Divisibility is a much stronger constraint than inclusion, and it is what makes the parity monotone.

The second truth is about centralisers in GL_n over a finite field. The centraliser of a single matrix is the unit group of the algebra of matrices commuting with it, and that algebra is a field exactly when the characteristic polynomial is irreducible, in which case the centraliser is $\mathbb{F}_{p^n}^{\times}$ of order $p^n - 1$. More generally, if the characteristic polynomial is a product of distinct irreducibles of degrees $d_1, \ldots, d_s$ then the centraliser has order $\prod_i (p^{d_i} - 1)$, which for $p = 2$ is a product of odd numbers whatever the degrees are. So the arithmetic of the whole problem funnels into the single question of whether numbers of the form $p^d - 1$ can be odd, which happens exactly when $p = 2$. That is the reason the answer depends on p and not on n, and you can see it before doing any of the work. It also explains why for $n = 3$, $p = 2$ there are elements with centralisers of order 3 as well as of order 7: the splitting types $(1, 2)$ and (3) both give odd products.

The discernment mechanism has three steps, and it is worth naming them because they transfer. First, when a game's rules constrain moves by compatibility with all previous moves, define the closure of the position and identify it algebraically; a rule you can name has theorems, a rule you cannot name has only cases. Second, ask whether the number of moves in a completed play is determined by the terminal position; if it is, abandon strategy and compute parity. Third, when a parity is controlled by a group order, look for the cheapest universal element forcing that parity: the centre for one direction, and an element generating a field for the other. Train this by replaying the same game on small groups: on any group of odd order the first player wins for free, and on any group with a central involution the second player wins for free, so the only interesting groups are those of even order with centre of odd order.

The examiners nudged you deliberately. They wrote rule (3) as "a player who cannot choose an element loses", which is normal play and is the phrasing that makes parity decisive; had they written a misere rule the analysis would change. They said the answer may depend on n and p, which is a hint that one of the two parameters is a red herring and invites you to find the reason n drops out. And they chose GL_n over $\mathbb{Z}/p\mathbb{Z}$ rather than an abstract group, because in an abstract group you would have no way to build an element with odd centraliser, whereas here the finite field structure hands it to you.

Revolving concepts

Centraliser and centre. For $S \subseteq G$, $Z(S) = \{h : hg = gh \ \forall g \in S\}$ is a subgroup; $Z(G)$ is the centre and lies inside every centraliser; $S \subseteq T$ implies $Z(T) \leq Z(S)$; and $Z(S) = Z(\langle S \rangle)$, so centralisers only see the subgroup generated by S. See Book 1, field: Abstract algebra, topic: centralisers and conjugacy.

Lagrange and Cauchy. Lagrange: the order of a subgroup divides the order of the group, and the order of an element divides the order of the group. Cauchy: if a prime q divides $|H|$ then H

contains an element of order q; for $q = 2$ there is also the one line pairing proof, since pairing each element with its inverse leaves the involutions and the identity unpaired. Both are the workhorses here. See Book 1, field: Abstract algebra, topic: Lagrange and Cauchy.

Order of the general linear group. $|GL_n(\mathbb{F}_p)| = \prod_{i=0}^{n-1}(p^n - p^i) = p^{n(n-1)/2}\prod_{i=1}^{n}(p^i - 1)$. For p odd this is even (already $p^1 - 1$ is even), and for $p = 2$ it is $2^{n(n-1)/2}\prod_{i=1}^{n}(2^i - 1)$, whose odd part is $\prod(2^i - 1)$. See Book 1, field: Matrices and determinants, topic: counting invertible matrices over a finite field.

Finite fields. For every prime p and every $n \geq 1$ there is a field $\mathbb{F}_{p^n}$ with p^n elements, unique up to isomorphism, and its multiplicative group is cyclic of order $p^n - 1$. Equivalently there exists an irreducible polynomial of degree n over $\mathbb{F}_p$, whose companion matrix is an element of $GL_n(\mathbb{F}_p)$ generating a copy of $\mathbb{F}_{p^n}$ inside the matrix algebra. This is the one genuine knowledge gap in the problem: without the existence of $\mathbb{F}_{p^n}$ and the cyclicity of its unit group there is no construction. See Book 1, field: Abstract algebra, topic: finite fields.

Singer cycles and non derogatory matrices. An element of $GL_n(\mathbb{F}_p)$ of order $p^n - 1$ arising as multiplication by a generator of $\mathbb{F}_{p^n}^{\times}$ is called a Singer cycle. More generally, if the minimal polynomial of g equals its characteristic polynomial then the matrices commuting with g are exactly the polynomials in g; the field case above is the special case where that algebra is a field. See Book 1, field: Linear algebra, topic: cyclic and non derogatory matrices.

Normal play, parity and the descending chain. The recursion in this problem is the descending chain $Z(S_0) \geq Z(S_1) \geq \cdots \geq Z(S_N) = S_N$, with the recursion step $Z(S_{t+1}) = Z(S_t) \cap Z(\{g_{t+1}\})$ for the newly chosen g_{t+1}. Generalised statement: play this game on any finite group G. The length of any completed play is $|Z(S_N)|$, so the first player wins if and only if some $g \in G$ has $|Z(\{g\})|$ odd. Two immediate corollaries: if $|G|$ is odd the first player always wins, and if $|Z(G)|$ is even the second player always wins. Solution method for the general recursion: never iterate it, instead find one invariant of the chain (here, the parity of the order) that is monotone, and identify the earliest move that commits it. See Book 1, field: Games and invariants, topic: normal play parity arguments.

Circuit Bank

- **WHEN** a game rule says a move must be compatible with every previous move **DO** define the position as the chosen set S, name the compatibility closure algebraically, and check whether S lies inside its own closure **BECAUSE** a named closure such as a centraliser or a span comes with theorems about its size, and the terminal condition becomes an equation rather than a case list
- **WHEN** the number of moves in a completed normal play game is determined by the terminal position **DO** abandon strategy entirely and compute the parity of that number **BECAUSE** the loser is whoever faces an empty move set, so the winner is a function of the length alone
- **WHEN** a chain of subgroups is forced to descend **DO** apply Lagrange to get divisibility, not merely inclusion, and look for a parity or a prime that can never come back **BECAUSE** divisibility makes coarse invariants monotone, which is what lets you decide the end of the game from its first move
- **WHEN** you need an element of order two inside a group of even order **DO** invoke Cauchy's theorem, or pair each element with its inverse and count the unpaired ones **BECAUSE** this is the standard bridge from a divisibility fact about the order to a concrete element you can actually play
- **WHEN** you need an element of $GL_n(\mathbb{F}_p)$ with the smallest possible centraliser **DO** identify $\mathbb{F}_p^n$ with $\mathbb{F}_{p^n}$ and take multiplication by a generator of the multiplicative group **BECAUSE** commuting with a generator upgrades to linearity over the larger field, and a one dimensional space has only scalar endomorphisms, so the centraliser collapses to $\mathbb{F}_{p^n}^{\times}$

Backtesting

1. First useful cue: rule (2) is a compatibility rule, so the legal moves form the centraliser $Z(S)$ and the position satisfies $S \subseteq Z(S)$.
2. Method the expression suggested: compute the length of the game rather than the strategy, because the game ends at $S = Z(S)$ and therefore lasts exactly $|Z(S_N)|$ moves.
3. Where you would hesitate, and the thought to reinforce: at the converse half of the criterion, where it feels as though Keeta needs a full strategy. She does not. She needs one involution, after which the parity is locked no matter how badly either player plays. Reinforce this: in a parity locked game, a single committing move replaces an entire strategy.
4. The trap the question sets: the invitation to compute Grundy values because the game is impartial. There is no sum decomposition here, so Grundy theory adds nothing and costs everything. The second trap is the phrase "may depend on n and p", which tempts you into an n dependent answer; n cancels because $p^n - 1$ is odd exactly when p is.
5. What makes you recognise it faster next time: the pattern "moves must commute with all previous moves" is the centraliser pattern, and the pattern "who cannot move loses" is the parity pattern. Seeing both in one problem should immediately produce the sentence "the length of the game is the order of a subgroup".
6. The reusable decision rule: when a game's move set is an algebraic closure of the position, compute the terminal condition as an equation and read off the length; then decide the winner with Lagrange.

Mirror drills

1. Play the same game on a finite group G of odd order. Who wins? *Answer:* the first player, always. Every completed play has length $|Z(S_N)|$, which divides $|G|$ and is therefore odd.
2. Play the same game on S_3. Who wins, and how long is the game? *Answer:* the first player, by opening with a 3 cycle σ, since $Z(\{\sigma\}) = \{e, \sigma, \sigma^2\}$ has order 3; the game then lasts exactly three moves. Opening with the identity or a transposition loses, since those centralisers have orders 6 and 2.
3. Show that if the centre of G has even order then the second player wins, and use this to re prove the case $p > 2$. *Answer:* $Z(G) \leq Z(S)$ for every S, so $|Z(S_N)|$ is always even and the game has even length; for $GL_n(\mathbb{Z}/p\mathbb{Z})$ with $p > 2$ the centre contains $-I$ of order two.

12.12 B6. Lipschitz plus integrality forces piecewise linearity

Let $f : [0,1] \to \mathbb{R}$ be a function for which there exists a constant $K > 0$ such that $|f(x) - f(y)| \leq K|x - y|$ for all $x, y \in [0,1]$. Suppose also that for each rational number $r \in [0,1]$, there exist integers a and b such that $f(r) = a + br$. Prove that there exist finitely many intervals $I_1, \ldots, I_n$ such that f is a linear function on each I_i and $[0,1] = \bigcup_{i=1}^{n} I_i$.

Read it out loud

Two hypotheses of completely different type are being pushed together, and the conclusion is a rigidity statement. The first hypothesis is analytic: f is Lipschitz, which is a bound on slopes and nothing else. The second is arithmetic: at every rational point the value of f lies in a lattice, and crucially the lattice depends on the point. The conclusion is that f is piecewise affine with finitely many pieces, that is, that a function satisfying both constraints has almost no freedom at all.

Whenever an analytic bound and an arithmetic quantisation are handed to you together, the intended mechanism is fixed: build a quantity that is simultaneously an integer and small, then conclude it is zero. That single sentence is the problem. Everything else is deciding which quantity, and the answer will be a slope, because slopes are the only thing the Lipschitz hypothesis knows about.

So convert the arithmetic hypothesis into something about denominators immediately. If $r = m/q$ in lowest terms and $f(r) = a + br$ with $a, b \in \mathbb{Z}$, then

$$qf(m/q) = aq + bm \in \mathbb{Z}.$$

In words: the denominator of $f(r)$ divides the denominator of r. That is the usable form, and you should write it inside thirty seconds, because it is the form in which slopes become integers. Indeed, for two rationals $m/q < m'/q'$ the slope of the chord is

$$\frac{f(m'/q') - f(m/q)}{m'/q' - m/q} = \frac{q'f(m'/q') \cdot q - qf(m/q) \cdot q'}{m'q - mq'},$$

whose numerator is an integer. So the slope is an integer as soon as $m'q - mq' = 1$. That condition, not smallness of $|r - r'|$, is the right notion of two rationals being adjacent, and recognising this is the whole discernment content of the problem.

The fields wired in are real analysis (Lipschitz continuity, density plus continuity), elementary number theory (Bezout, Farey neighbours, mediants, the Stern and Brocot refinement), and a small combinatorial termination argument to get from "locally affine" to "finitely many pieces".

Approaches, ranked worst to best

1. **Differentiate almost everywhere.** *Cue:* Lipschitz is the exact hypothesis of Rademacher's theorem. *Circuit:* reduce a regularity claim to a statement about the derivative. You get that f is differentiable almost everywhere with $|f'| \leq K$, and with some work that f' takes integer values at points where you can compute it. It fails, and it is important to know why rather than to feel vaguely unsatisfied: there exist 1 Lipschitz functions with $f' \in \{0,1\}$ almost everywhere that are not piecewise linear. Take a closed nowhere dense $C \subseteq [0,1]$ of positive measure and set $f(x) = \int_0^x \mathbf{1}_{[0,1]\setminus C}$. Then $f' = 1$ on the complement of C and $f' = 0$ at every density point of C, yet f is affine on no interval meeting C in positive measure. An almost everywhere statement is simply too weak; the hypothesis at every rational is much stronger and must be used pointwise.

2. **Prove local linearity at each point, then use compactness.** *Cue:* the conclusion is local plus finite. *Circuit:* cover $[0,1]$ by intervals of linearity and extract a finite subcover. This is the

alternative route recorded in the official solution, using the continued fraction convergents of each irrational x to approximate it by rationals with controlled denominators. It works, and it costs more: you need convergent estimates at every point, and then a separate compactness argument, and the final count of pieces is not explicit.

3. **Interpolate on the Farey sequence and show the interpolants stabilise.** *Cue:* denominators bounded by N give a finite grid. *Circuit:* define f_N to agree with f on the Farey sequence of order N and prove $f_N = f_{N-1}$ once $N > 2K$. This is the official solution, and it is good. It requires the Farey neighbour theorem as an input (consecutive fractions of order N satisfy $m'q - mq' = 1$ and $q + q' > N$), proved by its own induction, and it phrases the key computation in terms of the interpolant rather than in terms of f.

4. **Mediant refinement with a termination bound.** *Cue:* unimodular pairs make chord slopes integers; the mediant is the unique fraction of smallest denominator in between. *Circuit:* split each interval at the mediant until the denominators are large enough that a bounded integer must vanish. The winner, and the cleanest, because it never needs the Farey sequence as an external input: it constructs the right subdivision itself, the induction measure is explicit, and the termination argument is two lines. It also gives, for free, that each affine piece has integer slope and integer intercept.

The ideal solution

Every rational $r \in [0,1]$ is written $r = m/q$ in lowest terms with $q \geq 1$; thus $0 = 0/1$ and $1 = 1/1$. Call an ordered pair $r = m/q < r' = m'/q'$ of rationals in $[0,1]$ a *unimodular pair* if

$$m'q - mq' = 1.$$

Note that $(0/1, 1/1)$ is a unimodular pair. Since f is Lipschitz it is continuous, a fact used only at the very end.

Lemma 12.12.1. *Let (r, r') be a unimodular pair. Then $\gcd(q, q') = 1$, $r' - r = \dfrac{1}{qq'}$, and the mediant*

$$\mu = \frac{m + m'}{q + q'}$$

is in lowest terms, satisfies $r < \mu < r'$, and both (r, μ) and (μ, r') are unimodular pairs.

Proof. Any common divisor of q and q' divides $m'q - mq' = 1$, and $r' - r = (m'q - mq')/(qq') = 1/(qq')$. Any common divisor of $m + m'$ and $q + q'$ divides $(m + m')q - m(q + q') = m'q - mq' = 1$, so the mediant is in lowest terms. Finally

$$(m + m')q - m(q + q') = 1, \qquad m'(q + q') - (m + m')q' = 1,$$

which say exactly that (r, μ) and (μ, r') are unimodular; positivity of these two quantities also gives $r < \mu < r'$. □

Lemma 12.12.2. *For every rational $r = m/q$ in $[0,1]$ we have $qf(r) \in \mathbb{Z}$. Consequently, if (r, r') is a unimodular pair then the chord slope*

$$s(r, r') := \frac{f(r') - f(r)}{r' - r} = qq'\big(f(r') - f(r)\big)$$

is an integer with $|s(r, r')| \leq K$.

Proof. Write $f(r) = a + bm/q$ with $a, b \in \mathbb{Z}$; then $qf(r) = aq + bm \in \mathbb{Z}$. For the second claim, use $r' - r = 1/(qq')$ from Lemma 1 and write

$$qq'\big(f(r') - f(r)\big) = q\big(q'f(r')\big) - q'\big(qf(r)\big) \in \mathbb{Z}.$$

The bound $|s| \le K$ is the Lipschitz hypothesis applied to r and r'. □

Lemma 12.12.3. *Let (r, r') be a unimodular pair and let $\rho = u/v$ in lowest terms satisfy $r < \rho < r'$. Then $v \ge q + q'$, with equality if and only if ρ is the mediant μ.*

Proof. Both $uq - mv$ and $m'v - uq'$ are positive integers, so

$$\rho - r = \frac{uq - mv}{vq} \ge \frac{1}{vq}, \qquad r' - \rho = \frac{m'v - uq'}{vq'} \ge \frac{1}{vq'}.$$

Adding and using $r' - r = 1/(qq')$,

$$\frac{1}{qq'} \ge \frac{1}{vq} + \frac{1}{vq'} = \frac{q + q'}{vqq'},$$

so $v \ge q + q'$. If $v = q + q'$ then both displayed inequalities are equalities, so $uq - mv = 1$; substituting $v = q + q'$ and $m'q - mq' = 1$ gives $uq = m'q + mq$, hence $u = m + m'$ and $\rho = \mu$. □

Lemma 12.12.4. *Let (r, r') be a unimodular pair with $q + q' > 2K$, and let μ be its mediant. Then $f(\mu)$ lies on the chord joining $(r, f(r))$ to $(r', f(r'))$.*

Proof. Write $s = s(r, r')$, $s_1 = s(r, \mu)$, $s_2 = s(\mu, r')$, all integers of absolute value at most K by Lemmas 1 and 2. From Lemma 1,

$$r' - r = \frac{1}{qq'}, \qquad \mu - r = \frac{1}{q(q + q')}, \qquad r' - \mu = \frac{1}{q'(q + q')}.$$

Adding $f(\mu) - f(r)$ and $f(r') - f(\mu)$ gives

$$\frac{s}{qq'} = \frac{s_1}{q(q + q')} + \frac{s_2}{q'(q + q')},$$

and multiplying through by $qq'(q + q')$,

$$s(q + q') = s_1 q' + s_2 q.$$

Put $\delta_1 = s_1 - s$ and $\delta_2 = s_2 - s$, both integers. The displayed identity becomes $\delta_1 q' + \delta_2 q = 0$. Since $\gcd(q, q') = 1$, q divides δ_1; writing $\delta_1 = tq$ with $t \in \mathbb{Z}$ gives $\delta_2 = -tq'$, hence

$$s_1 - s_2 = t(q + q').$$

But $|s_1|, |s_2| \le K$, so $|t|(q + q') \le 2K < q + q'$, forcing $t = 0$. Thus $s_1 = s$, which says precisely that $f(\mu) = f(r) + s(\mu - r)$. □

Proposition 12.12.5. *Let (r, r') be a unimodular pair with $q + q' > 2K$, and let L be the affine function with $L(r) = f(r)$ and $L(r') = f(r')$. Then $f = L$ on $[r, r']$.*

Proof. First we prove $f(\rho) = L(\rho)$ for every rational $\rho \in [r, r']$. If $\rho \in \{r, r'\}$ this holds by definition, so let $\rho = u/v$ in lowest terms with $r < \rho < r'$; by Lemma 3 the integer $d := v - (q + q')$ is nonnegative. We argue by strong induction on d, the statement being quantified over all unimodular pairs with denominator sum exceeding $2K$ simultaneously.

If $d = 0$ then $\rho = \mu$ by Lemma 3, and Lemma 4 gives $f(\mu) = L(\mu)$.

If $d > 0$ then $\rho \neq \mu$. By Lemma 4, $f(\mu) = L(\mu)$. Suppose $r < \rho < \mu$; the case $\mu < \rho < r'$ is identical with the roles of the two ends exchanged. By Lemma 1, (r, μ) is a unimodular pair, with denominator sum $q + (q + q') > 2K$, and the affine function matching f at its two endpoints is L itself, since both $(r, f(r))$ and $(\mu, f(\mu))$ lie on L. Applying Lemma 3 to the pair (r, μ) gives $v \geq q + (q + q')$, so the induction parameter for ρ relative to this new pair is $v - (2q + q') = d - q$, which is nonnegative and strictly smaller than d. The inductive hypothesis therefore gives $f(\rho) = L(\rho)$.

Now f and L are continuous on $[r, r']$ and agree on $\mathbb{Q} \cap [r, r']$, which is dense in $[r, r']$; hence $f = L$ on $[r, r']$. □

Theorem 12.12.6. *There are finitely many closed intervals covering $[0, 1]$ on each of which f is affine.*

Proof. Fix an integer $N \geq 2K$. Build a finite binary tree of unimodular pairs as follows. The root is $(0/1, 1/1)$. A node (r, r') with $q + q' \leq N$ has two children, namely (r, μ) and (μ, r'), which are unimodular pairs by Lemma 1; a node with $q + q' > N$ is a leaf.

The tree is finite. Along any path from the root, the denominator sums form a strictly increasing sequence of positive integers, since the children of a node of sum $\sigma = q + q'$ have sums $\sigma + q$ and $\sigma + q'$ with $q, q' \geq 1$. Every internal node has sum at most N, and the root has sum 2, so no path contains more than $N - 1$ edges. A binary tree of bounded depth is finite.

The closed intervals $[r, r']$ attached to the leaves cover $[0, 1]$, because splitting a node partitions its interval into the two subintervals of its children and the root interval is $[0, 1]$. Each leaf is a unimodular pair with $q + q' > N \geq 2K$, so the Proposition applies to it and f is affine on the corresponding closed interval. Taking $I_1, \ldots, I_n$ to be the leaf intervals completes the proof. □

Remark 12.12.7. The pieces are integral, not merely affine. On a leaf interval, $f = a + bx$ with $b = s \in \mathbb{Z}$ by Lemma 2; moreover $qa = qf(r) - sm \in \mathbb{Z}$ and $q'a = q'f(r') - sm' \in \mathbb{Z}$, and $\gcd(q, q') = 1$, so Bezout gives $a \in \mathbb{Z}$. So the hypothesis "for each rational r there are integers a, b with $f(r) = a + br$" upgrades to a single pair of integers valid on a whole interval, which is the real content of the problem.

Remark 12.12.8. The leaves of the tree are exactly the pairs of consecutive elements of the Farey sequence of order N: every denominator appearing in the tree is the denominator sum of some earlier node and hence at most N, every leaf has $q + q' > N$, and unimodular pairs with $q, q' \leq N < q + q'$ are precisely the consecutive pairs of that Farey sequence. So the refinement tree reconstructs the Farey sequence rather than assuming it, and the number of pieces is $|F_N| - 1$, which is of order N^2, hence of order K^2.

Layer 1: WHAT

Block 1. Rewrite the arithmetic hypothesis in denominator form: $qf(m/q) \in \mathbb{Z}$, that is, the denominator of the value divides the denominator of the point.

Block 2. Ask which pairs of rationals have integer chord slope. Answer: unimodular pairs, because then $r' - r = 1/(qq')$ exactly and the slope is $q(q'f(r')) - q'(qf(r))$.

Block 3. Note that the Lipschitz bound caps that integer: $|s| \leq K$. An integer confined to a bounded range is the raw material of every rigidity argument.

Block 4. Introduce the mediant and compare the two half slopes with the whole slope: $\delta_1 q' + \delta_2 q = 0$, so $s_1 = s + tq$ and $s_2 = s - tq'$ for one integer t.

Block 5. Squeeze: $|s_1 - s_2| = |t|(q + q') \leq 2K$, so once $q + q' > 2K$ the integer t is zero and the mediant sits on the chord.

Block 6. Descend: every rational strictly inside a unimodular pair has denominator at least $q + q'$, with equality only for the mediant, so iterated mediants reach every rational and the quantity $v - (q + q')$ is a well founded induction measure.

Block 7. Close the interval by continuity: f agrees with an affine map on a dense set, hence everywhere.

Block 8. Make the cover finite: refine from $(0/1, 1/1)$ only while $q + q' \leq N$; the denominator sums strictly increase along every branch, so the tree has depth at most $N - 1$ and finitely many leaves.

Juxtaposition. The nominal route reaches for Rademacher because "Lipschitz" is a hypothesis with a famous theorem attached, gets f' integer valued almost everywhere, and then cannot finish, because almost everywhere information cannot see the rationals. The bifurcation point is the choice of what "two nearby rationals" means. The zombie route measures nearness by $|r - r'|$; the ideal route measures it by $m'q - mq' = 1$. Everything downstream follows from that single choice, because unimodularity is exactly the condition that makes $r' - r$ the reciprocal of the product of the denominators and hence makes the slope an integer.

Layer 2: HOW

Block 1. **Cue.** The hypothesis names the integers a and b but never uses them; they depend on r and are otherwise uncontrolled. **Why natural.** Data you cannot control should be eliminated. Multiplying by q removes both a and b and leaves one clean statement about f alone. This is the standard first move whenever a hypothesis is phrased as "there exist integers such that", and it converts an existential statement into a divisibility statement. **Link to the next block.** Now every value $f(r)$ has a known denominator, so differences of values have known denominators, so slopes are computable objects rather than arbitrary reals.

Block 2. **Cue.** A slope is a difference of values divided by a difference of points, and both differences now have denominators you control. **Why natural.** The denominator of $f(r') - f(r)$ divides qq', so the slope is an integer as soon as $r' - r$ equals $1/(qq')$ exactly, and $r' - r = (m'q - mq')/(qq')$ makes that condition literally $m'q - mq' = 1$. You are not choosing unimodularity for elegance; you are solving for the condition that makes the quotient integral. **Link to the next block.** The correct notion of adjacency on $\mathbb{Q} \cap [0, 1]$ is therefore arithmetic, and its combinatorics (mediants, the Stern and Brocot tree) becomes the scaffolding of the whole proof.

Block 3. **Cue.** Lipschitz is a bound and integrality is a quantisation. **Why natural.** This is the engine of the problem, and it should be named before it is used: a bounded integer lies in a finite set, and an integer of absolute value less than 1 is zero. Every rigidity theorem of this shape, from the rational root theorem to Liouville's approximation theorem, runs on the same fuel. **Link to the next block.** So the plan must be to manufacture an integer that the Lipschitz bound forces below 1. The slopes themselves are only bounded by K, which is not small, so you need a difference of slopes.

Block 4. **Cue.** Three collinear conditions and one obvious intermediate point: the mediant is the unique fraction of smallest denominator strictly between r and r'. **Why natural.** If you want to compare the chord over $[r, r']$ with the chords over its two halves, you must split at a point whose denominator you know, and the mediant is the only such point. The resulting identity $s(q + q') = s_1 q' + s_2 q$ is forced; there is no choice in it. **Link to the next block.** The identity says $\delta_1 q' = -\delta_2 q$ with $\gcd(q, q') = 1$, so a single integer parameter t controls both discrepancies. This is where the arithmetic does real work: coprimality converts one linear relation into a complete parametrisation.

Block 5. **Cue.** The two discrepancies are tq and $-tq'$, so their difference is $t(q + q')$, and both slopes are bounded by K. **Why natural.** The parameter t is amplified by $q + q'$, while the analytic bound on it stays fixed at $2K$. Large denominators therefore crush t to zero. This is exactly the manufactured small integer promised in Block 3, and it explains why the threshold in the problem is $2K$ and not K. **Link to the next block.** Once $t = 0$ for all sufficiently deep pairs, the function is pinned on the mediant, and so on every rational reachable by iterated mediants.

Block 6. **Cue.** An induction that must reach every rational in an interval, with no obvious ordering. **Why natural.** Iterating mediants is a recursion, and every recursion needs a well founded measure. Lemma 3 supplies one for free: any rational strictly inside a unimodular pair has denominator at least $q + q'$, so $v - (q + q')$ is a nonnegative integer, and descending into a half interval increases the denominator sum, hence strictly decreases the measure. Never write "and so on by repeated bisection"; name the measure. **Link to the next block.** With the measure in hand the induction is finite for each individual rational, and gives $f = L$ on all rationals of the interval.

Block 7. **Cue.** Two continuous functions agreeing on a dense set. **Why natural.** This is the only place the Lipschitz hypothesis is used qualitatively rather than quantitatively, and it must be used, because the conclusion is about all reals while every hypothesis so far was about rationals. Skipping this line is a real gap, not a formality. **Link to the next block.** It remains only to produce finitely many intervals of the required type, which is a combinatorial question about the refinement, not an analytic one.

Block 8. **Cue.** The Proposition needs $q + q' > 2K$, and the root pair $(0/1, 1/1)$ has $q + q' = 2$. **Why natural.** Refine only where the hypothesis fails, and stop as soon as it holds. Termination is then automatic from the same inequality that drove Block 5: sums strictly increase along a branch and refinement stops above N, so branches have length at most $N - 1$. This is the standard way to convert a local statement into a finite decomposition, and it is what the word "finitely" in the problem is testing. **Link to the next block.** The leaf intervals are the required $I_1, \ldots, I_n$, and as a bonus the leaves turn out to be exactly the Farey intervals of order N.

Layer 3: WHY

The mathematical truth beneath the route is that $\mathbb{Q} \cap [0, 1]$ carries a natural tree structure, the Stern and Brocot tree, in which the children of an interval $[m/q, m'/q']$ are obtained by splitting at the mediant, and in which every interval that appears is unimodular. This tree is not a convenience; it is the unique refinement scheme with the property that the endpoints of every cell have coprime denominators and the cell has length exactly $1/(qq')$. Those two facts are what make chord slopes integral, and integrality is what interacts with the Lipschitz bound. Metric subdivision, say into intervals of length $1/N$, has neither property and is therefore useless here, however natural it looks to an analyst.

The second truth is the quantisation principle. Two independent constraints act on the same quantity: arithmetic says the chord slope is an integer, analysis says it is at most K in absolute value. Neither alone says anything; together they confine the slope to a finite set. Then the mediant identity manufactures a second integer, t, whose analytic bound $|t| \leq 2K/(q + q')$ tends to zero as the denominators grow, and an integer that is forced below 1 is zero. That single mechanism, "a nonzero integer has absolute value at least one", is the load bearing idea in this problem exactly as it is in the rational root theorem, in Liouville's theorem on approximation of algebraic numbers, and in the standard proof that a polynomial taking small integer values on many points is zero.

The discernment mechanism is a two question filter that you should run on any hypothesis pair of this shape. First: what does each hypothesis constrain, and is there a single quantity constrained by both? Here the answer is the chord slope, which is the only thing Lipschitz continuity can see and, after the denominator reformulation, the only thing the arithmetic hypothesis can see. Second: in what index of locality does the arithmetic hypothesis behave well? Not in $|r - r'|$, because the denominator of a rational is not a continuous function of the rational; but unimodular adjacency is exactly the relation under which the arithmetic is clean. Train this by re deriving once, and then never forgetting, the three identities $r' - r = 1/(qq')$, $\mu - r = 1/(q(q + q'))$ and $r' - \mu = 1/(q'(q + q'))$; they are the entire toolkit.

The examiners nudged in two ways. They wrote the arithmetic hypothesis in the awkward form "there exist integers a and b such that $f(r) = a + br$" rather than the clean form "$qf(m/q) \in \mathbb{Z}$", because the translation is the first test. And they asked for finitely many intervals rather than for local linearity, because the finiteness is where the constant K must be used a second time, converting the qualitative statement into a bound on how deep the refinement must go.

Revolving concepts

Lipschitz continuity. $|f(x) - f(y)| \le K|x - y|$ implies uniform continuity and, by Rademacher's theorem, differentiability almost everywhere with $|f'| \le K$. The almost everywhere statement is a trap here: a 1 Lipschitz function can have $f' \in \{0, 1\}$ almost everywhere without being piecewise linear, as the fat Cantor construction above shows. See Book 1, field: Real analysis, topic: Lipschitz and absolutely continuous functions.

Density and continuity. If two continuous functions on an interval agree on a dense subset they agree everywhere. This is the bridge from rational information to real conclusions and must be stated explicitly. See Book 1, field: Real analysis, topic: density arguments.

Unimodular pairs, mediants and the Stern and Brocot tree. For $m/q < m'/q'$ in lowest terms with $m'q - mq' = 1$ one has $\gcd(q, q') = 1$ and $r' - r = 1/(qq')$; the mediant $(m + m')/(q + q')$ is automatically in lowest terms, lies strictly between, and forms a unimodular pair with each endpoint. Starting from $(0/1, 1/1)$ and splitting repeatedly at mediants generates every rational in $(0, 1)$ exactly once; the path spelling out a given rational is its continued fraction expansion in run length form. See Book 1, field: Number theory, topic: Farey fractions and continued fractions.

Farey sequences. F_N is the increasing list of rationals in $[0, 1]$ with denominator at most N. Consecutive elements are precisely the unimodular pairs with $q, q' \le N < q + q'$, and $|F_N| = 1 + \sum_{k \le N} \varphi(k)$, which is asymptotically $3N^2/\pi^2$. In this problem F_N appears as the set of leaves of the refinement tree rather than as an input. See Book 1, field: Number theory, topic: Farey sequences.

Bezout. $\gcd(q, q') = 1$ gives integers x, y with $xq + yq' = 1$, used to upgrade "qa and $q'a$ are integers" to "a is an integer", and used implicitly whenever coprimality converts a divisibility into a parametrisation. See Book 1, field: Number theory, topic: Bezout and linear Diophantine equations.

The bounded integer principle. If a quantity is an integer and is proved to satisfy $|t| < 1$, then $t = 0$; if it is an integer bounded by K then it takes at most $2K + 1$ values. See Book 1, field: Number theory, topic: integrality arguments.

The recursion, generalised. The refinement is the recursion $I = [r, r'] \mapsto \{[r, \mu], [\mu, r']\}$ on unimodular pairs, with associated data $\sigma(I) = q + q'$ satisfying $\sigma(\text{child}) = \sigma(I) + q$ or $\sigma(I) + q'$. Its general solution method has two halves that recur everywhere. To prove a statement for all rationals in I, induct on the nonnegative integer $v - \sigma(I)$, which strictly decreases as you descend, by Lemma 3. To prove finiteness of a refinement, exhibit a strictly increasing integer quantity along each branch that is bounded by the stopping rule; the depth bound then follows and with it the leaf count. The same pair of arguments settles any statement of the form "a function on $\mathbb{Q} \cap [0, 1]$ whose values have denominators dividing those of their arguments and whose slopes are bounded is piecewise affine", for instance with the Lipschitz hypothesis replaced by a bound on second differences. See Book 1, field: Number theory, topic: Stern and Brocot induction.

Circuit Bank

- **WHEN** a hypothesis reads "for each x there exist integers a, b with $f(x) = a + bx$" **DO** clear denominators to get one uncontaminated statement, here $qf(m/q) \in \mathbb{Z}$ **BECAUSE** the existential integers are uncontrolled and cannot be used, whereas a divisibility statement about the value is usable at every later step

- **WHEN** you have an analytic bound and an arithmetic quantisation on the same object **DO** build the quantity that is simultaneously an integer and small, and force it to be zero **BECAUSE** a nonzero integer has absolute value at least one, which is the only mechanism by which soft analytic bounds produce exact algebraic conclusions
- **WHEN** you must move between nearby rationals and the arithmetic must stay clean **DO** use unimodular pairs and mediants, not equally spaced points **BECAUSE** unimodularity is exactly the condition making $r' - r = 1/(qq')$, which is exactly the condition making chord slopes integral
- **WHEN** a claim must be proved for every rational in an interval by repeated subdivision **DO** name an explicit nonnegative integer measure that strictly decreases, here $v - (q + q')$, and induct on it **BECAUSE** "repeat this process" is not a proof; a well founded measure converts an informal descent into a finite induction and also tells you exactly when the descent stops
- **WHEN** a local property must be upgraded to a finite decomposition **DO** refine only while a stated inequality fails, and show a bounded integer strictly increases along each branch **BECAUSE** that is what makes the tree finite, and the word "finitely" in the question is precisely the request for this argument

Backtesting

1. First useful cue: the arithmetic hypothesis names uncontrolled integers, so clear denominators and read it as "$qf(m/q) \in \mathbb{Z}$".
2. Method the expression suggested: compute chord slopes between unimodular pairs, since those are integers bounded by K, and then compare a chord with its two mediant halves.
3. Where you would hesitate, and the thought to reinforce: at the choice of which rationals count as adjacent. Reinforce this: on $\mathbb{Q}$, adjacency is arithmetic, not metric; two rationals are close for the purposes of this problem when $m'q - mq' = 1$, however far apart they are on the line.
4. The trap the question sets: Rademacher's theorem. Lipschitz is exactly its hypothesis, so the reflex fires, and it leads to an almost everywhere statement that cannot be upgraded, since a 1 Lipschitz function with derivative in $\{0, 1\}$ almost everywhere need not be piecewise linear.
5. What makes you recognise it faster next time: the joint appearance of a slope bound and a denominator condition. That pair means "integer forced to be small", and the only question left is which combination of the data is that integer.
6. The reusable decision rule: when arithmetic and analysis constrain the same quantity, subdivide along the Stern and Brocot tree until the analytic bound beats one, then propagate by induction and close by density.

Mirror drills

1. Suppose $g : [0, 1] \to \mathbb{R}$ is continuous and $g(r) \in \mathbb{Z}$ for every rational r. Show g is constant. *Answer:* $\mathbb{Z}$ is closed and $\mathbb{Q}$ is dense, so $g([0, 1]) \subseteq \mathbb{Z}$; a continuous image of a connected set is connected, and a connected subset of $\mathbb{Z}$ is a point. No Lipschitz hypothesis is needed.
2. In the problem above, suppose additionally $K < 1$. Show f is a constant integer. *Answer:* the root pair $(0/1, 1/1)$ already has $q + q' = 2 > 2K$, so the Proposition applies to the whole of $[0, 1]$; the slope is an integer of absolute value less than 1, hence zero, and $f \equiv f(0) \in \mathbb{Z}$.
3. Prove that if $m/q < m'/q'$ are consecutive in the Farey sequence F_N then $q + q' > N$. *Answer:* their mediant has denominator $q + q'$ and lies strictly between them, so it is not in F_N; by Lemma 3 the smallest denominator strictly between them is exactly $q + q'$, so $q + q' > N$.

Chapter 13

Putnam 2013

13.1 A1. Icosahedron face labels force a repeat

Recall that a regular icosahedron is a convex polyhedron having 12 vertices and 20 faces, the faces congruent equilateral triangles. On each face of a regular icosahedron is written a nonnegative integer such that the sum of all 20 integers is 39. Show that there are two faces that share a vertex and have the same integer written on them.

Read it out loud

Say the data back before touching a pencil to it: twenty faces, one nonnegative integer on each, total 39, and you must prove a coincidence, two faces at a common corner sharing a value. A fixed total attached to a claim of forced repetition is the signature of a counting argument by contradiction: you do not construct the pair, you suppose it is absent and squeeze the total until it cannot fit.

The only geometry you need is the incidence pattern, not lengths or angles. Within thirty seconds you should have written down the two numbers that describe it: three vertices per face, free, since a face is a triangle, and less obviously, five faces per vertex. That second number comes from Euler's formula, $V - E + F = 2$ together with $3F = 2E$ for a triangulated solid, giving $V = 12$, $E = 30$, $F = 20$, so every vertex has degree $2E/V = 5$.

Three and five are the whole field embedded in the question: solid geometry reduced to a bipartite incidence structure between twelve vertices and twenty faces. A local relation of size five meeting a global total of twenty across that structure is exactly what a double count is for.

Approaches, ranked worst to best

1. **search the vertex figures by hand for the repeated pair.** *Cue:* the conclusion promises a specific pair exists. *Circuit:* none, this is zombie maths. Twelve vertices with five faces each give no symmetry to cut the casework, and the number 39 never enters the search in any structural way. It does not terminate inside an exam hour, and it never becomes a proof, only a hunt.

2. **bound how often a single label can repeat.** *Cue:* the conclusion is about a repeated label, so bound the label directly. *Circuit:* maximum size of a pairwise vertex disjoint family of faces. If no two same labelled faces may share a vertex, five faces carrying one label would need $5 \times 3 = 15$ distinct vertices against only 12 available, so every label appears on at most four faces; twenty faces then need at least five distinct labels, forcing a minimum total of $4(0 + 1 + 2 + 3 + 4) = 40 > 39$. A genuine proof, but it pays an extra step to pin down the cheapest label multiset, a step the winning route below never needs.

3. **sum the labels around each vertex and compare to the global total.** *Cue:* the conclusion names the vertex as the site of the coincidence. *Circuit:* double counting incidences between two levels of a polyhedron. Negate the conclusion, bound the five distinct labels at each vertex below by $0 + 1 + 2 + 3 + 4 = 10$, sum that bound over the twelve vertices, and compare the result against the same sum evaluated the other way, by counting each face three times. One inequality, no optimisation to justify. This is the route the problem was built for.

The ideal solution

Suppose, for contradiction, that no two faces sharing a vertex carry the same integer. Fix a vertex v. Exactly five faces meet at v: with $V = 12$, $E = 30$, $F = 20$ from Euler's formula and $3F = 2E$, the sum of all vertex degrees is $2E = 60$, and by the regularity of the solid every vertex carries the same degree, so each of the twelve vertices has degree $60/12 = 5$. By assumption the five faces at v carry five pairwise distinct nonnegative integers, so, sorted increasingly, their labels are termwise at least $0, 1, 2, 3, 4$, and their sum is at least $0 + 1 + 2 + 3 + 4 = 10$.

Sum this bound over all twelve vertices:

$$\sum_{v} \Big(\text{sum of the labels of the faces meeting } v\Big) \;\geq\; 12 \times 10 = 120.$$

Evaluate the left side exactly by counting the other way. Every face is a triangle, so it meets exactly three vertices, and its label is counted exactly three times as v ranges over all twelve vertices, once per corner. Writing $T = 39$ for the total of all twenty face labels,

$$\sum_{v} \Big(\text{sum of the labels of the faces meeting } v\Big) = 3T = 117.$$

So $117 \geq 120$, false. The assumption fails: there exist two faces sharing a vertex with the same integer written on them. ■

Layer 1: WHAT

Block 1. Negate the conclusion into a per vertex distinctness statement, and extract the incidence numbers three and five from Euler's formula. Block 2. Bound the minimum sum of five distinct nonnegative integers at a single vertex: at least 10. Block 3. Sum that local bound over all twelve vertices: at least 120. Block 4. Evaluate the same total the other way, face by face, using that each face is counted three times: exactly 117. Block 5. Collide $117 \geq 120$, false.

The nominal route (approach 1) hunts for the pair directly and never produces a number to set against 39. The bifurcation point against it sits at Block 1: the moment "two faces agree" is negated into "every vertex's five faces are pairwise distinct", the problem stops being a search and becomes an arithmetic bound, because a distinctness statement about nonnegative integers has a computable minimum sum attached to it for free. Against approach 2, the bifurcation is which object to sum first: summing by label needs an extra extremal step to find the cheapest multiset; summing by vertex, the object the conclusion itself names, needs nothing beyond $0 + 1 + 2 + 3 + 4 = 10$.

Layer 2: HOW

Block 1. **Cue.** The conclusion is phrased as two faces that share a vertex, naming the vertex, not the face, as the site of the coincidence. **Why natural.** The object you sum local data over should always be the object the conclusion names. You need the degree of a vertex, five, because that is the size of the distinct set you are about to bound. **Link to the next block.** Once every vertex sees five faces, the negated conclusion has an immediate numerical consequence.

Block 2. **Cue.** Five pairwise distinct nonnegative integers, no upper bound stated. **Why natural.** Nonnegative integers with no other constraint have exactly one extremal fact worth knowing: sorted, k distinct such integers are termwise at least $0, 1, \ldots, k-1$, because nothing smaller fits. This is the generic minimum $\binom{k}{2}$, and it should fire the instant "distinct" meets "nonnegative integers" with a sum nearby. **Link to the next block.** A per vertex minimum of 10 is only useful once aggregated, since the problem's one numerical fact is a sum over all twenty faces, not over one vertex.

Block 3. **Cue.** Twelve vertices, each contributing at least 10. **Why natural.** Summing a true inequality over a finite index set preserves it; naming this explicitly turns twelve local facts into one global fact. **Link to the next block.** You now have a global lower bound, 120, to set against the global total the problem handed you, 39.

Block 4. **Cue.** The sum just built is a sum over vertices of a sum over incident faces, secretly a sum over incident (vertex, face) pairs. **Why natural.** That same set of pairs sums the other way too: fix a face and ask how many vertices see it. Because a face is a triangle, the answer is always three, so every face's label contributes exactly three times. **Link to the next block.** The vertex sum, computed this second way, is exactly $3 \times 39 = 117$.

Block 5. **Cue.** A lower bound of 120 from Block 3 against an exact value of 117 from Block 4. **Why natural.** A quantity cannot be at least 120 and equal 117. The gap is not incidental: the setter chose 39 so that 3×39 sits three below 12×10, tight enough that no weaker bound would have closed it. **Link to the next block.** No further block; the contradiction closes the proof.

Layer 3: WHY

The route works because a triangulated closed surface offers two dual ways to sum over its own incidence structure, by vertex or by face, both counting exactly the same set of (vertex, face) pairs. Whenever a fixed global quantity sits on a structure that also carries a local adjacency relation, that duality is free, the same mechanism as the handshake lemma, $\sum_v \deg(v) = 2E$, lifted one dimension from edges and vertices to faces and vertices.

The discernment mechanism: the moment a problem gives one fixed global number and asks for a local coincidence, ask what quantity you can compute two ways, once by summing local bounds and once by summing the global data through the same incidences. If the answer is not obvious, sum the negated conclusion's consequence over exactly the object the conclusion names. The examiners nudged you here by giving the total, 39, no other role in the problem beyond being compared against a bound you must build, and by tuning the gap to three, tight enough to rule out a lazier bound (bounding by four distinct values instead of five, for instance, only reaches $4 \times 6 = 24$ against eleven, and does not close).

Revolving concepts

Double counting. Definition: when a finite incidence relation connects two finite sets X and Y, the total number of incident pairs can be summed from X's side or from Y's side, and the two totals agree. Orbiting results: the handshake lemma, Euler's formula $V - E + F = 2$, and the identity $\sum_f (\text{vertices of } f) = \sum_v (\text{faces at } v)$ used here. See Book 1, field: Combinatorics, topic: double counting and incidence structures.

Extremal sum of distinct nonnegative integers. Definition: k pairwise distinct nonnegative integers sum to at least $\binom{k}{2}$, with equality only for $\{0, 1, \ldots, k-1\}$. The single fact that turns "distinct" into a number. See Book 1, field: Combinatorics, topic: extremal counting.

Regular polyhedron combinatorics. Definition: for a triangulated closed polyhedral surface, $V - E + F = 2$ with $3F = 2E$ pins the combinatorics down once V or F is known, and regularity forces every vertex to degree $2E/V$. See Book 1, field: Geometry, topic: polyhedral combinatorics and Euler's formula.

Both ranked approaches above are instances of one general threshold, worth stacking once since it recurs on every fixed total, forced coincidence problem: if n labelled objects split into m overlapping groups of size t, each object in exactly d groups, and the total of all labels is T, two objects in a common group are forced to share a label whenever $T < \dfrac{mt(t-1)}{2d}$. Here $n = 20$, $m = 12$, $t = 5$, $d = 3$, threshold $\dfrac{12 \times 5 \times 4}{2 \times 3} = 40$, and the given total sits one below it. Writing this once with letters is what lets you recognise the dodecahedron and octahedron versions on sight.

Circuit Bank

- **WHEN** a fixed global sum is given and you must prove two local, adjacent objects share a value **DO** negate the conclusion, bound the minimum sum of a maximal set of pairwise distinct local values, and sum that bound over the whole structure to compare against the given total **BECAUSE** negation turns distinct into a computable minimum, and only a summed minimum set against a summed exact value makes a single fixed total load bearing
- **WHEN** you need the minimum possible sum of k pairwise distinct nonnegative integers **DO** write $0 + 1 + \cdots + (k-1) = \binom{k}{2}$ immediately, without rederiving it **BECAUSE** any k distinct nonnegative integers, sorted, are termwise at least $0, 1, \ldots, k-1$
- **WHEN** a polyhedron's incidence numbers, faces per vertex and vertices per face, are not stated outright **DO** extract them from $V - E + F = 2$ plus the degree sum $2E = \sum \deg(v)$, using regularity to divide evenly **BECAUSE** most polyhedral counting problems are unusable until these two numbers are known, and deriving them costs three lines, not a guess
- **WHEN** a counting inequality closes with an uncomfortably small gap between the two sides **DO** trust the gap was chosen deliberately and look for exactly the bound producing it, rather than a looser bound that will not close **BECAUSE** setters calibrate constants so the natural bound barely fails, and a small nonzero gap is itself a sign you found the intended route

Backtesting

1. First useful cue: the conclusion names the vertex, two faces that share a vertex, telling you which object to sum local data over.
2. Method the expression suggested: a fixed global total plus a local adjacency claim suggested double counting before any specific number was examined.
3. Where you would hesitate, and the thought to reinforce: choosing between summing by label multiplicity or by vertex; reinforce that summing over the object the conclusion names is shorter, since it needs no auxiliary optimisation.
4. The trap the question sets: the number 39 looks arbitrary until you compute $3 \times 39 = 117$ against $12 \times 10 = 120$; stopping before that numerical comparison hides why 39 was chosen at all.
5. What makes you recognise it faster next time: a solid described purely by face count, vertex count and a fixed label total, with no length or angle used anywhere, tells you the problem lives entirely in the incidence structure.
6. The reusable decision rule: fixed total plus a local repetition claim on an incidence structure means negate, bound the local minimum, sum, and compare against the exact global value from counting the same incidences the other way.

Mirror drills

1. A regular dodecahedron, 12 pentagonal faces, 20 vertices, each vertex meeting exactly 3 faces, carries a nonnegative integer on each face, total S. Find the largest S that still forces two faces sharing a vertex to carry the same integer. *Answer:* each vertex bound is $0 + 1 + 2 = 3$ over 20 vertices, at least 60; each face is counted 5 times, so $5S \geq 60$ needs $S \geq 12$ to escape, hence every $S \leq 11$ forces the repeat.
2. A regular octahedron, 8 triangular faces, 6 vertices, each vertex meeting exactly 4 faces, carries a nonnegative integer on each face, total S. Find the threshold. *Answer:* each vertex bound is

$0+1+2+3=6$ over 6 vertices, at least 36; each face is counted 3 times, so $3S \geq 36$ needs $S \geq 12$ to escape, hence $S \leq 11$ forces the repeat, matching $\frac{mt(t-1)}{2d} = \frac{6\times4\times3}{2\times3} = 12$.

3. Twelve committees of five people are drawn from twenty people, each person on exactly three committees; every person gets a nonnegative integer score, total 39. Show two people on a common committee share a score. *Answer:* the identical structure with committee in place of vertex and person in place of face, the same double count gives $117 < 120$.

13.2 A2. Injectivity of the minimal square completion map

Let S be the set of all positive integers that are not perfect squares. For n in S, consider choices of integers $a_1, a_2, \ldots, a_r$ such that $n < a_1 < a_2 < \cdots < a_r$ and $n \cdot a_1 \cdot a_2 \cdots a_r$ is a perfect square, and let $f(n)$ be the minimum of a_r over all such choices. For example, $2 \cdot 3 \cdot 6$ is a perfect square, while $2 \cdot 3$, $2 \cdot 4$, $2 \cdot 5$, $2 \cdot 3 \cdot 4$, $2 \cdot 3 \cdot 5$, $2 \cdot 4 \cdot 5$, and $2 \cdot 3 \cdot 4 \cdot 5$ are not, and so $f(2) = 6$. Show that the function f from S to the integers is one to one.

Read it out loud

Strip the dressing off first. Underneath, an integer k is a perfect square exactly when every prime in its factorisation occurs to an even power, so "the product of these integers is a perfect square" is a parity statement, a vanishing sum, once you track each integer by the exponents of its primes read modulo two. That reading is not needed to solve the problem, but it explains why the problem is solvable at all: perfect squares are closed under multiplication, and closure under an operation is the one thing every proof below leans on.

The command term is minimum. $f(n)$ is not a value you compute from a formula, it is defined as the smallest top term over an entire family of valid completions of n into a square. Whenever a quantity is handed to you as a minimum over a family, and you are asked to prove a property of it, the property is never going to fall out of algebra on $f(n)$ alone, because you have no closed form for f; the worked example already warns you of this, $f(2) = 6$ has no visible pattern to extrapolate from a single data point. The only lever a minimum gives you for free is this: if you can build a valid member of the family whose top term is smaller than the claimed minimum, you have a contradiction. That is the whole architecture you should have written down within the first thirty seconds: to prove f one to one, suppose $f(n) = f(m)$ for $n \neq m$, and manufacture, from that coincidence, a completion of the smaller of the two with a strictly smaller top term than the one you assumed was minimal.

The second cue sits in the hypothesis itself: you are handed two separate instances of the same closed-under-multiplication property, one for n and one for m. Two instances of an identical algebraic hypothesis are an invitation to combine them by the operation they are instances of.

Approaches, ranked worst to best

1. **hunt for a closed formula for f, or assume monotonicity.** *Cue:* a single worked value, $f(2) = 6$, tempts you to look for a pattern. *Circuit:* pattern extrapolation from one data point. There is no such formula to find. The remark attached to the official solution calls the behaviour of f erratic, and no argument for injectivity can rest on an order relationship between $f(n)$ and $f(m)$ that is never established. Every hour spent computing $f(3), f(5), f(6), \ldots$ by hand looking for structure is ink spent on a route that was never going to close.

2. **reformulate every integer as its squarefree kernel's exponent vector over $\mathbb{F}_2$, and argue with linear algebra.** *Cue:* "product is a perfect square" is a vanishing sum of parity vectors. *Circuit:* closure of a vanishing sum under vector addition, cancellation from $w + w = 0$. This is the honest underlying structure and it produces a correct proof: two vanishing sums, added, still vanish, and a vector appearing twice cancels itself. It costs more ink than it needs to, though, because you must first set up the map from integers to $\mathbb{F}_2^{(\text{primes})}$, verify it is a homomorphism, and translate "perfect square" into "kernel element" before you can say anything, machinery the next approach gets for free from a single line about exponents.

3. **multiply the two witnessing products together and cancel every repeated factor.** *Cue:* two hypotheses of the identical shape, product is a perfect square, sitting side by side. *Circuit:* closure of perfect squares under multiplication, plus a one line valuation lemma for cancellation.

Multiply the witnessing product for n by the witnessing product for m: still a perfect square. The two top terms are forced equal by $f(n) = f(m)$, so they are guaranteed to be among the values that repeat, and cancelling repeats in pairs preserves the perfect square property by a single valuation check, no vector space needed. What survives is a strictly smaller completion of n. This is the route the problem was built for.

The ideal solution

Suppose, for contradiction, that $f(n) = f(m)$ for some $n, m \in S$ with $n \neq m$; without loss of generality $n < m$. Let $n < a_1 < \cdots < a_r$ witness $f(n)$, so $n\,a_1 \cdots a_r$ is a perfect square and $a_r = f(n)$ is minimal among all such top terms for n; likewise let $m < b_1 < \cdots < b_s$ witness $f(m)$, with $m\,b_1 \cdots b_s$ a perfect square and $b_s = f(m)$ minimal for m. Since $r, s \geq 1$ (an empty product would leave n, respectively m, itself a perfect square, contrary to $n, m \in S$), we have $a_r \geq a_1 > n$ and $b_s \geq b_1 > m$. Write $M = a_r = b_s$ for the common value forced by the hypothesis $f(n) = f(m)$; then $M > m > n$.

Consider the product of the two witnessing relations,

$$N \;=\; \big(n\,a_1 \cdots a_r\big)\big(m\,b_1 \cdots b_s\big),$$

a product of two perfect squares, hence itself a perfect square. Use repeatedly the following one line fact: if $N = v^2 N'$ for a positive integer v and N is a perfect square, then N' is a perfect square, because for every prime p the exponent $v_p(N) = 2v_p(v) + v_p(N')$ is even, so $v_p(N')$ is even too. The top terms a_r and b_s are both equal to M and appear as literal factors of N, so $N = M^2 \cdot N_1$ where N_1 is the product of the remaining $r + s$ factors $n, a_1, \ldots, a_{r-1}, m, b_1, \ldots, b_{s-1}$; by the fact just stated, N_1 is a perfect square. If any further value repeats within this remaining list, for instance some $a_i = b_j$ with $i < r, j < s$, or conceivably m coinciding with some a_i, cancel that pair too, by the same fact, and continue until the surviving multiset D consists of pairwise distinct positive integers whose product is a perfect square.

Two facts about D finish the proof. First, $n \in D$: n is strictly smaller than every other original factor ($n < a_1 \leq \cdots$, and $n < m < b_1 \leq \cdots$), so n is never the larger element of a cancelled pair and is never removed. Second, every element of $D \setminus \{n\}$ is strictly between n and M: each is one of $a_1, \ldots, a_{r-1}, m, b_1, \ldots, b_{s-1}$, all of which lie strictly below M (the unique largest value of the original multiset of $r + s + 2$ factors, since $a_r = b_s = M$ dominates every a_i, b_j, and exceeds both n and m), while M itself has been cancelled out of D entirely.

If $D = \{n\}$, its product is n itself, forced to be a perfect square by the cancellation fact, contradicting $n \in S$ outright. Otherwise sort $D \setminus \{n\}$ increasingly as $a'_1 < \cdots < a'_t$; then $n < a'_1 < \cdots < a'_t < M$ and $n\,a'_1 \cdots a'_t$, being the product of D, is a perfect square. This is a valid completion of n with top term $a'_t < M = f(n)$, contradicting the minimality of $f(n)$. Either way the assumption $f(n) = f(m)$, $n \neq m$, is impossible, and f is one to one. ■

Layer 1: WHAT

Block 1. Read "minimum" as an extremal-contradiction architecture: suppose $f(n) = f(m)$ for $n < m$, and commit to manufacturing a smaller completion of n. Block 2. Multiply the two witnessing products; the result is a perfect square because perfect squares are closed under multiplication. Block 3. Identify the one forced coincidence the hypothesis actually supplies, the two top terms are literally equal, so they repeat inside the combined product. Block 4. Cancel every repeated value, not only the engineered one, using a one line valuation lemma, leaving a product of distinct integers that is still a perfect square. Block 5. Read off the extremes of what survives, minimum n, maximum strictly below M, and recognise this as a smaller completion of n, contradicting minimality.

The nominal route, formula hunting, never reaches a comparison between $f(n)$ and $f(m)$ at all, because it tries to pin down f itself rather than use the one structural fact the definition guarantees.

The bifurcation point sits at Block 1: the instant you stop trying to compute f and instead ask what a violation of minimality would look like, the problem becomes constructive rather than exploratory. Against the vector space approach, the bifurcation is representation only, both routes cancel a repeated coincidence to shrink a witness; working with the integers directly and a one line valuation check reaches the same cancellation without first building the homomorphism.

Layer 2: HOW

Block 1. **Cue.** $f(n)$ is defined as a minimum over a family, and the target claim, injectivity, is a property of that minimum, not its value. **Why natural.** Minimality is inert until you build a competitor; the only proof technique that uses a minimum definition directly is contradiction via a smaller witness, because no formula for f is available to reason about algebraically. **Link to the next block.** To build a competitor for n, you need to use the assumed coincidence $f(n) = f(m)$ together with the witnessing data for m, not n alone.

Block 2. **Cue.** Two hypotheses of the identical shape, a product of integers is a perfect square, one for n and one for m. **Why natural.** Perfect squares are closed under multiplication, so two instances of the same closed property combine for free; this is the natural way to bring two pieces of given structure into contact with each other. **Link to the next block.** The combined product is one large perfect square containing every original factor from both sequences, including both top terms.

Block 3. **Cue.** $f(n) = f(m)$ says the two top terms a_r, b_s are the same number, not merely that they are equal in size by coincidence, they are one number playing two roles. **Why natural.** This is the one fact the hypothesis $f(n) = f(m)$ actually contributes beyond two independent square-product facts, so it has to be what the rest of the proof leans on; a value appearing as a literal factor twice inside a product is exactly the situation a cancellation lemma wants. **Link to the next block.** Cancelling that repeated value out of the perfect square, by the valuation fact, produces a strictly smaller perfect square built from the remaining factors.

Block 4. **Cue.** Beyond the engineered repeat at the top, the remaining list $(n, a_1, \ldots, a_{r-1}, m, b_1, \ldots, b_{s-1})$ may hold further coincidental repeats. **Why natural.** The cancellation lemma does not care why two factors are equal, only that they are, so leaving any repeat unresolved keeps D from being the strictly increasing sequence the definition of f requires. **Link to the next block.** After exhausting every repeat, D is a set of pairwise distinct positive integers whose product is a perfect square, ready to be read as a candidate completion.

Block 5. **Cue.** D contains n as its minimum (nothing else was ever small enough to be n) and every other surviving value below M (the unique original maximum, now cancelled away). **Why natural.** This is precisely the data the definition of $f(n)$ asks for: a strictly increasing sequence starting above n whose product with n is a perfect square, only with a top term smaller than the value you assumed was minimal. **Link to the next block.** There is no further block: $f(n)$ cannot be M and also have a smaller completion, so the assumed coincidence is impossible.

Layer 3: WHY

The reason the multiply-and-cancel move is available at all is that perfect squares are the kernel of a homomorphism: send each positive integer k to the vector of its prime exponents read modulo two, and k is a perfect square exactly when this vector is zero. Multiplication of integers becomes addition of vectors under this map, and addition over $\mathbb{F}_2$ is its own inverse, a vector added to itself is zero, which is exactly why a value repeated in the combined product can be struck out without disturbing the perfect square property. You never need to name this vector space to run the proof,

because the one line valuation lemma delivers the same cancellation directly on the integers, but recognising the vector space underneath is what tells you, before you have written a single line, that cancellation will work here, rather than having to discover it by trial.

The discernment mechanism for reaching for this move: whenever a problem hands you two hypotheses of the identical algebraic shape, closed under some operation (a product that is a perfect square, a sum with fixed parity, a linear combination that vanishes), and asks you to relate the two instances, combine them by that operation and look for what must repeat as a consequence of whatever equality the problem's hypothesis forces. Here the forced equality was handed to you directly, $f(n) = f(m)$ makes $a_r = b_s$ literally, so the repeat you need is not something you search for, it is guaranteed the moment you accept the hypothesis you are trying to contradict. The setters chose to define f as a minimum, rather than asking you to compute it, precisely so that this architecture, and no other, would be the one that closes; the worked example $f(2) = 6$ exists to stop you from hunting for a formula before you have exhausted the structural route.

Revolving concepts

Squarefree kernel and the $\mathbb{F}_2$ reading of perfect squares. Definition: every positive integer k factors uniquely as $k = u^2 d$ with d squarefree, and k is a perfect square exactly when the exponent of every prime in k is even; recording those exponents modulo two turns multiplication of integers into vector addition over $\mathbb{F}_2$, with perfect squares as the zero vector. See Book 1, field Number theory, topic squarefree decomposition and quadratic residues.

p-adic valuation and the cancellation lemma. Definition: $v_p(k)$ is the exponent of the prime p in k, valuations are additive over products, and k is a perfect square iff $v_p(k)$ is even for every p; consequently if v^2 divides a perfect square N exactly, N/v^2 is again a perfect square. This one line fact is the entire engine of the cancellation step. See Book 1, field Number theory, topic p-adic valuations.

Extremal choice via a minimality contradiction. Definition: to prove a property of a quantity defined as the least (or greatest) member of a family satisfying some condition, assume the property fails and construct, from that failure, a strictly smaller (or larger) member of the family, contradicting the defining extremal property. This is the master architecture of the whole problem and recurs whenever a definition contains the word minimum or maximum. See Book 1, field Number theory, topic extremal principle and well ordering.

Closure of perfect squares under multiplication. Definition: the product of two perfect squares is a perfect square, since exponents add and the sum of two even numbers is even. A one sentence fact, but it is the only justification needed for combining the two witnessing relations in the first place, and it is worth naming explicitly rather than treating as obvious, since the whole proof is an exploitation of exactly this closure. See Book 1, field Number theory, topic multiplicative structure of the integers.

Quadratic sieve and factor bases, for context rather than necessity. Definition: integer factoring algorithms search for subsets of a fixed factor base whose product is a perfect square, by exactly the linear algebra over $\mathbb{F}_2$ described above; the source solution's own remark points here directly. Recognising the connection is not required for this proof, but it tells you the technique you have just installed is not a one-off competition device, it is the working method behind a real factoring algorithm. See Book 1, field Number theory, topic quadratic sieve.

Circuit Bank

- **WHEN** a quantity is defined as the minimum (or maximum) over a family of valid constructions, and you must prove a property of it **DO** assume the property fails and build, from that failure, a strictly smaller (or larger) member of the family **BECAUSE** a minimum is otherwise inert, it only becomes usable once a competitor exists to violate it

- **WHEN** you are given two separate instances of the same closed-under-an-operation hypothesis and must relate them **DO** combine the two instances by that operation and cancel whatever the problem's hypothesis forces to repeat **BECAUSE** closure guarantees the combination still satisfies the hypothesis, and a forced repeat is exactly what a cancellation argument needs to produce something strictly smaller
- **WHEN** you must justify that removing a repeated factor from a perfect square leaves a perfect square **DO** check it with one line of valuations, $v_p(N) = 2v_p(v) + v_p(N')$, rather than reaching for heavier machinery **BECAUSE** evenness of $v_p(N)$ and of $2v_p(v)$ forces evenness of $v_p(N')$ for every prime, which is the entire content of what needs proving
- **WHEN** a competition remark or a single worked example suggests a function has no closed form **DO** abandon formula hunting immediately and ask what structural property, closure, minimality, symmetry, the definition actually guarantees **BECAUSE** ink spent guessing at a nonexistent pattern is ink unavailable for the structural argument that was always going to be the real route

Backtesting

1. First useful cue: the word minimum inside the definition of $f(n)$, flagging an extremal-contradiction proof before any number theory is even needed.
2. Method the expression suggested: two instances of product is a perfect square suggested multiplying them together, since perfect squares are multiplicatively closed.
3. Where you would hesitate, and the thought to reinforce: justifying that cancelling a repeated factor preserves the perfect square property; reinforce the one line valuation check rather than reaching for a heavier vector space argument.
4. The trap the question sets: assuming the only repeat you need to handle is the engineered one, $a_r = b_s$, and forgetting that other coincidental repeats among the remaining factors must also be cancelled before D is genuinely a strictly increasing witness.
5. What makes you recognise it faster next time: any definition of the form let $f(n)$ be the minimum of such that, attached to a request to prove a global property (injectivity, surjectivity, an inequality) of f, rather than to compute a value.
6. The reusable decision rule: two hypotheses of an identical multiplicatively (or additively) closed shape, tied by an equality the problem's own hypothesis forces, combine by that operation and cancel, and the survivor contradicts whatever extremal definition you started from.

Mirror drills

1. Compute $f(3)$ directly from the definition and confirm it differs from $f(2) = 6$, consistent with injectivity. *Answer:* $f(3) = 8$, witnessed by $3 \cdot 6 \cdot 8 = 144 = 12^2$; checking every subset of $\{4, 5, 6, 7\}$ shows no completion with top term at most 7 exists.
2. If f were redefined using perfect cubes instead of perfect squares, explain why multiplying two witnessing products for a collision no longer immediately finishes the proof. *Answer:* cancelling a value appearing twice divides out a square factor, not a cube factor, so the survivor need not stay a perfect cube; the fix needs valuations tracked modulo three, forcing triple, not double, coincidences.
3. Restate the original proof entirely in the language of exponent vectors over $\mathbb{F}_2$: what does "the two top terms cancel" become? *Answer:* the vector of a_r and the vector of b_s are identical, so

their sum is zero, exposing a shorter vanishing sum, the vector-space shadow of the smaller completion of n.

13.3 A3. Geometric series bridge from sum to polynomial

Suppose that the real numbers $a_0, a_1, \ldots, a_n$ and x, with $0 < x < 1$, satisfy

$$\frac{a_0}{1-x} + \frac{a_1}{1-x^2} + \cdots + \frac{a_n}{1-x^{n+1}} = 0.$$

Prove that there exists a real number y with $0 < y < 1$ such that

$$a_0 + a_1 y + \cdots + a_n y^n = 0.$$

Read it out loud

Name the two objects before touching either of them. There is x, a single fixed real number, handed to you already satisfying one given equation, data, not an unknown. And there is y, ranging freely over $(0, 1)$, the unknown a conclusion must produce, attached not to the given fractional sum at all but to a different expression, $P(y) = a_0 + a_1 y + \cdots + a_n y^n$, that shares its coefficients with the hypothesis and nothing else. Confusing these two variables, treating x as though it were free to vary, is the single most expensive misreading available here, and the problem is built to invite exactly that misreading, since both expressions use the same letters $a_0, \ldots, a_n$ and both concern the interval $(0, 1)$.

Within the first thirty seconds you should have written two things down. First, the conclusion is an existence statement about a continuous function, P, having a zero, and you were never asked to find y, only to prove it exists, which is the signature of the intermediate value theorem run by contradiction: assume no zero, derive that P is one sign throughout, then break that sign. Second, the denominators $1 - x^{k+1}$ with $0 < x < 1$ are not incidental, $1/(1-r)$ for $|r| < 1$ is the sum of a convergent geometric series, and a denominator of exactly this shape, with an explicit constraint guaranteeing convergence sitting right next to it, is a standing invitation to expand it.

The field embedded in the question is therefore small and sharp: continuity and the intermediate value theorem on one side, geometric series and the interchange of a finite sum with a convergent infinite one on the other, meeting because the coefficients a_k are shared between the reciprocal sum in x and the polynomial in y.

Approaches, ranked worst to best

1. **clear every denominator and manipulate the resulting single equation in x.** *Cue:* a sum of fractions equalling zero invites clearing denominators. *Circuit:* none, this stalls immediately. Multiplying through by $(1-x)(1-x^2)\cdots(1-x^{n+1})$ produces one polynomial equation satisfied by the single fixed number x. There is no second unknown here to extract a root from, and no route from a true statement about one fixed x to an existence statement about a different polynomial P in a different variable. This treats x as though it carried freedom it does not have.

2. **apply the intermediate value theorem directly to the given sum, varying x.** *Cue:* "prove existence" plus a sum equal to zero pattern matches to IVT. *Circuit:* sign change search, aimed at the wrong function. This correctly senses that IVT is the right family of tool, and for that reason it is not simply wrong, only misaimed: x is one fixed number satisfying one given equation, not a variable ranging over an interval, so there is no family of values of the given sum to search for a sign change in. The existential object the problem wants, y, belongs to P, not to this sum.

3. **expand each $1/(1-x^{k+1})$ as a geometric series and resum by evaluating P along the orbit $x, x^2, x^3, \ldots$.** *Cue:* the same coefficients a_k sit on reciprocal powers of x in the hypothesis and on monomial powers of y in the conclusion. *Circuit:* generating function bridge, plus

IVT by contradiction applied to P itself. Assume P has no root in $(0,1)$; by IVT it is one sign throughout, say positive. Expanding every denominator as a geometric series and reorganising turns the given sum, term by term, into $\sum_{m\geq 0} x^m P(x^m)$, an infinite sum whose every term is nonnegative and whose $m = 1$ term alone is strictly positive, forcing the total strictly positive and contradicting the hypothesis value 0. Minimal ink, no wrong variable, closes cleanly.

The ideal solution

Suppose, for contradiction, that $P(y) := a_0 + a_1 y + \cdots + a_n y^n$ has no zero for $y \in (0,1)$. P is a polynomial, hence continuous on all of $\mathbb{R}$, and $(0,1)$ is an interval; if P took a positive value at one point of $(0,1)$ and a negative value at another, the intermediate value theorem would force a zero of P strictly between them, inside $(0,1)$, contradicting the assumption. So P has one constant sign throughout $(0,1)$. Replacing every a_k by $-a_k$ if necessary, negating P throughout, negating the hypothesis sum termwise, and so preserving both the fact that it equals zero and whether P has a root, we may assume

$$P(y) > 0 \qquad \text{for every } y \in (0,1).$$

For each integer $m \geq 1$, $x^m \in (0,1)$ since $0 < x < 1$, so $P(x^m) > 0$; as $x^m > 0$ too,

$$x^m P(x^m) \;=\; \sum_{k=0}^{n} a_k x^{m(k+1)} \;>\; 0 \qquad (m \geq 1).$$

For $m = 0$, $x^0 P(x^0) = P(1) = a_0 + \cdots + a_n$, which is not covered directly by the open interval bound, but continuity closes the gap: letting $y \to 1^-$ through $(0,1)$, where $P(y) > 0$ throughout, gives $P(1) = \lim_{y\to 1^-} P(y) \geq 0$.

Now sum $x^m P(x^m)$ over $m = 0, 1, 2, \ldots$. For each fixed $k \in \{0, \ldots, n\}$ the series $\sum_{m=0}^{\infty} x^{m(k+1)}$ is geometric with ratio $x^{k+1} \in (0,1)$, converging to $1/(1-x^{k+1})$; since k ranges over only the $n+1$ values $0, \ldots, n$, summing this finite index outside the convergent infinite series over m is a plain rearrangement of a finite sum of convergent series, needing nothing beyond that convergence itself:

$$\sum_{m=0}^{\infty} x^m P(x^m) \;=\; \sum_{m=0}^{\infty} \sum_{k=0}^{n} a_k x^{m(k+1)} \;=\; \sum_{k=0}^{n} a_k \sum_{m=0}^{\infty} x^{m(k+1)} \;=\; \sum_{k=0}^{n} \frac{a_k}{1 - x^{k+1}}.$$

By hypothesis the right side is 0. But the left side is $P(1) \geq 0$ plus infinitely many strictly positive terms $x^m P(x^m)$, $m \geq 1$; the single term at $m = 1$ already forces the total strictly positive. So $\sum_{m=0}^{\infty} x^m P(x^m) > 0$, contradicting that it equals 0. The assumption fails: P has a zero $y \in (0,1)$. ■

Layer 1: WHAT

Block 1. Separate the fixed datum x from the unknown y, and locate the existential target correctly: a zero of P on $(0,1)$, not a property of the given sum. Block 2. Assume P has no zero on $(0,1)$; by IVT it is one sign there; fix the sign to positive by a uniform relabelling. Block 3. Expand each $1/(1-x^{k+1})$ as a geometric series and reorganise into $\sum_m x^m P(x^m)$. Block 4. For $m \geq 1$, $x^m \in (0,1)$ forces each term strictly positive; for $m = 0$, a one sided limit gives $P(1) \geq 0$. Block 5. The resummed total equals the hypothesis's left side, forced strictly positive, contradicting the given value 0.

The nominal route (approach 1) never separates the two variables at all, and stalls the instant it tries to extract a statement about y from an equation purely about x. The bifurcation point against it sits at Block 1. The IVT-on-x route (approach 2) does separate correctly enough to reach for IVT, but aims the theorem at the given sum, which has no freedom to vary; the same Block 1 bifurcation applies, now phrased as which function the existential variable belongs to. Only once P, not the fractional sum, is recognised as the object IVT must act on does Block 2 onward become available, and Block 3 is what actually connects the fixed number x back to P: the geometric expansion is the bridge, not IVT itself.

Layer 2: HOW

Block 1. **Cue.** The conclusion names $y \in (0,1)$ and a specific different expression, $P(y)$, while the hypothesis concerns a fixed x and a sum of reciprocals. **Why natural.** An existential conclusion always names the function and the domain the missing object must satisfy; treating x as though it shared that freedom, or treating the hypothesis's sum as the function needing a root, is not a small slip, it removes the only variable the proof can actually move. **Link to the next block.** Once P on $(0,1)$ is correctly identified as the target, the standard route to proving a continuous function has a root without exhibiting it is the intermediate value theorem, run by contradiction.

Block 2. **Cue.** P is a polynomial, hence continuous everywhere, and $(0,1)$ is a single interval. **Why natural.** A continuous function with no zero on a connected set cannot change sign there, since a sign change would itself force a zero by IVT; this converts "no root" into the much stronger, much more usable statement "one constant sign throughout." **Link to the next block.** A constant, but as yet unknown, sign is only useful once fixed; negating every a_k uniformly flips the sign of P everywhere without disturbing whether the hypothesis sum is zero or whether P has a root, so WLOG the sign is positive.

Block 3. **Cue.** The denominators $1 - x^{k+1}$ carry the explicit constraint $0 < x < 1$ sitting right beside them. **Why natural.** $1/(1-r)$ for $|r| < 1$ is a convergent geometric series in disguise, and an explicit convergence guarantee attached to a term of exactly that shape is the reflex trigger to expand it; this is the one manipulation of the given hypothesis that introduces a free summation index to work with. **Link to the next block.** Expanding and reorganising by that new index exposes the fixed number x raised to every power m, which is exactly the sequence of points at which P can now be evaluated.

Block 4. **Cue.** Every x^m for $m \geq 1$ lies inside $(0,1)$, precisely the domain on which the assumed sign was established, while $x^0 = 1$ sits on the boundary, outside it. **Why natural.** The boundary point is exactly the kind of gap a full marks solution has to close explicitly, not wave past; continuity supplies a one sided limit from inside $(0,1)$ out to 1, weakening strict positivity to nonnegativity only there, which is all the rest of the argument needs. **Link to the next block.** With every term accounted for, one nonnegative and the rest strictly positive, the infinite sum is forced strictly positive.

Block 5. **Cue.** The reorganised infinite sum equals, term for term, the hypothesis's left hand side, which is given to be exactly 0. **Why natural.** A quantity proved strictly positive cannot also be given as zero; this contradiction is the entire reason the geometric expansion was worth doing, since it is what makes the fixed, single equation about x finally say something about P's values. **Link to the next block.** No further block: P must have a zero in $(0,1)$.

Layer 3: WHY

A finite linear relation among terms of the shape $1/(1-x^{k+1})$ is not really one equation, it is a compressed infinite family of equations, one for every power of x simultaneously, because $1/(1-x^{k+1})$ packages the entire sequence $1, x^{k+1}, x^{2(k+1)}, \ldots$ into a single closed form. Expanding it unpacks that family, and reorganising by the shared power m regroups the family into a sequence of evaluations of one fixed polynomial, P, along the geometric orbit $x^0, x^1, x^2, \ldots$. This is why the route works: it does not manufacture a connection between the hypothesis and P, it reveals a connection that was already there, hidden inside the reciprocal notation.

The discernment mechanism: whenever a hypothesis is stated as a finite linear combination of terms $1/(1-r_k)$ with each $|r_k| < 1$ made explicit, and the conclusion concerns a related expression

built from the same coefficients, expand every reciprocal as a geometric series and look for the related expression to reappear, evaluated at whatever the ratios reduce to under the shared index. Here every r_k was the same base x raised to a power matching the polynomial's own degree pattern exactly, a_k paired with x^{k+1} in the denominator and with y^k in the polynomial, an index shift of exactly one, which is the examiners' clearest nudge that a substitution linking x's powers to P's argument was intended, not a coincidence of notation. Training this discernment means drilling $1/(1-x^{k+1})$, under an explicit $0 < x < 1$, into a reflex for generating function language exactly as automatic as reading $\binom{n}{k}$ as "choose."

Revolving concepts

Geometric series and its recursive origin. Definition: for $|r| < 1$, $\sum_{m=0}^{\infty} r^m = 1/(1-r)$. The sequence of terms $c_m = r^m$ is the general solution of the first order linear recursion $c_{m+1} = r\,c_m$, $c_0 = 1$; every such recursion, for any constant ratio r, has closed form $c_m = r^m c_0$, and it is exactly this closed form that licenses summing the series and that produces the geometric orbit $x, x^2, x^3, \ldots$ used throughout the solution. See Book 1, field Sequences and recurrences, topic geometric series and first order linear recursions.

Intermediate value theorem and sign constancy on a connected interval. Definition: a continuous real function on an interval that takes both a positive and a negative value must take the value zero somewhere between them; equivalently, a continuous function with no zero on an interval has one constant sign throughout it. See Book 1, field Real analysis, topic IVT and sign constancy.

One sided limits extending an open interval bound to a closed endpoint. Definition: if f is continuous and $f(y) \geq 0$ for every y in an interval approaching a boundary point c, then $\lim_{y \to c} f(y) \geq 0$, though the strict inequality need not survive the limit. This is the entire justification for the $m = 0$ term of the sum, and skipping it silently is the most common way a candidate loses the point here. See Book 1, field Real analysis, topic continuity and limits.

Generating function bridge between reciprocal power sums and polynomial evaluation. Definition: when the same coefficient sequence $a_0, \ldots, a_n$ appears both as $\sum a_k/(1-x^{k+1})$ and as $\sum a_k y^k$, expanding the first as a family of geometric series and resumming by the shared exponent expresses the first sum exactly as $\sum_{m \geq 0} x^m P(x^m)$; this is the general device linking any reciprocal sum of this shape to evaluations of its coefficient polynomial along a geometric orbit, and it reappears whenever a problem pairs a rational function in one variable with a polynomial sharing its coefficients. See Book 1, field Generating functions, topic coefficient extraction and series to polynomial bridges.

Interchange of a finite sum with a convergent infinite series. Definition: summing a finite collection of convergent series termwise, then adding across the finite index, always agrees with adding across the finite index first and then summing each resulting series; no absolute convergence or Fubini style hypothesis is needed once one of the two indices ranges over a finite set, in contrast to swapping two genuinely infinite sums, which does require absolute convergence to justify reordering. See Book 1, field Real analysis, topic rearrangement of series and finite versus infinite index swaps.

Circuit Bank

- **WHEN** a hypothesis fixes one data variable while the conclusion asks you to find a different unknown for a related but different expression **DO** name both objects explicitly before choosing a technique, and attach any continuity or IVT argument to the function the unknown actually belongs to **BECAUSE** an existential conclusion is only ever about the function and domain it names, and no amount of manipulation of the fixed data variable can substitute for that
- **WHEN** you must prove a continuous function has a zero on an interval without finding it **DO** assume it has none, invoke IVT to force one constant sign throughout, then fix that sign by

whatever symmetry the problem allows (uniform negation, relabelling) **BECAUSE** a function with no zero on a connected interval cannot change sign there, and a constant but unfixed sign is not yet usable until pinned down

- **WHEN** a term of the shape $1/(1-r)$ appears with an explicit hypothesis guaranteeing $|r| < 1$ **DO** expand it immediately as the geometric series $\sum r^m$, and look for a shared index to resum by **BECAUSE** this is the one manipulation of such a term guaranteed to converge and to introduce a free summation variable, and an explicit convergence hypothesis sitting next to it is rarely accidental
- **WHEN** an infinite sum is built from an open interval sign bound but includes a boundary index (here $m = 0$, evaluating at $x^0 = 1$) outside that interval **DO** close the gap with a one sided limit from inside the interval, weakening strict positivity to nonnegativity only at that one boundary term **BECAUSE** a full marks solution accounts for every term explicitly; silently assuming the boundary term shares the strict inequality is exactly the kind of gap a grader is trained to find

Backtesting

1. First useful cue: the conclusion names a different variable, y, for a different expression, $P(y)$, sharing only coefficients with the hypothesis, flagging that x is fixed data, not the unknown.
2. Method the expression suggested: $1/(1-x^{k+1})$ with $0 < x < 1$ stated explicitly suggested a geometric series expansion before any other manipulation was considered.
3. Where you would hesitate, and the thought to reinforce: handling the $m = 0$ boundary term, which sits at $y = 1$, outside the open interval where the assumed sign was established; reinforce closing it with a one sided limit, never by silently extending the strict inequality.
4. The trap the question sets: the shared letters $a_0, \ldots, a_n$ and the shared interval $(0, 1)$ invite treating x as though it were free to vary and applying IVT directly to the given sum, which stalls because x is a single fixed number.
5. What makes you recognise it faster next time: seeing the same coefficient sequence attached to reciprocal powers of one variable in the hypothesis and to plain powers of a different variable in the conclusion, an index shift of exactly one between the two.
6. The reusable decision rule: fixed data variable plus an existence claim about a different, coefficient sharing expression means expand every reciprocal term as a geometric series and resum by the shared exponent to relate the two.

Mirror drills

1. Real numbers $b_0, \ldots, b_n$ and z with $0 < z < 1$ satisfy $b_0/(1+z) + b_1/(1+z^2) + \cdots + b_n/(1+z^{n+1}) = 0$. Does $Q(y) = b_0 + b_1 y + \cdots + b_n y^n$ still have a root in $(0, 1)$? *Hint:* expand $1/(1+z^{k+1}) = \sum_m (-z^{k+1})^m$; the resummed total is an alternating sum of $Q(z^m)$ terms, so the single sign argument needs pairing consecutive m, not the direct route above.
2. If real numbers $a_0, \ldots, a_n$ satisfy $\sum_{k=0}^{n} a_k/(k+1) = 0$, show $P(y) = a_0 + \cdots + a_n y^n$ has a root in $(0, 1)$. *Answer:* since $1/(k+1) = \int_0^1 y^k \, dy$, the hypothesis says $\int_0^1 P(y) \, dy = 0$; if P had no root in $(0, 1)$ it would be one sign throughout, forcing the integral to share that sign, a contradiction, the continuous analogue of the geometric orbit argument.
3. Real numbers $c_1, \ldots, c_5$ and a fixed t with $1 < t < 2$ satisfy $c_1 t + c_2 t^2 + \cdots + c_5 t^5 = 0$. Does this force a root in $(0, 1)$ of $Q(s) = c_1 + c_2 s + \cdots + c_5 s^4$? *Answer:* no general conclusion follows; there is no reciprocal-of-a-power structure to expand and $t \notin (0, 1)$, so the bridge circuit correctly does not fire here.

13.4 A4. Balanced circles and the density of zeros

A finite collection of digits 0 *and* 1 *is written around a circle. An* arc *of length* $L \geq 0$ *consists of* L *consecutive digits around the circle. For each arc* w*, let* $Z(w)$ *and* $N(w)$ *denote the number of* 0*'s in* w *and the number of* 1*'s in* w*, respectively. Assume that* $|Z(w) - Z(w')| \leq 1$ *for any two arcs* w, w' *of the same length. Suppose that some arcs* $w_1, \ldots, w_k$ *have the property that*

$$Z = \frac{1}{k}\sum_{j=1}^{k} Z(w_j) \quad \textit{and} \quad N = \frac{1}{k}\sum_{j=1}^{k} N(w_j)$$

are both integers. Prove that there exists an arc w *with* $Z(w) = Z$ *and* $N(w) = N$*.*

Read it out loud

Say the structure out loud before you touch anything. This is a cyclic binary word, a function from $\mathbb{Z}/m\mathbb{Z}$ to $\{0,1\}$, and every object in the question is a window statistic: slide a window of length L around the circle and count the zeros inside it. The hypothesis says that for each fixed length the count is almost constant. The conclusion hands you a target as an *average* of counts you already have. So the whole problem reads: an average of window counts is itself a window count, provided it is an integer. The fields wired in are combinatorics on words (balanced, that is Sturmian, words), double counting, and the integrality half of the pigeonhole principle.

The first useful cue is the phrase "any two arcs of the same length". A hypothesis that compares two objects of the same size and bounds the gap by 1 is never about those two objects; it is about a hidden real number that both of them are rounding. Integers with range at most 1 occupy two consecutive values, so such a family is pinned down completely once you know its *mean*. That is the reaction circuit: range at most one, so compute the mean and you own the whole family.

Within thirty seconds you should have written m for the number of digits, positions in $\mathbb{Z}/m\mathbb{Z}$, q for the total number of zeros, $z_L(p)$ for the zeros in the arc of length L starting at p, and $L := Z + N$, because $Z(w) + N(w)$ is the length of w, so the arc you must produce has its length already determined. After that the problem is one dimensional: find a window of length L with exactly Z zeros.

Approaches, ranked worst to best

1. **Classify the circles first.** *Cue:* the hypothesis looks like a structure theorem waiting to happen. *Circuit:* recognise the balanced condition, quote the Sturmian classification, then compute. A balanced cyclic word is a rotation of a mechanical word, its zeros sitting at positions $\lfloor jm/q \rfloor$, after which every window count is a floor function calculation. True, and a real theorem, but you would install an entire classification to extract one corollary of it, with $q = 0$ and $q = m$ to handle separately. Ranked last: the ink is enormous and none of it is needed.

2. **Slide one window and use the discrete intermediate value theorem.** *Cue:* "arcs of the same length" and a circle. *Circuit:* moving a window one step changes the count by at most one, so no integer value is skipped. Since $z_L(p+1) - z_L(p) \in \{-1, 0, 1\}$, the count attains every integer between its minimum and its maximum, so it would remain to prove $\min z_L \leq Z \leq \max z_L$. That is genuinely half the problem and worth having, but alone it stalls: your data are counts for arcs of *other* lengths L_j, and you cannot compare those with arcs of length L until you find the common currency.

3. **Re-cut the given arcs into arcs of the common length.** *Cue:* the lengths sum to kL, which is exactly k arcs of length L. *Circuit:* translate each w_j so that the translates lie end to end, then re-partition that run of digits into k arcs of length L. The official route, and correct. Translating costs at most 1 zero per arc, and the first arc need not move, so the total drift is at most $k - 1$;

but if no arc of length L had exactly Z zeros, every re-cut piece would miss Z by at least 1, forcing a drift of at least k. Second only because of the bookkeeping: you must check that the concatenation is a legitimate run of digits and that $L \le m$, and keep track of which side of the inequality each estimate lives on.

4. **Density: every window count is a rounding of length times density.** *Cue:* range at most 1, and the mean over all starting points is forced by double counting. *Circuit:* sum z_L over the m starting positions to get qL, so the mean is Lq/m; a family of integers of range at most 1 with that mean is exactly the integers within distance 1 of it. Best. Every count in the problem, given or wanted, becomes the same statement: an integer within distance 1 of its length times the global density q/m. That model is *linear* in the length, so it survives averaging for free, and integrality of Z closes the problem in one line. No wrap-around bookkeeping, no auxiliary arcs, and the exact set of achievable values comes out as a bonus.

The ideal solution

Let $m \ge 1$ be the number of digits around the circle, index the positions by $\mathbb{Z}/m\mathbb{Z}$, and let $d : \mathbb{Z}/m\mathbb{Z} \to \{0,1\}$ be the digit at each position. Let $q = \#\{p : d(p) = 0\}$ and let $\rho = q/m$ be the density of zeros. For $L \ge 0$ and $p \in \mathbb{Z}/m\mathbb{Z}$ let

$$z_L(p) = \#\{\, t : 0 \le t \le L-1,\ d(p+t) = 0 \,\}$$

be the number of zeros in the arc of length L beginning at p. Every arc of length L is of this form and has $L - z_L(p)$ ones, and arc lengths satisfy $0 \le L \le m$.

Lemma 13.4.1 (Counting). *For every $L \ge 0$ we have* $\displaystyle\sum_{p \in \mathbb{Z}/m\mathbb{Z}} z_L(p) = qL.$

Proof. Count the pairs (p,t) with $0 \le t \le L-1$ and $d(p+t) = 0$ in the two possible orders. Summing over p first gives $\sum_p z_L(p)$. Summing over t first gives $\sum_{t=0}^{L-1} \#\{p : d(p+t) = 0\} = \sum_{t=0}^{L-1} q = qL$, because $p \mapsto p+t$ is a bijection of $\mathbb{Z}/m\mathbb{Z}$ for each fixed t. □

Lemma 13.4.2 (Rounding). *Fix L with $0 \le L \le m$ and let $V_L = \{z_L(p) : p \in \mathbb{Z}/m\mathbb{Z}\}$ be the set of zero counts achieved by arcs of length L. Then V_L is exactly the set of integers n with $|n - L\rho| < 1$. In particular every arc of length L satisfies $|z_L(p) - L\rho| < 1$, and conversely every integer within distance 1 of $L\rho$ is achieved.*

Proof. Write $c = \min_p z_L(p)$ and $M = \max_p z_L(p)$. The hypothesis applied to two arcs of length L gives $M - c \le 1$, so $V_L \subseteq \{c, c+1\}$, and by the counting lemma the mean of z_L over the m starting positions is exactly $L\rho$, whence $c \le L\rho \le M$.

If $M = c$ then z_L is constant, $V_L = \{c\}$ and $L\rho = c$, and the only integer within distance 1 of c is c itself. If $M = c+1$ then both values occur, say r positions carry $c+1$ with $1 \le r \le m-1$, so $L\rho = c + r/m$ lies strictly between c and $c+1$; the integers within distance 1 of such a number are exactly c and $c+1$, and both are achieved. In both cases $V_L = \{n \in \mathbb{Z} : |n - L\rho| < 1\}$. □

Now the problem itself. Let $L_j = Z(w_j) + N(w_j)$ be the length of w_j, so $0 \le L_j \le m$, and set

$$L = Z + N = \frac{1}{k}\sum_{j=1}^{k}\big(Z(w_j) + N(w_j)\big) = \frac{1}{k}\sum_{j=1}^{k} L_j.$$

Then L is an integer, being the sum of the integers Z and N, and $0 \le L \le m$ because L is an average of numbers lying in $[0, m]$. By the rounding lemma applied to each w_j we have $|Z(w_j) - L_j\rho| < 1$ for every j, so by the triangle inequality

$$|Z - L\rho| = \left|\frac{1}{k}\sum_{j=1}^{k}\big(Z(w_j) - L_j\rho\big)\right| \le \frac{1}{k}\sum_{j=1}^{k}\big|Z(w_j) - L_j\rho\big| < \frac{1}{k}\cdot k = 1.$$

Thus Z is an integer with $|Z - L\rho| < 1$, so by the second half of the rounding lemma there is a position p with $z_L(p) = Z$. Let w be the arc of length L starting at p. Then $Z(w) = Z$ and $N(w) = L - Z(w) = (Z + N) - Z = N$, as required. ■

Two edge cases, handled and not swept away. If $L = 0$ then every w_j is empty, $Z = N = 0$, and the empty arc works, the argument covering it unchanged since $z_0 \equiv 0$; if $q = 0$ or $q = m$ then every count equals $L\rho$ exactly. Note also that the counting lemma never used $L \leq m$, so the proof is indifferent to whether your reading of the problem permits arcs that wrap past the full circle; the bound $L \leq m$ is needed only to know that an arc of length L exists.

Layer 1: WHAT

Block 1. Coordinatise: positions in $\mathbb{Z}/m\mathbb{Z}$, digits d, total zeros q, density ρ, window statistic $z_L(p)$.

Block 2. Read off the target length: $L = Z + N$ is forced, since length equals zeros plus ones. The problem is now one dimensional.

Block 3. Double count the pairs (starting position, offset): $\sum_p z_L(p) = qL$, so the mean of the window statistic is $L\rho$ for every L.

Block 4. Combine mean with range: a family of integers of range at most 1 with mean $L\rho$ is exactly the integers within distance 1 of $L\rho$. Every count is now a rounding of a linear function of length.

Block 5. Average: the model $L \mapsto L\rho$ is linear, so the averaged datum Z is within distance 1 of $L\rho$ by the triangle inequality alone.

Block 6. Cash in integrality: an integer within distance 1 of $L\rho$ is achieved. The ones count is free because the length is fixed.

Juxtaposition. The nominal route reads "arcs of the same length" and starts comparing the given arcs with each other: re-cut, translate, concatenate, track a drift of at most $k - 1$ against a deficit of at least k. That route works, but every step of it is bookkeeping about *positions*. The ideal route never moves an arc, because it notices that positions are irrelevant once the hypothesis has collapsed each length class to two consecutive values.

The bifurcation point is the sentence "the given arcs have different lengths, so I cannot compare them directly". The zombie response is to make the lengths equal by cutting and pasting. The ideal response is to find a scalar that all lengths report on, and there is exactly one candidate: the density q/m, the only number attached to the circle as a whole.

Layer 2: HOW

Block 1. Cue. A finite object written around a circle, with statistics attached to consecutive runs. **Why natural.** A circle has no distinguished start, so the honest coordinate system is $\mathbb{Z}/m\mathbb{Z}$ and the honest statistics are those invariant under rotation or averaged over it. Writing $z_L(p)$ makes both parameters explicit: the length you control, and the position you will quantify away. **Link to the next block.** The hypothesis constrains p uniformly, which is the hint that p is the variable to average out.

Block 2. Cue. $Z(w) + N(w)$ is not a new quantity: it is the length of w. **Why natural.** The problem gives two averages, and the second is redundant information dressed as extra data. The pair (Z, N) looks like a two dimensional target, but its second coordinate is determined by the first once the length is known. **Link to the next block.** With the length pinned to $L = Z + N$, everything afterwards works inside one length class.

Block 3. Cue. A quantity indexed by a group, with a hypothesis uniform over that group. **Why natural.** When a hypothesis is uniform over a group you learn nothing from any single element, so sum over the group. Summing $z_L(p)$ over p is a double count of pairs (position, offset) whose

inner sum is the same for every offset because rotation is a bijection. It is the cheapest calculation available and it produces the only global invariant, the density. **Link to the next block.** The mean is now exact and the hypothesis controls the spread; mean plus spread determines a two valued family completely.

Block 4. **Cue.** "Any two differ by at most 1", applied to integers. **Why natural.** Integers of range at most 1 occupy two consecutive values, and such a family is determined by its mean: constant if the mean is an integer, otherwise exactly the two neighbours of the mean, both occurring. No third case exists, so the qualitative hypothesis becomes an exact description. **Link to the next block.** You need that description in both directions: forwards for the data you are given, backwards for the arc you must produce.

Block 5. **Cue.** The conclusion is stated as an average of the hypotheses. **Why natural.** Averaging destroys almost every property but preserves membership of an interval and commutes with linear functions. The approximation $z_L(p) \approx L\rho$ has error below 1 uniformly and its main term is linear in L, so averaging the counts averages the main terms and the triangle inequality controls the errors. Had the expected count been a nonlinear function of L, this step would fail and the problem would be false. **Link to the next block.** You now hold a strict inequality about an integer, which is a pigeonhole waiting to be cashed.

Block 6. **Cue.** "are both integers". Examiners do not say "integer" for decoration. **Why natural.** The bound $|Z - L\rho| < 1$ says nothing about a real number, but for an integer it isolates one or two candidates, and the rounding lemma says exactly those candidates are achieved. The integrality hypothesis is not a convenience; it is the whole content of the final step. **Link to the next block.** Fixing the length gives the ones count for free, so the proof ends here.

Layer 3: WHY

The truth underneath is that a balanced cyclic word carries no local information beyond its density: the hypothesis forces $z_L(p) \in \{\lfloor L\rho \rfloor, \lceil L\rho \rceil\}$ for all p and L, which says the zeros are as equidistributed as integers permit, and is the defining property of mechanical (Sturmian) words, whose zeros sit at a Beatty type sequence $\lfloor j/\rho \rfloor$. So the circle carries one parameter, and the numbers $Z(w_1), \ldots, Z(w_k)$ are k noisy measurements of it, each with error below one unit. The conclusion is a statement about noise under averaging, and it holds because the signal $L\rho$ is linear in L.

The discernment mechanism, stated for reuse: when a hypothesis bounds the *spread* of a family of integers, immediately demand its *mean*, because spread plus mean is a complete description when the spread is at most one. The mean is almost always obtainable by double counting, since a mean over a group orbit is a sum in disguise. Thereafter treat every member of the family as "linear model plus error below one", and every later manipulation is arithmetic rather than combinatorics. Train it by writing, on sight of "differ by at most one", the words two consecutive values, mean determines everything; and on sight of a window sliding around a cycle, sum over all positions equals window length times total. Those two reflexes solved this in five lines, and they are the same reflexes behind equidistribution and three distance arguments.

The examiners nudged three times. They wrote the hypothesis as a comparison of two arcs rather than as "the count takes at most two values", which is the same statement dressed to look local. They handed you Z and N separately although only their sum matters for the length. And they inserted "are both integers", which is the pigeonhole flag. Miss all three and you can still find the concatenation proof, which is why this sits at A4 and not at A6.

Revolving concepts

Cyclic binary word. A function $d : \mathbb{Z}/m\mathbb{Z} \to \{0,1\}$, that is, a necklace. Orbiting it: the window statistics z_L, the density $\rho = q/m$, and the rotation action of $\mathbb{Z}/m\mathbb{Z}$ on itself. Anything rotation invariant is expressible in the density and the window statistics. See Book 1, field: Combinatorics, topic: cyclic words and necklaces.

Balanced word. Balanced means any two factors of the same length have zero counts differing by at most 1. Theorem (Morse and Hedlund): the aperiodic infinite balanced words are exactly the Sturmian words, and the periodic ones are exactly the rotations of the mechanical words $d(n) = \lfloor (n+1)\rho + \theta \rfloor - \lfloor n\rho + \theta \rfloor$. You do not need the classification, and quoting it would be a knowledge gap in disguise; the one consequence used, $z_L(p) \in \{\lfloor L\rho \rfloor, \lceil L\rho \rceil\}$, was proved from scratch. See Book 1, field: Combinatorics, topic: balanced and Sturmian words.

Double counting over a group orbit, with the recursion generalised. The identity $\sum_p \sum_{t=0}^{L-1} F(p+t) = L \sum_p F(p)$ holds for any F on $\mathbb{Z}/m\mathbb{Z}$: the moving sum of any function over a cycle has mean equal to the window length times the mean of the function. The moving sums obey the first order recursion $S_{L+1}(p) = S_L(p) + F(p+L)$, whose solution is the telescoping sum; summing that recursion over p gives $\sum_p S_{L+1} = \sum_p S_L + \sum_p F$, an arithmetic progression in L with common difference $\sum_p F$. That is the counting lemma, now visible as the solution of a linear recursion, and it is why the model is linear in L. See Book 1, field: Combinatorics, topic: double counting.

Discrete intermediate value theorem. If $u : \mathbb{Z}/m\mathbb{Z} \to \mathbb{Z}$ has $|u(p+1) - u(p)| \leq 1$ then u attains every integer between its extremes; here $z_L(p+1) - z_L(p) = [d(p+L) = 0] - [d(p) = 0]$. This gives achievability without the balanced hypothesis and is worth carrying separately, since many problems supply a step condition but no range condition. See Book 1, field: Combinatorics, topic: discrete continuity arguments.

Integrality and rounding. For real x and integer n, $|n - x| < 1$ forces $n \in \{\lfloor x \rfloor, \lceil x \rceil\}$, and $n = x$ if x is an integer. With the triangle inequality this is the engine of the last step. See Book 1, field: Number theory, topic: floor and ceiling.

Circuit Bank

- **WHEN** a family of integers is assumed to have range at most 1 **DO** compute the family's mean, usually by double counting, and declare the family to be exactly the integers within distance 1 of that mean **BECAUSE** integers of range at most 1 occupy two consecutive values, so the mean pins down both the values and their achievability
- **WHEN** a statistic is attached to a sliding window on a cycle **DO** sum it over all starting positions, using that the sum equals window length times total **BECAUSE** rotation is a bijection of the index group, so the double count collapses in one line and produces the only rotation invariant available
- **WHEN** a hypothesis says some quantity "is an integer" **DO** plan for a final step of the form "an integer strictly within 1 of x is a rounding of x" **BECAUSE** examiners insert integrality exactly where a strict inequality must be upgraded to an identification
- **WHEN** the conclusion is an average of the given data **DO** check whether the quantity has an exact linear model plus a uniformly bounded error **BECAUSE** averaging commutes with the linear part and the triangle inequality controls the errors, so linear models survive averaging and nothing else does
- **WHEN** two target statistics are given but their sum is a parameter you control **DO** eliminate one of them immediately **BECAUSE** a two dimensional target lying on a known line is one dimensional, and carrying the spare coordinate costs ink and invites case analysis

Backtesting

1. First useful cue: "any two arcs of the same length differ by at most 1", the range at most one flag, with its partner "are both integers", the pigeonhole flag.
2. Method the expression suggested: fix the length class, compute the mean of the window count by double counting, and describe the class exactly as the roundings of length times density.
3. Where you would hesitate, and the thought to reinforce: at the moment you notice the given arcs have different lengths. Do not cut and paste. Compare each arc to the linear model $L\rho$ instead, and the differing lengths become a strength, because the model is linear and averages cleanly.
4. The trap the question sets: presenting Z and N as a pair, tempting you to chase two conditions; and hiding the strictness of $|Z(w_j) - L_j\rho| < 1$, which must stay strict, since k errors of size exactly 1 would break the final bound.
5. What makes you recognise it faster next time: "an average of window statistics is itself a window statistic" is an equidistribution statement, and those are always proved by exhibiting the linear model and bounding the discrepancy.
6. The reusable decision rule: range at most one plus known mean equals complete description, after which every step is an error estimate.

Mirror drills

1. A balanced circle carries 100 digits of which exactly 30 are zeros. How many zeros are in an arc of length 10, and in an arc of length 7? *Answer:* exactly 3 for length 10, since $10 \cdot 30/100 = 3$ is an integer and a two valued family with integer mean is constant; either 2 or 3 for length 7, since $7 \cdot 30/100 = 2.1$, and both values occur.
2. Let $F : \mathbb{Z}/m\mathbb{Z} \to \mathbb{R}$ and $S_L(p) = \sum_{t=0}^{L-1} F(p+t)$. Prove $\max_p S_L(p) \geq \frac{L}{m} \sum_p F(p) \geq \min_p S_L(p)$, and deduce that some arc of length L contains at least $\lceil Lq/m \rceil$ zeros. *Answer:* average S_L over p, use $\sum_p S_L(p) = L \sum_p F(p)$, and a maximum is at least the mean; then take $F = [d = 0]$ and round, the count being an integer.
3. Show that $u : \mathbb{Z}/m\mathbb{Z} \to \mathbb{Z}$ with $|u(p+1) - u(p)| \leq 1$ attains every integer between its extremes, and give a circle and a length for which z_L takes three values, so that the balanced hypothesis is genuinely needed. *Answer:* walk from an argmin to an argmax one step at a time, skipping nothing; and take $d = 0, 0, 1, 1$ with $L = 2$, where z_2 runs $2, 1, 0, 1$, a word that is not balanced.

13.5 A5. Area definite lists and average shadows

For $m \geq 3$, a list of $\binom{m}{3}$ real numbers a_{ijk} $(1 \leq i < j < k \leq m)$ is said to be area definite *for $\mathbb{R}^n$ if the inequality*

$$\sum_{1 \leq i < j < k \leq m} a_{ijk} \cdot \text{Area}(\Delta A_i A_j A_k) \geq 0$$

holds for every choice of m points $A_1, \ldots, A_m$ in $\mathbb{R}^n$. For example, the list of four numbers $a_{123} = a_{124} = a_{134} = 1$, $a_{234} = -1$ is area definite for $\mathbb{R}^2$. Prove that if a list of $\binom{m}{3}$ numbers is area definite for $\mathbb{R}^2$, then it is area definite for $\mathbb{R}^3$.

Read it out loud

Read the quantifiers first, because they are the problem. The hypothesis is universal over planar configurations, the conclusion universal over spatial ones: you are given information about a small class of objects and asked for information about a larger class containing it. Exactly one shape of argument does that. Take an object of the large class, manufacture objects of the small class from it, apply the hypothesis, reassemble. So this is a transport problem, not a geometry problem, and the only canonical way to manufacture a planar configuration from a spatial one is orthogonal projection. The fields are vector geometry in $\mathbb{R}^3$ (the vector product and vector area), integral geometry (Cauchy's projection formula), and the decisive triviality that integration against a nonnegative measure preserves ≥ 0.

What is really being tested is whether you will *stop choosing*. Every solver projects; most then hunt for a good direction, because choosing is what we are trained to do. The direction cannot be chosen, since the distortion factor depends on the triangle and different triangles lie in different planes. The move is to project in every direction at once and average, and the average behaves only because the averaging measure can be taken rotation invariant.

Within thirty seconds you should have written the vector area $\mathbf{V}_{ijk} = \frac{1}{2}(A_j - A_i) \times (A_k - A_i)$, noted $\text{Area}(\Delta A_i A_j A_k) = |\mathbf{V}_{ijk}|$, and asked what the shadow of the triangle along a unit vector $\hat{u}$ has for its area. That is one cross product away from $|\mathbf{V}_{ijk} \cdot \hat{u}|$, and the moment you see the absolute value of a linear functional you should smell the plan: a norm is about to be recovered as an average of absolute values of linear functionals.

Approaches, ranked worst to best

1. **Classify the area definite lists for $\mathbb{R}^2$.** *Cue:* the hypothesis is a condition on the numbers a_{ijk}. *Circuit:* find the extreme rays of a convex cone, then verify the conclusion on each. The area definite lists form a closed convex cone, so its extreme rays would suffice. Nobody has a usable description of that cone for general m, so you would replace a two page problem with an open ended one. Ranked last: it attacks the coefficients when the quantifier that matters is over configurations.

2. **Squash the picture with one clever projection.** *Cue:* the conclusion is about $\mathbb{R}^3$ and the hypothesis about $\mathbb{R}^2$. *Circuit:* choose a projection direction adapted to the configuration. A single projection multiplies each area by the cosine of the angle between its plane and the screen, and those cosines differ from triangle to triangle, so the projected inequality is a differently weighted sum and says nothing about the unweighted one. It fails informatively: the obstruction is the triangle dependence of one number.

3. **Project onto the three coordinate planes and add.** *Cue:* three orthogonal screens ought to see everything. *Circuit:* sum the three projected inequalities. With $\hat{n}$ a unit normal, the shadows have areas $\text{Area} \cdot |n_1|$, $\text{Area} \cdot |n_2|$, $\text{Area} \cdot |n_3|$, so adding gives $\sum a_{ijk} \text{Area}_{ijk}(|n_1| + |n_2| + |n_3|) \geq 0$. The weight lies in $[1, \sqrt{3}]$ but still depends on the triangle. Ranked third: the correct idea

with the wrong index set, and its failure names the requirement, that the total weight must be blind to orientation.

4. **Work with vector areas and linear algebra.** *Cue:* $\mathbf{V}_{ijk}$ is additive over a surface, so perhaps the sum is algebraic. *Circuit:* replace areas by bivectors and use linearity. The vector area is beautifully linear, but the quantity in the problem is the norm $|\mathbf{V}_{ijk}|$, and norms are not. With mixed signs a_{ijk} no triangle inequality manipulation produces the result. Ranked fourth: right objects, but you need a device that linearises the norm, and that device is the integral in the next approach.

5. **Average the shadow over all directions.** *Cue:* the distortion depends on the triangle, so refuse to choose a direction and integrate over all of them. *Circuit:* the identity $|\mathbf{V}| = \frac{1}{2\pi}\int_{S^2} |\mathbf{V}\cdot\hat{u}|\,d\sigma(\hat{u})$, which is Cauchy's projection formula: a Euclidean norm is an average of absolute values of linear functionals. Best, and once seen, forced. For each fixed $\hat{u}$ the projected points form a genuine planar configuration, so the hypothesis gives a nonnegative number; the projected areas are $|\mathbf{V}_{ijk}\cdot\hat{u}|$; integrating against the rotation invariant measure turns each into $2\pi|\mathbf{V}_{ijk}|$, the same constant for every triangle precisely because the measure is rotation invariant. Divide by $2\pi > 0$.

The ideal solution

For $P, Q, R \in \mathbb{R}^3$ write $\mathbf{V}(P,Q,R) = \frac{1}{2}(Q-P)\times(R-P)$ for the *vector area* of the triangle PQR, so $\mathrm{Area}(\Delta PQR) = |\mathbf{V}(P,Q,R)|$, with $\mathbf{V} = \mathbf{0}$ exactly when the points are collinear. For a unit vector $\hat{u}$ let $\Pi_{\hat{u}} = \hat{u}^{\perp}$ and $\pi_{\hat{u}}(x) = x - (x\cdot\hat{u})\hat{u}$.

Lemma 13.5.1 (Shadow area)**.** *For all $P, Q, R \in \mathbb{R}^3$ and every unit vector $\hat{u}$,*

$$\mathrm{Area}\big(\Delta\,\pi_{\hat{u}}(P)\,\pi_{\hat{u}}(Q)\,\pi_{\hat{u}}(R)\big) = \big|\mathbf{V}(P,Q,R)\cdot\hat{u}\big|.$$

Proof. Put $u = Q-P$, $v = R-P$; by linearity $\pi_{\hat{u}}(Q) - \pi_{\hat{u}}(P) = \pi_{\hat{u}}(u)$ and likewise for v. Both lie in $\Pi_{\hat{u}}$, so their vector product is parallel to $\hat{u}$, say $\pi_{\hat{u}}(u)\times\pi_{\hat{u}}(v) = \lambda\hat{u}$ with $\lambda = (\pi_{\hat{u}}(u)\times\pi_{\hat{u}}(v))\cdot\hat{u}$. Expanding,

$$\pi_{\hat{u}}(u)\times\pi_{\hat{u}}(v) = u\times v - (v\cdot\hat{u})\,(u\times\hat{u}) - (u\cdot\hat{u})\,(\hat{u}\times v),$$

the term $(u\cdot\hat{u})(v\cdot\hat{u})\,\hat{u}\times\hat{u}$ vanishing. Taking the scalar product with $\hat{u}$ kills the last two terms, since $(u\times\hat{u})\cdot\hat{u} = 0 = (\hat{u}\times v)\cdot\hat{u}$, so $\lambda = (u\times v)\cdot\hat{u}$ and the shadow area is $\frac{1}{2}|\lambda| = |\mathbf{V}(P,Q,R)\cdot\hat{u}|$. If the points are collinear both sides vanish, so the identity holds in every case. □

Lemma 13.5.2 (Sphere average)**.** *Let σ denote the standard surface measure on $S^2 \subset \mathbb{R}^3$. For every $\mathbf{V} \in \mathbb{R}^3$,*

$$\int_{S^2} \big|\mathbf{V}\cdot\hat{u}\big|\,d\sigma(\hat{u}) = 2\pi\,|\mathbf{V}|.$$

Proof. Both sides vanish for $\mathbf{V} = \mathbf{0}$ and both are homogeneous of degree 1, so assume $\mathbf{V}$ is a unit vector and choose $R \in SO(3)$ with $Re_3 = \mathbf{V}$. Since σ is rotation invariant and $Re_3\cdot R\hat{v} = e_3\cdot\hat{v}$,

$$\int_{S^2} |\mathbf{V}\cdot\hat{u}|\,d\sigma(\hat{u}) = \int_{S^2} |e_3\cdot\hat{v}|\,d\sigma(\hat{v}) = \int_0^{2\pi}\!\!\int_0^{\pi} |\cos\theta|\sin\theta\,d\theta\,d\varphi = 2\pi\int_{-1}^{1}|c|\,dc = 2\pi. \qquad \square$$

Now let $\{a_{ijk}\}$ be area definite for $\mathbb{R}^2$, let $A_1,\dots,A_m \in \mathbb{R}^3$, write $\mathbf{V}_{ijk} = \mathbf{V}(A_i,A_j,A_k)$ and set

$$F(\hat{u}) = \sum_{1\le i<j<k\le m} a_{ijk}\,\big|\mathbf{V}_{ijk}\cdot\hat{u}\big| \qquad (\hat{u}\in S^2).$$

Claim: $F(\hat{u}) \ge 0$ for every $\hat{u} \in S^2$. Fix $\hat{u}$. The plane $\Pi_{\hat{u}}$ is a two dimensional Euclidean space, so choose a linear isometry $\Phi : \Pi_{\hat{u}} \to \mathbb{R}^2$; isometries preserve areas of triangles. Apply the hypothesis

to the points $B_i = \Phi(\pi_{\hat{u}}(A_i)) \in \mathbb{R}^2$, which is legitimate because area definiteness is assumed for *every* choice of m points of $\mathbb{R}^2$, with no requirement that they be distinct or in general position:

$$0 \le \sum_{i<j<k} a_{ijk} \operatorname{Area}(\Delta B_i B_j B_k) = \sum_{i<j<k} a_{ijk} \operatorname{Area}\big(\Delta\, \pi_{\hat{u}}(A_i)\pi_{\hat{u}}(A_j)\pi_{\hat{u}}(A_k)\big) = F(\hat{u}),$$

the last equality by the shadow lemma. This proves the claim.

F is continuous on S^2, being a finite linear combination of the continuous functions $\hat{u} \mapsto |\mathbf{V}_{ijk} \cdot \hat{u}|$, so it is integrable and the integral of the finite sum is the sum of the integrals. Integrating the claim and applying the sphere average lemma to each term,

$$0 \le \int_{S^2} F(\hat{u})\, d\sigma(\hat{u}) = \sum_{i<j<k} a_{ijk} \int_{S^2} \big|\mathbf{V}_{ijk} \cdot \hat{u}\big|\, d\sigma(\hat{u}) = 2\pi \sum_{i<j<k} a_{ijk} \operatorname{Area}(\Delta A_i A_j A_k).$$

Dividing by $2\pi > 0$ gives the required inequality, and since $A_1, \dots, A_m$ were arbitrary, the list is area definite for $\mathbb{R}^3$. ■

One remark on rigour, in case your reading of the definition insists on *m distinct* points: projections can collide, so replace B_i by $B_i + \varepsilon C_i$ for fixed generic $C_i \in \mathbb{R}^2$ and small $\varepsilon > 0$, which are distinct, and let $\varepsilon \to 0^+$; the sum of areas is continuous in the points, so the non strict inequality survives.

Layer 1: WHAT

Block 1. Read the quantifiers: hypothesis universal over planar configurations, goal universal over spatial ones. Build planar configurations from a spatial one, so project.

Block 2. Quantify the damage: with the vector area, the shadow of a triangle has area $|\mathbf{V} \cdot \hat{u}|$. One cross product line.

Block 3. Observe the damage is triangle dependent, so no single direction can work.

Block 4. Refuse to choose: integrate over all directions, demanding that $\int |\mathbf{V} \cdot \hat{u}|\, d\sigma$ be a fixed multiple of $|\mathbf{V}|$. Rotation invariance plus homogeneity deliver exactly that.

Block 5. Compute the constant, 2π, and note it is positive.

Block 6. Apply the hypothesis for each fixed $\hat{u}$, integrate the pointwise inequality, exchange the finite sum with the integral, divide.

Juxtaposition. The nominal route projects, notices the cosine factors, and then hunts: for a direction making them equal, for bounds on the cosines, for a decomposition into coplanar groups. All of that is choosing, and all of it costs pages. The ideal route projects once, symbolically, in a direction that is never specified.

The bifurcation point is the line "the factor $|\hat{n}_{ijk} \cdot \hat{u}|$ depends on ijk". Zombie response: choose $\hat{u}$ to control it. Ideal response: quantify over $\hat{u}$ and average, because that dependence is only through a direction, and directions are exactly what a rotation invariant average forgets.

Layer 2: HOW

Block 1. Cue. Two universal statements, one over a subcategory and one over the whole. **Why natural.** A hypothesis is useful only where it applies, so you must produce objects it applies to. Your points are genuinely spatial, so you cannot restrict; you must map. Orthogonal projection is the only map onto a plane that is canonical, linear and area controlling. **Link to the next block.** Once you commit to projecting, the only remaining question is what the map costs you.

Block 2. Cue. Areas of triangles under a linear map. **Why natural.** Area in $\mathbb{R}^3$ is $\frac{1}{2}|u \times v|$ and $\pi_{\hat{u}}$ is linear, so push the projection inside. The computation is two lines because $\pi(u)$ and $\pi(v)$ are both perpendicular to $\hat{u}$, forcing their product parallel to $\hat{u}$, so only one scalar is unknown. Writing the answer as $|\mathbf{V} \cdot \hat{u}|$ rather than Area $\cdot\, |\hat{n} \cdot \hat{u}|$ is the lazy choice in the good sense: it removes the sign

ambiguity of the normal and handles collinear triples silently. **Link to the next block.** The shadow area is the absolute value of a linear functional at the vector area, which is a strong structural cue.

Block 3. **Cue.** The factor depends on the triangle. **Why natural.** Different triangles lie in different planes, so one screen flatters some and crushes others; and the coefficients must have mixed signs, or the statement would be trivial, so a weighted inequality with unknown weights carries no information about the unweighted one. **Link to the next block.** The failure names its own cure: you need a family of screens whose total effect is blind to orientation.

Block 4. **Cue.** "Blind to orientation" is the definition of rotation invariance. **Why natural.** Set $I(\mathbf{V}) = \int_{S^2} |\mathbf{V} \cdot \hat{u}| \, d\sigma(\hat{u})$. If σ is rotation invariant then $I(R\mathbf{V}) = I(\mathbf{V})$, and I is homogeneous of degree 1, so $I(\mathbf{V}) = c|\mathbf{V}|$ because $SO(3)$ is transitive on each sphere. An invariant homogeneous positive function of a vector must be a multiple of the norm. You need only $c > 0$, which is clear since the integrand is positive off a great circle. **Link to the next block.** With c existing and positive the proof is finished up to arithmetic, and computing it costs two lines.

Block 5. **Cue.** A spherical integral of $|\cos\theta|$. **Why natural.** Rotate so that $\mathbf{V} = e_3$, legitimate by the invariance already invoked, then integrate in spherical coordinates; $c = \cos\theta$ turns it into $2\pi \int_{-1}^{1} |c| \, dc = 2\pi$. Equivalently the mean of $|\hat{n} \cdot \hat{u}|$ is $\frac{1}{2}$: the average shadow of a plate is half its area. **Link to the next block.** Only the exchange of sum and integral remains, and the sum is finite.

Block 6. **Cue.** An inequality holding pointwise in a parameter you are about to integrate over. **Why natural.** Integration against a nonnegative measure preserves ≥ 0, which is why the plan works at all, and it is linear, which is why the fixed coefficients pass through untouched. A supremum would preserve the inequality too but is not linear. **Link to the next block.** Divide by the positive constant and you have the definition of area definite for $\mathbb{R}^3$.

Layer 3: WHY

Three truths sit under this route. The first is Cauchy's projection formula: for a convex body in $\mathbb{R}^3$ the average shadow area is a quarter of the surface area, and for a flat plate, whose shadow is counted once rather than twice, it is half the area. The identity $\int_{S^2} |\mathbf{V} \cdot \hat{u}| \, d\sigma = 2\pi|\mathbf{V}|$ is that plate case, applied triangle by triangle.

The second, and the reusable one, is that the Euclidean norm is an average of absolute values of linear functionals, $|\mathbf{V}| = \frac{1}{2\pi} \int_{S^2} |\mathbf{V} \cdot \hat{u}| \, d\sigma(\hat{u})$. A norm admitting such a representation is called a zonoid norm, and the representation is the standard way to convert a nonlinear quantity into something that commutes with sums. That is exactly the obstruction that killed the algebraic approach, linearised at the cost of one integral. The third is that integration against a nonnegative measure preserves inequalities and finite sums pass through it: positivity is stable under averaging, and not under differentiation, substitution or rearrangement.

The discernment mechanism: when the hypothesis lives on a small class and the goal on a large one, find a map from large to small, compute the distortion factor, and ask whether averaging over a group acting transitively on the relevant data makes that factor object independent. Here the datum is the direction of the vector area, the group is $SO(3)$, it is transitive on directions, and the invariant measure on S^2 is unique up to scale, so the averaging measure is forced rather than chosen. Whenever you can name the group and check transitivity, the constant exists before you compute it. Train the reflex by asking at every projection or restriction step: does this distortion depend on the object, and if so, what group moves the objects around? That single question covers integral geometry, symmetrisation, averaging over cyclic shifts, and averaging over a finite group, which is the same circuit with counting measure.

The examiners nudged deliberately: a sample list with mixed signs, which kills any term by term argument; the dimensions 2 and 3, where projection is familiar and the vector product exists; and the words "for every choice of m points" rather than "in general position", which is what licenses feeding a collapsed projected configuration into the hypothesis.

Revolving concepts

Vector area. $\mathbf{V} = \frac{1}{2}(Q-P)\times(R-P)$; its magnitude is the area, its direction the normal. Vector areas add over a triangulated surface and sum to zero over a closed one, which is the divergence theorem for constant fields in disguise. See Book 1, field: Geometry, topic: vector product and vector area.

Orthogonal projection. $\pi_{\hat{u}}(x) = x - (x\cdot\hat{u})\hat{u}$, linear, idempotent and self adjoint with image $\hat{u}^{\perp}$. The area distortion of a linear map restricted to a plane W is $|\det(T|_W)|$ in orthonormal bases, here the cosine of the angle between W and $\hat{u}^{\perp}$. See Book 1, field: Linear algebra, topic: projections and determinants as volume scaling.

Cauchy's projection formula. The mean shadow of a convex body in $\mathbb{R}^3$ is a quarter of its surface area; in $\mathbb{R}^n$ the mean volume of the projection onto a hyperplane is a dimensional constant times the surface area. Beneath it sit mixed volumes and Steiner's formula. Knowing the name is how you recognise the technique later. See Book 1, field: Geometry, topic: integral geometry and convex bodies.

Invariant measure. The rotation invariant Borel measure on S^2 is unique up to a positive scalar, being Haar measure transported to $S^2 = SO(3)/SO(2)$; uniqueness is why the constant is forced. The finite analogue, averaging over a finite group with counting measure, is the same tool. See Book 1, field: Geometry, topic: invariant measures and homogeneous spaces.

Generalisation to all dimensions, with the recursion made explicit. Area definite for $\mathbb{R}^2$ implies area definite for $\mathbb{R}^n$ for every $n \geq 2$, by induction on n: project along a uniformly distributed unit vector onto $\hat{u}^{\perp} \cong \mathbb{R}^{n-1}$; a triangle spanning the plane W has its area multiplied by $\sqrt{1-\|P_W\hat{u}\|^2}$, whose average over S^{n-1} is independent of W because $O(n)$ is transitive on two dimensional subspaces and the measure is invariant. The recursion is $\mathcal{A}_{n-1} \Rightarrow \mathcal{A}_n$ with the same one step argument, and the general method is: identify the group acting on the fibre data, check transitivity, average. See Book 1, field: Geometry, topic: Grassmannians and invariant integration.

Circuit Bank

- **WHEN** a hypothesis is known in a small class and needed in a bigger one **DO** map the big object into the small class and compute the distortion factor of the quantity in question **BECAUSE** hypotheses travel along maps and nowhere else, so the only work is measuring what the map costs
- **WHEN** a distortion factor depends on the object **DO** stop choosing the map and integrate over the whole family of maps against an invariant measure **BECAUSE** an invariant homogeneous positive function of a vector is a constant times its norm, so object dependence is exactly what invariance forgets
- **WHEN** you meet $|\hat{n}\cdot\hat{u}|$ averaged over the unit sphere in $\mathbb{R}^3$ **DO** write down the value 2π, equivalently mean value $1/2$ **BECAUSE** it is Cauchy's projection formula for a flat plate, and knowing the constant saves re-deriving it under time pressure
- **WHEN** an inequality holds for every value of a parameter and you want a statement about a weighted sum **DO** integrate the inequality over the parameter rather than optimising over it **BECAUSE** integration against a nonnegative measure preserves ≥ 0 and is linear, so fixed coefficients pass straight through, while suprema are not linear

- **WHEN** a nonlinear quantity such as a norm or an area blocks a linearity argument **DO** look for a representation of it as an average of absolute values of linear functionals **BECAUSE** that representation is what makes norms compatible with sums, and every Euclidean norm has one

Backtesting

1. First useful cue: the mismatch of quantifiers, hypothesis over $\mathbb{R}^2$ and goal over $\mathbb{R}^3$, which forces a map from $\mathbb{R}^3$ into a plane.
2. Method the expression suggested: areas as $|\mathbf{V}_{ijk}|$, shadows as $|\mathbf{V}_{ijk} \cdot \hat{u}|$, and the norm recovered as the sphere average of the linear functional.
3. Where you would hesitate, and the thought to reinforce: when the cosine factor turns out to depend on the triangle. A factor depending only on a direction is annihilated by averaging over a transitive group, so an object dependent constant is an instruction to integrate, not a defeat.
4. The trap the question sets: the temptation to pick a good direction, and the smaller trap that projected points may coincide, harmless only because the definition quantifies over every choice of m points.
5. What makes you recognise it faster next time: "property in dimension n implies property in dimension $n+1$" is an integral geometry flag, and the answer is always project, compute the distortion, average it away.
6. The reusable decision rule: when the distortion depends on the object, average over the symmetry group with its invariant measure; the constant is positive and may be computed later, or never.

Mirror drills

1. Show that the average area of the shadow of a fixed triangle $T \subset \mathbb{R}^3$ on a plane with uniformly random normal direction is $\frac{1}{2}\text{Area}(T)$. *Answer:* the shadow lemma gives shadow area $|\mathbf{V} \cdot \hat{u}|$, and the sphere average lemma gives mean $\frac{1}{4\pi} \cdot 2\pi|\mathbf{V}|$.
2. Deduce Cauchy's surface area formula for a convex polyhedron: the mean shadow area is one quarter of the surface area. *Answer:* triangulate and apply the previous drill to each face; for a convex body almost every line in direction $\hat{u}$ meets the boundary twice, so the shadow is covered exactly twice by the shadows of the faces, giving $\frac{1}{2} \cdot \frac{1}{2}S = \frac{1}{4}S$.
3. Prove that a list area definite for $\mathbb{R}^2$ is area definite for $\mathbb{R}^n$ for every $n \geq 2$. *Answer:* induct on n, projecting along a uniformly random unit vector onto a hyperplane; the area factor for a triangle spanning W is $\sqrt{1 - \|P_W \hat{u}\|^2}$, whose average over S^{n-1} is independent of W since $O(n)$ is transitive on two dimensional subspaces, so divide by that positive constant.

13.6 A6. A positive definite weight on the integer lattice

Define a function $w : \mathbb{Z} \times \mathbb{Z} \to \mathbb{Z}$ as follows. For $|a|, |b| \leq 2$, let $w(a,b)$ be as in the table shown; otherwise, let $w(a,b) = 0$.

$w(a,b)$	$b=-2$	$b=-1$	$b=0$	$b=1$	$b=2$
$a=-2$	-1	-2	2	-2	-1
$a=-1$	-2	4	-4	4	-2
$a=0$	2	-4	12	-4	2
$a=1$	-2	4	-4	4	-2
$a=2$	-1	-2	2	-2	-1

For every finite subset S of $\mathbb{Z} \times \mathbb{Z}$, define

$$A(S) = \sum_{(\mathbf{s},\mathbf{s}') \in S \times S} w(\mathbf{s} - \mathbf{s}').$$

Prove that if S is any finite nonempty subset of $\mathbb{Z} \times \mathbb{Z}$, then $A(S) > 0$. (For example, if $S = \{(0,1),(0,2),(2,0),(3,1)\}$, then the terms in $A(S)$ are $12, 12, 12, 12, 4, 4, 0, 0, 0, 0, -1, -1, -2, -2, -4, -4$.)

Read it out loud

Strip the table away and say what is left. You have a double sum over $S \times S$ whose summand depends only on the *difference* of the two indices. That single feature is the entire problem: a kernel depending only on $\mathbf{s} - \mathbf{s}'$ is translation invariant, the associated quadratic form is a convolution, and translation invariant operators on $\mathbb{Z}^2$ are diagonalised by the characters $\mathbf{u} \mapsto e^{i\langle \mathbf{u}, \theta \rangle}$. So the sentence to say out loud is: this is a Toeplitz quadratic form, and "is it positive?" is the question "is its symbol nonnegative?".

The fields are Fourier analysis on the lattice $\mathbb{Z}^2$ and its dual torus, trigonometric identities of Chebyshev type, and a two variable polynomial inequality on the square $[-1,1]^2$. The table is a smokescreen: all that matters is that it is supported on a 5×5 box, that it is symmetric under $a \mapsto -a$ and under $b \mapsto -b$, and that its entries were manufactured by expanding a nonnegative expression. Somebody chose a nonnegative function on the torus and printed its Fourier coefficients; your job is to run the construction in reverse.

What is tested is whether $w(\mathbf{s} - \mathbf{s}')$ triggers a circuit. If it does, this is thirty minutes of careful routine work. If it does not, you will spend two hours charging the negative entries against the 12s and you will fail, because the failure is structural: the entries sum to zero, so for large S the positive mass is cancelled almost exactly and $A(S)$ is governed by the boundary of S, not its size. Within thirty seconds write the generating function $f(\theta) = \sum_{\mathbf{u} \in S} e^{i\langle \mathbf{u}, \theta \rangle}$, the symbol $g(\theta) = \sum_{\mathbf{c}} w(\mathbf{c}) e^{i\langle \mathbf{c}, \theta \rangle}$, and the guess $A(S) = \mathcal{M}\left[|f|^2 g\right]$. Note also the word "nonempty" and the strict inequality: strictness will not come from the positivity argument, so plan a separate final step.

Approaches, ranked worst to best

1. **Induct on the size of S.** *Cue:* a statement about all finite sets. *Circuit:* add one point and control the increment. Adding $\mathbf{p}$ changes A by $12 + 2\sum_{\mathbf{s} \in S} w(\mathbf{p} - \mathbf{s})$. Take S containing the four neighbours $\mathbf{p} \pm (1,0)$, $\mathbf{p} \pm (0,1)$: the increment is $12 + 2(-4-4-4-4) = -20 < 0$. So A is not monotone and the induction has no hypothesis to run on. Ranked last, and worth killing in one line so that you stop believing in it.

2. **Local charging.** *Cue:* a big positive diagonal entry against small negative ones. *Circuit:* pay each negative term from a nearby positive one. The diagonal contributes $12|S|$ and you try to distribute it. For S an $n \times n$ box, $A(S) = 32(n-1)$ while $|S| = n^2$, so $A(S)/|S| \to 0$: the positivity is a boundary effect, invisible to any local and uniform argument. Instructive, because that computation tells you the form is degenerate in the direction of constants.

3. **Write w as an autocorrelation.** *Cue:* if $w = \varphi \star \tilde{\varphi}$ then $A(S) = \sum_{\mathbf{u}} \left| \sum_{\mathbf{s} \in S} \varphi(\mathbf{s} + \mathbf{u}) \right|^2 \geq 0$ automatically. *Circuit:* exhibit the quadratic form as a sum of squares in the data. The right kind of idea, since this is what positive semidefiniteness means concretely, but searching for φ is blind guesswork and no finitely supported φ exists here: in two variables a nonnegative trigonometric polynomial need not be a squared modulus, which is exactly the failure of the Fejer and Riesz theorem beyond one variable.

4. **Diagonalise the quadratic form as a matrix.** *Cue:* $A(S) = \mathbf{1}_S^{\mathsf{T}} W \mathbf{1}_S$ with $W_{\mathbf{uv}} = w(\mathbf{u} - \mathbf{v})$. *Circuit:* a symmetric matrix is positive definite iff its eigenvalues are positive. Right, and since the eigenvectors of a translation invariant matrix are the characters, this approach *is* the next one in linear algebra clothing. Ranked here only because stated this way it invites finite dimensional eigenvalue estimates such as Gershgorin, which fail: the diagonal entry is 12 against an absolute off diagonal row sum of 36.

5. **Fourier transform on $\mathbb{Z}^2$.** *Cue:* the summand depends only on $\mathbf{s} - \mathbf{s}'$. *Circuit:* a translation invariant kernel is a convolution, characters diagonalise it, and positivity of the form is nonnegativity of the symbol, the easy half of the Herglotz theorem. Best, and forced by the shape of the sum. Write $A(S)$ as the average of $|f|^2 g$ over the torus, compute $g = 16G(\cos s, \cos t)$ from the double symmetry of the table, prove $G \geq 0$ on $[-1,1]^2$ by pairing each negative monomial with the positive monomial that divides it, and take the strictness from the equality analysis, which the same identities hand you for free.

The ideal solution

Write $\theta = (s,t)$, $\langle \mathbf{u}, \theta \rangle = u_1 s + u_2 t$ for $\mathbf{u} \in \mathbb{Z}^2$, and $e_{\mathbf{u}}(\theta) = e^{i\langle \mathbf{u}, \theta \rangle}$. Let

$$\mathcal{M}[F] := \frac{1}{4\pi^2} \int_0^{2\pi} \int_0^{2\pi} F(s,t)\, ds\, dt.$$

The orthogonality relation $\mathcal{M}[e_{\mathbf{u}}] = 1$ if $\mathbf{u} = \mathbf{0}$ and 0 otherwise follows at once from $\int_0^{2\pi} e^{ins} ds = 0$ for $n \neq 0$. For a finite nonempty $S \subseteq \mathbb{Z}^2$ put

$$f(\theta) = \sum_{\mathbf{u} \in S} e_{\mathbf{u}}(\theta), \qquad g(\theta) = \sum_{\mathbf{c} \in \mathbb{Z}^2} w(\mathbf{c})\, e_{\mathbf{c}}(\theta),$$

the second sum being finite since w vanishes off the 5×5 box.

Step 1 (the form is the average of $|f|^2 g$). Since the coefficients of f are real, $\overline{f(\theta)} = \sum_{\mathbf{v} \in S} e_{-\mathbf{v}}(\theta)$, so $|f|^2 = \sum_{\mathbf{u},\mathbf{v} \in S} e_{\mathbf{u}-\mathbf{v}}$ and

$$\mathcal{M}\big[|f|^2 g\big] = \sum_{\mathbf{u},\mathbf{v} \in S} \sum_{\mathbf{c} \in \mathbb{Z}^2} w(\mathbf{c})\, \mathcal{M}\big[e_{\mathbf{u}-\mathbf{v}+\mathbf{c}}\big] = \sum_{\mathbf{u},\mathbf{v} \in S} w(\mathbf{v} - \mathbf{u}) = A(S),$$

the only surviving terms being those with $\mathbf{c} = \mathbf{v} - \mathbf{u}$. All sums are finite, so the interchanges are trivial. Taking $g \equiv 1$ in the same computation gives the fact needed at the end,

$$\mathcal{M}\big[|f|^2\big] = \#\{(\mathbf{u},\mathbf{v}) \in S \times S : \mathbf{u} = \mathbf{v}\} = |S|.$$

Step 2 (compute the symbol). The table satisfies $w(a,b) = w(-a,b) = w(a,-b)$, so grouping the 25 terms into the classes $(|a|,|b|)$ and using $e^{i\alpha} + e^{-i\alpha} = 2\cos\alpha$ gives, with $z = \cos s$ and $y = \cos t$,

$$\begin{aligned} g &= w(0,0) + 2w(1,0)\cos s + 2w(2,0)\cos 2s + 2w(0,1)\cos t + 2w(0,2)\cos 2t \\ &\quad + 4w(1,1)\cos s\cos t + 4w(1,2)\cos s\cos 2t + 4w(2,1)\cos 2s\cos t + 4w(2,2)\cos 2s\cos 2t \\ &= 12 - 8z + 4(2z^2-1) - 8y + 4(2y^2-1) + 16zy - 8z(2y^2-1) - 8y(2z^2-1) - 4(2z^2-1)(2y^2-1). \end{aligned}$$

The constants cancel as $12-4-4-4=0$, the terms in z alone as $-8z+8z=0$, likewise those in y, and collecting the rest gives

$$g(s,t)=16\,G(\cos s,\cos t),\qquad G(z,y)=z^2+y^2+zy-z^2y-zy^2-z^2y^2.$$

Step 3 ($G\geq 0$ on the square). Every negative monomial of G is a positive monomial multiplied by a factor of modulus at most 1, and all three carry a factor zy, so the grouping to use is decided by the sign of zy. For $z,y\in[-1,1]$ we have two identities, each verified by expanding:

$$\text{(i)}\qquad G(z,y)=zy(1-zy)+z^2(1-y)+y^2(1-z),$$
$$\text{(ii)}\qquad G(z,y)=(z+y)^2-zy(1+z)(1+y).$$

If $zy\geq 0$, use (i): each of zy, $1-zy$, $1-y$, $1-z$, z^2, y^2 is nonnegative. If $zy\leq 0$, use (ii): $(z+y)^2\geq 0$ while $-zy\geq 0$, $1+z\geq 0$, $1+y\geq 0$. Hence $g\geq 0$ everywhere on the torus, and by Step 1, $A(S)\geq 0$.

Step 4 (where G vanishes). The same identities give the equality cases at no extra cost. If $zy\geq 0$ and $G=0$, all three terms of (i) vanish: $z^2(1-y)=0$ gives $z=0$ or $y=1$, and $y^2(1-z)=0$ gives $y=0$ or $z=1$. If $z=0$ the second forces $y=0$; if instead $y=1$ then, as $y\neq 0$, the second forces $z=1$, and $zy(1-zy)=0$ holds. So the solutions are $(0,0)$ and $(1,1)$. If $zy\leq 0$ and $G=0$, then $(z+y)^2=0$ and $zy(1+z)(1+y)=0$; the first gives $y=-z$ and substituting into the second gives $z^2(1-z^2)=0$, so the solutions are $(0,0)$, $(1,-1)$ and $(-1,1)$.

Hence G vanishes on $[-1,1]^2$ exactly at $(0,0)$, $(1,1)$, $(1,-1)$, $(-1,1)$. Since $\cos s=0$ has two solutions in $[0,2\pi)$ while each of $\cos s=1$ and $\cos s=-1$ has one, the zero set $N=\{\theta\in[0,2\pi)^2:g(\theta)=0\}$ consists of exactly seven points, and in particular is finite.

Step 5 (strict positivity). Suppose $A(S)=0$. The function $|f|^2g$ is continuous and, by Step 3, nonnegative, with average $A(S)=0$; a continuous nonnegative function with vanishing integral is identically zero, so $|f|^2g\equiv 0$. For $\theta\notin N$ we have $g(\theta)>0$, hence $f(\theta)=0$; so f vanishes off a finite set and therefore, being continuous on a space where the complement of a finite set is dense, vanishes identically. But then $|S|=\mathcal{M}\big[|f|^2\big]=0$, contradicting nonemptiness. Therefore $A(S)>0$. ■

Two remarks worth banking. The argument never used that the coefficients of f are 0 or 1: replacing f by $\sum_{\mathbf{u}}c_{\mathbf{u}}e_{\mathbf{u}}$ for finitely supported $c:\mathbb{Z}^2\to\mathbb{C}$ proves that $\sum_{\mathbf{u},\mathbf{v}}c_{\mathbf{u}}\overline{c_{\mathbf{v}}}w(\mathbf{u}-\mathbf{v})>0$ whenever $c\not\equiv 0$, that is, w is a strictly positive definite function on $\mathbb{Z}^2$. And $A(S)$ is an integer, so in fact $A(S)\geq 1$; the true minimum is 12, attained by a single point, while the $n\times n$ box gives $A(S)=32(n-1)$ for $n\geq 2$, so no bound $A(S)\geq c|S|$ can hold. The reason is visible in the symbol: $g(0,0)=16G(1,1)=0$, the entries of the table summing to zero, so the form is degenerate on the trivial character and large sets are cheap.

Layer 1: WHAT

Block 1. Read the kernel: $w(\mathbf{s}-\mathbf{s}')$ depends only on the difference, so A is a translation invariant quadratic form, that is, a convolution.

Block 2. Diagonalise: characters diagonalise convolutions, so pass to the torus, $A(S)=\mathcal{M}\big[|f|^2g\big]$.

Block 3. Reduce: since $|f|^2\geq 0$, it suffices to prove $g\geq 0$ pointwise. The set S has now left the problem entirely.

Block 4. Compute g: the double symmetry converts exponentials into cosines, and $\cos 2\alpha=2\cos^2\alpha-1$ converts those into a polynomial $16G(z,y)$ on the square.

Block 5. Prove $G\geq 0$ on the square by pairing each negative monomial with a positive monomial it divides, splitting on the sign of zy because all three negative monomials carry that factor.

Block 6. Read the equality cases off the same identities: four points, hence a finite zero set on the torus.

Block 7. Upgrade to strict: a continuous nonnegative function of zero average vanishes identically, $g > 0$ off a finite set, so $f \equiv 0$, contradicting $\mathcal{M}\left[|f|^2\right] = |S| \geq 1$.

Juxtaposition. The nominal route sees sixteen terms in the worked example and a big 12 on the diagonal, and starts accounting: 12 per point, -4 per edge neighbour, $+4$ per diagonal neighbour. That is zombie maths of the purest kind, and the $n \times n$ box shows why it can never close, since the leading terms cancel identically and the truth lives on the boundary.

The bifurcation point is the very first line: how you read $w(\mathbf{s} - \mathbf{s}')$. Read as "a number attached to each pair", you go to accounting. Read as "the kernel depends only on the difference", you go to Fourier. The two routes never meet again.

Layer 2: HOW

Block 1. **Cue.** The summand is a function of $\mathbf{s} - \mathbf{s}'$ alone. **Why natural.** Difference kernels are exactly the kernels invariant under translating S, and a quadratic form invariant under a group action should be studied in a basis adapted to the group. For $\mathbb{Z}^2$ acting on itself, that basis is the characters, because translation multiplies each character by a scalar. This is the principle that diagonalises a circulant matrix, and it is not special to this problem: convolution becomes multiplication, always. **Link to the next block.** You now need the dictionary between a set and a function on the torus, which is the generating function.

Block 2. **Cue.** A sum over $S \times S$ of a difference kernel, plus a wish to see squares. **Why natural.** $\sum_{\mathbf{u},\mathbf{v} \in S} e_{\mathbf{u}-\mathbf{v}}$ is exactly $|f|^2$, the identity behind every "count pairs with a given difference" argument. Multiplying by g and averaging kills every character but the trivial one, selecting $\mathbf{c} = \mathbf{v} - \mathbf{u}$. Equivalently $A(S)$ is the constant coefficient of the Laurent polynomial $f(x,y)f(x^{-1},y^{-1})g(x,y)$, and averaging over the two unit circles is how one extracts a constant coefficient. **Link to the next block.** The factor $|f|^2$ is nonnegative whatever S is, so all the content has been pushed into g.

Block 3. **Cue.** A product of a nonnegative factor with an unknown one, integrated. **Why natural.** Proving $g \geq 0$ proves the theorem for all S simultaneously. This is why the Fourier route dominates every combinatorial route: it replaces infinitely many statements, one per S, by one statement about a fixed function of two real variables. The price is a two variable inequality, which is finite work. **Link to the next block.** So compute g, in a form you can actually test for nonnegativity.

Block 4. **Cue.** The table is symmetric in $a \mapsto -a$ and in $b \mapsto -b$. **Why natural.** Symmetry pairs the exponentials into cosines, so g is real valued, which it had better be if it is to be nonnegative; then $\cos 2\alpha = 2\cos^2 \alpha - 1$, the first nontrivial case of the Chebyshev recursion, turns everything into a polynomial in z and y. Crucially the substitution also records the constraint $z, y \in [-1, 1]$, which is not a technicality but the whole reason the inequality is true, since $G(2,2) = -20 < 0$. **Link to the next block.** You now need a polynomial inequality on a box, and boxes have canonical nonnegative atoms.

Block 5. **Cue.** A polynomial with a few negative monomials, to be shown nonnegative on $[-1,1]^2$. **Why natural.** On the square the manifestly nonnegative quantities are $1 \pm z$, $1 \pm y$, $1 - zy$, squares and their products, so any certificate is assembled from those. Write each negative monomial as a positive monomial times a factor known to lie in $[-1,1]$: $z^2 y = z^2 \cdot y$, $zy^2 = y^2 \cdot z$, $z^2y^2 = zy \cdot zy$. That produces (i) at once, valid whenever the leftover zy is nonnegative. When $zy \leq 0$ the term $zy(1 - zy)$ turns against you, so regroup: $z^2 + y^2 + zy = (z+y)^2 - zy$ is comfortable when $-zy \geq 0$, and the remaining $-zy(1 + z + y + zy)$ factors as $-zy(1+z)(1+y)$, which is (ii). Neither identity is a rabbit from a hat; each is the natural grouping in its regime. **Link to the next block.** Both identities are sums of products of nonnegative factors, so their equality cases are free.

Block 6. **Cue.** You proved ≥ 0 but the problem said > 0. **Why natural.** A sum of nonnegative terms vanishes only if every term does, so the zero set follows from two or three small equations in each regime. Do it now, while the identities are in front of you: the alternative, quoting that the zero set of a nonzero real analytic function has measure zero, is heavier machinery for a weaker conclusion. **Link to the next block.** A finite zero set makes the closing continuity argument elementary.

Block 7. **Cue.** Strict inequality from a nonstrict one, with an integral in hand. **Why natural.** A continuous nonnegative function with zero integral is identically zero; this is where the continuity of $|f|^2 g$ earns its keep. Then f vanishes off a finite set, hence everywhere, hence S is empty. The identity $\mathcal{M}\left[|f|^2\right] = |S|$ is Parseval, and it is the only place nonemptiness is used, exactly as it should be, since the statement is false for $S = \emptyset$. **Link to the next block.** Done; the remarks on integrality and the box are for the Circuit Bank, not the proof.

Layer 3: WHY

The governing truth is the Herglotz theorem, the $\mathbb{Z}^d$ case of Bochner's theorem: a function w on $\mathbb{Z}^d$ is positive definite exactly when it is the Fourier transform of a nonnegative measure on the torus, and for finitely supported w that measure has the symbol g as its density. This problem asks for the easy direction of that equivalence, applied to a w constructed by running the theorem backwards. Once you know the theorem exists the problem is not a puzzle: it is a Fourier transform and a polynomial inequality on a box.

Two subsidiary truths carry weight. Convolution operators are diagonalised by characters, which is why S can be eliminated at Step 3. And positivity on a semialgebraic set is certified by combinations of the defining inequalities: on $[-1,1]^2$ those are $1 \pm z \geq 0$ and $1 \pm y \geq 0$, so every certificate is a positive combination of products of these with squares. That is the systematic reason the two identities look as they do, and the reason a split on the sign of zy is the natural first attempt, since zy is the one product of atoms whose sign the square does not fix.

The discernment mechanism: whenever a quantity is a double sum over a set with a summand depending only on the difference of the indices, you are looking at $\mathcal{M}\left[|f|^2 g\right]$ in disguise, and the question about all sets collapses to a question about one function. The recognition cue is syntactic, not semantic; you need not understand what w means, only see the subtraction inside its argument. Train it by rewriting every "count the pairs at difference $\mathbf{d}$" problem as a constant coefficient extraction until the translation is automatic.

The examiners' nudges are clear once you know the answer. They printed a symmetric table so the symbol would be real; they chose entries summing to zero so that no crude size based bound could work; they gave a worked example with sixteen listed terms as bait for the accounting route; and they asked for a strict inequality, warning that the natural argument gives only ≥ 0. A6 status is honest here: without the Fourier circuit installed there is no reasonable route in, and the knowledge gap is real rather than a matter of ingenuity.

Revolving concepts

Characters and the dual torus. The characters of $\mathbb{Z}^d$ are $\mathbf{u} \mapsto e^{i\langle \mathbf{u}, \theta \rangle}$ for θ in $(\mathbb{R}/2\pi\mathbb{Z})^d$, and the orthogonality relation $\mathcal{M}[e_{\mathbf{u}}] = [\mathbf{u} = \mathbf{0}]$ is the whole computational content of the subject at this level. See Book 1, field: Discrete structures and graphs, topic: Fourier analysis on abelian groups.

Convolution and Toeplitz forms. A kernel $K(\mathbf{u}, \mathbf{v}) = w(\mathbf{u} - \mathbf{v})$ gives the form $\langle w * c, c \rangle$ with symbol $\hat{w} = g$, and the values of g are the eigenvalues of the operator. The finite one dimensional analogue is a circulant matrix, whose eigenvalues are the discrete Fourier transform of its first row. See Book 1, field: Matrices and determinants, topic: circulant and Toeplitz matrices.

Herglotz and Bochner. Positive definite on $\mathbb{Z}^d$ iff the Fourier transform is a nonnegative measure. Positive definite functions also arise as autocorrelations $\varphi \star \tilde{\varphi}$; in one variable Fejer and Riesz gives $g = |\varphi|^2$ for a trigonometric polynomial φ, but this fails in two or more variables, which is exactly why the autocorrelation approach has no chance here. See Book 1, field: Real analysis, topic: positive definite functions.

Chebyshev polynomials, with the recursion generalised. $\cos(n\alpha) = T_n(\cos\alpha)$ where $T_0 = 1$, $T_1 = z$ and $T_{n+1}(z) = 2zT_n(z) - T_{n-1}(z)$. Any trigonometric polynomial even in α is a polynomial in $\cos\alpha$, obtained by running that linear recursion, and the substitution automatically encodes $\cos\alpha \in [-1, 1]$; in several variables you run it once per variable, which is Step 2. Solving the recursion by its characteristic equation $\lambda^2 - 2z\lambda + 1 = 0$, whose roots are $e^{\pm i\alpha}$ with $z = \cos\alpha$, gives the closed form $T_n(z) = \cos(n \arccos z)$; every constant coefficient linear recursion is handled the same way. See Book 1, field: Polynomials, topic: Chebyshev polynomials.

Certificates of positivity on a box. To prove $P \geq 0$ on $[-1, 1]^2$, write P as a positive combination of products of squares with the atoms $1 - z$, $1 + z$, $1 - y$, $1 + y$; such a representation always exists for strictly positive P (Positivstellensatz), and finding one by hand means matching each negative monomial with a positive monomial that divides it. See Book 1, field: Inequalities, topic: sums of squares and positivity certificates.

From nonnegative to positive. If $F \geq 0$ is continuous on a compact space with $\int F = 0$ then $F \equiv 0$; combined with an explicit description of where one factor vanishes, this converts any "≥ 0" proof into a "> 0" proof. See Book 1, field: Real analysis, topic: vanishing integrals.

Circuit Bank

- **WHEN** a double sum over $S \times S$ has a summand depending only on $\mathbf{s} - \mathbf{s}'$ **DO** write it as $\mathcal{M}\left[|f|^2 g\right]$ with f the generating function of S and g the symbol of the kernel **BECAUSE** difference kernels are convolutions, characters diagonalise convolutions, and the set then disappears from the problem
- **WHEN** you must prove a quadratic form in an arbitrary finite set is nonnegative **DO** compute the symbol and prove it is nonnegative pointwise **BECAUSE** by Herglotz that condition is necessary and sufficient, so you are guaranteed not to be wasting effort
- **WHEN** a trigonometric polynomial is invariant under $s \mapsto -s$ and $t \mapsto -t$ **DO** substitute $z = \cos s$, $y = \cos t$ through the Chebyshev recursion and remember the constraint $z, y \in [-1, 1]$ **BECAUSE** the problem becomes a polynomial inequality on a box, and the box constraint is usually essential rather than cosmetic
- **WHEN** you must certify a polynomial is nonnegative on $[-1, 1]^2$ **DO** match each negative monomial with a positive monomial that divides it, and split on the sign of the one product whose sign is undetermined **BECAUSE** the certificate must be built from the atoms $1 \pm z$, $1 \pm y$ and squares, and this matching is the mechanical way to find it
- **WHEN** your argument yields ≥ 0 but the problem demands > 0 **DO** extract the equality cases from the certificate you already wrote, show the zero set is small, and use that a continuous nonnegative function of zero integral vanishes identically **BECAUSE** the equality analysis is free once the proof is a sum of nonnegative terms, and it is far cheaper than any analytic continuation argument

Backtesting

1. First useful cue: the subtraction inside $w(\mathbf{s} - \mathbf{s}')$, announcing a translation invariant kernel and therefore Fourier; secondary cues are the double symmetry of the table and its entries summing to zero.

2. Method the expression suggested: express $A(S)$ as the constant coefficient of $f(x,y)f(x^{-1},y^{-1})g(x,y)$, extracted by averaging over the two unit circles, and then work only with g.
3. Where you would hesitate, and the thought to reinforce: at the two variable inequality, where the negative monomials look unmanageable. On a box every negative monomial is a positive monomial times a factor of modulus at most one, so pair them off and split only on the sign the box does not fix.
4. The trap the question sets: the worked example with its sixteen terms, inviting local accounting; and the strict inequality, which the main argument does not give, so that stopping at $A(S) \geq 0$ loses marks after the hard part is done.
5. What makes you recognise it faster next time: any weight depending on a difference of lattice points is a Fourier problem, full stop; and any table symmetric in both coordinates was printed from a cosine polynomial.
6. The reusable decision rule: convolution structure implies pass to the symbol; positivity of the form is nonnegativity of the symbol; strictness comes from the size of the symbol's zero set.

Mirror drills

1. Let $v(0) = 2$, $v(\pm 1) = -1$ and $v = 0$ elsewhere on $\mathbb{Z}$. Show $\sum_{s,s' \in S} v(s - s') > 0$ for every finite nonempty $S \subseteq \mathbb{Z}$, in two ways. *Answer:* the symbol is $2 - 2\cos s = 4\sin^2(s/2) \geq 0$, vanishing only at $s = 0$, so the argument above applies verbatim; combinatorially the sum is $2|S| - 2e$ with e the number of pairs of consecutive integers in S, and $e \leq |S| - 1$ since those pairs form a union of paths, giving at least 2.
2. Prove that $\mathcal{M}\left[e^{i(ns+mt)}\right] = 1$ if $n = m = 0$ and 0 otherwise, and deduce that $A(S)$ is the constant coefficient of $f(x,y)f(x^{-1},y^{-1})g(x,y)$. *Answer:* integrate in each variable separately using $\int_0^{2\pi} e^{ins} ds = 0$ for $n \neq 0$; the deduction is Step 1 read backwards.
3. Compute $A(S)$ for S the $n \times n$ box and say what the answer demands of any correct proof. *Answer:* the number of ordered pairs of difference (a,b) is $(n - |a|)(n - |b|)$, so $A(S) = \sum_{a,b} w(a,b)(n - |a|)(n - |b|) = 32n - 32$ for $n \geq 2$, using $\sum w = 0$, $\sum w(a,b)(|a| + |b|) = -32$ and $\sum w(a,b)|a||b| = -32$. Since $A(S)$ grows like the perimeter while $|S|$ grows like the area, any proof must be global, which rules out every local charging scheme.

13.7 B1. Telescoping a parity-keyed sign recursion

For positive integers n, let the numbers $c(n)$ be determined by the rules $c(1) = 1$, $c(2n) = c(n)$, and $c(2n+1) = (-1)^n c(n)$. Find the value of

$$\sum_{n=1}^{2013} c(n)\, c(n+2).$$

Read it out loud

Say what c is before you touch the sum. It is defined on positive integers by how the integer is built from a smaller one by doubling, so it is a recursion keyed on parity, not on subtraction by one. Two rules, one for each parity class of the argument: $c(2n) = c(n)$ carries no sign, while $c(2n+1) = (-1)^n c(n)$ injects a sign that depends on the smaller index n, not on the number $2n+1$ itself. From $c(1) = 1$ and two multiplications by ± 1 at every step you get $c(n) \in \{-1, 1\}$ for every n: this is a sign sequence, and the question asks for the sum of 2013 products of signs taken two apart.

The cue that should fire first is the shift. You are not asked for $\sum c(n)c(n+1)$, you are asked for $\sum c(n)c(n+2)$, and the recursion that defines c is itself about halving the argument. So ask what a shift of 2 in n does to the halved argument $\lfloor n/2 \rfloor$: it always adds exactly 1, whichever parity n has, because $\lfloor (n+2)/2 \rfloor = \lfloor n/2 \rfloor + 1$ for every integer n. A shift of 1 would not have this property, since it changes the halved index for one parity class and leaves it alone for the other. The problem has chosen the one shift value that lines the two parity classes up against each other. That is the guiding light: split the sum by the parity of the summation index n, not by unwinding each $c(n)$ down to $c(1)$ one integer at a time.

Within thirty seconds you should have written the two rules again, side by side, and asked what $c(2k)c(2k+2)$ and $c(2k+1)c(2k+3)$ become separately. Both reduce to a product $c(k)c(k+1)$, up to sign, and comparing those two signs is the entire problem.

Approaches, ranked worst to best

1. **Direct unwinding.** *Cue:* a finite, explicit recursion and a finite sum. *Circuit:* compute $c(n)$ for every n from 1 to 2015 by repeated halving, then add the 2013 products. (worst. It is correct and it terminates, but it treats 2013 as an instruction to do 2013 separate pieces of work rather than as a number whose parity structure you should exploit once. No step explains itself to the next step, so a single sign slip anywhere is invisible until the final total is already wrong. This is grinding, not reading.)

2. **Binary digit count.** *Cue:* $c(n)$ has a closed form from the binary expansion of n. *Circuit:* write n in binary as $b_k \cdots b_0$, use $c(n) = (-1)^{b_0 b_1 + b_1 b_2 + \cdots + b_{k-1} b_k}$, the parity of the count of adjacent 11 pairs, and analyse how adding 2 to n changes that count. (middle. The closed form is true and worth knowing, but adding 2 to n in binary can trigger a carry of unbounded length whenever the bit above the units digit already sits inside a run of ones, so comparing the 11-counts of n and $n+2$ needs a case split on how long that run is. It gets there, but it costs more ink than the recursion already hands you for free.)

3. **Pair by the recursion's own parity.** *Cue:* the recursion is stated separately for even and odd arguments. *Circuit:* split $\sum_n c(n)c(n+2)$ into its even-n and odd-n parts, reduce each part to $c(k)c(k+1)$ using the defining rules directly, and compare the two resulting signs. (best. No case split on binary digits, no carries: the two defining rules already do the case split for you, and each reduction is one line.)

The ideal solution

Every value of c is ± 1: $c(1) = 1$, and each defining rule multiplies a previously constructed value by ± 1, so by strong induction $c(n) \in \{-1, 1\}$ for all $n \geq 1$.

Separate the one term the pairing below cannot reach from the rest:

$$\sum_{n=1}^{2013} c(n)c(n+2) = c(1)c(3) + \sum_{n=2}^{2013} c(n)c(n+2).$$

For the remaining sum, group the index n into consecutive pairs $(2k, 2k+1)$ for $k = 1, \ldots, 1006$. These pairs run through $n = 2, 3, 4, \ldots, 2013$ exactly once each: $2 \cdot 1 = 2$ is the first even value in range, $2 \cdot 1006 + 1 = 2013$ is the last odd value in range, and the 1006 pairs cover the 2012 integers from 2 to 2013.

Fix $k \geq 1$ and reduce both halves of the pair using the two defining rules.

Even half: $c(2k) = c(k)$ and $c(2k+2) = c(2(k+1)) = c(k+1)$, so

$$c(2k)\, c(2k+2) = c(k)\, c(k+1).$$

Odd half: $c(2k+1) = (-1)^k c(k)$, and applying the odd rule again with $n = k+1$ gives $c(2k+3) = c(2(k+1)+1) = (-1)^{k+1} c(k+1)$, so

$$c(2k+1)\, c(2k+3) = (-1)^k (-1)^{k+1} c(k)c(k+1) = (-1)^{2k+1} c(k)c(k+1) = -c(k)c(k+1),$$

since $2k+1$ is odd. Adding the two halves,

$$c(2k)c(2k+2) + c(2k+1)c(2k+3) = c(k)c(k+1) - c(k)c(k+1) = 0$$

for every $k \geq 1$. Summing this identity over $k = 1, \ldots, 1006$ gives

$$\sum_{n=2}^{2013} c(n)c(n+2) = \sum_{k=1}^{1006} \Big(c(2k)c(2k+2) + c(2k+1)c(2k+3) \Big) = 0.$$

It remains to compute the boundary term. From the odd rule with $n = 1$, $c(3) = (-1)^1 c(1) = -1$, so $c(1)c(3) = 1 \cdot (-1) = -1$. Therefore

$$\sum_{n=1}^{2013} c(n)c(n+2) = -1 + 0 = \boxed{-1}.$$

Layer 1: WHAT

Block 1: recognise that c is built by a doubling recursion, so the natural way to organise any sum over c is by the same doubling structure, not by evaluating each term separately.

Block 2: notice the shift is 2, and that $\lfloor (n+2)/2 \rfloor = \lfloor n/2 \rfloor + 1$ for every n, so pairing the summation index by parity makes both halves of a pair reduce to the same smaller pair $(k, k+1)$.

Block 3: reduce the even half, $c(2k)c(2k+2)$, to $c(k)c(k+1)$ with no sign, directly from $c(2n) = c(n)$.

Block 4: reduce the odd half, $c(2k+1)c(2k+3)$, to $c(k)c(k+1)$ with a sign of $(-1)^{2k+1} = -1$, from $c(2n+1) = (-1)^n c(n)$ applied twice.

Block 5: compare the two reductions, same magnitude and opposite sign, so the pair sums to zero for every $k \geq 1$.

Block 6: sum the zero identity over all 1006 pairs, covering $n = 2$ to $n = 2013$, and separately evaluate the one term the pairing cannot reach, $n = 1$.

The nominal route computes $c(n)$ for $n = 1, \ldots, 2015$ one integer at a time, unwinding each down to $c(1)$ by repeated halving, then adds 2013 signed products. The bifurcation point is Block 1: the ideal route regroups the summation index by parity before evaluating a single $c(n)$, because the recursion defining c is itself parity-keyed, so any structure in a sum over c has to come from matching that same key. The zombie route evaluates first and looks for structure second, by which point every value is a bare number with no visible relationship to its neighbours.

Layer 2: HOW

Block 1. **Cue.** Two separate defining rules, one for even arguments and one for odd arguments. **Why natural.** A function defined by cases on parity has no reason to behave uniformly across an arbitrary sum. The only sums of it you can control are ones that respect the same parity split, so you regroup the index before computing anything. **Link to the next block.** Once you commit to a parity split of n, you need to know exactly which shift in n gives a clean, uniform shift in the halved index $\lfloor n/2 \rfloor$.

Block 2. **Cue.** The problem's shift is $c(n)c(n+2)$, exactly 2. **Why natural.** $\lfloor (n+2)/2 \rfloor = \lfloor n/2 \rfloor + 1$ holds for every n, of either parity, whereas a shift of 1 does not have this property. The shift 2 is the smallest value for which halving turns it into a uniform shift of 1, which is exactly what makes pairing $n = 2k$ with $n = 2k+1$ reduce both to the same smaller pair $(k, k+1)$. **Link to the next block.** Evaluate the even half of the pair first, since it uses the rule with no sign.

Block 3. **Cue.** $c(2k)$ and $c(2k+2)$ both have even argument. **Why natural.** $c(2n) = c(n)$ applies directly to both, with no sign to track, giving $c(2k)c(2k+2) = c(k)c(k+1)$ in one line. **Link to the next block.** The odd half uses the other rule, the one that carries a sign, so that sign is where the content of the problem is hiding.

Block 4. **Cue.** $c(2k+1)$ and $c(2k+3)$ both have odd argument, and the odd rule is $c(2n+1) = (-1)^n c(n)$. **Why natural.** Applying the odd rule once at $n = k$ and once at $n = k+1$ multiplies in $(-1)^k$ and $(-1)^{k+1}$, opposite signs for consecutive integers, so their product is always -1 regardless of k. The sign rule was written as $(-1)^n$, not as a fixed -1, precisely so that it flips between consecutive values of the reduced index. **Link to the next block.** Compare this signed reduction with the unsigned one from Block 3; they share the same magnitude $c(k)c(k+1)$.

Block 5. **Cue.** Even half equals $c(k)c(k+1)$, odd half equals $-c(k)c(k+1)$. **Why natural.** Two quantities that are negatives of each other sum to zero. This holds for every $k \geq 1$ with no exceptions, so it is safe to sum over all of them at once. **Link to the next block.** Sum this zero identity over the full range of k and account for exactly which values of n that range covers.

Block 6. **Cue.** The pairs $(2k, 2k+1)$ for $k = 1, \ldots, 1006$ cover $n = 2$ through $n = 2013$, leaving $n = 1$ untouched. **Why natural.** $n = 1$ has no partner, because the pairing needs both $2k$ and $2k+1$ inside the range, and $2k = 1$ has no integer solution. $n = 1$ sits outside the block structure, not inside it as a special case. **Link to the next block.** Evaluate $c(1)c(3)$ directly from the definition and add it to the remainder, now known to be zero.

Layer 3: WHY

The truth underneath the cancellation is that the recursion for c is first order in the halved index: everything about $c(n)$ is determined by $c(\lfloor n/2 \rfloor)$ and the parity of n. Consecutive integers $2k$ and $2k+1$ share the same halved index k, so any product of two values of c taken at arguments shifted by two collapses to a product at the halved indices, and the only thing that can differ between the even case and the odd case is a sign contributed by the odd rule. The shift of exactly 2 is what makes the halved indices of the two paired terms match (Layer 2, Block 2); this is not a coincidence of the problem, it is the reason the problem is solvable in closed form at all.

The discernment mechanism to install: whenever a sequence is defined by a recursion keyed on a modulus m (here $m = 2$, halving), and you must sum an expression built from that sequence at arguments shifted by a fixed d, check first whether d is a multiple of m. If it is, the shift descends to a uniform shift on the reduced index regardless of residue class, and pairing residues is likely to

telescope. If it is not, the reduced shift depends on the residue class itself, and you should expect casework rather than a clean pairing. Here $d = 2$ and $m = 2$, so the test passes immediately, which is exactly why Approach 3 costs so much less ink than Approach 2, an attempt to run the same idea in the wrong coordinate (binary digits of n directly, rather than the halved index).

What the examiners nudged toward: writing the recursion as two rules, one keyed on $2n$ and one on $2n + 1$, is itself an instruction to think in halved indices; asking for a shift of 2, rather than 1 or 3, is an instruction that this particular shift behaves uniformly. Every number the problem gives you is load-bearing.

More generally, for any integer $e \geq 1$ and $k \geq 1$, the same two reductions give $c(2k)c(2k + 2e) = c(k)c(k + e)$ and $c(2k + 1)c(2k + 2e + 1) = (-1)^e c(k)c(k + e)$, so the paired sum $c(2k)c(2k + 2e) + c(2k + 1)c(2k + 2e + 1)$ equals $c(k)c(k + e)\big(1 + (-1)^e\big)$: it is 0 whenever e is odd, exactly the cancellation above, and it is $2c(k)c(k + e)$ whenever e is even, which does not cancel but instead reduces a shift-$2e$ sum to twice a shift-e sum of the same type, a smaller instance of the same problem. The shift $d = 2$ actually asked for is the case $e = 1$, the smallest odd e, which is exactly why the problem telescopes to a single boundary term rather than to a smaller sum of the same shape.

Revolving concepts

Sequences defined by a doubling recursion. A sequence $a(n)$ with $a(2n) = g(n, a(n))$ and $a(2n + 1) = h(n, a(n))$ is determined entirely by the binary expansion of n read one digit at a time; such sequences are called 2-automatic, produced by a finite automaton reading the base-2 digits of n. A question about $a(n)$ for a specific large n, or a sum of a over a range, should be attacked by descending along this digit structure rather than by treating n as an opaque integer. See Book 1, field: Discrete structures and graphs, topic: automatic sequences.

The Rudin and Shapiro sequence. The sequence here, with $c(1) = 1$, $c(2n) = c(n)$, $c(2n + 1) = (-1)^n c(n)$, is exactly the Rudin and Shapiro sequence, with closed form $c(n) = (-1)^{b_0 b_1 + b_1 b_2 + \cdots + b_{k-1} b_k}$ for binary expansion $b_k \cdots b_0$: the exponent counts adjacent pairs of 1's. The sequence is known for partial sums $\left|\sum_{n \leq N} c(n)\right|$ that grow only like $\sqrt{N}$, far slower than a typical ± 1 sequence, and for producing ± 1 coefficient polynomials with the smallest possible maximum modulus on the unit circle. See Book 1, field: Number theory, topic: digit sum sequences.

Telescoping by index-shift pairing. Grouping a sum's index into blocks whose contribution collapses is the method of differences applied to a pair rather than a single term. The generalisation above, where an even e doubles rather than cancels, shows the technique is really a reduction, of which cancellation is the special case e odd. See Book 1, field: Sequences and recurrences, topic: telescoping sums.

Circuit Bank

- **WHEN** a sequence is defined by separate rules for even and odd arguments **DO** regroup any sum over that sequence by the parity of the summation index before evaluating a single term **BECAUSE** the sequence's own structure is keyed on parity, so only a parity-respecting regrouping can expose it
- **WHEN** you must sum a shifted product $a(n)a(n + d)$ of a base-m-keyed recursion **DO** check whether m divides d before choosing a method **BECAUSE** when it does, the shift descends to a uniform shift on the reduced index for every residue class, which is what makes pairing and telescoping possible
- **WHEN** a pairing argument covers a range of the summation index but not all of it **DO** isolate and separately evaluate every boundary term the pairing cannot reach, and state explicitly why it cannot reach them **BECAUSE** silently absorbing a boundary term into a cancellation

you only proved for the interior is the most common way to get the right method and the wrong final number

- **WHEN** a recursive sign rule is written as $(-1)^n$ rather than a fixed constant **DO** treat the dependence on n as load-bearing and check what happens to it at consecutive values of n **BECAUSE** a sign that flips with n is what produces cancellation across consecutive terms, a constant sign would not

Backtesting

1. First useful cue: the recursion is stated as two separate rules keyed on the parity of the argument, so any sum over c should be regrouped by parity before evaluation.
2. Method the expression suggested: pair the summation index as $(2k, 2k+1)$, because the shift of 2 descends to a uniform shift of 1 on the halved index for both parities.
3. Where you would hesitate, and the thought to reinforce: at the range bookkeeping, whether the pairing covers $n = 1$. Reinforce by checking explicitly which n would require $2k = 1$ or $2k + 1 = 1$ for an integer $k \geq 1$, and noticing neither has a solution, so $n = 1$ sits outside the paired block by construction.
4. The trap the question sets: assuming the whole sum telescopes to zero and forgetting the one boundary term, worth -1, which is the entire final answer.
5. What makes you recognise it faster next time: seeing a recursion keyed on doubling together with a summation shift that is itself even. Test immediately whether the shift is a multiple of the recursion's modulus.
6. The reusable decision rule: when a sum's shift is a multiple of a recursion's base, pair residue classes of the index; the two rules of the recursion applied to consecutive reduced indices differ by at most a controllable sign.

Mirror drills

1. Show that $\sum_{n=1}^{N} c(n)c(n+2) = -1$ for every odd $N \geq 3$, not only $N = 2013$. *Answer:* the same pairing covers $n = 2, \ldots, N$ in blocks of two whenever N is odd, so the interior always cancels and only $c(1)c(3) = -1$ survives, independent of N.
2. Using the generalised identity for shift $2e$, reduce $\sum_{n=1}^{2013} c(n)c(n+4)$ to a sum of shift-2 type. *Hint:* here $e = 2$ is even, so each interior pair contributes $2c(k)c(k+2)$ rather than 0; separate the boundary terms $n = 1, 2, 3$ by hand and rewrite the rest as twice a shift-2 sum, which you already know how to telescope.
3. Define $d(1) = 1$, $d(2n) = d(n)$, $d(2n+1) = -d(n)$, a constant sign rather than $(-1)^n$, and compute $\sum_{n=1}^{7} d(n)d(n+2)$. *Answer:* -3. Here the odd rule applied twice gives $(-d(k))(-d(k+1)) = +d(k)d(k+1)$, the same sign as the even half, so interior pairs now double instead of cancel. This is the precise reason the original problem needed a sign rule depending on n, not a fixed -1.

13.8 B2. Maximising a cosine polynomial at zero

Let $C = \bigcup_{N=1}^{\infty} C_N$, where C_N denotes the set of those cosine polynomials of the form

$$f(x) = 1 + \sum_{n=1}^{N} a_n \cos(2\pi n x)$$

for which:

(i) $f(x) \geq 0$ *for all real* x*, and*

(ii) $a_n = 0$ *whenever* n *is a multiple of* 3*.*

Determine the maximum value of $f(0)$ *as* f *ranges through* C*, and prove that this maximum is attained.*

Read it out loud

Notice first that the quantity you are asked to maximise is not hidden: $f(0) = 1 + \sum_{n=1}^{N} a_n \cos(0) = 1 + \sum_{n=1}^{N} a_n$, so the whole problem is really how large the coefficients of a nonnegative cosine polynomial can sum to, given that every third coefficient is forced to be zero. Two hypotheses are in tension. Hypothesis (i) is a global statement, true at every real x at once, and on its own it is hard to use directly, since you cannot check infinitely many points one by one. Hypothesis (ii) is a local, arithmetic statement about the index n, tied to the number 3.

The cue that should fire is exactly that number 3. A cosine polynomial is a finite sum of functions $\cos(2\pi n x)$, each periodic in x with period $1/n$, and evaluating the whole sum at a rational point $x = j/m$ makes every term's value depend only on n mod m, because $\cos(2\pi n \cdot j/m)$ is unchanged when n is shifted by a multiple of m. Hypothesis (ii) hands you the modulus $m = 3$ for free. So the global hypothesis (i), unusable as stated, becomes usable the instant you specialise it to $x = 1/3$: it turns into a single numeric inequality involving only $\sum a_n$, which is exactly the quantity $f(0)$ needs.

Within thirty seconds you should have written $f(1/3) \geq 0$ and started computing $\cos(2\pi n/3)$ for the values of n that are still allowed to be nonzero.

Approaches, ranked worst to best

1. **Continuous optimisation over all coefficients.** *Cue:* maximise a linear function subject to $f \geq 0$ everywhere. *Circuit:* treat $f \geq 0$ for all x as an infinite family of linear constraints on $(a_1, \ldots, a_N)$ and reach for Lagrange multipliers or duality. (worst. This is the technically general method and it is not wrong, but the constraint set is infinite dimensional, one inequality per real x, and N itself is free. Nothing in the raw optimisation tells you to look at $x = 1/3$ specifically, so you would have to rediscover the binding constraint before the machinery does any work, which means the machinery is not actually doing the work.)

2. **Trial and error on small N.** *Cue:* the extremal example, once seen, is small. *Circuit:* try $N = 1$, find $f(0) \leq 2$ is forced, push to $N = 2$ and search numerically for coefficients that do better. (middle. This can stumble onto the right construction by luck, but it produces no proof that 3 is an upper bound for every N, only a candidate value. A full solution still needs the argument below to close the gap between having found something that gives 3 and knowing that nothing gives more.)

3. **Evaluate at the point the hypothesis names.** *Cue:* the vanishing condition is stated modulo 3. *Circuit:* compute $f(1/3)$ directly, using that $\cos(2\pi n/3)$ takes only one value among the surviving coefficients. (best. One evaluation converts an infinite dimensional hypothesis into a single linear inequality in $\sum a_n$, which is exactly $f(0) - 1$.)

The ideal solution

Upper bound. Fix $f \in C_N \subset C$. For any n with $3 \nmid n$, the residue of n modulo 3 is 1 or 2, and

$$\cos\left(\tfrac{2\pi n}{3}\right) = \cos\left(\tfrac{2\pi}{3}\right) = \cos\left(\tfrac{4\pi}{3}\right) = -\tfrac{1}{2},$$

since $\cos(4\pi/3) = \cos(2\pi - 2\pi/3) = \cos(2\pi/3)$. So every n with $3 \nmid n$ contributes $a_n \cos(2\pi n/3) = -\frac{1}{2}a_n$, while every n with $3 \mid n$ has $a_n = 0$ by hypothesis (ii) and so contributes 0 regardless of its cosine value. Hence

$$f\left(\tfrac{1}{3}\right) = 1 + \sum_{n=1}^{N} a_n \cos\left(\tfrac{2\pi n}{3}\right) = 1 - \tfrac{1}{2}\sum_{n=1}^{N} a_n,$$

where the sum on the right is the full sum $\sum_{n=1}^{N} a_n$, since the terms with $3 \mid n$ vanish in it too. By hypothesis (i), $f(1/3) \geq 0$, so $\sum_{n=1}^{N} a_n \leq 2$. Therefore

$$f(0) = 1 + \sum_{n=1}^{N} a_n \leq 3.$$

This holds for every $f \in C$, so 3 is an upper bound for $f(0)$ on all of C.

A tight example. Equality above forces $f(1/3) = 0$. Since $f \geq 0$ on all of $\mathbb{R}$ and f is smooth, a zero at the interior point $x = 1/3$ is a global minimum, so it is not a simple sign-changing zero: nonnegativity forbids that. The natural object to try is a perfect square vanishing exactly where $\cos(2\pi x) = -\frac{1}{2}$, namely a multiple of $\big(2\cos(2\pi x) + 1\big)^2$. Set

$$f(x) := \tfrac{1}{3}\big(2\cos(2\pi x) + 1\big)^2,$$

manifestly nonnegative for every x, since it is $\frac{1}{3}$ times a real square. Expand it with the double angle identity $\cos(4\pi x) = 2\cos^2(2\pi x) - 1$:

$$\big(2\cos(2\pi x) + 1\big)^2 = 4\cos^2(2\pi x) + 4\cos(2\pi x) + 1 = 2\big(2\cos^2(2\pi x) - 1\big) + 4\cos(2\pi x) + 3 = 2\cos(4\pi x) + 4\cos(2$$

so

$$f(x) = 1 + \tfrac{4}{3}\cos(2\pi x) + \tfrac{2}{3}\cos(4\pi x).$$

This is a cosine polynomial with $N = 2$, $a_1 = \frac{4}{3}$, $a_2 = \frac{2}{3}$. Neither index 1 nor 2 is a multiple of 3, so hypothesis (ii) holds vacuously, and hypothesis (i) holds because f is a nonnegative multiple of a square. Hence $f \in C_2 \subset C$, and

$$f(0) = 1 + \tfrac{4}{3} + \tfrac{2}{3} = 3.$$

The upper bound is attained, so the maximum value of $f(0)$ on C is $\boxed{3}$.

Layer 1: WHAT

Block 1: read $f(0) = 1 + \sum a_n$; the target of the problem is literally the sum of the coefficients.

Block 2: read hypothesis (ii) as an instruction of where to evaluate, namely $x = 1/3$, since it names the modulus 3.

Block 3: compute $\cos(2\pi n/3) = -\frac{1}{2}$ for every surviving n, so $f(1/3)$ collapses to a single linear expression in $\sum a_n$.

Block 4: use $f(1/3) \geq 0$ to get $\sum a_n \leq 2$ and hence $f(0) \leq 3$.

Block 5: read the equality condition, $f(1/3) = 0$, as forcing a global minimum of a nonnegative function, which motivates a perfect square ansatz vanishing there.

Block 6: build $\frac{1}{3}(2\cos 2\pi x + 1)^2$, expand via the double angle identity, read off a_1, a_2, check both hypotheses explicitly, and confirm $f(0) = 3$.

The nominal route sets up a Lagrangian or a blind coefficient search over an unspecified N, treating the divisibility condition as bookkeeping to carry through the optimisation rather than as an instruction. The bifurcation point is Block 2: the ideal route reads the number 3 in the hypothesis as telling you where to plug in, the zombie route treats it as just an algebraic side condition on the index set.

Layer 2: HOW

Block 1. **Cue.** $f(0)$ is asked for, and $\cos(0) = 1$. **Why natural.** Every $\cos(2\pi n \cdot 0)$ term equals 1, so the target collapses to 1 plus the coefficient sum with no cosine values to compute at all. Knowing exactly what you are chasing before you touch the constraint tells you what form the constraint needs to be squeezed into. **Link to the next block.** You need a way to bound $\sum a_n$ using hypothesis (i), which as stated says nothing about that sum directly.

Block 2. **Cue.** The vanishing condition is stated as a divisibility by 3. **Why natural.** A coefficient constraint tied to a modulus m is an invitation to evaluate the cosine polynomial at $x = j/m$, since that is exactly the substitution that makes every surviving term's phase depend only on n mod m, collapsing a sum over many different cosines into a sum over few. **Link to the next block.** Compute what that substitution actually does to the terms that are still allowed to be nonzero.

Block 3. **Cue.** The surviving residues modulo 3 are 1 and 2, and $\cos(2\pi/3) = \cos(4\pi/3) = -\frac{1}{2}$. **Why natural.** Cosine is even and 2π-periodic, so residues r and $m - r$ modulo m always give equal values at $x = 1/m$; with $m = 3$ the only two nonzero residues, 1 and 2, are exactly this pair, so every surviving coefficient is multiplied by the same number. **Link to the next block.** A single evaluation now gives a genuine linear inequality in $\sum a_n$, which you can combine with nonnegativity.

Block 4. **Cue.** $f(1/3) \geq 0$ by hypothesis (i). **Why natural.** This is the only place hypothesis (i) is used in the bound: one point suffices, because the earlier step arranged for that one point to see every surviving coefficient at once. **Link to the next block.** The bound $f(0) \leq 3$ is now proved; existence of an example achieving it is a separate task.

Block 5. **Cue.** Equality in $f(0) \leq 3$ forces $f(1/3) = 0$ exactly, together with $f \geq 0$ everywhere. **Why natural.** A nonnegative smooth function that touches zero at an interior point has that point as a global minimum, hence a critical point, and cannot cross through zero there; the natural way to build such a function is as a square, or a positive multiple of one, vanishing at the required point. **Link to the next block.** Choose the simplest expression, linear in $\cos(2\pi x)$, that vanishes at $\cos(2\pi x) = -\frac{1}{2}$, and square it.

Block 6. **Cue.** $2\cos(2\pi x) + 1$ vanishes exactly when $\cos(2\pi x) = -\frac{1}{2}$. **Why natural.** Squaring it and rescaling by $\frac{1}{3}$ so the constant term is 1 turns the candidate into an actual member of C; expanding with the double angle identity is the only way to see its coefficients a_1, a_2 explicitly and check both hypotheses against them. **Link to the next block.** Nothing remains: every hypothesis has been checked and the target value computed directly.

Layer 3: WHY

The mechanism beneath the bound is a roots-of-unity style collapse: evaluating a cosine polynomial at a rational point $x = j/m$ replaces $2\pi n x$ by an angle depending only on n mod m. In general this splits the surviving coefficients into as many distinct cosine values as there are residue classes paired by $r \leftrightarrow m - r$ that are not otherwise excluded. What makes $m = 3$ special is that it is the

smallest modulus with exactly one such pair covering every nonzero residue: 1 and $2 = 3 - 1$. A modulus like 4 does not have this property, since its nonzero residues split into the pair $\{1, 3\}$, giving $\cos(\pi/2) = 0$, and the self-paired residue $\{2\}$, giving $\cos(\pi) = -1$, two different values rather than one, so a single evaluation would not collapse everything there. The problem's choice of divisibility by 3 is exactly what makes one evaluation enough.

The construction step rests on a deeper fact worth naming rather than treating as luck: the Fejér and Riesz theorem states that every trigonometric polynomial nonnegative on the whole real line is $|P(e^{2\pi ix})|^2$ for some polynomial P, so restricting a search for an extremal nonnegative cosine polynomial to perfect squares is not a guess, it is a complete parametrisation. Once you also know the extremiser must vanish at $x = 1/3$, the square is forced up to scale, since the only degree-one factor of $\cos(2\pi x)$ vanishing there is a multiple of $2\cos(2\pi x) + 1$.

The discernment mechanism: whenever a family of Fourier or cosine coefficients is constrained to vanish on an arithmetic progression modulo m, and you are asked to bound a value of the function at a specific point, test evaluation at rational points with denominator m before attempting a general optimisation. Separately, whenever a linear bound derived from one inequality is tight, look for the extremal object as a square (or sum of squares) forced to vanish exactly where the bound was evaluated, since Fejér and Riesz guarantees this search is exhaustive for nonnegative trigonometric polynomials.

What the examiners nudged toward: writing the coefficient condition as divisibility by 3, rather than simply "$a_1 = a_2 = 0$" or some other finite list, is what points at a modulus to substitute rather than a fixed set of indices to avoid; asking for $f(0)$, the value at the one point where every cosine is 1, sets up the exact mirror of the evaluation at $x = 1/3$, where every surviving cosine is a single other value.

Revolving concepts

The Fejér and Riesz theorem. Every trigonometric polynomial $g(\theta) = \sum_{k=-N}^{N} c_k e^{ik\theta}$ that is nonnegative for all real θ can be written as $|P(e^{i\theta})|^2$ for a polynomial P of degree N. Consequently every nonnegative cosine polynomial is a positive multiple of a perfect square of a lower cosine polynomial, which is why the construction in this problem is a systematic search, not a guess. See Book 1, field: Real analysis, topic: nonnegative trigonometric polynomials.

Collapsing a sum by evaluating at a rational point. For a cosine polynomial, evaluating at $x = j/m$ makes every term depend only on n mod m, and $\cos(2\pi r/m) = \cos(2\pi(m - r)/m)$ for every residue r, so residues r and $m - r$ always coincide in value. This pairing covers every nonzero residue with a single value only for small m, notably $m = 3$; for larger m several distinct values survive and a single point evaluation is not enough. See Book 1, field: Number theory, topic: roots of unity filters.

Chebyshev polynomials. $\cos(n\theta) = T_n(\cos\theta)$ for the Chebyshev polynomial T_n of the first kind, with $T_2(t) = 2t^2 - 1$. This is the identity used to expand $(2\cos(2\pi x) + 1)^2$ back into a genuine cosine polynomial in x, and it generalises: T_n satisfies $T_{n+1}(t) = 2tT_n(t) - T_{n-1}(t)$ with $T_0 = 1$, $T_1 = t$, the same recursion that $\cos(n\theta)$ obeys as a function of n for fixed θ. See Book 1, field: Polynomials, topic: Chebyshev polynomials.

The Fejér kernel and unbounded coefficient sums. $F_N(y) = 1 + \sum_{j=1}^{N} 2\left(1 - \frac{j}{N+1}\right)\cos(2\pi jy)$ is nonnegative for every y and has $F_N(0) = N + 1$, unbounded as $N \to \infty$. Substituting $y = 3x$ turns F_N into a nonnegative cosine polynomial supported only on multiples of 3, showing that the reversed divisibility condition, requiring $a_n = 0$ *unless* $3 \mid n$, has no finite maximum at all. This is why excluding multiples of 3, rather than restricting to them, is essential to the finite answer 3. See Book 1, field: Real analysis, topic: summability kernels.

Circuit Bank

- **WHEN** a Fourier or cosine coefficient constraint vanishes on an arithmetic progression modulo m **DO** evaluate the function at $x = j/m$ for a small j **BECAUSE** the surviving cosines collapse to values depending only on n mod m, turning a global nonnegativity constraint into a single numeric inequality
- **WHEN** a linear bound obtained from one evaluation is tight **DO** search for the extremal object as a perfect square, or a positive multiple of one, vanishing exactly at the point the bound used **BECAUSE** the Fejér and Riesz theorem guarantees every nonnegative trigonometric polynomial has such a squared form, so the search is complete, not a guess
- **WHEN** you need to convert between $\cos(n\theta)$ and powers of $\cos\theta$ **DO** use the double angle or Chebyshev identity $T_n(\cos\theta) = \cos(n\theta)$ directly **BECAUSE** it is faster than expanding through complex exponentials and it keeps every quantity manifestly real

Backtesting

1. First useful cue: hypothesis (ii) names a modulus, 3, which is an instruction to evaluate the cosine polynomial at $x = 1/3$.
2. Method the expression suggested: a single evaluation collapses every surviving coefficient's cosine to $-\frac{1}{2}$, turning the global constraint $f \geq 0$ into the single inequality $\sum a_n \leq 2$.
3. Where you would hesitate, and the thought to reinforce: whether the constructed extremal example genuinely satisfies hypothesis (ii). Reinforce by checking the indices actually used, here only 1 and 2, against the divisibility condition explicitly, rather than assuming a square automatically qualifies.
4. The trap the question sets: attempting a full Lagrangian optimisation over an unspecified N with infinitely many pointwise constraints, which is technically sound but never terminates without first finding the binding point.
5. What makes you recognise it faster next time: a coefficient hypothesis phrased as divisibility by a fixed modulus, paired with a nonnegativity hypothesis; test rational evaluation points with that modulus as denominator before anything else.
6. The reusable decision rule: once a linear bound is proved tight at a specific point, build the extremal example as a square vanishing there rather than searching coefficient space unguided.

Mirror drills

1. Cosine polynomials with $a_n = 0$ whenever n is even, otherwise unconstrained, with $f \geq 0$; find the maximum of $f(0)$. *Answer:* evaluate at $x = 1/2$, where $\cos(n\pi) = -1$ for every surviving odd n, giving $f(0) \leq 2$; attained by $f(x) = 1 + \cos(2\pi x)$, already nonnegative since $\cos \geq -1$.
2. Same setup as the original problem but with $a_n = 0$ required whenever $4 \mid n$ instead of $3 \mid n$. Why does evaluating at $x = 1/4$ fail to reproduce the same clean bound? *Hint:* the surviving residues modulo 4 are $1, 2, 3$, and $\cos(2\pi n/4)$ equals 0 for $n \equiv 1, 3$ but -1 for $n \equiv 2$, two distinct values rather than one, so a single evaluation only bounds the coefficients with $n \equiv 2 \pmod 4$, not the full sum.
3. Suppose instead $a_n = 0$ is required *unless* $3 \mid n$ (the reversed condition). Show $f(0)$ has no finite maximum on the resulting family. *Hint:* substitute $x \mapsto 3x$ into the Fejér kernel F_N;

the result is nonnegative, supported only on multiples of 3, and has value $N+1$ at $x=0$, unbounded as $N \to \infty$.

13.9 B3. Additivity of a modular set function

Let $\mathcal{P}$ be a nonempty collection of subsets of $\{1, \ldots, n\}$ such that:

(i) if $S, S' \in \mathcal{P}$, then $S \cup S' \in \mathcal{P}$ and $S \cap S' \in \mathcal{P}$, and

(ii) if $S \in \mathcal{P}$ and $S \neq \emptyset$, then there is a subset $T \subset S$ such that $T \in \mathcal{P}$ and T contains exactly one fewer element than S.

Suppose that $f : \mathcal{P} \to \mathbb{R}$ is a function such that $f(\emptyset) = 0$ and

$$f(S \cup S') = f(S) + f(S') - f(S \cap S') \text{ for all } S, S' \in \mathcal{P}.$$

Must there exist real numbers $f_1, \ldots, f_n$ such that

$$f(S) = \sum_{i \in S} f_i$$

for every $S \in \mathcal{P}$?

Read it out loud

Name the structure before you name a strategy. A collection of subsets closed under union and intersection is closed under both lattice operations of the Boolean lattice at once, so $(\mathcal{P}, \subseteq)$ is a lattice in its own right, sitting inside $2^{\{1,\ldots,n\}}$. The identity $f(S \cup S') = f(S) + f(S') - f(S \cap S')$ is not an arbitrary equation, it is inclusion and exclusion for two sets, imposed as an axiom on an abstract real valued f rather than derived from f being a size or a measure. A function on a lattice satisfying exactly this identity has a name, a valuation, and the question is asking whether every valuation on this kind of lattice is automatically the simplest possible kind: a sum of weights attached to single elements.

Hypothesis (ii) is the second cue, and it is doing different work from (i). It says that below every nonempty set of the family sits another set of the family exactly one element smaller. That is a graded, covering condition: it lets you strip one element at a time and never leave $\mathcal{P}$, all the way down. Whenever a hypothesis hands you a covering structure like this, the shape of the proof is almost forced: define something using one covering step, check it does not depend on which step you used, then induct along the covering structure.

Within thirty seconds you should have written: for an element i that appears in some set of $\mathcal{P}$, look at a pair $S = T \sqcup \{i\}$ with both $S, T \in \mathcal{P}$, call $f(S) - f(T)$ the weight of i, and worry immediately about whether two different such pairs could disagree.

Approaches, ranked worst to best

1. **Assume singletons are available.** *Cue:* the target formula $f(S) = \sum_{i \in S} f_i$ looks like it wants $f_i = f(\{i\})$. *Circuit:* set $f_i := f(\{i\})$ whenever $\{i\} \in \mathcal{P}$ and try to build S up from singletons using the modularity identity. (worst, and not merely inelegant: it is not even well posed in general. Hypothesis (ii) guarantees a subset one element smaller than any nonempty S, not that every element has its own singleton in $\mathcal{P}$; a family can satisfy (i) and (ii) while never containing a single one-element set. Defining f_i this way can leave some elements with no definition at all.)

2. **Explore special cases.** *Cue:* small, symmetric families are easy to compute by hand. *Circuit:* check the claim when $\mathcal{P}$ is the full power set, or when $\mathcal{P}$ is a chain totally ordered by inclusion. (middle. This is genuinely useful for building intuition and for discovering the right definition of f_i, and it is where you should start scratch work. But a chain or a power set both happen to contain every singleton, so success there proves nothing about a general $\mathcal{P}$; stopping here leaves the real difficulty, families without singletons, completely untouched.)

3. **Define an invariant intrinsically and induct along the covering structure.** *Cue:* the two hypotheses are exactly a lattice identity and a graded covering condition. *Circuit:* define the weight of i as a difference across any covering step containing it, prove that difference is independent of the step using the modularity identity itself, then induct on $|S|$ using hypothesis (ii) for both the base case and the inductive step. (best. Every hypothesis is used exactly where it is needed, nothing about singletons is assumed, and the argument is the same length regardless of how far from a chain or a power set $\mathcal{P}$ happens to be.)

The ideal solution

Lemma 13.9.1. *If some $S \in \mathcal{P}$ contains i, then there exist $S, T \in \mathcal{P}$ with $S = T \sqcup \{i\}$, that is S is T together with i, disjointly.*

Proof. Among all sets in $\mathcal{P}$ containing i, choose S of minimum cardinality; such S exists since the family of sets in $\mathcal{P}$ containing i is nonempty and finite. Since $S \neq \emptyset$, hypothesis (ii) gives $T \subset S$ with $T \in \mathcal{P}$ and $|T| = |S| - 1$, so $T = S \setminus \{x\}$ for the unique element $x \in S \setminus T$. If $x \neq i$ then $i \in T$, and $T \in \mathcal{P}$ has smaller cardinality than S while still containing i, contradicting the choice of S. So $x = i$, that is $T = S \setminus \{i\}$, as claimed. □

Lemma 13.9.2. *Suppose $S_1 = T_1 \sqcup \{i\}$ and $S_2 = T_2 \sqcup \{i\}$ with $S_1, S_2, T_1, T_2 \in \mathcal{P}$. Then*

$$f(S_1) - f(T_1) = f(S_2) - f(T_2).$$

Proof. Since $i \notin T_2$, distributing intersection over union gives $S_1 \cap T_2 = (T_1 \cup \{i\}) \cap T_2 = T_1 \cap T_2$, and $S_1 \cup T_2 = T_1 \cup T_2 \cup \{i\}$. Apply the hypothesis to the pair $(S_1, T_2) \in \mathcal{P} \times \mathcal{P}$:

$$f(T_1 \cup T_2 \cup \{i\}) = f(S_1 \cup T_2) = f(S_1) + f(T_2) - f(S_1 \cap T_2) = f(S_1) + f(T_2) - f(T_1 \cap T_2).$$

Symmetrically, since $i \notin T_1$, $T_1 \cap S_2 = T_1 \cap T_2$ and $T_1 \cup S_2 = T_1 \cup T_2 \cup \{i\}$, the same set as before, and applying the hypothesis to the pair (T_1, S_2):

$$f(T_1 \cup T_2 \cup \{i\}) = f(T_1 \cup S_2) = f(T_1) + f(S_2) - f(T_1 \cap S_2) = f(T_1) + f(S_2) - f(T_1 \cap T_2).$$

The left sides agree, so $f(S_1) + f(T_2) = f(T_1) + f(S_2)$, which rearranges to the claim. □

Define $f_1, \ldots, f_n$ as follows. If no $S \in \mathcal{P}$ contains i, set $f_i = 0$. Otherwise, by Lemma 1 choose any $S, T \in \mathcal{P}$ with $S = T \sqcup \{i\}$, and set $f_i := f(S) - f(T)$; by Lemma 2 this value does not depend on which such pair was chosen.

Claim 13.9.3. *$f(S) = \sum_{i \in S} f_i$ for every $S \in \mathcal{P}$.*

Proof. First, $\emptyset \in \mathcal{P}$: since $\mathcal{P}$ is nonempty, take any $S_0 \in \mathcal{P}$ and, if it is nonempty, apply hypothesis (ii) repeatedly to descend through sets of $\mathcal{P}$ of strictly decreasing size, necessarily reaching $\emptyset \in \mathcal{P}$ after finitely many steps. The hypothesis $f(\emptyset) = 0$ is therefore meaningful, and it is the base case of an induction on $|S|$ for $S \in \mathcal{P}$: the empty sum on the right agrees with $f(\emptyset) = 0$.

For the inductive step, let $S \in \mathcal{P}$ be nonempty and suppose the claim holds for every set of $\mathcal{P}$ with fewer elements than S. By hypothesis (ii), there is $T \in \mathcal{P}$, $T \subset S$, with $|T| = |S| - 1$; write $S = T \sqcup \{j\}$ for the unique $j \in S \setminus T$. This is exactly a covering pair isolating j, so by the definition of f_j, applied to this particular pair,

$$f(S) - f(T) = f_j.$$

By the induction hypothesis, $f(T) = \sum_{i \in T} f_i$. Adding,

$$f(S) = f(T) + f_j = \sum_{i \in T} f_i + f_j = \sum_{i \in T \cup \{j\}} f_i = \sum_{i \in S} f_i,$$

since T and $\{j\}$ are disjoint with union S. This completes the induction. □

So yes: such real numbers $f_1, \ldots, f_n$ always exist.

Layer 1: WHAT

Block 1: name the two hypotheses' separate jobs. Closure under union and intersection makes the modularity identity a genuine, checkable algebraic fact about $\mathcal{P}$; the one-smaller subset condition is a covering condition letting you strip elements one at a time and stay inside $\mathcal{P}$.

Block 2: since the target is additivity over single elements, define, for each element i that occurs somewhere in $\mathcal{P}$, its own marginal contribution as the jump in f across a pair that adds exactly i. Hypothesis (ii) is exactly what produces such a pair.

Block 3: worry immediately that different covering pairs for the same i might give different jumps, and resolve this by feeding two candidate pairs into the modularity identity at once, using closure under union and intersection to force the two resulting expressions to agree.

Block 4: with every f_i pinned down, or set to zero, prove $f(S) = \sum_{i \in S} f_i$ by strong induction on $|S|$, with hypothesis (ii) supplying both the base case, that $\emptyset \in \mathcal{P}$, and the inductive step, a size-one-smaller predecessor of S inside $\mathcal{P}$.

Block 5: conclude that such $f_1, \ldots, f_n$ always exist.

The nominal route tries to define f_i as $f(\{i\})$ and build every set up from singletons using the modularity identity, or attempts an induction over an arbitrary enumeration of $\mathcal{P}$ not respecting inclusion. The bifurcation point is Block 2: the moment you ask what quantity, defined intrinsically from f and not assuming any singleton lies in $\mathcal{P}$, deserves to be called f_i. Recognising that a difference across a covering step, not a value at a literal singleton, is the right invariant is the fork between an argument that stalls the moment $\mathcal{P}$ has no singletons and one that never needed them.

Layer 2: HOW

Block 1. **Cue.** A family of subsets closed under both union and intersection, together with a graded covering condition. **Why natural.** Closure under the two lattice operations is precisely what makes $(\mathcal{P}, \subseteq)$ a lattice, and a graded covering condition is precisely what a rank-based induction needs; naming the structure up front tells you which two tools to reach for before you have written a single line of proof. **Link to the next block.** The target claim, additivity over elements, asks for a per-element quantity, so the next question is how to define one without assuming more than the hypotheses give you.

Block 2. **Cue.** The target formula sums over $i \in S$, so it wants a value f_i attached to each element. **Why natural.** The only elements the hypotheses let you isolate are ones for which hypothesis (ii) supplies a covering pair; that pair is a legitimate way to isolate i's contribution even when $\{i\}$ itself is nowhere in $\mathcal{P}$. **Link to the next block.** A quantity defined by an arbitrary choice is not yet well defined; the choice of covering pair must be shown not to matter.

Block 3. **Cue.** Two different covering pairs, $S_1 = T_1 \sqcup \{i\}$ and $S_2 = T_2 \sqcup \{i\}$, both isolating the same element i. **Why natural.** Apply the governing identity to (S_1, T_2) and separately to (T_1, S_2): because i sits outside both T_1 and T_2, both applications compute f of the very same set, $T_1 \cup T_2 \cup \{i\}$, on the union side, forcing the two right-hand expressions to be equal. This is not a special property of the particular pairs chosen, it works for any two covering pairs at once. **Link to the next block.** With f_i now well defined for every occurring element, the additivity claim becomes a statement you can try to prove by induction on the size of S.

Block 4. **Cue.** The target is a statement for every $S \in \mathcal{P}$, and hypothesis (ii) supplies a smaller set inside $\mathcal{P}$ for every nonempty one. **Why natural.** Strong induction on $|S|$ is the natural vehicle whenever a hypothesis is phrased in terms of cardinality directly, as (ii) is; the covering pair used to descend from S to a smaller set is exactly a covering pair isolating one element, so it plugs directly

into the definition of f_j from Block 3. **Link to the next block.** The only remaining gap is the base case, since $f(\emptyset) = 0$ is only useful if $\emptyset \in \mathcal{P}$ in the first place.

Block 5. Cue. $\mathcal{P}$ is nonempty, but nothing yet says $\emptyset$ is one of its members. **Why natural.** Hypothesis (ii), applied repeatedly starting from any element of $\mathcal{P}$, strictly decreases cardinality at each step and must terminate, and the only place it can terminate is at $\emptyset$; this is the same covering hypothesis doing a second job, securing the floor of the induction rather than a step of it. **Link to the next block.** The induction is now complete in both its base case and its step, so the claim holds for every $S \in \mathcal{P}$.

Layer 3: WHY

The truth underneath the proof is that $\mathcal{P}$, closed under union and intersection, is a finite distributive lattice, and by Birkhoff's representation theorem every finite distributive lattice arises exactly this way, as a ring of subsets of some ground set. The hypothesis on f is the definition of a valuation on that lattice. Hypothesis (ii) additionally makes the lattice graded by cardinality with every element covering another, and a valuation on a graded lattice of this kind is reconstructible from its jumps across covering pairs in exactly the way a function with path-independent increments along the edges of a graph is reconstructible from a single value plus those increments: you need the increments to be well defined regardless of path, which is Lemma 2, and you need every point to be reachable from the base point by a path of covering steps, which is hypothesis (ii). The word valuation is not decoration: it tells you the exact theorem this problem is a disguised instance of, and it tells you where to look if a harder version of this problem ever drops hypothesis (ii) and asks what survives.

The discernment mechanism: whenever a hypothesis gives you closure under two lattice-like operations together with an inclusion-exclusion identity, and a second hypothesis gives you a graded, one-step-at-a-time covering property, look for a per-element invariant defined as a jump across one covering step, prove it path-independent using the SAME identity applied to two candidate steps simultaneously, then induct along the covering relation for the base case and the general claim together. The reason equality, not merely an inequality, is essential here is worth isolating: Lemma 2 needs the identity in both directions at once to force $f(S_1) - f(T_1) = f(S_2) - f(T_2)$ exactly; a submodular inequality, $f(S \cup S') + f(S \cap S') \leq f(S) + f(S')$, of the kind matroid rank functions satisfy, only gives one-directional information and does not pin down a well defined per-element jump. The two parallel-element matroid on $\{1, 2\}$, with $\text{rank}(\{1, 2\}) = 1$ but $\text{rank}(\{1\}) = \text{rank}(\{2\}) = 1$, is exactly such a submodular function that is not additive over elements, and it satisfies hypotheses (i) and (ii) trivially on the full power set of $\{1, 2\}$, confirming the equality hypothesis is where the additivity genuinely comes from, not the lattice structure alone.

What the examiners nudged toward: phrasing hypothesis (ii) as "exactly one fewer element", rather than merely "a proper subset", is the unmistakable signature of a graded poset with a rank function, and it is the direct instruction to induct on cardinality; writing the identity on f with a union and an intersection rather than, say, only unions of disjoint sets, is the instruction that the natural setting is a lattice closed under both operations, not merely a union-closed family.

Revolving concepts

Distributive lattices as rings of sets. A family of subsets of a finite set closed under union and intersection forms a distributive lattice under inclusion, and by Birkhoff's representation theorem every finite distributive lattice arises this way up to relabelling, so $\mathcal{P}$ here is a completely general finite distributive lattice, not a special case of one. See Book 1, field: Discrete structures and graphs, topic: lattice theory.

Valuations on a lattice. A function f on a lattice with $f(a \vee b) + f(a \wedge b) = f(a) + f(b)$ for all a, b is a valuation, or modular function; the hypothesis in this problem is exactly this identity specialised to a ring of sets, with $\vee, \wedge$ read as union and intersection. See Book 1, field: Discrete structures and graphs, topic: lattice valuations.

Graded posets and induction along covering relations. Hypothesis (ii) says $(\mathcal{P}, \subseteq)$ is graded by cardinality, with every nonempty element covering another element of $\mathcal{P}$ exactly one size smaller; such a poset supports strong induction on rank, with the base case obtained by iterating the covering property down to the minimum element. The same pattern, reconstruct a valuation from its edge jumps by inducting along covering relations, applies to any graded lattice with a modular rank function, not only rings of sets: for instance the flats of a matroid under a modular rank function, or the subgroup lattice of a finite group restricted to a modular chain. See Book 1, field: Combinatorics, topic: graded posets.

Matroid rank as a submodular, non additive contrast. A matroid rank function satisfies $\text{rank}(A \cup B) + \text{rank}(A \cap B) \leq \text{rank}(A) + \text{rank}(B)$, submodularity, with strict inequality possible: two parallel elements $1, 2$ have $\text{rank}(\{1,2\}) = 1 \neq \text{rank}(\{1\}) + \text{rank}(\{2\}) = 2$. This shows the equality in this problem's hypothesis, not merely an inequality, is what forces additivity over elements. See Book 1, field: Combinatorics, topic: matroid rank functions.

Two-set inclusion and exclusion. For a counting or measure function μ, $\mu(A \cup B) = \mu(A) + \mu(B) - \mu(A \cap B)$ is inclusion and exclusion for two sets; this problem imposes exactly that identity as an axiom on an abstract real valued f, which is what makes the lattice reading of the problem natural rather than invented. See Book 1, field: Combinatorics, topic: inclusion and exclusion.

Circuit Bank

- **WHEN** a hypothesis gives closure under two lattice-like operations together with an inclusion-exclusion style identity **DO** recognise the structure as a valuation on a lattice **BECAUSE** this tells you immediately which theorems, such as Birkhoff's representation theorem and the theory of modular functions, are relevant, and reframes "prove additivity" as "prove the valuation is determined by its covering-edge jumps"
- **WHEN** a covering-type hypothesis says there is always a legal object exactly one smaller or bigger, still in the family **DO** set up strong induction on size or rank **BECAUSE** it supplies both the base case, by iterating down to the bottom of the family, and the inductive step, one legal predecessor, for free
- **WHEN** you want a per-element quantity from a set function, but individual elements' singleton sets might not lie in your family **DO** define it as a difference across any legal covering step containing that element, then immediately prove independence of the step using the same functional identity applied to two candidate steps at once **BECAUSE** assuming singletons exist without proof is exactly the trap this kind of problem is built to catch

Backtesting

1. First useful cue: closure under union and intersection together with an inclusion-exclusion style identity names f as a valuation on a lattice, which tells you the target is really "is every valuation on this lattice additive over ground set elements".
2. Method the expression suggested: define a per-element weight as a jump across a covering pair supplied by hypothesis (ii), prove it well defined by feeding two candidate pairs into the modularity identity at once, then induct on cardinality.
3. Where you would hesitate, and the thought to reinforce: whether $\emptyset \in \mathcal{P}$ can be assumed. Reinforce by deriving it explicitly from hypothesis (ii), iterating the one-smaller condition down from any starting set, rather than treating $f(\emptyset) = 0$ as proof that $\emptyset$ is a member.

4. The trap the question sets: defining $f_i := f(\{i\})$, which silently assumes every element has its own singleton set inside $\mathcal{P}$, a fact hypothesis (ii) never actually guarantees.
5. What makes you recognise it faster next time: seeing closure under both union and intersection together with a "one fewer element" covering condition; that pairing is the signature of a graded lattice valuation problem, not a generic functional equation.
6. The reusable decision rule: when a functional identity is imposed as an axiom on a lattice-closed family together with a graded covering hypothesis, reconstruct the function from jumps across covering steps rather than trying to build every element up from singletons.

Mirror drills

1. Suppose additionally that $\mathcal{P}$ is exactly the full power set of $\{1, \ldots, n\}$. Show directly that $f_i = f(\{i\})$, matching the general construction. *Answer:* the pair $(\{i\}, \emptyset)$ is a legal covering pair for every i, since $\emptyset \in \mathcal{P}$ here trivially, so by the well-definedness lemma the general construction's f_i equals $f(\{i\}) - f(\emptyset) = f(\{i\})$.
2. Show that $\mathcal{P} = \{\{1\}, \{1,2\}\}$, for $n \geq 2$, is closed under union and intersection but does not contain $\emptyset$, and identify exactly which part of hypothesis (ii) this violates. *Answer:* $\{1\} \cup \{1,2\} = \{1,2\}$ and $\{1\} \cap \{1,2\} = \{1\}$, both in $\mathcal{P}$, so (i) holds; but (ii) applied to $S = \{1\}$ demands a subset of size 0 inside $\mathcal{P}$, forcing $\emptyset \in \mathcal{P}$, which this family does not contain, so (ii) fails and this is precisely the kind of family (ii) is written to exclude.
3. Replace the hypothesis by the inequality $f(S \cup S') + f(S \cap S') \leq f(S) + f(S')$ (submodularity) under the same (i) and (ii). Show the conclusion can fail. *Hint:* take $\mathcal{P}$ the full power set of $\{1,2\}$ and f = rank of the matroid where $1, 2$ are parallel: $\text{rank}(\emptyset) = 0$, $\text{rank}(\{1\}) = \text{rank}(\{2\}) = \text{rank}(\{1,2\}) = 1$; this satisfies the inequality but $\text{rank}(\{1,2\}) \neq \text{rank}(\{1\}) + \text{rank}(\{2\})$, so no such f_1, f_2 exist.

13.10 B4. Variance of a product of continuous functions

For any continuous real-valued function f defined on the interval $[0,1]$, let

$$\mu(f) = \int_0^1 f(x)\,dx, \quad \mathrm{Var}(f) = \int_0^1 (f(x) - \mu(f))^2\,dx,$$
$$M(f) = \max_{0\le x\le 1} |f(x)|.$$

Show that if f and g are continuous real-valued functions defined on the interval $[0,1]$, then

$$\mathrm{Var}(fg) \le 2\,\mathrm{Var}(f)M(g)^2 + 2\,\mathrm{Var}(g)M(f)^2.$$

Read it out loud

Say the structure before you touch a pen. The interval $[0,1]$ with dx is a probability space, because its total measure is 1. So μ is expectation, Var is variance, and M is the supremum norm. The claim is a statement about two random variables: the variance of a product is controlled by the variance of each factor weighted by the size of the other factor. Nothing about $[0,1]$ matters except the total mass 1, and nothing about continuity matters except that it makes $M(f)$ an attained maximum (extreme value theorem) and every integral finite. This is L^2 geometry on a probability space, plus one elementary inequality. There is no integral to evaluate, no special function, no differentiation. What is being tested is whether you know what variance *is*, as opposed to what its formula looks like.

The second cue is the constant 2 sitting in front of both terms on the right. A clean factor of 2 in an inequality almost always comes from one place: $(A+B)^2 \le 2A^2 + 2B^2$. Read the right side backwards and it is telling you what to do. It wants the quantity inside $\mathrm{Var}(fg)$ split into exactly two pieces, one piece carrying $\mathrm{Var}(f)$ and $M(g)$, the other carrying $\mathrm{Var}(g)$ and $M(f)$. Two terms, two bounds, one factor of 2 each. The examiners have told you the shape of the proof in the statement.

The third cue decides the problem. Var is invariant under adding a constant, but M is not. So the reflex "without loss of generality $\mu(f) = 0$" is illegal here: centring f leaves the left side and the Var factors alone but destroys your control of $M(f)$. That one observation is the entire difficulty. You may not normalise the means away, so you must carry $\mu(f)$ and $\mu(g)$ through the argument and kill them only at the last line with $|\mu(f)| \le M(f)$. Every candidate who lost time on this problem lost it by trying to centre and then not knowing what to do with the M terms.

Within thirty seconds you should have written three lines: $f_0 = f - \mu(f)$ and $g_0 = g - \mu(g)$; the algebraic identity $fg - \mu(f)\mu(g) = f_0 g + \mu(f) g_0$; and the fact that $\mathrm{Var}(h) \le \int_0^1 (h(x) - c)^2\,dx$ for *every* constant c, not just $c = \mu(h)$. Those three lines are the solution.

Approaches, ranked worst to best

1. **Bound everything by sup norms.** *Cue:* $M(f)$ and $M(g)$ appear on the right. *Circuit:* crude domination: replace each factor by its maximum. (worst, and it fails outright. You write $\mathrm{Var}(fg) \le \int_0^1 f^2 g^2 \le M(f)^2 M(g)^2$ and then discover the right side of the problem does not dominate that. Test it: take g constant. Then $\mathrm{Var}(fg) = g^2\mathrm{Var}(f) = \mathrm{Var}(f)M(g)^2$, while your bound gives $M(f)^2 M(g)^2$, which is far too big. The right side of the problem vanishes whenever f and g are both constant; any proof built from quantities that do not vanish there is dead on arrival. Kill this route in ten seconds by testing constants.)

2. **Full expansion.** *Cue:* everything in sight is an explicit integral. *Circuit:* expand $\mathrm{Var}(fg) = \int f^2 g^2 - (\int fg)^2$, substitute $f = f_0 + \mu(f)$ and $g = g_0 + \mu(g)$, and grind. (this is the route the official solution essentially takes, in an organised form. It works. It costs seven or eight terms,

several of which cancel only because $\int f_0 = \int g_0 = 0$, and it needs an auxiliary inequality at the end to absorb the surviving cross term $2\mu(f)\mu(g)\int f_0 g_0$. Two thirds of a page of correct algebra with no structural guidance, and every line is a chance for a skipped-calculation error. Zombie maths: correct, expensive, and it teaches you nothing about why the 2 is there.)

3. **Variance of a sum.** *Cue:* the right side is a sum of two terms with a 2 each. *Circuit:* write fg as a sum of two functions and use $\mathrm{Var}(A+B) \le 2\mathrm{Var}(A) + 2\mathrm{Var}(B)$. (good, and morally the ideal route. $fg = f_0 g + \mu(f)g$, so $\mathrm{Var}(fg) \le 2\mathrm{Var}(f_0 g) + 2\mu(f)^2\mathrm{Var}(g)$, then $\mathrm{Var}(f_0 g) \le \int f_0^2 g^2 \le M(g)^2\mathrm{Var}(f)$. It ranks just below the next one only because it makes you quote the sub-additivity of variance as a separate fact rather than exhibiting it.)

4. **Distance to the constants.** *Cue:* Var is a squared distance, and the problem forbids centring. *Circuit:* $\mathrm{Var}(h) = \min_c \int_0^1 (h-c)^2$, so choose the constant yourself. (best. You never compute $\mu(fg)$ at all: you bound the minimum by one well-chosen competitor, $c = \mu(f)\mu(g)$, which factors the difference as $f_0 g + \mu(f)g_0$ in one line. Then $(A+B)^2 \le 2A^2 + 2B^2$ pointwise and two sup-norm extractions finish it. Four lines, no cancellation, and it hands you the strictly stronger inequality in the remark below for free.)

The ideal solution

Write $\mu_f = \mu(f)$, $\mu_g = \mu(g)$, $f_0 = f - \mu_f$ and $g_0 = g - \mu_g$, so that

$$\int_0^1 f_0 = \int_0^1 g_0 = 0, \qquad \int_0^1 f_0^2 = \mathrm{Var}(f), \qquad \int_0^1 g_0^2 = \mathrm{Var}(g).$$

All of these integrals exist because f and g are continuous on a compact interval, and for the same reason $M(f)$ and $M(g)$ are attained and finite.

Lemma 13.10.1. *For any continuous h on $[0,1]$ and any real constant c,*

$$\mathrm{Var}(h) \le \int_0^1 (h(x) - c)^2 \, dx,$$

with equality exactly when $c = \mu(h)$.

Proof. Put $h_0 = h - \mu(h)$. Then $h - c = h_0 + (\mu(h) - c)$, and since $\int_0^1 h_0 = 0$ the cross term integrates to zero:

$$\int_0^1 (h-c)^2 = \int_0^1 h_0^2 + 2(\mu(h)-c)\int_0^1 h_0 + (\mu(h)-c)^2 = \mathrm{Var}(h) + (\mu(h)-c)^2.$$

The last term is nonnegative and vanishes only for $c = \mu(h)$. □

Apply the lemma to $h = fg$ with the constant $c = \mu_f \mu_g$. The point of this choice is the identity

$$f(x)g(x) - \mu_f\mu_g = \big(f(x) - \mu_f\big)g(x) + \mu_f\big(g(x) - \mu_g\big) = f_0(x)g(x) + \mu_f g_0(x),$$

which is checked by expanding the middle expression: the two copies of $\mu_f g(x)$ cancel. Hence

$$\mathrm{Var}(fg) \le \int_0^1 \big(f_0(x)g(x) + \mu_f g_0(x)\big)^2 dx.$$

Now use $(A+B)^2 \le 2A^2 + 2B^2$, valid for all reals because it rearranges to $(A-B)^2 \ge 0$, at each point x:

$$\mathrm{Var}(fg) \le 2\int_0^1 f_0(x)^2 g(x)^2 \, dx + 2\mu_f^2 \int_0^1 g_0(x)^2 \, dx.$$

For the first integral, $g(x)^2 \le M(g)^2$ for every x and $f_0(x)^2 \ge 0$, so

$$\int_0^1 f_0^2 g^2 \le M(g)^2 \int_0^1 f_0^2 = \operatorname{Var}(f) M(g)^2.$$

The second integral is $\operatorname{Var}(g)$. Finally

$$|\mu_f| = \left| \int_0^1 f(x)\,dx \right| \le \int_0^1 |f(x)|\,dx \le M(f),$$

so $\mu_f^2 \le M(f)^2$. Putting the three estimates together,

$$\operatorname{Var}(fg) \le 2\operatorname{Var}(f) M(g)^2 + 2\operatorname{Var}(g) M(f)^2,$$

which is the required inequality.

Remark 13.10.2. The same three lines give something stronger at no extra cost. Writing $\|u\|_2 = (\int_0^1 u^2)^{1/2}$, the triangle inequality in L^2 (Minkowski) applied to $f_0 g + \mu_f g_0$ gives

$$\sqrt{\operatorname{Var}(fg)} \le \|f_0 g\|_2 + |\mu_f|\,\|g_0\|_2 \le M(g)\sqrt{\operatorname{Var}(f)} + M(f)\sqrt{\operatorname{Var}(g)},$$

and squaring recovers the stated inequality through $(u+v)^2 \le 2u^2 + 2v^2$. So the constant 2 is exactly the slack in one application of the quadratic mean against the arithmetic mean. The 2 is not optimal: taking $g \equiv 1$ gives $\operatorname{Var}(fg) = \operatorname{Var}(f)$ and $\operatorname{Var}(g) = 0$, $M(g) = 1$, so no constant below 1 can work. The true optimal constant lies somewhere in $[1, 2]$ and the problem sensibly does not ask for it.

Layer 1: WHAT

Block 1. Read $[0, 1]$ as a probability space and Var as a squared L^2 distance from h to the constant functions.

Block 2. Notice that you cannot normalise $\mu(f)$ to 0, because $M(f)$ is not translation invariant. Conclusion: do not centre the inputs; centre the *output* instead.

Block 3. Bound $\operatorname{Var}(fg)$ by $\int (fg - c)^2$ for a constant c of your choosing, and choose $c = \mu(f)\mu(g)$ because it makes $fg - c$ factor as $f_0 g + \mu(f) g_0$.

Block 4. Split with $(A + B)^2 \le 2A^2 + 2B^2$. This is where the 2 in the problem is spent.

Block 5. Extract sup norms: $g^2 \le M(g)^2$ inside the first integral, and $\mu(f)^2 \le M(f)^2$ by the triangle inequality for integrals.

The nominal route bifurcates from the ideal at Block 3. The zombie writes $\operatorname{Var}(fg) = \int f^2 g^2 - \left(\int fg\right)^2$, because that is the definition unpacked, and then substitutes $f = f_0 + \mu_f$, $g = g_0 + \mu_g$ into both terms and expands. That produces the correct answer after eight terms and one extra auxiliary inequality. The ideal route refuses to compute $\mu(fg)$ at all, on the grounds that a minimum only ever needs to be bounded above by one competitor. **The bifurcation point is the line where you decide whether to expand** $\operatorname{Var}(fg)$ **or to over-estimate it by a distance to a constant you get to name.** Every saving in this problem comes from that one decision.

Layer 2: HOW

Block 1. Cue. The three definitions given in the question are exactly expectation, variance and supremum norm on a measure space of total mass 1. The question defines them rather than naming them, which is the examiner telling you that the probabilistic reading is available but not assumed. **Why natural.** A definition-level reading always beats a formula-level reading. $\operatorname{Var}(h) = \int (h - \mu(h))^2$ is a formula; "the squared L^2 distance from h to the constant function nearest it" is what the object is. The second reading immediately supplies a family of upper bounds, one for each constant, and upper bounds are what you need. **Link to the next block.** Once variance is a distance to the constants, the next question asks itself: which constant should I compare fg to?

Block 2. **Cue.** $\mathrm{Var}(f+c) = \mathrm{Var}(f)$ but $M(f+c) \neq M(f)$. **Why natural.** The standard reflex on any variance problem is to assume the mean is zero. Here the right side of the inequality contains M, which is not translation invariant, so that reflex would be changing the problem. Spotting a forbidden normalisation early is worth more than any amount of algebra later; it tells you that $\mu(f)$ and $\mu(g)$ are genuine actors in the proof and must survive to the last line. **Link to the next block.** If the means cannot be removed at the start, they must be removed at the end, and the only tool for that is $|\mu(f)| \leq M(f)$. Now you know the last line of the proof before you know the middle.

Block 3. **Cue.** You need a constant c near fg. The only constants in the problem are $\mu(f)$, $\mu(g)$ and their combinations. **Why natural.** Take $c = \mu_f \mu_g$ because it is the unique choice that makes $fg - c$ split into two products, each with one centred factor: $fg - \mu_f\mu_g = f_0 g + \mu_f g_0$. This is the same telescoping move as $ab - a'b' = (a-a')b + a'(b-b')$, which you already use for products of sequences and for the product rule. The right side of the problem has two terms, so you need exactly two pieces, so you need exactly one telescoping step. **Link to the next block.** The expression is now a sum of two functions, and the target is a bound on the square of that sum.

Block 4. **Cue.** A squared sum, and a target with a 2 in front of each term. **Why natural.** $(A+B)^2 \leq 2A^2 + 2B^2$ is the cheapest possible way to break a squared sum, and the presence of the 2s in the statement is confirmation that the intended proof breaks the sum exactly once and crudely. If the intended constant had been 1 you would have had to be delicate about the cross term; the 2 is the examiners paying you to be lazy. **Link to the next block.** Two integrals remain, each with one centred factor and one uncentred factor.

Block 5. **Cue.** $\int f_0^2 g^2$ has a factor bounded by a constant and a factor whose integral is a named quantity. **Why natural.** This is Hoelder with the exponent pair $(1, \infty)$, which in plain language is: pull the maximum out of the integral. It costs one line and it is the only estimate available that keeps $\mathrm{Var}(f)$ intact. The second integral is $\mathrm{Var}(g)$ on the nose, and the coefficient μ_f^2 is converted by the triangle inequality for integrals, which is where the promise made in Block 2 is honoured. **Link to the next block.** Nothing remains. Assemble.

Layer 3: WHY

The mathematical truth underneath is that $\sqrt{\mathrm{Var}}$ is a seminorm. Consider the space of continuous functions on $[0,1]$ modulo the constants: $\mathrm{Var}(h)$ is the squared L^2 norm of the class of h in that quotient. Two consequences follow immediately and both are used. First, a seminorm satisfies the triangle inequality, so $\sqrt{\mathrm{Var}(A+B)} \leq \sqrt{\mathrm{Var}(A)} + \sqrt{\mathrm{Var}(B)}$, which is exactly the sub-additivity that makes the two-term split legal. Second, a squared distance to a subspace is a minimum over that subspace, so it is bounded above by the value at any single point of it. Almost all of the ink saved in this solution is saved by the second consequence: you never compute the foot of the perpendicular, you just name a nearby point.

The discernment mechanism is a general one and you should install it as a habit. *Inspect the degenerate cases of the inequality before choosing a method, and build the proof only from quantities that degenerate the same way.* Here the right side vanishes when f and g are both constant, and more sharply, the right side is unchanged when a constant is added to f while the left side is not. That asymmetry says the proof must treat $\mu(f)$ as a coefficient, never as something to be removed. A candidate who tests g constant in the first minute learns three things: crude sup-norm bounds are dead, the inequality is nearly tight in that case, and the factor 2 is slack. Three facts for ten seconds of work.

What the examiners nudged. They wrote the definitions out instead of saying "variance", which invites the probabilistic reading without requiring it. They put a 2 in front of both terms, which is a

public announcement that a single crude split is intended. They used $M(f) = \max |f|$ rather than $\max f - \min f$, which forces the appearance of $|\mu(f)| \leq M(f)$ and thereby forces you to confront the translation issue. Every one of those choices is a guiding light, and none of them is decoration.

To train the discernment: for the next month, whenever you meet $\int (h - \text{something})^2$, write the sentence "this is the squared distance from h to the set of somethings" and then ask what the nearest something is and whether you can afford to guess it. That reflex converts a large family of analysis inequalities (variance bounds, best approximation, least squares, Bessel and Parseval, the Poincare inequality) into one move.

Revolving concepts

Variance as a squared distance. For $h \in L^2$ of a probability space, $\mathrm{Var}(h) = \min_{c \in \mathbb{R}} \int (h-c)^2 = \int h^2 - \left(\int h\right)^2$, the minimum being attained at $c = \mu(h)$. Orbiting facts: $\mathrm{Var}(ah+b) = a^2\mathrm{Var}(h)$; $\mathrm{Var}(h) = 0$ if and only if h is constant (for continuous h); $\mathrm{Cov}(u,v) = \int u_0 v_0$ and $\mathrm{Cov}(u,v)^2 \leq \mathrm{Var}(u)\mathrm{Var}(v)$, which is Cauchy and Schwarz applied to the centred functions. See Book 1, field: Real analysis, topic: L^2 geometry.

Minkowski and Hoelder. $\|u+v\|_p \leq \|u\|_p + \|v\|_p$ for $p \geq 1$, and $\int |uv| \leq \|u\|_p \|v\|_q$ for $1/p + 1/q = 1$. The two special cases used here are $p = q = 2$ (Cauchy and Schwarz) and the pair $(1, \infty)$, that is $\int |uv| \leq \|v\|_\infty \int |u|$, which is the formal name of "pull the maximum out of the integral". See Book 1, field: Inequalities, topic: L^p norms.

The two-term power bound and its generalisation. $(A+B)^2 \leq 2(A^2+B^2)$ generalises to $(A_1 + \cdots + A_m)^2 \leq m(A_1^2 + \cdots + A_m^2)$, which is Cauchy and Schwarz against the all-ones vector, and further to $(\sum A_i)^p \leq m^{p-1} \sum A_i^p$ for $p \geq 1$ by convexity of $t \mapsto t^p$. If a problem hands you a constant m in front of an m-term right side, that is the intended tool. There is no recursion in this problem, but the m-term version is proved by induction on m from the two-term version, and the general convexity statement is Jensen. See Book 1, field: Inequalities, topic: power mean.

Popoviciu and the range bound. A companion to the lemma: choosing c to be the midpoint of the range gives $\mathrm{Var}(h) \leq \frac{1}{4}(\max h - \min h)^2$. This is a second, sharper way to convert a sup bound into a variance bound and it is often the missing ingredient in Putnam variance problems. See Book 1, field: Inequalities, topic: variance bounds.

Extreme value theorem. A continuous real function on a closed bounded interval attains its maximum. This is what makes $M(f)$ well defined as a maximum rather than a supremum, and it is the only place continuity is used beyond integrability. See Book 1, field: Real analysis, topic: compactness.

Circuit Bank

- **WHEN** you see $\int (h - \mu(h))^2$ or the word variance **DO** write $\mathrm{Var}(h) \leq \int (h-c)^2$ for a constant c you will choose **BECAUSE** variance is the squared distance from h to the constants, so any constant gives an upper bound and you never have to compute the true mean
- **WHEN** an inequality has a factor m in front of an m term right side **DO** split the left side into m pieces and apply $(\sum A_i)^2 \leq m \sum A_i^2$ **BECAUSE** that constant is the examiner telling you the intended proof splits once and crudely
- **WHEN** a supremum norm multiplies an integral **DO** bound the squared factor pointwise by its maximum and pull it outside **BECAUSE** this is Hoelder with the pair $(1, \infty)$, the cheapest estimate that leaves the other factor intact
- **WHEN** you want to normalise a mean to zero **DO** first check every quantity in the target for translation invariance **BECAUSE** Var is translation invariant and M is not, so centring silently changes the problem

- **WHEN** you must compare a product fg with a product ab of constants **DO** telescope as $fg - ab = (f - a)g + a(g - b)$ **BECAUSE** one telescoping step turns a product difference into a sum of two terms, each with exactly one centred factor

Backtesting

1. First useful cue: the definitions given are expectation, variance and supremum norm on a probability space, so $\mathrm{Var}(h)$ is a squared distance from h to the constant functions.
2. Method the expression suggested: bound that distance by the distance to one named constant, $\mu(f)\mu(g)$, then split the resulting square into two pieces because the right side has two terms with a 2 each.
3. Where you would hesitate, and the thought to reinforce: at the temptation to set $\mu(f) = 0$. Reinforce by checking translation invariance of every symbol in the target before normalising anything. M moves, Var does not, so the means stay.
4. The trap the question sets: expanding $\mathrm{Var}(fg) = \int f^2 g^2 - (\int fg)^2$ and grinding. It is correct and it costs eight terms plus one extra inequality to absorb the surviving cross term $2\mu(f)\mu(g)\int f_0 g_0$.
5. What makes you recognise it faster next time: the pattern "squared deviation from a mean" plus a constant multiplier on the right. Test the degenerate case (g constant) in the first ten seconds, which simultaneously kills the crude bounds and shows the 2 is slack.
6. The reusable decision rule: a minimum is bounded above by any single competitor, so never compute the optimal one when you only need an upper bound.

Mirror drills

1. Show that $\mathrm{Var}(f+g) \le 2\mathrm{Var}(f) + 2\mathrm{Var}(g)$ for continuous f, g on $[0,1]$, and state when equality holds. *Answer:* $\mathrm{Var}(f+g) = \int (f_0+g_0)^2 \le 2\int f_0^2 + 2\int g_0^2$; equality requires $f_0 = g_0$ almost everywhere, that is f and g differ by a constant.
2. Prove that $\mathrm{Var}(f) \le \frac{1}{4}(\max f - \min f)^2$ for continuous f on $[0,1]$. *Answer:* take c to be the midpoint of $[\min f, \max f]$ in the lemma; then $|f - c| \le \frac{1}{2}(\max f - \min f)$ pointwise, so $\mathrm{Var}(f) \le \int (f-c)^2 \le \frac{1}{4}(\max f - \min f)^2$.
3. Prove the sharper bound $\sqrt{\mathrm{Var}(fg)} \le M(g)\sqrt{\mathrm{Var}(f)} + M(f)\sqrt{\mathrm{Var}(g)}$, deduce the problem's inequality from it, and show that the constant 2 in the problem cannot be replaced by any constant smaller than 1. *Answer:* Minkowski in L^2 on $f_0 g + \mu(f) g_0$, then $(u+v)^2 \le 2u^2 + 2v^2$; for the lower bound on the constant take $g \equiv 1$, so the left side is $\mathrm{Var}(f)$ and the right side is $C\,\mathrm{Var}(f)$.

13.11 B5. Counting maps that funnel into a set

Let $X = \{1, 2, \ldots, n\}$, and let $k \in X$. Show that there are exactly $k \cdot n^{n-1}$ functions $f : X \to X$ such that for every $x \in X$ there is a $j \geq 0$ such that $f^{(j)}(x) \leq k$. (Here $f^{(j)}$ denotes the jth iterate of f, so that $f^{(0)}(x) = x$ and $f^{(j+1)}(x) = f(f^{(j)}(x))$.)

Read it out loud

This is a labelled counting problem about a discrete dynamical system. A function $f : X \to X$ on a finite set is nothing but a directed graph in which every vertex has out-degree exactly 1, the functional graph. Following $x, f(x), f^{(2)}(x), \ldots$ you must eventually repeat a value, so every orbit runs down a tail and then loops forever around a cycle. The condition imposed is that every orbit passes through the target block $K = \{1, \ldots, k\}$. Read that once more in graph language and it says: every vertex can reach K, that is, K absorbs the whole system. The fields wired in are enumerative combinatorics of labelled structures (trees, forests, Cayley's formula) and induction with a binomial recursion.

Two features of the answer are shouting at you. First, $k \cdot n^{n-1}$ is linear in k. A count that is linear in a parameter almost never comes out of a case analysis; it comes out of a factorisation or a telescoping, so look for a place where k enters exactly once. Second, n^{n-1} and n^{n-2} are the Cayley numbers: n^{n-2} labelled trees on n vertices, n^{n-1} labelled rooted trees on n vertices, $q\, n^{n-q-1}$ labelled rooted forests on n vertices with a prescribed set of q roots. An answer of that shape is the examiners saying "this object is a forest in disguise". Third, and most useful of all, divide by the total: $kn^{n-1} = \frac{k}{n} \cdot n^n$. So the claim is equivalent to saying that a uniformly random function $X \to X$ funnels into K with probability exactly k/n. A clean probability like that is never a coincidence, and it is worth keeping in view because it tells you the final identity will be exact and simple, with no floor functions or case splits.

Now the cue that starts the work, and you should have it within twenty seconds. For $x \leq k$ the condition is satisfied by $j = 0$: nothing is being asked about the points of K at all. So the values $f(1), \ldots, f(k)$ ought to be completely unconstrained, giving a free factor n^k. That is a guess until you prove it, because $f(1)$ could in principle be needed on somebody else's orbit. It is not, and the reason is the only genuinely delicate observation in the problem: once an orbit reaches K it has already succeeded, so the part of the orbit that matters never consults a value of f at a point of K. Write that down, pull out n^k, and you are left with counting maps from $Y = \{k+1, \ldots, n\}$ into X that are absorbed by K.

Approaches, ranked worst to best

1. **Inclusion and exclusion over the bad set.** *Cue:* count objects avoiding a global condition. *Circuit:* complement plus sieve. (worst. The complementary condition is "some cycle of f avoids K", and sieving over which subsets carry those cycles produces sums over cycle types with factorial weights. Nothing closes, and there is no sign of kn^{n-1} at the end of it.)

2. **Induction on n with k fixed.** *Cue:* a formula in two integer parameters. *Circuit:* peel off the point n and relate to the same problem one size down. (stalls. Deleting a point of Y changes the ambient set that everything maps into, and deleting a point of K changes the target. The recursion you get relates the count for (n, k) to counts with a different, structurally new target set, so with k held fixed it never closes. The failure is informative: it is telling you which parameter must be freed.)

3. **Exponential generating functions for functional graphs.** *Cue:* labelled objects, so use the species machinery. *Circuit:* $T = xe^T$ for rooted trees, functional graphs are sets of cycles of trees. (viable and heavy. You must install Lagrange inversion to extract coefficients, and the

marked block K of fixed size k breaks the plain symmetry of the species, so you end up with a two variable generating function anyway. In the arena this costs three times what the ideal route costs.)

4. **Quote the rooted forest formula.** *Cue:* n^{n-1} is Cayley shaped. *Circuit:* a map absorbed by K is a forest of trees hanging off K, and forests on n vertices with q prescribed roots number qn^{n-q-1}. (fast and correct if the formula is installed. Two risks: you must justify the bijection between absorbed maps and forests carefully, and you are leaning on the matrix tree theorem or on Pruefer codes, which a grader may reasonably ask you to prove. Excellent as a sanity check on the answer, slightly exposed as a write-up.)

5. **Induction on $n-k$ through the new cycle.** *Cue:* the answer increases by exactly n^{n-1} when k increases by 1. *Circuit:* count the maps absorbed by $\{1,\dots,k+1\}$ but not by $\{1,\dots,k\}$, which are those with a unique cycle inside $\{k+1,\dots,n\}$, necessarily through $k+1$. (this is the official route and it works. The count for fixed cycle length ℓ is a product of falling factorials, and the resulting sum only collapses after you rewrite $k+\ell$ as $n-(n-k-\ell)$ to make two telescoping halves. Correct, but the collapse is the sort of step you will not find under time pressure unless you already know it.)

6. **Peel the first entry layer.** *Cue:* the condition is about first arrival at a set, so condition on the set of points that arrive in one step. *Circuit:* strengthen the statement to two free parameters, sink size and non-sink size, then induct. (best. The recursion is forced by the structure, the induction closes because the sink is a variable, and the binomial sum collapses in one line using $j\binom{p}{j} = p\binom{p-1}{j-1}$ and the binomial theorem. It also proves Cayley's formula and the forest formula as by-products, so you are never quoting anything you cannot derive.)

The ideal solution

Throughout, $K=\{1,\dots,k\}$ and $Y=\{k+1,\dots,n\}$, so $X=K\cup Y$ with $|K|=k\geq 1$ and $|Y|=n-k\geq 0$. Call $f:X\to X$ *good* if for every $x\in X$ there is $j\geq 0$ with $f^{(j)}(x)\in K$.

Lemma 13.11.1. *Whether f is good depends only on the restriction $f|_Y$. Consequently the number of good functions is n^k times the number of good restrictions.*

Proof. Let f and $\tilde{f}$ agree on Y, and suppose f is good. Fix $x\in X$. If $x\in K$ then $j=0$ works for $\tilde{f}$ as well. If $x\in Y$, let j be minimal with $f^{(j)}(x)\in K$; then $j\geq 1$ and $f^{(i)}(x)\in Y$ for $0\leq i<j$ by minimality. Since f and $\tilde{f}$ agree on Y, an induction on i gives $\tilde{f}^{(i)}(x)=f^{(i)}(x)$ for $0\leq i\leq j$, so $\tilde{f}^{(j)}(x)\in K$. Hence $\tilde{f}$ is good. The values $f(1),\dots,f(k)$ are therefore free, contributing the factor n^k. □

Now fix the general shape of what remains. Let Y and S be disjoint finite sets with $|Y|=p\geq 0$ and $|S|=q\geq 1$, and let $h:Y\to Y\cup S$. Define the orbit of $y\in Y$ by $y_0=y$ and $y_{i+1}=h(y_i)$, continued only while $y_i\in Y$. Call h *absorbing* (for the pair (Y,S)) if the orbit of every $y\in Y$ reaches S.

Lemma 13.11.2. *The number of absorbing $h:Y\to Y\cup S$ is $q\,(p+q)^{p-1}$.*

Proof. Induct on p. If $p=0$ there is exactly one map, the empty map, and it is vacuously absorbing; the formula gives $q\cdot q^{-1}=1$, as required.

Let $p\geq 1$ and let h be absorbing. Put $T=h^{-1}(S)\subseteq Y$. First, $T\neq\emptyset$: pick any $y\in Y$ and let i be minimal with $y_i\in S$; then $i\geq 1$ and $y_{i-1}\in Y$ satisfies $h(y_{i-1})\in S$, so $y_{i-1}\in T$.

Fix a subset $T\subseteq Y$ with $|T|=j\geq 1$. I claim that the absorbing maps h with $h^{-1}(S)=T$ are in bijection with the pairs consisting of an arbitrary map $T\to S$ and an absorbing map $Y\setminus T\to (Y\setminus T)\cup T$ for the pair $(Y\setminus T,T)$. Indeed, $h^{-1}(S)=T$ says exactly that h maps T into S and maps $Y\setminus T$ into $Y=(Y\setminus T)\cup T$, so h splits into those two pieces and any pair of pieces can

be recombined. It remains to see that h is absorbing for (Y, S) if and only if $h|_{Y\setminus T}$ is absorbing for $(Y \setminus T, T)$. If h is absorbing and $y \in Y \setminus T$, its orbit reaches S, and the point immediately before it does so lies in T; every earlier point of the orbit lies in $Y \setminus T$, since a point of T is sent straight into S. So the orbit reaches T while staying inside $Y \setminus T$ beforehand, which is precisely absorption for $(Y \setminus T, T)$. Conversely, if every $y \in Y \setminus T$ reaches T, then one further step sends it into S, and points of T enter S in one step, so h is absorbing.

Since $|Y \setminus T| = p - j < p$ and $|T| = j \geq 1$, the induction hypothesis counts the second piece as $j\,\big((p-j)+j\big)^{p-j-1} = j\,p^{p-j-1}$, while the first piece may be chosen in q^j ways. Summing over j and over the $\binom{p}{j}$ choices of T,

$$\#\{h \text{ absorbing}\} = \sum_{j=1}^{p} \binom{p}{j} q^j\, j\, p^{p-j-1}.$$

Use $j\binom{p}{j} = p\binom{p-1}{j-1}$, which is the identity "choose a committee of j from p and then its chair" counted the other way round. The factor p cancels one power of p:

$$\sum_{j=1}^{p} p\binom{p-1}{j-1} q^j p^{p-j-1} = \sum_{j=1}^{p} \binom{p-1}{j-1} q^j p^{p-j} = q \sum_{i=0}^{p-1} \binom{p-1}{i} q^i p^{(p-1)-i},$$

where $i = j - 1$. The last sum is the binomial expansion of $(p+q)^{p-1}$, so the total is $q(p+q)^{p-1}$. □

Apply the second lemma with $S = K$, so $q = k$ and $p = n - k$ and $p + q = n$: the number of good restrictions $f|_Y$ is $k\,n^{n-k-1}$. Note this is exactly the statement that $f|_Y$ is absorbing, which by the definition of good is what "good" says once the points of K are known to be free. By the first lemma the number of good functions is

$$n^k \cdot k\, n^{n-k-1} = \boxed{k\, n^{n-1}}.$$

For $k = n$ the set Y is empty, every function is good, and the formula correctly returns $n \cdot n^{n-1} = n^n$.

Remark 13.11.3. Two reformulations worth carrying away. Since every orbit is eventually periodic, f is good if and only if every cycle of f meets K: a cycle avoiding K traps its own points forever, and conversely every orbit runs into a cycle. And dividing by n^n, the theorem says that a uniformly random $f : X \to X$ is good with probability exactly k/n, the fraction of the ground set occupied by K.

Remark 13.11.4. The second lemma with $q = 1$ says: the number of maps h from a p element set Y into $Y \cup \{r\}$ all of whose orbits reach r is $(p+1)^{p-1}$. Drawing the edge $y \to h(y)$, such maps are exactly the trees on the $p + 1$ vertices $Y \cup \{r\}$, rooted at r with all edges directed towards the root. With $N = p + 1$ vertices that count is N^{N-2}, which is Cayley's formula. So the ideal route proves the tool it would otherwise have had to quote.

Layer 1: WHAT

Block 1. Read the condition on the functional graph: every vertex reaches K. Notice it is vacuous on K itself.

Block 2. Prove that goodness reads only $f|_Y$, by taking the *first* entry time into K, and pull out the free factor n^k.

Block 3. The residual count is an absorption count. Recognise that an induction will change the sink set, so promote the sink from the fixed set K to a free parameter: state the two parameter lemma.

Block 4. Peel the first entry layer $T = h^{-1}(S)$. Prove it is nonempty, and that the rest of h is an absorbing map with sink T.

Block 5. Assemble the binomial recursion, kill the stray factor j with $j\binom{p}{j} = p\binom{p-1}{j-1}$, and recognise the binomial theorem.

Block 6. Multiply the two factors and check $k = n$.

The nominal route sets k up as a fixed constant, tries to induct on n, and drowns. The bifurcation point is Block 3: **the line where you decide that the target set is a variable of the problem and not a constant.** Everything downstream is forced once the statement has two free parameters, because the peel then lands you back inside the same statement rather than outside it. This is the classic strengthen the induction hypothesis moment, and the tell is precisely that the naive recursion produced an object of the same kind but the wrong size.

Layer 2: HOW

Block 1. **Cue.** "For every x there is $j \geq 0$ with $f^{(j)}(x) \leq k$" is a reachability statement, and $j = 0$ is allowed. **Why natural.** Allowing $j = 0$ is not decoration; it is the examiner telling you that the points of K are exempt. When a condition is vacuous on part of the domain, the corresponding coordinates are candidates to be free, and free coordinates should always be factored out before any counting begins, because they shrink the problem without costing anything. **Link to the next block.** You now need to prove the values on K really are free, which means proving that nobody else's orbit depends on them.

Block 2. **Cue.** The word "there is j" hides a minimal j, and minimal first entry times are what make orbit arguments finite and clean. **Why natural.** Take the least j with $f^{(j)}(x) \in K$. Then the whole path $x, f(x), \ldots, f^{(j-1)}(x)$ lies in Y, so it only ever reads values of f at points of Y. That is the definition-level reason the factor n^k exists: success is decided before the orbit ever touches K. Never argue with an arbitrary j when a minimal one is available. **Link to the next block.** With n^k removed, the object left is a map $Y \to X$ absorbed by K, and K now behaves like a single exit door rather than a set of ordinary points.

Block 3. **Cue.** The naive induction wants to remove a point, and removing one changes the sink. **Why natural.** If a recursion produces the same kind of object with a different parameter, the statement you are inducting on must have that parameter free. Here the peel will replace the sink K by a subset of Y of some new size j, so a lemma phrased with sink size exactly k can never be applied to its own consequences. Freeing the parameter costs one line of notation and buys the entire induction. **Link to the next block.** Once the lemma reads "sink of size q, ground of size p", you may look for the natural way to cut Y.

Block 4. **Cue.** "Every orbit reaches S" is a first arrival condition, and the natural slice of a first arrival condition is the set of points that arrive in exactly one step. **Why natural.** Set $T = h^{-1}(S)$. It is nonempty (take any orbit and look one step before it lands), it is canonically determined by h (no choices, so no overcounting), and by construction the rest of h maps $Y \setminus T$ into Y and must funnel into T. So the peel is exactly the same problem with sink T: the structure regenerates itself. Compare this with the official route, which peels a cycle instead; a cycle needs to be chosen and ordered, which is why falling factorials appear there and binomials appear here. **Link to the next block.** Summing over the size and the choice of T produces a binomial recursion.

Block 5. **Cue.** A binomial sum with a stray linear factor j in it. **Why natural.** The identity $j\binom{p}{j} = p\binom{p-1}{j-1}$ exists precisely to remove such factors, and after removing it the summand is $\binom{p-1}{i} q^i p^{(p-1)-i}$, which is a term of $(p+q)^{p-1}$. The fact that everything collapses to the plain binomial theorem is the confirmation that the peel was the right cut: a wrong decomposition leaves

you with a sum that needs an Abel type identity instead. **Link to the next block.** The answer $q(p+q)^{p-1}$ with $p+q=n$ is the missing factor.

Block 6. **Cue.** Two independent factors, one free and one counted. **Why natural.** Multiply, then test the extreme case $k=n$ where Y is empty. Checking a boundary case at the end costs five seconds and catches an off by one in the exponent, which is the most likely error in a formula with n^{n-k-1} in it. **Link to the next block.** Done.

Layer 3: WHY

The structural truth is that an absorbing map is a rooted forest. Draw an edge from y to $h(y)$ for each $y \in Y$. Absorption says no cycle lies inside Y, so the digraph on $Y \cup S$ is a forest of trees whose roots are the elements of S, edges directed towards the roots. The second lemma is therefore the classical count of labelled rooted forests on $p+q$ vertices with a prescribed set of q roots, namely $q(p+q)^{p-1}$, and the case $q=1$ is Cayley's formula. That is why the answer had to look like kn^{n-1}: nothing else of that shape exists in the labelled world.

Why is the count linear in q? Because the true statement is a probability. Of all $(p+q)^p$ maps $Y \to Y \cup S$, exactly the fraction $q/(p+q)$ are acyclic on Y. In the original problem this is the statement that a random $f : X \to X$ funnels into K with probability $|K|/|X|$. That is the same phenomenon as the cycle lemma of Dvoretzky and Motzkin, which computes the probability that a random cyclic arrangement never dips below a level and returns an answer linear in the parameter for the same reason: there is a $(p+q)$ to 1 symmetry acting on all maps under which exactly q of every $p+q$ related objects are good. Knowing that a clean fraction is in play tells you before you start that the final identity will be an exact binomial collapse, not an estimate.

The discernment mechanism has three parts, and all three transfer. First, *separate free coordinates from constrained coordinates before counting anything*: any coordinate the condition never reads contributes a bare power, and removing it shrinks the problem for free. Second, *when a decomposition regenerates your problem with a parameter you were treating as constant, promote that parameter*. A statement is easier to prove when it is stronger, if the extra strength is exactly what the recursion consumes. Third, *cut along a canonical set, never along a chosen one*. The first entry layer $h^{-1}(S)$ is determined by h, so summing over it partitions the objects with no overcounting and no ordering factors; choosing a cycle, or a vertex, or an edge, introduces a symmetry you must divide out later.

To train it: take any counting problem whose answer contains n^{n-1}, n^{n-2} or $(n+1)^{n-1}$ and force yourself to name the forest. Then take any condition of the form "every point eventually reaches A" and write down $h^{-1}(A)$ before doing anything else.

Revolving concepts

Functional graphs. For finite X, a map $f : X \to X$ drawn as the digraph $x \to f(x)$ has every out-degree 1. Each connected component consists of a unique directed cycle with rooted trees hanging off its vertices (the rho shape). Consequences used here: every orbit is eventually periodic; f is good if and only if every cycle meets K; the maps with no cycle inside Y are exactly the forests rooted at Y^c. See Book 1, field: Discrete structures and graphs, topic: functional graphs.

Cayley's formula and rooted forests. The number of labelled trees on N vertices is N^{N-2}; the number rooted at a prescribed vertex is also N^{N-2}; the number of rooted forests on N vertices whose roots are a prescribed set of q vertices is qN^{N-q-1}. Proofs available: Pruefer sequences (bijective), the matrix tree theorem of Kirchhoff (linear algebraic), Joyal's doubly rooted trees (bijective), and the first entry peel used above (recursive). See Book 1, field: Discrete structures and graphs, topic: labelled trees.

The generalised recursion. The recursion extracted here is

$$F(p,q) = \sum_{j=1}^{p} \binom{p}{j} q^j F(p-j,j), \qquad F(0,q) = 1,$$

with solution $F(p,q) = q(p+q)^{p-1}$. Its general form is a binomial convolution $a_p = \sum_j \binom{p}{j} b_j a_{p-j}$, which in exponential generating function language is $A(x) = B(x)A(x)$ style multiplication, that is, the sequence of a labelled structure decomposed into a first layer and a remainder. Whenever a labelled counting problem is cut into a canonical layer plus a smaller copy of itself, you get a binomial convolution, and the two standard finishing moves are (a) recognise the binomial theorem after clearing a linear factor with $j\binom{p}{j} = p\binom{p-1}{j-1}$, or (b) pass to exponential generating functions and use Lagrange inversion when the convolution is genuinely nonlinear. The same peel counts surjections (layer = preimage of a chosen element), parking functions, and forests by number of trees. See Book 1, field: Combinatorics, topic: binomial convolutions.

Binomial identities. $j\binom{p}{j} = p\binom{p-1}{j-1}$ (committee and chair); $\binom{p}{j}\binom{j}{i} = \binom{p}{i}\binom{p-i}{j-i}$ (subset of a subset); and the binomial theorem $\sum_i \binom{m}{i} q^i p^{m-i} = (p+q)^m$, which is the target shape you steer every such sum towards. See Book 1, field: Combinatorics, topic: binomial coefficients.

Random mappings. A uniformly random $f : X \to X$ has about $\sqrt{\pi n/2}$ cyclic points and expected rho length of order $\sqrt{n}$; the statement proved here, that the probability of being absorbed by a k set is exactly k/n, is one of the few facts in this area that is exact rather than asymptotic. See Book 1, field: Probability, topic: random mappings.

Circuit Bank

- **WHEN** a counting condition never reads certain coordinates **DO** factor out the corresponding bare power immediately, after proving the independence with a minimal first entry time **BECAUSE** free coordinates shrink the problem at zero cost, and the proof of freeness is always the same: the successful path stops before it touches them
- **WHEN** a condition says every point eventually reaches the set S **DO** condition on the layer $h^{-1}(S)$, prove it is nonempty, and recurse with S replaced by that layer **BECAUSE** the layer is canonical, so summing over it partitions the objects with no overcounting, and the remainder is the same problem one size down
- **WHEN** your recursion regenerates the problem with a parameter you were holding fixed **DO** promote that parameter to a variable and prove the stronger two parameter statement **BECAUSE** an induction can only close if the hypothesis is broad enough to cover what the recursion produces
- **WHEN** the answer contains n^{n-1}, n^{n-2} or qn^{n-q-1} **DO** look for labelled trees and rooted forests **BECAUSE** those are the only standard labelled counts of that shape, and identifying the structure tells you what the proof must construct
- **WHEN** a binomial sum carries a stray linear factor j **DO** apply $j\binom{p}{j} = p\binom{p-1}{j-1}$ and re-index **BECAUSE** clearing the linear factor is what turns the sum into a plain binomial expansion, and if it does not, your decomposition was the wrong one

Backtesting

1. First useful cue: $j = 0$ is permitted, so the condition is vacuous on $\{1,\dots,k\}$ and those k values of f should be free.
2. Method the expression suggested: pull out n^k, then count absorbed maps $Y \to X$ by peeling the layer that enters K in one step, with the sink size promoted to a free parameter.

3. Where you would hesitate, and the thought to reinforce: at the claim that $f(1), \ldots, f(k)$ are free. Reinforce with the minimal first entry time: the winning path lies entirely in Y before it lands, so it never reads a value on K.
4. The trap the question sets: holding k fixed and inducting on n. The recursion changes the target set, so the hypothesis must carry a free sink size or the induction cannot be applied to its own output.
5. What makes you recognise it faster next time: the shape $kn^{n-1} = \frac{k}{n}n^n$. An exact rational probability plus a Cayley number means forests plus a canonical peel, every time.
6. The reusable decision rule: for reachability conditions, cut at the first entry layer, because it is canonical and it reproduces the problem exactly.

Mirror drills

1. Deduce Cayley's formula from the absorbing lemma. *Answer:* take $q = 1$ and $p = N - 1$; absorbing maps $Y \to Y \cup \{r\}$ are exactly the trees on $Y \cup \{r\}$ rooted at r with edges pointing to the root, so their number is $1 \cdot N^{N-2}$.
2. Count the functions $f : \{1, \ldots, n\} \to \{1, \ldots, n\}$ such that every x has some iterate equal to 1, and check your answer by hand for $n = 2$. *Answer:* n^{n-1}; for $n = 2$ the good maps are $f(2) = 1$ with $f(1)$ either value, so $2 = 2^1$.
3. Prove directly that $\sum_{j=1}^{p} \binom{p}{j} j\, p^{p-j-1} q^j = q(p+q)^{p-1}$. *Answer:* replace $j\binom{p}{j}$ by $p\binom{p-1}{j-1}$, cancel one power of p, substitute $i = j - 1$ and apply the binomial theorem to $(p+q)^{p-1}$.

13.12 B6. Stone game on a line of odd length

Let $n \geq 1$ be an odd integer. Alice and Bob play the following game, taking alternating turns, with Alice playing first. The playing area consists of n spaces, arranged in a line. Initially all spaces are empty. At each turn, a player either

- *places a stone in an empty space, or*
- *removes a stone from a nonempty space s, places a stone in the nearest empty space to the left of s (if such a space exists), and places a stone in the nearest empty space to the right of s (if such a space exists).*

Furthermore, a move is permitted only if the resulting position has not occurred previously in the game. A player loses if he or she is unable to move. Assuming that both players play optimally throughout the game, what moves may Alice make on her first turn?

Read it out loud

Impartial game, normal play, two players, a finite board, and one clause that changes everything: no position may be repeated. Read the pieces separately. The move set is symmetric under left and right reflection, the board has odd length, and the initial position is symmetric. So the only reflection invariant first move is the central space, and your first guess, made in ten seconds and before any analysis, is that the answer is exactly that. Guessing the answer early is not cheating; it tells you what you are trying to prove and therefore which invariant to hunt for.

Now the fields. Combinatorial game theory supplies the vocabulary (a position is an N position if the player to move wins, a P position if the previous player wins), but the Sprague and Grundy theorem is of no use here because this game is not a disjunctive sum of independent components: a type 2 move at one end can reach across the whole board. What is genuinely being tested is whether you can build a numerical invariant on positions such that from every non target value exactly one move restores the target. That is the Nim mechanism stripped of Nim.

The no-repetition clause is the real obstacle and it must be dealt with first, not last. A game whose legal moves depend on the whole history is not a game on positions at all; it is a game on histories, and none of the standard machinery applies to it. So the opening job is to prove that the clause is inert: find a quantity that never decreases, show that it forces every position to be new, and confine the damage to the small set of positions where the quantity is maximal.

Within thirty seconds you should have tabulated what a type 2 move does to the number of stones. Take the maximal block of consecutive stones containing s, and let i and j be the empty spaces flanking it. The move deletes the stone at s and fills i and j, so it changes the count by $+1$ if both flanks are on the board, by 0 if exactly one flank falls off the end, and by -1 if both do, which happens only when the board is full. That table is the whole skeleton of the solution: the number of stones never decreases until the board is nearly full.

Approaches, ranked worst to best

1. **Grundy values on positions with history.** *Cue:* impartial game, so compute Grundy values. *Circuit:* Sprague and Grundy. (worst. The state is a pair consisting of a position and the set of all positions used so far, so the state space has size around 2^{2^n}, and the game is not a sum of components, so the Grundy value of a state tells you nothing you could not read off from a win or lose label. The machinery is real and it is the wrong machinery.)

2. **Brute force small cases and guess.** *Cue:* the answer should be a simple description. *Circuit:* compute $n = 1, 3, 5$ and extrapolate. (this correctly produces the answer, the central space, in a few minutes of hand or machine work. As a solution it is worth nothing, and as reconnaissance

it is worth a great deal: it tells you the answer is unique, which means the invariant you are looking for separates one move from all $n-1$ others.)

3. **Induction on the board length.** *Cue:* a family of games indexed by n. *Circuit:* relate the game on n spaces to smaller boards. (stalls badly. A type 2 move at the end of a block can move stones across a gap, so the two halves of the board never become independent games and there is no smaller board hiding inside.)

4. **Mirror strategy.** *Cue:* odd board, symmetric rules, symmetric start. *Circuit:* take the centre, then reflect every reply. (good and genuinely half of the truth. Alice takes the centre and thereafter restores left and right symmetry; the symmetric positions with the centre occupied are exactly her safe positions. Two defects: proving that the reflected reply is always legal takes real care, especially against a type 2 move that straddles the centre, and the strategy says nothing about why the other $n-1$ first moves lose. A mirror argument proves one move wins; the question asks for all winning moves.)

5. **Level by level descending induction with two pile Nim.** *Cue:* the count of stones is nearly monotone. *Circuit:* analyse P_{n-2} as two pile Nim on the extremal blocks, then induct downwards on the number of stones. (this is the official route and it is correct. The extremal block moves really are Nim moves, and the induction does close. It costs a separate analysis of one level plus an induction, and it obscures the fact that the modular invariant alone already settles every level at once.)

6. **Weight modulo $n+1$ plus the end zone.** *Cue:* the board has n spaces and the type 2 move sometimes loses a stone off the end. *Circuit:* label the spaces 1 to n, add ghost spaces 0 and $n+1$ whose labels are both $\equiv 0$, and take the sum of the occupied labels modulo $n+1$. (best. The ghost convention makes one formula cover all four cases of the move, the increments then realise every nonzero residue exactly once, and the whole game collapses to a single congruence with no induction on levels at all. It also explains the mirror strategy: the symmetric positions with the centre occupied are precisely the ones of weight $(n+1)/2$.)

The ideal solution

The answer: Alice's only winning first move is to place a stone in the central space $\frac{n+1}{2}$.

Number the spaces 1 to n from left to right, and adopt two permanently empty *ghost* spaces labelled 0 and $n+1$ at the two ends. Identify a position with its set of occupied spaces, and let P_i denote the set of positions with i stones. For a stone at s, its *block* is the maximal set of consecutive occupied spaces containing s; write i and j for the empty spaces immediately to the left and right of the block, where $i = 0$ or $j = n+1$ if the block runs to an end of the board. A type 2 move at s deletes the stone at s and places a stone at i and at j, ignoring any ghost. Define the *weight*

$$w(x) \equiv \sum_{s \in x} s \pmod{n+1}, \qquad T = \frac{n+1}{2},$$

and call x *balanced* if $w(x) \equiv T$. Set $Z = P_{n-1} \cup P_n$, the *end zone*, so $\#Z = n+1$.

Assume $n \geq 3$; the case $n = 1$ is treated at the end.

Lemma 13.12.1. *The number of stones never decreases except from the full board, where it drops by one. Precisely: a type 1 move adds a stone; a type 2 move adds a stone if both flanks are on the board, keeps the count if exactly one flank is a ghost, and removes a stone if both flanks are ghosts, which happens only from the full board.*

Proof. Immediate from the description: one stone is removed and one stone is placed for each non-ghost flank. □

Lemma 13.12.2. *Let x have at most $n-2$ stones, and let a and b be the lengths of the maximal blocks of stones at the left and right ends of the board (possibly 0). The moves from x that do not increase the stone count are exactly the type 2 moves at a stone of one of those two extremal blocks. Such a move at $s \le a$ replaces a by $s-1$ and changes nothing else; symmetrically at the right. In particular $a+b$ strictly decreases, so no position can repeat while the stone count stays fixed.*

Proof. By the first lemma, a move keeps the count only if it is type 2 with exactly one ghost flank, that is, at a stone of an extremal block. Take $s \le a$, so the block is $\{1,\dots,a\}$, i is the ghost 0 and $j = a+1$. The new position is x with s emptied and $a+1$ filled, so it agrees with x outside $\{s, a+1\}$ and its left extremal block is $\{1,\dots,s-1\}$, of length $s-1 < a$. Since x has at least two empty spaces, some empty space $e \ge a+2$ survives the move, so the occupied spaces beyond e, and hence the right extremal block, are untouched. □

Lemma 13.12.3. *Every move from a position with at most $n-2$ stones is legal, so such a position always has exactly n legal moves, one for each space.*

Proof. By the first lemma the stone count is nondecreasing until the board is full, and the board is full only after passing through P_{n-1}; so while the count is at most $n-2$ it is nondecreasing. Let $x \in P_i$ with $i \le n-2$ be the current position. Every earlier position has at most i stones, and those with exactly i form a consecutive run ending at x along which $a+b$ strictly decreases, by the second lemma. A move to P_{i+1} produces a position with more stones than any that has occurred; a move inside P_i produces a position with smaller $a+b$ than any that has occurred at this level. Either way the resulting position is new. Finally there are exactly n moves available, one type 1 move per empty space and one type 2 move per occupied space. □

Lemma 13.12.4. *From any position of Z, the positions reachable in one move are exactly the other n positions of Z. Consequently, once the game enters Z it stays there and every move consumes a fresh element of Z, so the player who first moves into Z loses.*

Proof. Write F for the full board. From F there is no type 1 move, and a type 2 move at s has both flanks ghosts, so it produces $F \setminus \{s\}$; as s varies this gives all of P_{n-1}. From $F \setminus \{e\}$ the type 1 move at e gives F, and a type 2 move at $s \ne e$ acts on the block $\{1,\dots,e-1\}$ or $\{e+1,\dots,n\}$, which has exactly one ghost flank and the other flank e, so it produces $F \setminus \{s\}$; as s varies over the $n-1$ occupied spaces this gives every element of P_{n-1} except the current one. So the move graph on Z is complete. By the first lemma no position of Z can occur before the game first enters Z, so from that moment the players simply enumerate distinct elements of Z until none is left. Since $\#Z = n+1$ is even and the entering player supplies the first, third and so on, the entering player is the one facing an exhausted end zone, and loses. □

Lemma 13.12.5. *Fix any position x. For each $t \in \{1,\dots,n\}$ there is exactly one move from x whose weight increment is t modulo $n+1$. No move leaves the weight unchanged.*

Proof. A type 1 move at an empty space e changes the weight by e. A type 2 move at s, with flanks i and j, changes it by $i+j-s$ modulo $n+1$: this is exact when both flanks are on the board, and it remains correct when a flank is a ghost, because the ghost labels 0 and $n+1$ are both $\equiv 0$. This is the reason the modulus is $n+1$ and not something else. Now fix a block $B = \{i+1,\dots,j-1\}$. The map $s \mapsto i+j-s$ sends B bijectively onto itself, so the type 2 moves at the stones of B realise exactly the increments lying in B, each once. The empty spaces and the blocks partition $\{1,\dots,n\}$, and the type 1 moves realise exactly the increments lying in the set of empty spaces, each once. Hence the n moves realise the n nonzero residues bijectively. □

Lemma 13.12.6. *$w(F) \equiv T$, and no position of P_{n-1} is balanced.*

Proof. $w(F) = 1+2+\cdots+n = n\cdot\frac{n+1}{2} \equiv (-1)T \equiv (n+1) - T = T \pmod{n+1}$, using $n \equiv -1$ and $2T = n+1$. Hence $w(F \setminus \{e\}) \equiv T - e$, and $1 \le e \le n$ forces $T - e \not\equiv T$. □

Theorem 13.12.7. *A position outside Z is a loss for the player to move if and only if it is balanced.*

Proof. Let x lie outside Z, so x has at most $n-2$ stones and every move from it is legal, and every resulting position has at most $n-1$ stones and so is not F.

Suppose x is unbalanced. By the increment lemma there is a unique move with increment $T - w(x) \not\equiv 0$; it produces a position y with $w(y) \equiv T$. Then $y \notin P_{n-1}$, since no position of P_{n-1} is balanced, and $y \neq F$, so y lies outside Z and is balanced. So from any unbalanced position outside Z there is a move to a balanced position outside Z.

Suppose x is balanced. Every move changes the weight, so every move produces an unbalanced position, which is either outside Z or in P_{n-1}.

Now let a player face an unbalanced position outside Z and always play the unique restoring move above. That player never moves into Z. The game terminates, because no position ever repeats and there are finitely many positions; and by the third lemma a player at a position outside Z always has a move, so the game can only end inside Z. Hence the opponent is the player who first moves into Z, and by the fourth lemma the opponent loses. So an unbalanced position outside Z is a win for the player to move. A balanced position outside Z hands the opponent either a position in Z, which means the mover entered Z first and loses, or an unbalanced position outside Z, which the opponent then wins; so it is a loss for the player to move. □

The initial position is empty, with $0 \leq n-2$ stones and weight $0 \not\equiv T$, so it is unbalanced and outside Z: Alice wins, and her winning first moves are exactly the moves to a balanced position. From the empty board every move is a type 1 move at a space e, with increment e, so the unique balancing move is $e = T$. Any other first move leaves an unbalanced position with one stone, which lies outside Z because $1 \leq n-2$, and therefore hands Bob the win. Alice's unique winning first move is to place a stone in the central space $\boxed{\frac{n+1}{2}}$.

For $n = 1$ the only first move is to fill the single space, which is the central space, and Bob's only candidate move returns the empty board, which has already occurred; Bob cannot move and loses. (This case needs the reading that the initial empty position counts as having occurred once the game has begun; that reading affects no other n.)

Remark 13.12.8. The invariant is the mirror strategy in disguise. A position symmetric about the centre with the centre occupied has weight T, because the off centre stones pair as s and $n+1-s$ with $s + (n+1-s) \equiv 0$, leaving only the central label T. Alice's rule "restore the symmetry" is a special case of "restore the weight", and the weight version has the advantage that it is defined for every position, needs no case analysis for moves straddling the centre, and shows the other first moves lose.

Layer 1: WHAT

Block 1. Guess the answer from the symmetry of the rules and the parity of n: the centre.

Block 2. Tabulate the effect of each move type on the number of stones, using ghost spaces at 0 and $n+1$ so that all four cases read off one description.

Block 3. Deduce that the count is nondecreasing away from the full board, that level preserving moves only shrink the two extremal blocks, and hence that no position can repeat outside the end zone. The no-repetition clause is now quarantined.

Block 4. Identify the end zone $Z = P_{n-1} \cup P_n$, show its move graph is complete, count $\#Z = n+1$, and conclude that the first player to move into Z loses.

Block 5. Introduce the weight modulo $n+1$ and prove the increment lemma: the n available moves realise the n nonzero increments bijectively.

Block 6. Compute the weights inside the end zone: the full board is balanced, nothing in P_{n-1} is.

Block 7. Conclude: balanced positions outside Z are exactly the losing ones, and the empty board has a unique balancing move, at the centre.

The nominal route reads "impartial game" and starts computing Grundy values, or reads "what may Alice do" and starts playing out small cases hoping for a pattern to prove itself. **The bifurcation point is Block 2: the decision to classify a type 2 move by the two empty spaces flanking its block rather than by the stone it moves.** Everything follows from that reindexing. It makes the stone count table trivial, it produces the formula $i + j - s$ for the weight increment, and it exposes the involution $s \mapsto i + j - s$ that gives the bijection in Block 5. A player who keeps thinking about the stone at s sees an unruly move; a player who thinks about the flanks sees a reflection inside an interval.

Layer 2: HOW

Block 1. **Cue.** Odd n, an empty symmetric start, and rules invariant under reflecting the board. **Why natural.** In a symmetric game the symmetric option is always the first candidate, because after taking it the second player's every reply can be answered by its mirror image. Guessing early is cheap and it fixes your target: you are now looking for a quantity that distinguishes the centre from every other space. **Link to the next block.** A mirror argument alone will not show the other moves lose, so you need an invariant that is defined at every position, not just at symmetric ones.

Block 2. **Cue.** The move description has two parenthetical clauses, "if such a space exists", which is the question flagging its own boundary cases. **Why natural.** When a rule has exceptions only at the ends of the board, invent objects that make the exceptions disappear. Ghost spaces at 0 and $n + 1$ let you say "the move deletes the stone at s and fills the two flanks", with the convention that filling a ghost does nothing. Now there is one rule and the case analysis is bookkeeping. The same instinct adds a point at infinity in projective geometry and an empty product in algebra. **Link to the next block.** With one uniform description, the effect on the stone count is a two line table.

Block 3. **Cue.** The table shows the count can only fall from the full board. **Why natural.** A quantity that never decreases is a potential function, and a potential function is exactly what kills a no-repetition rule: positions at different levels are automatically distinct, so repetition can only happen inside a level. Then look at what a level preserving move does, discover that it only shortens one extremal block, and repetition dies there too. That is the whole reason the clause is harmless. **Link to the next block.** Only the levels where the count can fall remain dangerous, and those are P_{n-1} and P_n.

Block 4. **Cue.** The moves inside the top two levels are unusually rich: from the full board every stone can be removed, and from a board with one gap every stone can be moved to the gap. **Why natural.** Once you notice every position of Z can reach every other in one move, the game inside Z has no structure left: it is the trivial game "name an unused element", whose outcome is decided by the parity of $\#Z$. Here $\#Z = n + 1$ is even, so entering Z is fatal. Notice that this is where the oddness of n enters the argument, and the only place it does. **Link to the next block.** You may now forget the repetition rule entirely and play the equivalent game: no repetitions, and moving into Z loses.

Block 5. **Cue.** A game where you must control who is forced to make the last kind of move, and a board whose spaces are naturally labelled 1 to n. **Why natural.** Sum the labels. Two facts make this the right functional. First, the ghosts are labelled 0 and $n + 1$, which are congruent modulo $n + 1$, so working in $\mathbb{Z}/(n + 1)$ makes the boundary cases obey the interior formula $\Delta \equiv i + j - s$. Second, the symmetric positions Alice wants all have the same weight T, because s pairs with $n + 1 - s$. So the modulus and the target are both forced by the guess in Block 1, and neither is pulled out of the air. Then compute the set of achievable increments and find the bijection: the map $s \mapsto i + j - s$

reflects a block onto itself, so the type 2 moves in a block hit exactly the labels of that block, and the type 1 moves hit exactly the labels of the empty spaces. **Link to the next block.** An invariant with surjective increments always produces a single winning residue, provided the terminal positions behave.

Block 6. Cue. Before declaring the target residue, check what it does at the boundary of your analysis. **Why natural.** The strategy "always restore weight T" is only safe if the restoring move never walks into the end zone. It does not, and for a sharp reason: the only balanced position with many stones is the full board, which cannot be reached from a position with at most $n-2$ stones. That one computation, $w(F) \equiv T$ and $w(F \setminus \{e\}) \equiv T - e$, is what makes the whole strategy legal. **Link to the next block.** All the pieces are in place.

Block 7. Cue. The empty board has weight 0, and $T \neq 0$. **Why natural.** Since the empty board is unbalanced, Alice wins, and the increment lemma says exactly one of her n moves balances it. From an empty board the increments are just the labels, so the balancing move is the type 1 move at T, the centre. Uniqueness is not an extra argument; it is the same lemma read once. **Link to the next block.** Handle $n = 1$ separately and stop.

Layer 3: WHY

Three truths are doing the work.

The first is the standard characterisation of losing positions in an impartial normal play game: a set $\mathcal{L}$ of positions is the set of losses for the player to move, provided every move from a position of $\mathcal{L}$ leaves $\mathcal{L}$, and from every position outside $\mathcal{L}$ some move enters $\mathcal{L}$, and every terminal position lies in $\mathcal{L}$. You never need to compute anything recursively if you can guess $\mathcal{L}$ and verify those three closure conditions. Here $\mathcal{L}$ is a single congruence class, and the verification is the increment lemma plus the end zone parity. This is exactly the mechanism of Nim, where $\mathcal{L}$ is the set of positions of zero binary digit sum, and of every other game that turns out to be governed by a modular invariant.

The second is that a no-repetition rule is only dangerous where the position graph has cycles. Find a potential function that never decreases, and cycles can only live inside a level; kill those by finding a second quantity that strictly decreases inside a level. What is left is a subgraph where cycles are unavoidable, here the end zone, and in a strongly connected terminal set of size N with no repetition the game is pure parity. Any game with a Ko style rule should be attacked in that order: potential, then level structure, then the residual cycles.

The third is the reason the modulus is $n+1$. A type 2 move conserves the sum of labels up to the two flanks, and the flanks are labelled 0 and $n+1$ when they fall off the board. A conservation law that is exact in the interior and off by $n+1$ at the boundary is an exact conservation law modulo $n+1$. Whenever a local rearrangement rule conserves something except for a fixed boundary defect, quotient by that defect and the rule becomes clean. That is the single most transferable idea in this problem.

The discernment mechanism, stated as a protocol you can run on any unfamiliar game. (1) Test the symmetry of the rules and the start; guess the answer. (2) Find a monotone potential and use it to quarantine any history dependent rule. (3) Identify the terminal region and settle it by counting, not by strategy. (4) Label the cells and compute the set of increments of the label sum from a general position; if the increments realise every nonzero residue modulo some m exactly once, the losing positions form one residue class and you only have to find which one. (5) Check that the winning move never falls into the terminal region.

Honest knowledge gaps. Before this problem is even approachable you need: the language of impartial games under normal play and the three closure conditions above; the fact that Sprague and Grundy applies to disjunctive sums and that this game is not one, so Grundy values buy

nothing; two pile Nim, which is where the official solution routes its level P_{n-2} analysis and which is worth knowing even though the route above avoids it; and the notion of a loopy game, meaning a game whose move graph has cycles, together with the fact that the standard theory assumes a directed acyclic graph and that the no-repetition rule is what restores acyclicity here. None of that is in the A level or Further Maths syllabus and none of it is hard; it is simply a body of definitions that must be installed before the reflexes can fire.

Revolving concepts

Impartial games, normal play, N and P positions. Both players have the same options; the player unable to move loses. $\mathcal{P}$ is the set of positions losing for the player to move, characterised by: terminal positions lie in $\mathcal{P}$; every option of a $\mathcal{P}$ position lies outside $\mathcal{P}$; every position outside $\mathcal{P}$ has an option in $\mathcal{P}$. Exhibiting $\mathcal{P}$ and checking these three is always cheaper than a recursion. See Book 1, field: Games and invariants, topic: normal play.

Nim and Sprague and Grundy. Nim with piles $a_1, \dots, a_r$ is a loss for the mover exactly when $a_1 \oplus \cdots \oplus a_r = 0$; every impartial game position has a Grundy value, and the value of a disjunctive sum is the binary digit sum of the values. Two pile Nim, the case used in the official route, has $\mathcal{P}$ equal to the equal pile positions. The relevance here is negative and worth internalising: this game is not a sum of independent components, so Grundy theory does not apply, and recognising that saves a wasted half hour. See Book 1, field: Games and invariants, topic: Nim.

Modular invariants with surjective increments. Suppose φ maps positions to $\mathbb{Z}/m$ and from every non terminal position each nonzero increment is realised by exactly one move. Then no move fixes φ, and every value is reachable from every other in one move, so $\mathcal{P}$ is a single class $\varphi^{-1}(T)$ as soon as the terminal behaviour selects T. This is the abstract engine of the solution, and of subtraction games, Mock Turtles style games, and many coin and token games on a labelled strip. See Book 1, field: Games and invariants, topic: modular invariants.

Potentials, loopy games and the no-repetition rule. A game whose move graph has cycles is called loopy and has draws in general; adding a no-repetition rule makes the state space the set of histories and restores finiteness. The technique used here is the standard one: find a monotone potential (stone count), then a strictly decreasing secondary quantity within a level ($a + b$), which together force the history graph to be acyclic except on a small terminal set. The residual recursion is not a sequence recursion but a descending induction over the levels $P_{n-2}, P_{n-3}, \dots$; the general form of that recursion is "grade the state space by a monotone potential, solve the top grade directly, and induct downwards", and the invariant here is exactly what lets you solve all grades at once instead. See Book 1, field: Games and invariants, topic: loopy games.

Terminal cliques and parity. If a set Z of positions is closed under moves, mutually reachable in one move, and subject to no repetition, the game inside Z is "name an unused element of Z" and is decided by the parity of $\#Z$: with $\#Z$ even the player who enters Z loses, with $\#Z$ odd that player wins. This is where the hypothesis that n is odd is consumed, and it is also why the even case is a genuinely different and still unresolved problem. See Book 1, field: Combinatorics, topic: parity arguments.

Symmetry and mirror strategies. On a board with a reflection symmetry fixing a unique cell, the first player takes the fixed cell and thereafter mirrors, provided mirrored replies are always legal. The weight invariant here is the arithmetic shadow of that strategy, since central symmetry with the centre occupied is equivalent to weight T in the cases where both make sense. See Book 1, field: Games and invariants, topic: strategy stealing and symmetry.

Circuit Bank

- **WHEN** a game forbids repeating a position **DO** find a quantity that never decreases, then a second quantity that strictly decreases whenever the first is constant **BECAUSE** together

they make the history graph acyclic, so the clause is inert except where the first quantity is maximal, and that residue is small enough to count

- **WHEN** a set of positions is closed under moves and mutually reachable **DO** stop doing strategy and count the set **BECAUSE** with no repetition the play inside such a set is a bare enumeration, so only the parity of its size matters
- **WHEN** a move rule has exceptions at the ends of the board **DO** introduce ghost cells at 0 and $n+1$ and choose the modulus $n+1$ **BECAUSE** the ghost labels are then congruent to zero, so one formula covers interior and boundary and the conservation law becomes exact
- **WHEN** the rules and the starting position are symmetric and the board has a unique fixed cell **DO** guess that the fixed cell is the winning move and look for an invariant that is constant on symmetric positions **BECAUSE** a mirror strategy proves that move wins, and its arithmetic shadow proves the others lose
- **WHEN** from every position each nonzero increment of a labelling is realised by exactly one move **DO** declare the losing positions to be one residue class and verify the terminal condition **BECAUSE** that is the Nim mechanism, and the only remaining work is to identify the class and to check the restoring move never lands in the terminal region

Backtesting

1. First useful cue: the two clauses "if such a space exists" flag the boundary, so index a type 2 move by the empty spaces flanking its block, with ghosts at 0 and $n+1$.
2. Method the expression suggested: the stone count table, then a weight given by the sum of occupied labels modulo $n+1$, with target $(n+1)/2$.
3. Where you would hesitate, and the thought to reinforce: at the no-repetition clause. Reinforce with the protocol potential first, level structure second, residual cycles third, so that the clause is quarantined before any strategy is attempted.
4. The trap the question sets: treating this as a disjunctive sum and reaching for Grundy values, or proving the mirror strategy and thinking the question has been answered when it asked for all winning first moves.
5. What makes you recognise it faster next time: a local rearrangement rule that conserves a label sum except for a fixed defect at the boundary is a conservation law modulo that defect.
6. The reusable decision rule: if each nonzero increment of an invariant is realised exactly once from every position, the losing positions form a single residue class; find it by evaluating the invariant on the terminal region.

Mirror drills

1. Play the game with $n = 3$ by hand and confirm that space 2 is the unique winning first move. *Answer:* after Alice plays 2 every one of Bob's three moves reaches a two stone position, and $Z = P_2 \cup P_3$ with $\#Z = 4$, so Bob has entered the end zone first and loses; after Alice plays 1 or 3, Bob replies at the centre by the same reasoning.
2. Show that any position symmetric about the centre with the central space occupied has weight $(n+1)/2$, and that any symmetric position with the centre empty does not. *Answer:* pair s with $n+1-s$, each pair contributing $n+1 \equiv 0$; the centre contributes $(n+1)/2$ if occupied and 0 if not.

3. Identify exactly where the argument fails when n is even, and say what is known. *Answer:* $\#Z = n + 1$ becomes odd, so the first player into the end zone now wins and every position of P_{n-2} is a win; the weight argument collapses and must be replaced by a different invariant. Savitt conjectured, on the evidence of $n = 2$ and $n = 4$, that Alice still wins with either central space, and only a partial analysis is known.

Chapter 14

Putnam 2012

14.1 A1. an acute triple forced by Fibonacci growth

Let $d_1, d_2, \ldots, d_{12}$ be real numbers in the open interval $(1, 12)$. Show that there exist distinct indices i, j, k such that d_i, d_j, d_k are the side lengths of an acute triangle.

Read it out loud

Twelve real numbers, each confined to $(1, 12)$, and a claim that some three of them, in some order, form an acute triangle. The command term is "there exist": an existence claim, not a "for all", which means proof by contradiction is on the table from the first read, since contradicting a universal negation is usually cheaper than constructing a witness by hand. The fields wired into this problem are elementary geometry (what makes a triangle acute, given only its side lengths) and sequences (a forced growth argument). What is really being tested is whether you notice that twelve objects and a range factor of twelve, next to each other, are not a coincidence.

Within thirty seconds you should have written down two things. First: acuteness only depends on the relative sizes of the three chosen numbers, so sorting the twelve values costs nothing and turns an unordered search into an ordered one. Second: twelve numbers means $\binom{12}{3} = 220$ possible triples, and no fluent solution checks 220 cases, so whatever route you take must collapse that count drastically, most likely by restricting attention to a much smaller, structured sub-family and letting an induction do the rest.

The number 12 doing double duty, once as the count of terms and once as the upper bound of the interval, is the loudest cue on the page. Anyone who has drilled sequences knows that $F_{12} = 144 = 12^2$ for the Fibonacci numbers $F_1 = F_2 = 1$, $F_{n+2} = F_n + F_{n+1}$; seeing a count of exactly twelve objects bounded by a range whose square is a suspiciously round number should fire the reaction circuit for a growth-outpaces-the-range contradiction before you have written a single inequality.

Approaches, ranked worst to best

1. **brute enumeration.** *Cue:* "there exist distinct indices i, j, k". *Circuit:* casework over all triples. (checking all 220 triples individually against the cosine rule offers no structural insight, costs enormous ink for a ten point problem, and is exactly the kind of grinding this problem is designed to punish).
2. **sort, then search without a growth bound.** *Cue:* sorting makes "largest side" well defined for every triple. *Circuit:* ordered casework. (sorting is a real improvement, since it tells you instantly which of the three sorted values plays the role of the longest side in any candidate triple, but without noticing the Fibonacci substructure you are still hunting through an unstructured set of comparisons with no closed reason to stop early; better than raw enumeration, still far more ink than necessary).
3. **sort, negate on consecutive triples only, induct.** *Cue:* the sorted list turns "no acute triple" into a two-term recursive inequality. *Circuit:* Fibonacci domination. (the ideal route: assume for contradiction that even the ten cheapest candidate triples, the consecutive ones in sorted order, all fail to be acute; this assumption is a Fibonacci-shaped recursion in disguise, and a two-line induction turns it into a single numerical contradiction against the given bounds).

The ideal solution

Relabel the twelve given numbers so that $d_1 \leq d_2 \leq \cdots \leq d_{12}$. This is legitimate because the conclusion only asks for the existence of three distinct indices among the original twelve whose values form an acute triangle; a permutation of labels changes nothing about which triples of values are available, so it suffices to find three consecutive positions in the sorted list.

First fix exactly what "acute triangle" costs you to verify. Suppose $a \le b \le c$ are positive reals with $c^2 < a^2 + b^2$. Since $(a+b)^2 = a^2 + 2ab + b^2 \ge a^2 + b^2 > c^2$ and $a + b > 0$, this already gives $a + b > c$, so the triangle inequality holds automatically; the other two instances of it, $a + c > b$ and $b + c > a$, are immediate since $a, b, c > 0$ and c is the largest. So a, b, c genuinely are the side lengths of a triangle. In that triangle, c is the longest side, hence opposite the largest angle γ; by the cosine rule, $\cos\gamma = \dfrac{a^2 + b^2 - c^2}{2ab} > 0$, so γ is acute. Since the three angles of a triangle sum to π, at most one of them can be at least $\pi/2$, and that one, if it existed, would have to be the largest; since the largest angle γ is already acute, all three are. Conversely if $c^2 \ge a^2 + b^2$ the triangle is right angled or obtuse at γ. So for positive reals $a \le b \le c$: they form an acute triangle exactly when $c^2 < a^2 + b^2$, one inequality, nothing more to check.

Suppose, for contradiction, that no three of $d_1, \ldots, d_{12}$ are the side lengths of an acute triangle. In particular no three consecutive terms of the sorted list are, so for every $i = 1, \ldots, 10$,

$$d_{i+2}^2 \ge d_i^2 + d_{i+1}^2.$$

Let $F_1 = F_2 = 1$ and $F_{n+2} = F_n + F_{n+1}$ be the Fibonacci numbers. Claim: $d_i^2 \ge F_i\, d_1^2$ for every $i = 1, \ldots, 12$. For $i = 1$ this reads $d_1^2 \ge d_1^2$. For $i = 2$, $d_2 \ge d_1 > 0$ gives $d_2^2 \ge d_1^2 = F_2 d_1^2$. For the inductive step, if $d_i^2 \ge F_i d_1^2$ and $d_{i+1}^2 \ge F_{i+1} d_1^2$ then

$$d_{i+2}^2 \ge d_i^2 + d_{i+1}^2 \ge (F_i + F_{i+1}) d_1^2 = F_{i+2} d_1^2,$$

so the claim holds for all $i \le 12$ by induction on i. Taking $i = 12$, and using $F_{12} = 144$,

$$d_{12}^2 \ge 144\, d_1^2.$$

Since $d_1 \in (1, 12)$, $d_1 > 1$ strictly, so $d_1^2 > 1$ and $d_{12}^2 > 144$, giving $d_{12} > 12$. But $d_{12} \in (1, 12)$ forces $d_{12} < 12$. This is a contradiction.

Hence the assumption fails: some three consecutive terms of the sorted list, and therefore some three of the original $d_1, \ldots, d_{12}$ at three distinct original indices i, j, k, are the side lengths of an acute triangle. ■

Layer 1: WHAT

1. Sort and relabel the twelve values into increasing order.
2. Reduce "acute triangle" to a single inequality: largest side squared less than the sum of the squares of the other two, with the triangle inequality automatic.
3. Assume, for contradiction, only the weakest necessary sub-claim: that no three *consecutive* sorted values are acute.
4. Recognise that assumption as a Fibonacci-shaped recursive inequality and induct to $d_{12}^2 \ge 144\, d_1^2$.
5. Plug in the two given numeric bounds, $d_1 > 1$ and $d_{12} < 12$, to close the contradiction.

The nominal route enumerates some or all of the $\binom{12}{3} = 220$ triples directly, checking the cosine-rule inequality case by case with no structural reduction. The ideal route restricts the negation to the ten consecutive triples in sorted order and lets a single two-line induction absorb all ten checks at once. The bifurcation point sits immediately after sorting: it is the decision to assume the negation only on consecutive triples, a strict sub-family of all triples, rather than on every triple. That one choice is what turns two hundred and twenty cases into one Fibonacci induction of length ten, since a contradiction obtained from the smaller sub-family already proves the larger claim false as well.

Layer 2: HOW

Block 1. **Cue.** Twelve unordered real numbers, and a target property (acuteness) that depends only on which of three chosen values is largest, not on their original positions. **Why natural.** Sorting costs nothing precisely because the property you are hunting for is invariant under relabelling; once sorted, any candidate triple instantly tells you which value plays the role of the longest side, without further casework. **Link to the next block.** With a sorted list in hand, the cheapest candidate triples to examine are neighbours, since they require no extra bookkeeping about which positions to compare.

Block 2. **Cue.** The problem hands you only side lengths and the word "acute", nothing about angles directly. **Why natural.** The cosine rule converts an angle condition into a polynomial inequality in the squared side lengths, and because the three angles of a triangle sum to π, only the largest angle can possibly fail to be acute; testing all three angles collapses to testing one inequality on the longest side. **Link to the next block.** With a single inequality test available, "no acute triple exists" translates directly into a chain of numerical inequalities that can be chased by induction.

Block 3. **Cue.** You need a contradiction from "no acute triple exists" among all 220 triples, but the sorted list hands you ten much cheaper candidates for free: the consecutive ones. **Why natural.** Assuming the negation on the smallest sufficient sub-family, rather than on every triple, is the ink-saving move: if a contradiction already follows from the weaker assumption (failure on consecutive triples only), the stronger claim (failure on all triples) is impossible too, for free, since consecutive triples are a subset of all triples. **Link to the next block.** Written out, the negation on ten consecutive triples is exactly the inequality $d_{i+2}^2 \geq d_i^2 + d_{i+1}^2$ for $i = 1, \ldots, 10$: a recursion in disguise.

Block 4. **Cue.** An inequality of the shape "term $i+2$ is at least the sum of the two before it", repeated ten times. **Why natural.** Whenever an inequality has this shape, the sequence of terms is forced to grow at least as fast as the Fibonacci numbers, by an immediate two-line induction; recognising the shape tells you the growth rate, and hence the bound you are heading toward, without inventing anything from scratch. **Link to the next block.** Once $d_{12}^2 \geq 144\, d_1^2$ is established, the only work left is to substitute the two numeric bounds the problem actually gave you on d_1 and d_{12}.

Block 5. **Cue.** The problem's interval is $(1, 12)$, and $144 = 12^2$ is exactly F_{12}. **Why natural.** This is not a coincidence: the numbers were chosen so the Fibonacci bound at $i = 12$ lands exactly on the boundary 12^2, leaving no room to spare, which is also why twelve numbers are needed and not eleven (the same argument with eleven numbers only forces $d_{11} > \sqrt{89} \approx 9.43$, no contradiction against an interval of width twelve). **Link to the next block.** This closes the proof: the assumption that no acute triple exists is false, so three of the original twelve numbers, at three distinct indices, form one.

Layer 3: WHY

The truth underneath the route is that any two-term recursive inequality $x_{n+2} \geq x_n + x_{n+1}$ among positive terms grows at least geometrically, at the golden ratio rate $\phi = (1 + \sqrt{5})/2$, because the Fibonacci numbers themselves grow like $\phi^n/\sqrt{5}$. This problem is one instance of a wide family: a local avoidance condition, assumed everywhere, forces a monotone quantity to grow at a fixed exponential rate; make the given range too short to sustain that growth for as many steps as you have terms, and avoidance becomes impossible.

The discernment mechanism is this: whenever a problem gives you n bounded quantities together with a numeric bound whose ratio (upper over lower, or upper alone against a natural

scale) is suspiciously close to n itself, suspect that n has been tuned to a named growth sequence, Fibonacci, powers of two, factorials, evaluated near n. The setter needs the forced growth to exactly outrun the given range at the given count, no earlier and no later, and that exactness is the strongest possible hint that a growth-versus-range contradiction is the intended engine, not a red herring. Train this by asking, whenever a problem says "distinct indices" and "show there exist" about n bounded reals: what happens if I sort, then assume the negation locally, on neighbours only, rather than on every subset; does that generate a recursion I recognise?

Revolving concepts

Acute triangle from side lengths: for positive reals $a \le b \le c$, they form an acute triangle exactly when $c^2 < a^2 + b^2$, via the cosine rule $c^2 = a^2 + b^2 - 2ab\cos\gamma$ and the fact that the largest angle is opposite the largest side; since at most one angle in a triangle can be at least $\pi/2$, one inequality on the longest side decides all three angles at once, and the triangle inequality is then automatic. See Book 1, field Geometry, topic triangle classification by side lengths.

Fibonacci recursion and its generalisation: $F_1 = F_2 = 1$, $F_{n+2} = F_n + F_{n+1}$, with closed form (Binet) $F_n = (\phi^n - \psi^n)/\sqrt{5}$ where $\phi, \psi = (1 \pm \sqrt{5})/2$ solve $t^2 = t + 1$, giving geometric growth at rate $\phi \approx 1.618$. The recursion generalises directly: for any inequality $x_{n+2} \ge \alpha x_n + \beta x_{n+1}$ with $\alpha, \beta > 0$ and positive initial terms, the same two-line induction shows x_n grows at least like the dominant positive root of $t^2 = \beta t + \alpha$; Fibonacci is the case $\alpha = \beta = 1$. Any linear recursion of higher order $x_{n+k} = c_1 x_{n+k-1} + \cdots + c_k x_n$ is handled the same way, via the characteristic polynomial $t^k = c_1 t^{k-1} + \cdots + c_k$ and its root of largest absolute value. See field Sequences and recurrences, topic linear recurrence lower bounds by induction.

Extremal ordering: sorting a finite unordered collection so an order-dependent or symmetric claim becomes tractable, a standing technique across combinatorics whenever a statement does not depend on the original labelling. See field Combinatorics, topic the extremal principle.

Sufficient sub-family reroute: to prove existence of a k-subset with a property among n objects, it suffices to show the property holds somewhere in a smaller, structured sub-family, here the consecutive windows of a sorted list, since that sub-family is contained in the full collection of k-subsets. See field Combinatorics, topic search-space reduction for existence proofs.

Circuit Bank

- **WHEN** a problem gives n bounded reals in an interval and asks you to find k of them satisfying an order-dependent inequality, such as the acute triangle condition **DO** sort them first and rename in increasing order **BECAUSE** the target property depends only on relative size, so sorting is free and turns "which k-subset" into "which window of the sorted list"
- **WHEN** you want to contradict "no such k-tuple exists" **DO** apply the assumption only to consecutive windows of the sorted list, not to every k-subset **BECAUSE** a contradiction from the smaller, structured sub-family already proves the full statement, and consecutive windows are exactly what turn into a one-parameter recursion
- **WHEN** a recursive inequality of the shape $x_{i+2} \ge x_i + x_{i+1}$ appears among positive terms, or any order-two linear recursion **DO** identify it immediately as at least Fibonacci growth and induct $x_i \ge F_i x_1$, or the analogous bound for general coefficients **BECAUSE** Fibonacci-type recursions grow geometrically at the golden ratio rate, and this rate is almost always the engine behind whatever numeric bound the problem handed you
- **WHEN** you must decide whether a triangle with given side lengths is acute **DO** compute only the longest side squared against the sum of the squares of the other two **BECAUSE** the largest angle is opposite the longest side, at most one angle in a triangle can be non-acute, and the triangle inequality follows automatically once that one inequality holds

Backtesting

1. First useful cue: twelve numbers, each confined to $(1, 12)$, a range factor of twelve sitting next to a count of twelve.
2. Method the expression suggested: an order-statistic argument, sort then bound growth, since the target property is entirely about relative size.
3. Where you would hesitate, and the thought to reinforce: the moment you decide to assume the negation only on consecutive triples rather than on all 220; reinforce that a contradiction from a sub-family is still a contradiction for the whole family.
4. The trap the question sets: spending a line verifying the triangle inequality separately from the acuteness inequality, when it is automatic once the longest side squared is less than the sum of the other two squared.
5. What makes you recognise it faster next time: a count of objects matching an interval's ratio bound exactly, twelve numbers against a ratio of twelve, is a fingerprint of a hidden Fibonacci or geometric-growth argument.
6. The reusable decision rule: when a bounded existence claim about k-tuples from n reals appears, sort, restrict the negation to consecutive windows, and compare the induced recursion's growth rate against the given numeric bound.

Mirror drills

1. Show that among any six positive reals in the open interval $(1, \sqrt{8})$, three form an acute triangle. *Answer:* identical argument with $n = 6$: sort, assume no three consecutive are acute, induct to $d_6^2 \geq F_6 d_1^2 = 8d_1^2 > 8$, contradicting $d_6 < \sqrt{8}$.
2. Explain in one line why the same conclusion fails to follow from only eleven numbers in $(1, 12)$. *Answer:* the induction only reaches $d_{11}^2 \geq F_{11} d_1^2 = 89 d_1^2 > 89$, giving $d_{11} > \sqrt{89} \approx 9.43$, which does not contradict $d_{11} < 12$.
3. Twelve reals lie in $(1, 12)$, all equal to 6. Does the argument produce a contradiction, and what does this tell you about the sharpness of the bound? *Answer:* no contradiction is needed; every triple is equilateral, hence acute, since equal side lengths trivially satisfy $c^2 < a^2 + b^2$; this shows the bound is not violated by a "flat" example, only genuinely spread-out, near-Fibonacci-ratio examples come close to avoiding an acute triple.

14.2 A2. forcing cancellation without identity or inverses

Let $*$ be a commutative and associative binary operation on a set S. Assume that for every x and y in S, there exists z in S such that $x * z = y$. (This z may depend on x and y.) Show that if a, b, c are in S and $a * c = b * c$, then $a = b$.

Read it out loud

No ambient structure is given beyond three axioms on an otherwise bare set S. "Commutative and associative" is a command term telling you that you may rebracket and reorder any expression built from $*$ freely, forever, without further comment, exactly as with ordinary multiplication. The solvability clause, for every x, y there exists z with $x * z = y$, is a disguised surjectivity statement: fixing x, the map $z \mapsto x * z$ hits every element of S. The field wired into this problem is abstract algebra, specifically the theory of cancellative and invertible structures sitting just below group theory.

What is really being tested is whether you can recognise that surjectivity of every one of these translation maps, together with associativity and commutativity, is already enough to force cancellation on one specific pair of elements, without first reconstructing the whole group structure the remark afterward promises exists. The goal equation, $a * c = b * c \Rightarrow a = b$, is literally asking you to cancel c. Since S has no named identity or inverses to cancel with directly, your only tool is the solvability axiom, so within thirty seconds you should have written down, symbolically, that solvability lets you solve $x * z = y$ for a fresh z whenever you choose x and y, and started hunting for the two applications of it that will bridge a and b.

One phrase deserves a second look: "this z may depend on x and y." That is a warning, not decoration. It tells you not to expect a single uniform witness reusable across every equation you write; each application of the axiom earns you one disposable helper element, valid only through the one equation it satisfies, and the proof must be built so that it never needs more from that helper than exactly that equation.

Approaches, ranked worst to best

1. **concrete guessing.** *Cue:* "binary operation on a set S". *Circuit:* testing familiar operations. (trying the claim on $\mathbb{Z}$ under addition or the positive reals under multiplication to guess the mechanism never becomes a proof for an abstract S and abstract $*$; every minute here is unrecoverable in the arena).

2. **build the full group structure first.** *Cue:* commutative, associative, always solvable reads like the group axioms minus a named identity. *Circuit:* the full group theorem. (prove an identity element exists and is unique, then that every element has an inverse, then cancel using ordinary group cancellation; this is valid, the remark confirms S is in fact an abelian group, but it proves a strictly stronger global theorem than the problem asks for, costing far more ink for a result that only needs the pair a, b reconciled).

3. **local bootstrap via two witnesses.** *Cue:* $a * c = b * c$ names a single shared value. *Circuit:* surjectivity of translations, applied exactly twice. (the ideal route: name the shared value $d = a * c = b * c$, solve $d * e = a$, build $f = c * e$, which is simultaneously a right identity for a and for b, then solve $g * a = b$ to close the loop; no global structure is ever proved, matching exactly what the problem asks and nothing more).

The ideal solution

Write d for the common value $a * c = b * c \in S$.

By hypothesis, applied to the pair $(x, y) = (d, a)$, there exists $e \in S$ with

$$d * e = a.$$

Set $f = c * e \in S$. Using associativity and then $a * c = d$,

$$a * f = a * (c * e) = (a * c) * e = d * e = a.$$

Using associativity again, and $b * c = d$, the same d as above,

$$b * f = b * (c * e) = (b * c) * e = d * e = a.$$

So f satisfies both $a * f = a$ and $b * f = a$.

By hypothesis, applied to the pair $(x, y) = (a, b)$, there exists $g \in S$ with

$$g * a = b.$$

Using $a = a * f$ from above, then associativity, then $g * a = b$,

$$b = g * a = g * (a * f) = (g * a) * f = b * f.$$

But $b * f = a$ was already shown. Hence $b = a$, as required. ■

It is worth repeating that this argument never constructs a global identity for S, only the single element f, built exclusively from c and the witness e, that happens to fix both a and b. A caveat worth clarifying: with more work one can show that the same f fixes every element of S, giving S a genuine two-sided identity, and that every element then has an inverse under $*$, so $(S, *)$ is in fact an abelian group; the problem only asks for the cheapest consequence of that structure, extracted without paying for the rest of it.

Layer 1: WHAT

1. Name the shared value $d = a * c = b * c$.
2. Solve $d * e = a$ using solvability: the first witness.
3. Build $f = c * e$; associativity gives $a * f = a$.
4. Run the identical computation with b in place of a, using $b * c = d$ as well, giving $b * f = a$.
5. Solve $g * a = b$ using solvability: the second witness.
6. Chain $b = g * a = g * (a * f) = (g * a) * f = b * f = a$.

The nominal route reconstructs the full group structure of S, a global identity followed by inverses for every element, then invokes ordinary group cancellation. The ideal route builds one local bridging element f tailored only to a, b, and c, using exactly two applications of the solvability axiom. The bifurcation point sits right after writing $d = a * c = b * c$: the choice either to chase a global identity valid for all of S, or to chase a local identity that only needs to work for the pair a, b. The ideal route is the realisation that the problem only ever needs one pair of elements reconciled, not the whole set.

Layer 2: HOW

Block 1. **Cue.** $a * c = b * c$: two different left-hand sides producing the identical element of S. **Why natural.** Naming that common output as a single symbol d is natural because everything you know about a and about b, for the rest of the proof, only enters through how each relates to d; collapsing two equations into one named quantity is what lets you treat a and b symmetrically in every later step. **Link to the next block.** Having a single target d, the next question is whether you can undo d back to a using the only tool the problem gives you, which is exactly the solvability axiom applied with $x = d, y = a$.

Block 2. **Cue.** The solvability axiom licenses a witness for any equation $x * z = y$ you choose to write down, for any $x, y \in S$; you have just produced a specific target, $d * z = a$. **Why natural.** Choosing $(x, y) = (d, a)$ is natural because d is exactly the quantity you want to strip away to recover a, and the axiom is licence to name a witness for that without knowing anything else about S. **Link to the next block.** The witness e only becomes useful once combined with the leftover factor c, since c is what is still attached to a and to b inside d.

Block 3. **Cue.** You have e with $d * e = a$, and c is the common factor inside $d = a * c = b * c$. **Why natural.** Forming $f = c * e$ and expanding $a * f = a * (c * e) = (a * c) * e = d * e = a$ uses only associativity, licensed freely by hypothesis; the reason to build f from exactly c and e, rather than any other combination, is that it is the unique combination that will let the identical computation run again with b in place of a. **Link to the next block.** Since f was built without ever referring to a specifically, only to c and e, the identical computation should run again verbatim with b.

Block 4. **Cue.** $b * c = d$ also, by the given hypothesis $a * c = b * c$. **Why natural.** Running the identical three-line computation with b in place of a, $b * f = b * (c * e) = (b * c) * e = d * e = a$, is natural precisely because f was deliberately built from c and e alone, so it treats a and b identically; this is the moment the shared hypothesis $a * c = b * c$ earns its keep a second time. **Link to the next block.** Two facts now sit side by side, $a * f = a$ and $b * f = a$; the goal $a = b$ will fall out the instant $b * f$ can be related back to b itself, not only to a.

Block 5. **Cue.** You want to turn "$b * f = a$" into "$b = a$", and your only tool remains the solvability axiom. **Why natural.** Applying it with $(x, y) = (a, b)$ produces g with $g * a = b$: a second, independent bridge, this time from a to b rather than from d to a; this pair is natural because b is the element you are trying to reach, and a is the element over which you already have full control, through f. **Link to the next block.** Substituting $a = a * f$ inside $g * a$ is the only way to bring f, and hence the fact $b * f = a$, into contact with $g * a = b$.

Block 6. **Cue.** $b = g * a$ on one hand, $a = a * f$ on the other. **Why natural.** Substituting the second into the first and re-associating, $b = g * a = g * (a * f) = (g * a) * f$, replaces the inner $g * a$ with b, by definition of g, purely through associativity, landing on $b * f$, a quantity already evaluated in Block 4 to equal a; the whole construction was aimed so that this final substitution has nowhere else to go but $a = b$. **Link to the next block.** None; this closes the proof.

Layer 3: WHY

The solvability hypothesis says every translation map $L_x : S \to S$, $L_x(z) = x * z$, is surjective. Associativity and commutativity make $*$-expressions manipulable exactly like ordinary multiplication, regrouping and reordering always licensed once stated. The proof shows that surjectivity of translations, together with associativity and commutativity, forces injectivity of translations too, which is exactly what cancellation means: L_c injective, even though injectivity was never assumed.

The discernment mechanism: whenever a hypothesis gives existence of a solution to an equation, without uniqueness, inside an abstract algebraic structure, and the goal is a cancellation or uniqueness statement, the standard reroute is to apply the existence hypothesis exactly twice, choosing the two input pairs so the resulting witnesses compose back into the two sides of the equation you are trying to equate. You are not solving for the witnesses explicitly, since in general you cannot; you are only ever using the one defining equation each witness satisfies. Train this by noticing the pattern in the problem: the goal equation ($a = b$) and the hypothesis equation ($a * c = b * c$) differ by exactly one factor of c, and you then look for a way to attach and detach that c using the existence axiom, which is precisely what e and f do. The phrase "this z may depend on x and y"

in the problem statement is a signpost that you must not expect uniform witnesses, only ad hoc ones for the specific pairs you choose, and that is exactly why the proof never names e or g except through the one equation each satisfies.

Revolving concepts

Translation maps and surjectivity in a magma: for $(S, *)$, $L_x(z) = x * z$; the hypothesis states L_x is surjective for every x. See field Abstract algebra, topic translations in a groupoid.

Generalised associativity and commutativity as bookkeeping licence: once associativity holds for three terms, any bracketing of a fixed sequence of elements under $*$ agrees, provable by induction on the number of terms; commutativity additionally frees the order. Used four times in this proof without further comment. See field Abstract algebra, topic semigroup bracketing.

Existential witnesses used opaquely: when a hypothesis only guarantees existence, the standard technique is to name the object and use only its defining equation afterward, never attempting to solve for it explicitly; this is how e and g are used throughout, neither ever computed. See field Logic and set theory, topic existential elimination in proofs.

Commutative, associative, everywhere solvable implies abelian group: construct an identity as any solution u of $u * a = a$ for a fixed a, show by an argument in the same spirit as this problem that u fixes every element of S, then produce inverses similarly. Cancellation, the content of this problem, is the cheapest consequence of that structure, extracted without paying for the rest of it. See field Abstract algebra, topic quasigroups, Latin squares, and when they are groups.

Circuit Bank

- **WHEN** a hypothesis guarantees a solution to $x * z = y$ exists for every x, y, but not uniquely **DO** name the witness with a fresh letter and use only the one equation it satisfies in every later step **BECAUSE** the witness cannot be solved for explicitly in an abstract structure, so the proof must stay entirely inside the algebra generated by the defining equations
- **WHEN** you must prove a cancellation law $a * c = b * c \Rightarrow a = b$ in a structure without a named identity or inverses **DO** look for an existence axiom and apply it exactly twice, choosing the two input pairs so the witnesses compose back into a and b **BECAUSE** cancellation is equivalent to injectivity of the translation map by c, and two applications of a solvability axiom typically suffice to convert surjectivity into injectivity for one specific element
- **WHEN** a problem's hypotheses read like group axioms minus a named identity **DO** resist proving the full group structure unless the problem actually asks for it **BECAUSE** proving only the minimal local fact the problem wants costs far less ink than reconstructing identity and inverses for the whole set, even when the stronger theorem is also true
- **WHEN** the same algebraic expression must be evaluated for two different elements sharing a piece of data, here c inside $d = a * c = b * c$ **DO** build one auxiliary object from only the shared data, then run the identical computation for each element **BECAUSE** an auxiliary object that never refers to either element individually is guaranteed to apply verbatim to both, which is exactly what lets $a * f = a$ and $b * f = a$ both fall out of a single construction

Backtesting

1. First useful cue: the goal $a * c = b * c \Rightarrow a = b$ is literally "cancel c", with no identity or inverse anywhere in sight.
2. Method the expression suggested: repeated, controlled use of the solvability axiom as an existence tool, combined with free rebracketing from associativity and commutativity.

3. Where you would hesitate, and the thought to reinforce: right after naming e with $d * e = a$, unsure what to do with it; reinforce that the next move is always to attach the leftover factor, here c, to the witness, building something usable on both a and b.
4. The trap the question sets: the temptation to prove S is a full abelian group before cancelling, which is true but far more expensive than the problem requires.
5. What makes you recognise it faster next time: reading "there exists z such that $x * z = y$, for every x, y" and translating it instantly into "every translation map is surjective", a single phrase naming exactly which tool you may invoke, and roughly how many times, here twice, for a two-element cancellation.
6. The reusable decision rule: when an abstract commutative, associative structure gives solvability but not identity or inverses, and asks for cancellation, apply solvability exactly twice, choosing input pairs that let the two witnesses recombine into the two sides of the target equation.

Mirror drills

1. Under the same hypotheses on $(S, *)$, show that for any $a \in S$ there is a unique $u_a \in S$ with $u_a * a = a$. *Answer:* existence is solvability with $(x, y) = (a, a)$; uniqueness follows from the cancellation law just proved, since $u * a = a = u' * a$ gives $u = u'$.
2. Show the element u_a from the previous drill does not depend on a, so S has a genuine two-sided identity. *Answer:* use solvability to find z with $a * z = b$; then $u_a * b = u_a * (a * z) = (u_a * a) * z = a * z = b$, so u_a also fixes b, and uniqueness forces $u_a = u_b$.
3. If $(S, *)$ satisfies the original hypotheses and S is finite, explain in one line why S is in fact a group. *Answer:* on a finite set, cancellativity makes each translation L_x injective, hence bijective since domain and codomain coincide, so every equation $x * z = y$ now has a unique solution, giving inverses as well as identity.

14.3 A3. pinning a functional equation at its fixed point

Let $f : [-1,1] \to \mathbb{R}$ be a continuous function such that

(i) $f(x) = \dfrac{2-x^2}{2} f\left(\dfrac{x^2}{2-x^2}\right)$ *for every x in $[-1,1]$,*

(ii) $f(0) = 1$, *and*

(iii) $\lim\limits_{x\to 1^-} \dfrac{f(x)}{\sqrt{1-x}}$ *exists and is finite.*

Prove that f is unique, and express $f(x)$ in closed form.

Read it out loud

This is a functional equation of the self-referential kind: f at x is expressed in terms of f at a different point, $\phi(x) = \dfrac{x^2}{2-x^2}$, together with an explicit algebraic prefactor $\dfrac{2-x^2}{2}$. The fields wired into this problem are functional equations, real analysis (continuity, limits, sequential arguments) and a touch of dynamics: ϕ is a self-map of $[-1,1]$, and the whole problem is really about where its iterates go. The command "prove that f is unique, and express $f(x)$ in closed form" is two separate tasks in disguise: an at-most-one-solution argument, and a genuine existence check, both required for a complete answer.

The first useful cue is $f(0) = 1$ sitting next to the prefactor $\frac{2-x^2}{2}$, which equals 1 exactly at $x = 0$. Substituting $x = \sin\theta$ mentally, or simply noticing that $\sqrt{1-x^2}$ is the one elementary function that is 1 at $x = 0$ and behaves well under this kind of rescaling, should be in your head within the first thirty seconds: the shape of the prefactor is screaming for a normalisation by $\sqrt{1-x^2}$, so that whatever multiplies f after you divide it out becomes an unadorned invariance.

The second cue, easy to miss on a first read, is that $\phi(0) = 0$: the argument transformation fixes the very point where hypothesis (ii) hands you a value. Whenever a functional equation's argument transformation has a fixed point, and you are independently given the function's value there, the fixed point is where the whole problem collapses; everything else, including the limit condition (iii) near the other end of the interval, should be treated with some suspicion until you have tried building the solution around that fixed point first.

Approaches, ranked worst to best

1. **guess and verify only.** *Cue:* $f(0) = 1$ suggests a natural closed form. *Circuit:* pattern matching. (guessing $f(x) = \sqrt{1-x^2}$ from the shape of the equation and checking it satisfies (i) to (iii) confirms existence, but the problem explicitly demands uniqueness as well; stopping here leaves open the possibility of other solutions and loses most of the marks despite having the right formula).

2. **boundary driven, via hypothesis (iii).** *Cue:* a limit condition is stated at $x \to 1^-$, so use it. *Circuit:* control f near the singular end and propagate inward. (a genuine route, and the one that uses every hypothesis on the page, which often feels reassuring; but it requires delicate estimates to relate values across the domain from a boundary point, and Elkies' own remark on this problem confirms condition (iii) is not even necessary, so this route pays for machinery the problem does not require).

3. **normalise, then iterate to the interior fixed point.** *Cue:* $\phi(0) = 0$ sits exactly where $f(0) = 1$ is given. *Circuit:* invariance under an attracting map, closed by continuity. (the ideal route: divide by $\sqrt{1-x^2}$ to strip the prefactor, reduce (i) to the bare invariance $g(x) = g(\phi(x))$, iterate ϕ from any starting point to its attracting fixed point at 0, and read off $g \equiv 1$ from continuity and $f(0) = 1$; only (ii) and the standing continuity of f are used).

The ideal solution

Uniqueness. Suppose f satisfies (i), (ii), (iii). Define $g : (-1,1) \to \mathbb{R}$ by

$$g(x) = \frac{f(x)}{\sqrt{1-x^2}},$$

well defined and continuous on $(-1,1)$ since f is continuous there and $\sqrt{1-x^2} > 0$ there.

Substitute $f(x) = g(x)\sqrt{1-x^2}$ into (i). Write $\phi(x) = \dfrac{x^2}{2-x^2}$. Since $2 - x^2 > 0$ on $[-1,1]$,

$$1 - \phi(x)^2 = \frac{(2-x^2)^2 - x^4}{(2-x^2)^2} = \frac{(2-2x^2)(2)}{(2-x^2)^2} = \frac{4(1-x^2)}{(2-x^2)^2},$$

using the difference of squares $(2-x^2)^2 - x^4 = \big((2-x^2)-x^2\big)\big((2-x^2)+x^2\big)$. For $x \in (-1,1)$, $1 - x^2 > 0$, so

$$\sqrt{1-\phi(x)^2} = \frac{2\sqrt{1-x^2}}{2-x^2}.$$

Hence (i) becomes

$$g(x)\sqrt{1-x^2} = \frac{2-x^2}{2} \cdot g(\phi(x)) \cdot \frac{2\sqrt{1-x^2}}{2-x^2} = g(\phi(x))\sqrt{1-x^2},$$

and dividing by $\sqrt{1-x^2} > 0$,

$$g(x) = g(\phi(x)) \qquad \text{for every } x \in (-1,1).$$

Fix $x \in (-1,1)$ and define a sequence by $a_1 = x$, $a_{n+1} = \phi(a_n) = \dfrac{a_n^2}{2-a_n^2}$. Since $|a_1| < 1$, an easy induction shows $a_n \in [0,1)$ for all $n \geq 2$: if $|a_n| < 1$ then $2 - a_n^2 > 1$, so $\phi(a_n) = \dfrac{a_n^2}{2-a_n^2} < a_n^2 < 1$, and $\phi(a_n) \geq 0$ always. In particular $|a_{n+1}| < a_n^2 \leq |a_n|$ for $a_n \neq 0$, so the sequence $\{|a_n|\}$ is decreasing, and in fact $|a_n| \leq |x|^n$ for all n: this holds for $n = 1$, and if $|a_n| \leq |x|^n$ then $|a_{n+1}| < a_n^2 \leq |x|^{2n} \leq |x|^{n+1}$ since $|x| < 1$. As $|x| < 1$, $|x|^n \to 0$, so $a_n \to 0$.

By the invariance just derived, $g(a_1) = g(a_2) = \cdots = g(x)$: the sequence $g(a_n)$ is constant, equal to $g(x)$, for every n. Since g is continuous at 0 and $a_n \to 0$, $g(a_n) \to g(0)$; a constant sequence converges only to its constant value, so

$$g(x) = g(0) = \frac{f(0)}{\sqrt{1-0}} = f(0) = 1.$$

This holds for every $x \in (-1,1)$, so $f(x) = \sqrt{1-x^2}$ throughout the open interval.

It remains to pin down $x = \pm 1$, and hypothesis (i) does this by itself, with no limiting argument at all. Hypothesis (i) is asserted for every $x \in [-1,1]$, including the endpoints, and $\phi(1) = \dfrac{1}{2-1} = 1$: the point $x = 1$ is itself a fixed point of ϕ. Substituting $x = 1$ into (i) directly,

$$f(1) = \frac{2-1}{2} f(\phi(1)) = \frac{1}{2} f(1),$$

a single algebraic equation in the one unknown $f(1)$, forcing $f(1) = 0$. Substituting $x = -1$, and using $\phi(-1) = \dfrac{1}{2-1} = 1$,

$$f(-1) = \frac{2-1}{2} f(\phi(-1)) = \frac{1}{2} f(1) = 0.$$

So $f(\pm 1) = 0 = \sqrt{1-(\pm 1)^2}$ as well, and $f(x) = \sqrt{1-x^2}$ on the whole of $[-1,1]$. This shows any f satisfying (i) to (iii) must equal this formula: at most one solution exists.

Existence. Direct substitution confirms $F(x) = \sqrt{1-x^2}$ satisfies all three hypotheses. For (i): with $\phi(x) = x^2/(2-x^2)$, the identity $1-\phi(x)^2 = 4(1-x^2)/(2-x^2)^2$ derived above gives

$$\frac{2-x^2}{2} F(\phi(x)) = \frac{2-x^2}{2} \cdot \frac{2\sqrt{1-x^2}}{2-x^2} = \sqrt{1-x^2} = F(x).$$

For (ii): $F(0) = \sqrt{1-0} = 1$. For (iii):

$$\lim_{x\to 1^-} \frac{F(x)}{\sqrt{1-x}} = \lim_{x\to 1^-} \frac{\sqrt{1-x^2}}{\sqrt{1-x}} = \lim_{x\to 1^-} \sqrt{1+x} = \sqrt{2},$$

which exists and is finite. So F is continuous on $[-1,1]$ and satisfies all of (i) to (iii): a solution exists.

Combining both halves, f is unique, and

$$\boxed{f(x) = \sqrt{1-x^2}}. \qquad \blacksquare$$

Layer 1: WHAT

1. Normalise by the natural candidate $\sqrt{1-x^2}$: set $g = f/\sqrt{1-x^2}$ on $(-1,1)$.
2. Simplify (i) after substitution, using a difference-of-squares identity, until the prefactor cancels and only the bare invariance $g(x) = g(\phi(x))$ survives.
3. Iterate $\phi(t) = t^2/(2-t^2)$ from any starting point: the orbit is squeezed toward its fixed point at 0.
4. Use continuity of g at 0 to pass from "g constant along the orbit" to "$g(x) = g(0) = f(0) = 1$" on the open interval.
5. Pin the two endpoints directly, by substituting $x = \pm 1$ into (i) itself, since 1 is also a fixed point of ϕ: no limiting argument is needed there at all.
6. Verify the resulting formula independently against all three hypotheses, to confirm a solution genuinely exists and not only that at most one could.

The nominal route reaches for hypothesis (iii) immediately, since it is the hypothesis that looks hardest to justify, and tries to control f near the boundary $x = 1$ and propagate values inward from there. The ideal route instead notices that the argument transformation ϕ already fixes the interior point $x = 0$, exactly where hypothesis (ii) hands over a known value, and drives every orbit there instead. The bifurcation point is the choice of which point in the domain to build the whole proof around: the boundary, where a hypothesis is visibly aimed at you, or the interior fixed point, which requires noticing a piece of hidden structure inside hypothesis (i) itself. The ideal route needs no limiting argument at the boundary either, since $x = 1$ turns out to be a second fixed point of ϕ, resolved by one substitution.

Layer 2: HOW

Block 1. **Cue.** The prefactor $\frac{2-x^2}{2}$ equals exactly 1 at $x = 0$, the same point where hypothesis (ii) gives $f(0) = 1$. **Why natural.** Dividing out a natural candidate function turns a multiplicative, inhomogeneous-looking functional equation into a plain invariance, the shape you know how to attack by iteration; guessing the normalising factor from the boundary data and the coefficient structure of (i) is a standard reroute whenever a functional equation mixes an unknown function with an explicit algebraic prefactor, and $\sqrt{1-x^2}$ is the simplest function matching both the value at 0 and the visible $(2-x^2)$ structure. **Link to the next block.** Once g is defined, every part of hypothesis (i) involving $\sqrt{1-x^2}$ should cancel if the guess was correct, so the next step is to perform that algebra and watch it collapse.

Block 2. **Cue.** After substitution, the surviving expression $1-\phi(x)^2$ is exactly "1 minus a squared fraction". **Why natural.** Factoring as a difference of two squares is the reflex whenever that shape appears, since it is cheaper than expanding directly, and here it produces a factor of $(2-x^2)$ in the numerator that exactly cancels the prefactor sitting outside f; watching that cancellation happen is the confirmation that $\sqrt{1-x^2}$ was the right choice of normalisation, not an assumption taken on faith. **Link to the next block.** With every $\sqrt{1-x^2}$ factor gone, what remains is an equation purely in g, of the form "value at x equals value at another specific point", an open invitation to iterate.

Block 3. **Cue.** The relation $g(x)=g(\phi(x))$ holds for every x in the domain, including $\phi(x)$ itself, and $\phi(\phi(x))$, and so on. **Why natural.** Whenever a functional equation asserts invariance under a self-map of the domain, the natural move is to iterate that map and ask where the orbit goes: if it converges to a point where the function's value is separately known, invariance forces the function to equal that value everywhere the orbit reaches. Here $\phi(t)=t^2/(2-t^2)$ visibly fixes 0, and the bound $\phi(t)<t^2$ for $|t|<1$ shows every orbit is squeezed toward 0, never toward the boundary. **Link to the next block.** Convergence of the orbit to 0 only becomes useful once that information can be transferred there through continuity, the next block.

Block 4. **Cue.** $g(a_n)=g(x)$ for every n, a constant sequence in disguise, and $a_n\to 0$. **Why natural.** Continuity is precisely the tool that turns "the same value along a convergent sequence" into "the same value at the limit point", and g is continuous at 0 because it is a quotient of the given continuous f by $\sqrt{1-x^2}$, itself continuous and strictly positive there. This is the single quantum leap of the problem: hypothesis (ii), $f(0)=1$, sits exactly at the attracting fixed point of the recursion hidden inside (i), so it, not the limit condition (iii), is what pins the function down on the open interval. **Link to the next block.** Having fixed $g\equiv 1$ on $(-1,1)$, only the two endpoints remain, and the domain of the closed interval $[-1,1]$ on which f is defined has not yet been fully covered.

Block 5. **Cue.** The interval is closed, $[-1,1]$, but the g-equation only ever made sense where $\sqrt{1-x^2}\neq 0$. **Why natural.** Hypothesis (i) itself, not merely its consequence for g, is asserted at every point of the closed interval, including $x=\pm 1$, and $\phi(1)=1$: the boundary point is itself a fixed point of ϕ, only a repelling one rather than an attracting one. Substituting $x=1$ into (i) directly therefore produces a closed algebraic equation in the single unknown $f(1)$, solving it outright with no limiting argument, cheaper even than invoking continuity of f at the boundary. **Link to the next block.** With $f(1)=0$ in hand, $x=-1$ resolves the same way in one further substitution, completing the domain without ever touching hypothesis (iii).

Block 6. **Cue.** The problem asks for both uniqueness and a closed form, which in a functional equation problem always means two separate half-proofs. **Why natural.** Skipping direct verification and treating the derived formula as automatically valid would leave open the possibility that no function satisfies (i) to (iii) at all, making the whole derivation vacuous; substituting $\sqrt{1-x^2}$ back into all three hypotheses is the cheapest possible way to rule that out. **Link to the next block.** None; both halves together give the complete answer the problem wants.

Layer 3: WHY

The truth underneath the route is a general one about functional equations built from an iterated self-map: whenever $f\circ(\text{normalisation})$ satisfies $g(x)=g(\phi(x))$ for a self-map ϕ with an attracting fixed point, and the value of g is independently known at that fixed point, continuity forces g to be constant on the whole basin of attraction. This is a recurring analytic technique, not specific to this problem: invariance under an attracting map, plus continuity, forces constancy.

The discernment mechanism is to examine every functional equation of the form $f(x) = c(x)f(\phi(x))$ for the fixed points of ϕ inside the domain, before reaching for any other given hypothesis. Compute ϕ at a couple of natural points, check whether a fixed point is attracting by a quick bound or a derivative, and check whether the function's value is already known there. If so, that fixed point is almost always the cheapest route, and every other hypothesis, however prominently placed, should be treated as possibly redundant until proven otherwise. Here, hypothesis (ii) is planted exactly at the point the hidden recursion collapses to, while hypothesis (iii), the limit near $x = 1$, is a genuine red herring for this route: Elkies observed precisely this, that (iii) is unnecessary, and the reason is now visible, $x = 1$ is a second fixed point of ϕ, but a repelling one, so it never receives an orbit from the interior; it only ever needs to be checked once, directly, by substitution, which needs no limit at all. Train the discernment by, on every functional equation with a visible argument transformation, computing where that transformation sends 0 and where it sends the endpoints of the domain, before touching any other hypothesis on the page.

Revolving concepts

Functional equations invariant under an iterated self-map: normalise by an explicit factor to remove any algebraic prefactor, reducing $f(x) = c(x)f(\phi(x))$ to the bare invariance $g(x) = g(\phi(x))$, then iterate ϕ toward a fixed point where the value is separately known. See field Functional equations, topic invariance under iteration and fixed points.

Attracting versus repelling fixed points of a real map: if $\phi(t_0) = t_0$ and ϕ is differentiable at t_0, the fixed point is attracting when $|\phi'(t_0)| < 1$, so nearby orbits converge to it, and repelling when $|\phi'(t_0)| > 1$, so nearby orbits move away. Here $\phi(t) = t^2/(2 - t^2)$ has $\phi'(t) = 4t/(2 - t^2)^2$, giving $\phi'(0) = 0$, an attracting, in fact superattracting, fixed point, and $\phi'(1) = 4$, a repelling one; this is exactly why iterating from any interior point drives the orbit to 0 and never toward 1. See field Sequences and recurrences, topic fixed points of iterated maps.

Self-referential boundary points resolve algebraically: if a functional equation $f(x) = c(x)f(\phi(x))$ is asserted at a fixed point t_0 of ϕ lying in the domain, substituting $x = t_0$ gives the single algebraic equation $f(t_0) = c(t_0)f(t_0)$, pinning $f(t_0)$ directly, to 0 whenever $c(t_0) \neq 1$, without any limiting or continuity argument; every fixed point of ϕ deserves this check, not only the attracting one used for the main argument. See field Functional equations, topic fixed-point substitution.

Difference of squares as a standing reflex: $1 - u^2 = (1 - u)(1 + u)$, used here to simplify $1 - \phi(x)^2$ into a form sharing a common factor with the prefactor $(2 - x^2)/2$. See field Polynomials, topic factoring in functional equation simplification.

Sequential continuity transfers a constant-along-a-sequence property to the limit: if h is continuous at L and $x_n \to L$, then $h(x_n) \to h(L)$; combined with $h(x_n)$ constant, this forces $h(L)$ to equal that constant. See field Real analysis, topic sequential continuity.

Recursion generalisation: the sequence $a_{n+1} = \phi(a_n)$ with $\phi(t) = t^2/(2 - t^2)$ is a first-order nonlinear recursion with no closed form for a_n; the general method for such "invariance under iteration" functional equations is to locate every fixed point of ϕ in the domain, classify each as attracting or repelling by $|\phi'|$ there, and note that any continuous function invariant under ϕ must be constant on the full basin of attraction of an attracting fixed point wherever a side condition pins its value there. This skeleton solves many Putnam functional equations of the shape $f(x) = c(x)f(\phi(x))$: first remove $c(x)$ by an integrating-factor-style substitution, as $\sqrt{1 - x^2}$ did here, then analyse ϕ alone. See field Sequences and recurrences, topic nonlinear first-order recursions and their fixed points.

Circuit Bank

- **WHEN** a functional equation has the shape $f(x) = c(x)f(\phi(x))$ for an explicit prefactor $c(x)$ and an explicit self-map ϕ **DO** guess a normalising function $h(x)$, often suggested by a boundary value or the shape of $c(x)$, so that $g = f/h$ satisfies the bare invariance $g(x) = g(\phi(x))$ **BECAUSE** stripping the explicit prefactor turns an inhomogeneous equation into a pure invariance statement solvable by iteration, and the correct h is the one for which the prefactor exactly cancels against $h(x)/h(\phi(x))$
- **WHEN** an iterated self-map ϕ appears inside a functional equation **DO** find its fixed points and classify each by $|\phi'|$ there against 1 **BECAUSE** an attracting fixed point is where every nearby orbit converges, and if the function's value is independently known there, continuity forces the function to equal that value on the whole basin
- **WHEN** you see 1 minus something squared inside a root or a simplification **DO** factor immediately as a difference of two squares **BECAUSE** it is always cheaper than expanding, and it very often exposes a common factor against another part of the expression
- **WHEN** a problem gives three hypotheses but your cleanest route only visibly needs two of them **DO** finish the proof with the two you needed, then state explicitly which hypothesis was not used and why the route did not require it **BECAUSE** full marks proofs must be honest about which hypotheses are load bearing, and naming the unused one shows the redundancy was noticed rather than missed

Backtesting

1. First useful cue: the coefficient $\frac{2-x^2}{2}$ multiplying f at a rescaled argument, sitting next to $f(0) = 1$, suggests dividing by a function equal to 1 at $x = 0$ that interacts cleanly with that coefficient.
2. Method the expression suggested: normalise by $\sqrt{1-x^2}$, reduce to a bare invariance $g(x) = g(\phi(x))$, then iterate ϕ toward its fixed point.
3. Where you would hesitate, and the thought to reinforce: uncertainty about whether hypothesis (iii) needs to appear somewhere; reinforce that a hypothesis not yet used is not automatically a hypothesis you are missing, check first whether the route is already complete and rigorous without it, and if so, say so explicitly.
4. The trap the question sets: the limit condition (iii) at a visually singular-looking boundary tempts you to build the whole proof around $x \to 1^-$; the actual pivot is the interior fixed point at 0, and even the boundary itself resolves by one substitution, not a limit.
5. What makes you recognise it faster next time: seeing a self-referential functional equation $f(x) = c(x)f(\phi(x))$ and immediately computing $\phi(0)$ and $\phi'(0)$ before touching anything else, since a fixed point at 0 paired with a given value $f(0)$ is the strongest structural cue on the page.
6. The reusable decision rule: when a functional equation expresses f at x in terms of f at $\phi(x)$ for an explicit self-map ϕ, normalise to remove the prefactor, locate every fixed point of ϕ and its attracting or repelling character, and pin the function at each fixed point, by iteration and continuity for an attracting one, by direct substitution for the rest.

Mirror drills

1. Suppose $h : [0,1] \to \mathbb{R}$ is continuous, $h(0) = 5$, and $h(x) = h(x/2)$ for every $x \in [0,1]$. Find h. *Answer:* iterating $x \mapsto x/2$ drives every orbit to 0; continuity forces $h \equiv 5$.

2. For $\phi(t) = t^2/(2-t^2)$, verify directly that $\phi'(0) = 0$ and $\phi'(1) = 4$, and explain in one line why this confirms orbits starting in $(-1, 1)$ head to 0, not to 1. *Answer:* $\phi'(t) = 4t/(2-t^2)^2$ by the quotient rule, so $\phi'(0) = 0$ and $\phi'(1) = 4$; a derivative below 1 in size at a fixed point means nearby points move closer under iteration, above 1 means they move away, so 0 attracts and 1 repels.

3. A continuous $f : [-1, 1] \to \mathbb{R}$ satisfies hypothesis (i) alone, with (ii) and (iii) both dropped. Show this is not enough to determine f, by exhibiting two different continuous functions satisfying it. *Answer:* $f \equiv 0$ works trivially, and $f(x) = \sqrt{1-x^2}$ also works, as verified in the ideal solution; indeed any $f(x) = k\sqrt{1-x^2}$ satisfies (i) for every constant k, so (i) alone has infinitely many solutions, and it is exactly the extra value $f(0) = 1$ from (ii) that selects $k = 1$.

14.4 A4. A small box forces the solutions into a progression

Let q and r be integers with $q > 0$, and let A and B be intervals on the real line. Let T be the set of all $b + mq$ where b and m are integers with b in B, and let S be the set of all integers a in A such that ra is in T. Show that if the product of the lengths of A and B is less than q, then S is the intersection of A with some arithmetic progression.

Read it out loud

Say the definitions back to yourself in one breath. T is the set of integers whose residue modulo q is realised by an integer lying in B, so membership of S unwinds to a single sentence: an integer a lies in S exactly when $a \in A$ and there exists an integer $b \in B$ with $ra \equiv b \pmod{q}$. One free variable a, one existential witness b, one congruence tying the two together, and a conclusion that mentions only a. The whole difficulty of the problem is that the witness b is quantified away in the statement of S but is the only thing that carries the information.

The fields wired into this are elementary number theory (linear congruences, reduction by $\gcd(q, r)$), the geometry of numbers (a congruence between two integer variables is a sublattice of $\mathbb{Z}^2$, and the modulus is its index), convexity (intervals, hence a rectangle, hence a convex region), and the combinatorial characterisation of an arithmetic progression.

Now the guiding light, and it is planted in plain sight. The hypothesis is not $\ell(A) < q$, and it is not $\ell(B) < q$. It is the *product* of two lengths, one measured in the world of a and one measured in the world of b, compared with q. A product of two lengths is an area. Nothing in the problem looks two-dimensional yet, so the examiners have told you exactly where to go and work: the plane of pairs (x, y), the rectangle $R = A \times B$ whose area is that product, and the set of admissible pairs

$$L = \{(x, y) \in \mathbb{Z}^2 : y \equiv rx \pmod{q}\}.$$

Within thirty seconds you should have written those two objects down together with the one line that reorganises the problem: S is the image of $L \cap R$ under the projection $(x, y) \mapsto x$. Everything after that is a comparison between the area of R and the index of L.

What is really being tested is whether you own that comparison. Lattices support exactly one quantitative statement of this shape, the one Minkowski built a subject on: a region that is too small for the lattice cannot contain many independent lattice points. Here "too small" is area less than q, and "not many" turns out to mean "no three of them off a single line".

Approaches, ranked worst to best

1. **Residue chase with no size control.** *Cue:* the word “modulo” and nothing else. *Circuit:* write S out as a set of residues and try to force structure by hand. You take two elements $a_1 < a_2$ of S, subtract, and try to prove $a_2 + (a_2 - a_1) \in S$. It stalls immediately and for an instructive reason: $r(2a_2 - a_1) \equiv 2b_2 - b_1 \pmod{q}$ is true, but $2b_2 - b_1$ has no reason to lie in B, and B is the only place a witness is allowed to live. Any route that forgets b is dead before it starts, because the hypothesis is a statement about B.

2. **Three distance theorem.** *Cue:* ra modulo q is an orbit of a rotation, and B is an arc. *Circuit:* return times of a rotation into an interval take at most three values. S is the set of visit times of the orbit $a \mapsto ra/q \bmod 1$ to an arc, so the gaps between consecutive elements of S take at most three values, and the two largest are the sum of the other two. To finish you still have to prove that the small area forces one value, which needs the same estimate as the winning route. You imported a heavy theorem and saved nothing.

3. **Euclidean descent on the pair (q, r).** *Cue:* $\gcd(q, r)$ appears the moment you try to solve the congruence. *Circuit:* continued fraction reduction, the engine behind the three distance

theorem. Replace r by $r \bmod q$, swap the roles of the two intervals, and induct on the size of r. The statement is genuinely stable under this descent, but you must carry two intervals, two lengths and the area hypothesis through every step, and the bookkeeping is where marks die.

4. **Congruence with sign control, then an arithmetic progression lemma.** *Cue:* "absolute value less than q and divisible by q" is the only way a congruence becomes an equation. *Circuit:* size beats modulus. Divide q, r and B by $\gcd(q,r)$; take three elements $a_1 < a_2 < a_3$ of S with witnesses b_1, b_2, b_3; two of the three differences $b_i - b_j$ share a sign, and for that pair the two products $(a_i - a_j)(b_k - b_l)$ and $(b_i - b_j)(a_k - a_l)$ are congruent modulo q, of the same sign and of absolute value at most $\ell(A)\ell(B) < q$, hence equal. That is collinearity of (a_i, b_i), from which $a_1 + a_3 - a_2 \in S$, and then a standalone lemma says a finite set closed under $a + c - b$ is an arithmetic progression. Completely elementary and completely correct, at the cost of one reduction, one sign discussion and one auxiliary lemma.

5. **Lattice and area.** *Cue:* "the product of the lengths", i.e. an area, compared with the modulus. *Circuit:* a linear congruence in two variables is a sublattice of index q; a region of area less than the index holds no lattice triangle. Three lines of setup, one determinant estimate, one triangle estimate, done. No reduction by $\gcd(q,r)$ is needed, because the index of L is q whatever r is, and no auxiliary progression lemma is needed, because collinear lattice points inside a convex set already are a progression. This is the route to write.

The ideal solution

Write $\ell(X)$ for the length of an interval X. Read "arithmetic progression" as a set $\{a_0 + kd : k \in K\}$ with d a fixed integer and K a set of consecutive integers, possibly finite, possibly empty. A caveat worth clarifying before any mathematics: if arithmetic progressions were required to be infinite the statement would be false. Take $q = 5$, $r = 1$, $A = [1,3]$, $B = [0,2]$, so that $\ell(A)\ell(B) = 4 < 5$. Then T is the set of integers congruent to 0, 1 or 2 modulo 5, so $S = \{1,2\}$, while any infinite progression containing 1 and 2 has common difference 1 and therefore meets A in $\{1,2,3\}$. With finite progressions allowed, and since $S \subseteq A$ forces $S = A \cap S$, the task is exactly this: prove that S is an arithmetic progression.

Degenerate lengths first. If $\ell(A) = 0$ then A contains at most one integer and S has at most one element, which is a progression. If $\ell(B) = 0$ then B contains at most one integer; if it contains none then $S = \emptyset$, and if it contains exactly one integer b then

$$S = \{a \in A \cap \mathbb{Z} : ra \equiv b \pmod{q}\},$$

and the solution set of that congruence is empty when $\gcd(q,r) \nmid b$ and otherwise a full residue class modulo $q/\gcd(q,r)$, an infinite arithmetic progression. So assume from now on that $\ell(A)$ and $\ell(B)$ are both positive; the hypothesis $\ell(A)\ell(B) < q$ then forces both to be finite, so $R := A \times B$ is a bounded rectangle of area less than q.

The lattice. Put $L = \{(x,y) \in \mathbb{Z}^2 : y \equiv rx \pmod{q}\}$. It is the kernel of the group homomorphism $\mathbb{Z}^2 \to \mathbb{Z}/q\mathbb{Z}$ sending (x,y) to $rx - y$, which is surjective because $(0,-1) \mapsto 1$. Hence L is a subgroup of $\mathbb{Z}^2$ of index exactly q, for every r. Unwinding the definitions, $a \in S$ if and only if $a \in A \cap \mathbb{Z}$ and there is an integer $b \in B$ with $ra - b \in q\mathbb{Z}$, that is, if and only if $(a,b) \in L \cap R$ for some b. Therefore

$$S = \pi(L \cap R), \qquad \pi(x,y) = x.$$

Lemma 14.4.1. *If $P_1, P_2, P_3 \in L$ are not collinear, the triangle $P_1P_2P_3$ has area at least $q/2$.*

Proof. Put $u = P_2 - P_1$ and $w = P_3 - P_1$; both lie in L and they are linearly independent. The lattice $M = \mathbb{Z}u + \mathbb{Z}w$ satisfies $M \subseteq L \subseteq \mathbb{Z}^2$, and the index of a full rank sublattice of $\mathbb{Z}^2$ equals the absolute value of the determinant of any of its bases (this is the standard fact from the geometry of

numbers, proved by Smith normal form: no A-level route to it exists, so install it as a named fact). Hence

$$|\det(u,w)| = [\mathbb{Z}^2 : M] = [\mathbb{Z}^2 : L]\,[L : M] = q\,[L : M],$$

a positive multiple of q. The area of the triangle is $\frac{1}{2}|\det(u,w)| \geq q/2$. □

Lemma 14.4.2. *A triangle contained in a rectangle with sides parallel to the axes of side lengths w and h has area at most $wh/2$.*

Proof. Label the vertices $P_i = (x_i, y_i)$ with $x_1 \leq x_2 \leq x_3$. If $x_1 = x_3$ the triangle is degenerate and has area 0. Otherwise the vertical line through P_2 meets the segment P_1P_3 in a point $Q = (x_2, y_Q)$, and the segment P_2Q cuts the triangle into two triangles with common base P_2Q and heights $x_2 - x_1$ and $x_3 - x_2$. Both y_2 and y_Q lie in the projection of the rectangle to the y-axis, an interval of length h, so $|P_2Q| = |y_2 - y_Q| \leq h$ and the total area is

$$\tfrac{1}{2}|y_2 - y_Q|\,(x_2 - x_1) + \tfrac{1}{2}|y_2 - y_Q|\,(x_3 - x_2) = \tfrac{1}{2}|y_2 - y_Q|(x_3 - x_1) \leq \tfrac{1}{2}hw.$$ □

Collinearity. Any three points of $L \cap R$ lie in the closed rectangle of side lengths $\ell(A)$ and $\ell(B)$, so by Lemma 2 they span a triangle of area at most $\ell(A)\ell(B)/2 < q/2$, and by Lemma 1 they must therefore be collinear. If $L \cap R$ has at most two points it lies on a line trivially; if it has at least three, fix two distinct points of it and note that every further point is collinear with those two, so the whole of $L \cap R$ lies on one line λ.

From a line to a progression. If $L \cap R = \emptyset$ then $S = \emptyset$. Otherwise fix $P_0 = (x_0, y_0) \in L \cap R$. The set $(L \cap \lambda) - P_0$ is a subgroup of $\mathbb{Z}^2$ contained in a one dimensional subspace, hence a discrete subgroup of a line, hence $\mathbb{Z}w_0$ for some $w_0 = (\alpha, \beta) \in \mathbb{Z}^2$ (with $w_0 = 0$ allowed). So $L \cap \lambda = \{P_0 + kw_0 : k \in \mathbb{Z}\}$. Since R is convex and $k \mapsto P_0 + kw_0$ is affine, $K := \{k \in \mathbb{Z} : P_0 + kw_0 \in R\}$ is the set of integers in an interval, that is, a set of consecutive integers. Consequently

$$S = \pi(L \cap R) = \{x_0 + k\alpha : k \in K\},$$

which is an arithmetic progression with common difference α (a single point if $\alpha = 0$, since then every point of $L \cap R$ shares its first coordinate). Finally $S \subseteq A$, so $S = A \cap S$ is the intersection of A with an arithmetic progression, as required. ■

Layer 1: WHAT

Block 1. Unwind S: $a \in S$ iff $a \in A \cap \mathbb{Z}$ and $ra \equiv b \pmod{q}$ for some integer $b \in B$. Refuse to discard b.

Block 2. Read the hypothesis as an area and promote the pair (a, b) to a point of the plane: $S = \pi(L \cap R)$ with $R = A \times B$ and L the solution set of the congruence.

Block 3. Identify L as a subgroup of $\mathbb{Z}^2$ of index q, by exhibiting the congruence as the kernel of a surjection onto $\mathbb{Z}/q\mathbb{Z}$.

Block 4. Compare the two areas. Three non-collinear points of L cost at least $q/2$; a triangle inside R costs at most $\ell(A)\ell(B)/2 < q/2$. So $L \cap R$ is collinear.

Block 5. Collinear points of a lattice form a coset of $\mathbb{Z}w_0$; intersecting with the convex set R keeps a consecutive run of them; projecting gives an arithmetic progression.

The nominal route runs: fix $a_1 < a_2 < a_3$ in S, write down the three congruences, cross multiply differences, argue about signs and sizes to turn a congruence into an equation, deduce $a_1 + a_3 - a_2 \in S$, then prove and apply a lemma that a finite set closed under $a + c - b$ is a progression. It is a correct route and it is what most solvers write, and it costs one gcd reduction, one sign case analysis and one extra lemma. The bifurcation point is a single decision taken in the first minute: when you unwind S, do you treat b as a nuisance to be quantified away, or as the second coordinate of a point? Keep b as a coordinate and the area hypothesis becomes a statement about a lattice; drop it and you spend the rest of the problem reconstructing that statement by hand in the language of inequalities.

Layer 2: HOW

Block 1. **Cue.** The definition of T is a two step definition (b in B, then translate by multiples of q), which always means: there is a hidden witness. **Why natural.** A set defined by an existential quantifier is never understood by staring at the quantified variable. You either eliminate the witness (needs gcd information and inequalities) or you keep it. The hypothesis mentions the length of B, and the only thing the length of B can constrain is the witness, so the witness must survive into your notation. **Link to the next block.** Once you have decided to keep b, you have a pair, and a pair wants a plane.

Block 2. **Cue.** "The product of the lengths of A and B is less than q". **Why natural.** A product of two lengths is an area, and an area is a two dimensional quantity, so the examiners have already told you that the intended arena is $A \times B$. The comparison object q is the modulus; and modulus is index. In other words the hypothesis reads: the area of the box is less than the index of the lattice. That sentence is the entire problem. **Link to the next block.** Now you must know what an index of a sublattice of $\mathbb{Z}^2$ controls, which is the determinant.

Block 3. **Cue.** A congruence linking two integer variables, linear in both. **Why natural.** $\{(x, y) : y \equiv rx \pmod{q}\}$ is the kernel of $(x, y) \mapsto rx - y$ into $\mathbb{Z}/q\mathbb{Z}$. Kernels of surjections have index equal to the size of the target, so the index is q, with no dependence on r at all. That independence is what lets you skip the $\gcd(q, r)$ reduction that the elementary route needs. **Link to the next block.** An index q sublattice has all its fundamental parallelograms of area q, so any two independent lattice vectors span area at least q.

Block 4. **Cue.** You want a conclusion of the form "the points are collinear", and the only currency you have is area. **Why natural.** Non-collinear means the determinant is nonzero, and for lattice vectors nonzero determinants are quantised in multiples of the index. So non-collinearity has a price, $q/2$ of area, while the box can only pay $\ell(A)\ell(B)/2 < q/2$. Any statement of the form "too small to contain independent lattice points" is proved this way; Minkowski's theorem is the same accounting with a factor of 4 and a convex symmetric body. **Link to the next block.** Collinearity is a structural conclusion, so now convert it into the structure the question asked for.

Block 5. **Cue.** The target word is "arithmetic progression", i.e. a one dimensional lattice coset. **Why natural.** The points of a lattice on a line form a coset of a discrete subgroup of that line, and every discrete subgroup of a line is cyclic. Convexity of R makes the surviving parameters consecutive, because an affine map pulls a convex set back to an interval. Projecting a progression of points to their first coordinates is again a progression. **Link to the next block.** Nothing remains except to notice $S \subseteq A$, so $S = A \cap S$ and the phrasing of the question is satisfied.

Layer 3: WHY

The mathematical truth beneath this route is the dictionary between congruences and lattices. A single linear congruence $y \equiv rx \pmod{q}$ in two integer unknowns is not an equation to be solved; it is a subgroup of $\mathbb{Z}^2$, and every fact about it is a fact about that subgroup. The index of the subgroup is the modulus. The determinant of any basis is the index. The area of a fundamental domain is the determinant. Once those four sentences are one object in your head, an inequality comparing an area with a modulus stops being a coincidence.

The discernment mechanism is the dimension count. Ask: how many free variables does the hypothesis constrain? Here the hypothesis constrains a product of two lengths, and lengths of different objects only multiply if you are measuring a region in a product space. A hypothesis of the form (size in A) times (size in B) less than (modulus) is always the geometry of numbers speaking.

Train it by collecting the family: Minkowski's convex body theorem, Thue's lemma (which produces $x \equiv ry \pmod{q}$ with $|x|, |y| < \sqrt{q}$ from exactly this area count and which is the engine of the two squares theorem), the Farey and continued fraction estimates, and Pick's theorem. Every one of them is the same accounting.

The examiners nudged twice. First they wrote "intervals" rather than sets, which hands you convexity, the property that turns "lattice points on a line" into "consecutive lattice points on a line". Second they compared a product of lengths with q rather than comparing each length with something, which is the plain instruction to build the rectangle. The rest of the wording, in particular the deliberately loose phrase "the intersection of A with some arithmetic progression", is there to let a finite progression count as an answer, and you should notice that this phrasing is doing exactly that job, because with infinite progressions only the statement is false.

Revolving concepts

Lattice in $\mathbb{Z}^2$. A subgroup $L \leq \mathbb{Z}^2$ of finite index. Orbiting facts: L has a basis u, w (structure theorem for finitely generated abelian groups); $[\mathbb{Z}^2 : L] = |\det(u, w)|$, independent of the basis; a fundamental domain has area $[\mathbb{Z}^2 : L]$. See Book 1, field: Number theory, topic: geometry of numbers.

Index and determinant. For $M \subseteq L \subseteq \mathbb{Z}^2$ of finite index, indices multiply: $[\mathbb{Z}^2 : M] = [\mathbb{Z}^2 : L][L : M]$. This is the mechanism that quantises the areas of lattice triangles, and it is exactly why a lattice triangle in $\mathbb{Z}^2$ itself has area at least $1/2$ (Pick's theorem in miniature). See Book 1, field: Abstract algebra, topic: index of a subgroup.

Minkowski's theorem. A convex region symmetric about the origin of area greater than $4 \det L$ contains a nonzero point of L. The statement used here is its cheaper twin: a region of area less than $\det L$ contains no three non-collinear points of L. Learn both, and learn that the first is proved by the pigeonhole principle on translates and the second by the multiplicativity of the index.

Linear congruence. $ra \equiv b \pmod{q}$ is solvable in a if and only if $\gcd(q, r) \mid b$, and then the solution set is a residue class modulo $q/\gcd(q, r)$, hence an infinite arithmetic progression. See Book 1, field: Number theory, topic: congruences.

Arithmetic progressions and the closure recursion. The elementary route needs this lemma: a finite set $S \subset \mathbb{Z}$ with at least three elements such that $a + c - b \in S$ whenever $a \leq b \leq c$ lie in S is an arithmetic progression. The recursion inside it generalises: let a_1 be the least element, d the gap to the second least, and let m be least with $a_1 + md \notin S$; if some element of S lay outside $\{a_1, \ldots, a_1 + (m-1)d\}$, the least such element n would give $n - d \in S$ by closure, again outside the run and smaller, contradicting minimality. The general pattern is: a set closed under a translation-type operation is a union of orbits of that translation, and minimal counterexample induction converts closure into exhaustion. The same recursion proves that a subgroup of $\mathbb{Z}$ is $d\mathbb{Z}$, that a discrete subgroup of $\mathbb{R}$ is cyclic (used above for $L \cap \lambda$), and that a numerical semigroup is eventually a progression.

Three distance theorem. For irrational θ, the points $\{\theta\}, \{2\theta\}, \ldots, \{N\theta\}$ cut the circle into arcs of at most three distinct lengths, the largest being the sum of the other two. It governs the general shape of the set of visit times to an arc, of which our S is an example; here the area hypothesis collapses the three possible gaps to one. See Book 1, field: Number theory, topic: equidistribution.

Circuit Bank

- **WHEN** a hypothesis multiplies a size in one variable by a size in another and compares it with a modulus **DO** build the rectangle in the product space and the lattice cut out by the congruence, then compare area with index **BECAUSE** a product of two lengths is an area, and the only theorems comparing areas with moduli are the geometry of numbers

- **WHEN** a set is defined by "there exists b in B with ..." **DO** keep b as a coordinate rather than quantifying it away **BECAUSE** the hypotheses that mention B can only act on the witness, so a proof that forgets the witness cannot use them
- **WHEN** you need to upgrade a congruence to an equation **DO** bound both sides in absolute value by less than the modulus, or force them to share a sign **BECAUSE** $q \mid N$ and $|N| < q$ give $N = 0$; this is the only elementary way a congruence becomes an equality
- **WHEN** lattice points are known to be collinear and you want a progression **DO** intersect the line with the lattice to get a coset of a cyclic group and use convexity to make the parameters consecutive **BECAUSE** a discrete subgroup of a line is cyclic, and the preimage of a convex set under an affine map is an interval
- **WHEN** a question says "some arithmetic progression" rather than "an arithmetic progression with difference d" **DO** check whether finite or empty progressions are being allowed, and test the extreme case **BECAUSE** loose phrasing in a competition statement is usually load bearing, and here the statement is false if progressions must be infinite

Backtesting

1. First useful cue: the hypothesis is a product of two lengths, that is an area, not two separate size conditions.
2. Method the expression suggested: promote (a, b) to a point of $\mathbb{Z}^2$, read the congruence as a sublattice of index q, and compare the area of $A \times B$ with q.
3. Where you would hesitate, and the thought to reinforce: at the claim that non-collinear lattice points span area at least $q/2$. Reinforce it by rebuilding the chain index equals determinant equals twice the triangle area, and by remembering that indices multiply down a chain of sublattices.
4. The trap the question sets: quantifying away the witness b, after which the length of B has nothing to act on and the problem looks impossible; and, at the end, believing that an arithmetic progression must be infinite, which makes the statement false.
5. What makes you recognise it faster next time: the phrase "the product of the lengths" together with a modulus. Also the shape of the conclusion, since an arithmetic progression is a one dimensional lattice coset, and lattice conclusions come from lattice hypotheses.
6. The reusable decision rule: congruence plus area equals geometry of numbers; keep every existential witness as a coordinate until the geometry has been extracted.

Mirror drills

1. Let $q = 100$, $r = 7$, $A = [0, 9]$, $B = [0, 9]$. Compute S and check the conclusion directly. *Answer:* $7a$ for $a = 0, \ldots, 9$ is $0, 7, 14, \ldots, 63$, and only 0 and 7 land in $[0, 9]$, so $S = \{0, 1\}$, a progression; the box has area $81 < 100$.
2. Prove Thue's lemma: for $\gcd(r, q) = 1$ there are integers x, y with $0 < \max(|x|, |y|) < \sqrt{q}$ and $x \equiv ry \pmod{q}$. *Answer:* the pairs (y, x) with $0 \le x, y \le \lfloor \sqrt{q} \rfloor$ number more than q, so two are congruent under $x - ry$; subtract, and the difference lies in the lattice L and in a box of side less than $\sqrt{q}$.
3. Show that a triangle contained in a w by h axis parallel rectangle has area at most $wh/2$, and that the bound is attained. *Answer:* split along the vertical through the middle vertex, as in Lemma 2; equality for a triangle with one full side of the rectangle as base and the opposite side containing the apex.

14.5 A5. When an affine map is one long cycle

Let $\mathbb{F}_p$ denote the field of integers modulo a prime p, and let n be a positive integer. Let v be a fixed vector in $\mathbb{F}_p^n$, let M be an $n \times n$ matrix with entries in $\mathbb{F}_p$, and define $G : \mathbb{F}_p^n \to \mathbb{F}_p^n$ by $G(x) = v + Mx$. Let $G^{(k)}$ denote the k-fold composition of G with itself, that is, $G^{(1)}(x) = G(x)$ and $G^{(k+1)}(x) = G(G^{(k)}(x))$. Determine all pairs p, n for which there exist v and M such that the p^n vectors $G^{(k)}(0)$, $k = 1, 2, \ldots, p^n$, are distinct.

Read it out loud

This is a dynamics question about an affine map of a finite vector space, and the command term is "determine all pairs", which tells you two things before you have done any mathematics. First, the answer is a short list, so you owe explicit constructions. Second, everything not on the list must be excluded by a single mechanism, so you owe one obstruction, not a case analysis over matrices.

Now read the condition itself. There are exactly p^n vectors in $\mathbb{F}_p^n$ and you are asked for p^n distinct iterates, so the orbit of 0 is the whole space. In other words G is a single cycle of length p^n, and therefore G is an element of order p^n in the group of invertible affine maps of $\mathbb{F}_p^n$. Hear that sentence properly: the order of the element is a power of the characteristic. In characteristic p that is a very strong statement about the linear part, because $M^{p^m} = I$ forces $(M - I)^{p^m} = 0$ by the Frobenius identity. So the linear part is unipotent, $N := M - I$ is nilpotent, and $N^n = 0$.

The fields are linear algebra over a finite field (unipotent matrices, nilpotency index, Cayley and Hamilton), group theory (order of an element of the affine group, cyclic permutations), and the p-adic arithmetic of binomial coefficients. The recursion at the heart is the affine one, $x_{k+1} = Mx_k + v$, and it is the same recursion you solved at A-level for $u_{n+1} = au_n + b$; you must be able to write its closed form instantly.

Within the first minute you should have written down $G^{(k)}(0) = (I + M + \cdots + M^{k-1})v$, tested $n = 1$ (where $G(x) = x + 1$ is obviously a p-cycle), and tested $p = n = 2$ by hand.

Approaches, ranked worst to best

1. **Search small cases and stop.** *Cue:* "determine all pairs". *Circuit:* numerical reconnaissance. Mandatory as reconnaissance and worthless as a solution. It does hand you the answer, because $n = 1$ works for every p and the only sporadic case you find is $(2, 2)$, but by itself it proves nothing about $n \geq 3$.

2. **Jordan form of M.** *Cue:* a matrix that is being iterated. *Circuit:* diagonalise or triangularise, then read off powers. You can push this through, but $\mathbb{F}_p$ is not algebraically closed, so you either move to an extension field or work with the rational canonical form, and then you must separately handle eigenvalue 1 blocks, other eigenvalues, and the inhomogeneous term. Enormous ink for a statement about a single element's order.

3. **Order of G in the affine group via the exponent of a Sylow p-subgroup.** *Cue:* the order of G must be p^n, a p-power. *Circuit:* a p-power order element lies in a Sylow p-subgroup, whose exponent is bounded. If $N^j = 0$ then $M^{p^s} = I$ as soon as $p^s \geq j$, so the order of M is at most $p^{\lceil \log_p n \rceil}$ and the order of G is at most p times that. This kills all large n at a stroke and is worth knowing, but the bound leaves $(p, 2)$ for every p and $(2, 3)$ standing, so you still have to compute. It gets you to a finite check and no further.

4. **Fitting decomposition.** *Cue:* $M - I$ is what governs fixed points of an affine map. *Circuit:* split the space into the part where $M - I$ is nilpotent and the part where it is invertible. Write $\mathbb{F}_p^n = V \oplus W$ with $M - I$ nilpotent on V and invertible on W; if $W \neq 0$, solve $(M - I)w = -v_2$ inside W and check that the orbit of w never leaves the coset $w + V$, whose size is smaller than p^n, contradicting transitivity. It is correct, and it reproves "$M - I$ is nilpotent" by hand.

5. **Order is a p-power, hence the linear part is unipotent, then binomial coefficients.** *Cue:* p^n distinct iterates, in characteristic p. *Circuit:* Frobenius, so $M^{p^m} = I$ is the same as $(M-I)^{p^m} = 0$; then expand $(I+N)^j$ and sum. Three lines to unipotency, one line to the closed form, one line to the kill. Write this.

The ideal solution

The pairs are

$$\boxed{n = 1 \text{ with } p \text{ arbitrary, together with } (p,n) = (2,2).}$$

These work. For $n = 1$ take $v = (1)$ and $M = (1)$, so that $G(x) = x + 1$ and $G^{(k)}(0) = k$; for $k = 1, \dots, p$ these run through all of $\mathbb{F}_p$. (If you ever see the pair $v = (1)$, $M = (0)$ quoted here, check it before you trust it: it gives the constant map $G(x) = 1$, whose iterates of 0 are all equal. The honest witness is the translation.) For $(p,n) = (2,2)$ take

$$v = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \qquad M = \begin{pmatrix} 1 & 1 \\ 0 & 1 \end{pmatrix},$$

and compute over $\mathbb{F}_2$:

$$G^{(1)}(0) = \begin{pmatrix} 0 \\ 1 \end{pmatrix}, \quad G^{(2)}(0) = \begin{pmatrix} 1 \\ 0 \end{pmatrix}, \quad G^{(3)}(0) = \begin{pmatrix} 1 \\ 1 \end{pmatrix}, \quad G^{(4)}(0) = \begin{pmatrix} 0 \\ 0 \end{pmatrix},$$

four distinct vectors, as required.

Nothing else works. Suppose the p^n vectors $G^{(k)}(0)$, $k = 1, \dots, p^n$, are distinct. They then exhaust $\mathbb{F}_p^n$, so $0 = G^{(k_0)}(0)$ for some k_0 with $1 \le k_0 \le p^n$. If $k_0 < p^n$ then $G^{(k+k_0)}(0) = G^{(k)}(0)$ for all $k \ge 0$, so the list takes at most $k_0 < p^n$ values, a contradiction. Hence

$$G^{(p^n)}(0) = 0.$$

Since $G^{(p^n)}$ commutes with G, it fixes $G^{(k)}(0)$ for every k, and these vectors are all of $\mathbb{F}_p^n$; therefore $G^{(p^n)}$ is the identity map.

Induction on k gives the closed form of the affine recursion,

$$G^{(k)}(x) = (I + M + \cdots + M^{k-1})v + M^k x.$$

Apply it with $k = p^n$. Putting $x = 0$ shows the constant term vanishes, and then $M^{p^n} x = x$ for every x, so

$$M^{p^n} = I.$$

Set $N = M - I$. The matrices M and I commute and lie in the commutative ring $\mathbb{F}_p[M]$, which has characteristic p, so the Frobenius identity $(a+b)^p = a^p + b^p$ applies and iterates:

$$N^{p^n} = (M-I)^{p^n} = M^{p^n} + (-I)^{p^n} = M^{p^n} - I = 0,$$

where for $p = 2$ one uses $-I = I$ and for odd p one uses $(-1)^{p^n} = -1$. So N is nilpotent, and a nilpotent $n \times n$ matrix has characteristic polynomial x^n, so Cayley and Hamilton give

$$N^n = 0.$$

Now expand. Since $N^i = 0$ for $i \ge n$,

$$M^j = (I+N)^j = \sum_{i=0}^{n-1} \binom{j}{i} N^i,$$

the integer binomial coefficients being read modulo p. Summing over $j = 0, 1, \ldots, k-1$ and using the telescoping identity $\sum_{j=0}^{k-1} \binom{j}{i} = \binom{k}{i+1}$ (immediate from $\binom{j+1}{i+1} - \binom{j}{i+1} = \binom{j}{i}$) gives the formula that decides the problem:

$$G^{(k)}(0) = \sum_{j=0}^{k-1} M^j v = \sum_{i=0}^{n-1} \binom{k}{i+1} N^i v. \tag{$*$}$$

Lemma 14.5.1. *If $m \geq 1$ and $1 \leq t \leq p^m - 1$ then $p \mid \binom{p^m}{t}$.*

Proof. From $t\binom{p^m}{t} = p^m \binom{p^m - 1}{t-1}$ we get $v_p\left(\binom{p^m}{t}\right) = m - v_p(t)$, where v_p is the exponent of p. Since $1 \leq t \leq p^m - 1$, p^m does not divide t, so $v_p(t) \leq m - 1$ and the exponent is at least 1. □

Suppose now $n \geq 2$ and $p^{n-1} > n$, and take $k = p^{n-1}$ in $(*)$. Every index $i+1$ occurring there satisfies $1 \leq i+1 \leq n \leq p^{n-1} - 1$, so by the lemma every coefficient $\binom{p^{n-1}}{i+1}$ vanishes in $\mathbb{F}_p$ and $G^{(p^{n-1})}(0) = 0$. But $1 \leq p^{n-1} < p^n$, contradicting the fact established above that p^n is the least positive k with $G^{(k)}(0) = 0$. Therefore

$$p^{n-1} \leq n.$$

For $n = 2$ this reads $p \leq 2$, so $p = 2$. For $n \geq 3$ we have $p^{n-1} \geq 2^{n-1} \geq n + 1 > n$ (the inequality $2^{n-1} > n$ for $n \geq 3$ is a one line induction), so no pair survives. Hence the only pairs with $n \geq 2$ are $(p, n) = (2, 2)$, and every pair with $n = 1$ occurs. ■

A remark worth keeping. Formula $(*)$ does not merely exclude; it constructs. For $n = 2$ it says $G^{(k)}(0) = \binom{k}{1} v + \binom{k}{2} N v$, so you need $\binom{k}{1}$ and $\binom{k}{2}$ modulo p to have joint period p^2. For odd p both vanish at $k = p$, which is why odd primes fail; for $p = 2$, $\binom{k}{2}$ modulo 2 has period 4, which is exactly the sporadic case. The obstruction and the example are the same computation read in two directions.

Layer 1: WHAT

Block 1. Distinctness of p^n iterates means the orbit of 0 is everything, hence $G^{(p^n)} = \mathrm{id}$ and no smaller positive power of G fixes 0.

Block 2. Compare linear parts in $G^{(p^n)} = \mathrm{id}$ using the closed form of the affine recursion: $M^{p^n} = I$.

Block 3. Characteristic p plus p-power order gives $(M - I)^{p^n} = 0$, so $N = M - I$ is nilpotent and $N^n = 0$.

Block 4. Expand $M^j = (I + N)^j$, sum with the hockey stick identity, and obtain $G^{(k)}(0) = \sum_{i=0}^{n-1} \binom{k}{i+1} N^i v$: the orbit is governed by n binomial coefficients only.

Block 5. Choose $k = p^{n-1}$, at which all n of those coefficients are divisible by p provided $p^{n-1} > n$; the orbit returns to 0 too early.

Block 6. Solve $p^{n-1} \leq n$ in positive integers, then exhibit the two families that survive.

The nominal route is: take M, put it in Jordan or rational canonical form, split off the eigenvalue 1 part, and grind the iteration block by block. It works, and it takes a page. The bifurcation point is Block 2: do you look at M as a matrix to be canonicalised, or as an element of a finite group whose order you already know? Choosing the second means the only structural fact you ever need about M is that its order is a power of p, and in characteristic p that single fact delivers unipotency for free.

Layer 2: HOW

Block 1. Cue. The list has length p^n and the space has size p^n. **Why natural.** A list of p^n distinct elements of a set of size p^n is the whole set. That forces 0 to appear in the list, and once a point of a forward orbit repeats, the orbit is periodic from the start, so the period must be the full p^n. Restating a distinctness hypothesis as a transitivity statement is always the right first move for

iteration problems, because transitivity is a group theoretic property and distinctness is not. **Link to the next block.** A permutation with a single cycle of length p^n has order p^n, so you now know the order of G as a group element.

Block 2. **Cue.** G is affine, so $G^{(k)}$ is affine, and two affine maps are equal exactly when their linear parts and their constant terms agree. **Why natural.** The closed form $G^{(k)}(x) = (I + \cdots + M^{k-1})v + M^k x$ is the matrix version of the A-level linear recurrence $u_{k+1} = au_k + b$; you should write it without thinking. Comparing the two halves of $G^{(p^n)} = \mathrm{id}$ separates the information: the constant term dies at $x = 0$, and the linear part gives $M^{p^n} = I$ cleanly. **Link to the next block.** An identity of the form $M^{p^m} = I$ over $\mathbb{F}_p$ is not an ordinary order condition, because the exponent is a power of the characteristic.

Block 3. **Cue.** Exponent p^m and characteristic p in the same line. **Why natural.** In any commutative ring of characteristic p the map $a \mapsto a^p$ is a ring homomorphism, so $(M - I)^{p^m} = M^{p^m} - I$. Hence order a power of p is exactly the same statement as $M - I$ nilpotent. This is the definition-level reason and it is why the answer depends on p at all: the problem is asking a question about the characteristic and answering it with the characteristic. Cayley and Hamilton then bounds the nilpotency index by n, which is the only place the dimension enters. **Link to the next block.** Nilpotency lets you expand powers of M as short binomial sums.

Block 4. **Cue.** $M = I + N$ with $N^n = 0$ is an invitation to the binomial theorem, and $G^{(k)}(0)$ is a sum over j of M^j. **Why natural.** Two commuting matrices obey the binomial theorem, and nilpotency truncates the expansion to n terms with coefficients $\binom{j}{i}$. Summing a polynomial in j over j is a hockey stick, so the sum stays inside the same family: $\sum_{j<k} \binom{j}{i} = \binom{k}{i+1}$. The whole orbit is therefore encoded by the residues of n binomial coefficients, which is a drastic reduction of a matrix problem to an arithmetic one. **Link to the next block.** You now control $G^{(k)}(0)$ for every k, so you may choose a k designed to make it vanish.

Block 5. **Cue.** You want an early return to 0, so you want all n coefficients $\binom{k}{i+1}$ to die modulo p at once. **Why natural.** The binomial coefficients $\binom{p^m}{t}$ with $0 < t < p^m$ are exactly the ones divisible by p; that is Kummer's theorem in miniature, and the one line proof is $t\binom{p^m}{t} = p^m\binom{p^m-1}{t-1}$. So $k = p^{n-1}$ kills every term provided the index range $1 \le t \le n$ stays below p^{n-1}. Choosing the largest p-power below p^n is the natural choice, because you need a contradiction with minimality of the period and you want the strongest possible divisibility. **Link to the next block.** The surviving cases are exactly those with $p^{n-1} \le n$, a small arithmetic inequality.

Block 6. **Cue.** A clean inequality in two integer variables. **Why natural.** p^{n-1} grows exponentially in n and n does not, so only the smallest values can survive; the check is $n = 1$ trivially, $n = 2$ giving $p \le 2$, and $2^{n-1} > n$ for $n \ge 3$ closing everything else. Then you must supply constructions, since an obstruction argument never proves existence. **Link to the next block.** The constructions are read straight off formula $(*)$.

Layer 3: WHY

The truth beneath this problem is that in characteristic p, the p-power map is a homomorphism, and therefore p-power order and unipotence are the same phenomenon. Every statement in the solution is a consequence: the linear part is $I + N$ with N nilpotent, the affine map lies in a p-group, and the orbit lengths are constrained by how fast binomial coefficients acquire factors of p. The competition is really asking whether you know that the affine group $\mathrm{AGL}_n(\mathbb{F}_p)$ contains an element of order p^n only in the two listed cases, and that the reason is arithmetic, not linear algebraic.

The discernment mechanism is a question you should ask on every iteration problem over a finite structure: what group does the map live in, and what does the hypothesis say about its order? Distinctness of a full orbit always means "single cycle", which always means "order equals the size of the space". Once the order is known, you interrogate the group, not the matrix. Train this by rebuilding the following chain from memory: order of an element divides the group order; the size of $\mathrm{GL}_n(\mathbb{F}_p)$ is $\prod_{i=0}^{n-1}(p^n - p^i)$; the Sylow p-subgroup of $\mathrm{GL}_n(\mathbb{F}_p)$ is the group of upper unitriangular matrices, of order $p^{n(n-1)/2}$; the exponent of that group is $p^{\lceil \log_p n \rceil}$. The last item alone would kill every $n \geq 4$ here.

The examiners nudged with the phrase "the p^n vectors $G^{(k)}(0)$, $k = 1, 2, \ldots, p^n$", which names the exact size of the space to make sure you convert distinctness into transitivity, and with the choice of an affine rather than a linear map, which guarantees that a fixed point argument alone cannot settle it. A purely linear G could never be transitive at all, since it fixes 0; the vector v is what makes the question non-trivial, and the vector v is exactly what the binomial formula tracks.

Revolving concepts

Affine map and its recursion. $G(x) = v + Mx$ generates $x_{k+1} = Mx_k + v$, whose closed form is $x_k = M^k x_0 + (I + M + \cdots + M^{k-1})v$. Generalised recursion: for any linear operator M and any forcing term, the solution of an inhomogeneous linear recursion is (homogeneous solution) plus (a convolution of the forcing with the powers of the operator), and when $I - M$ is invertible the second term collapses to $(I - M)^{-1}(I - M^k)v$, which is the matrix geometric series. Here $I - M$ is nilpotent, which is precisely why the geometric series must be handled by binomial expansion instead. See Book 1, field: Sequences and recurrences, topic: linear recurrences with constant coefficients.

Unipotent and nilpotent matrices. N is nilpotent when $N^j = 0$ for some j; then its only eigenvalue is 0, its characteristic polynomial is x^n, and $N^n = 0$ by Cayley and Hamilton. M is unipotent when $M - I$ is nilpotent, equivalently when M has order a power of p in characteristic p. See Book 1, field: Linear algebra, topic: nilpotency and canonical forms.

Frobenius. In a commutative ring of characteristic p, $(a + b)^p = a^p + b^p$, because $p \mid \binom{p}{i}$ for $0 < i < p$. Iterating gives $(a + b)^{p^m} = a^{p^m} + b^{p^m}$. This is the single most used fact in finite field arithmetic and it is beyond A-level; install it with its proof. See Book 1, field: Abstract algebra, topic: finite fields.

Binomial coefficients modulo p. $v_p\left(\binom{p^m}{t}\right) = m - v_p(t)$; Kummer's theorem (the exponent of p in $\binom{a+b}{a}$ is the number of carries when adding a and b in base p); Lucas's theorem (the residue of $\binom{a}{b}$ is the product of the residues of the base p digit binomials). See Book 1, field: Number theory, topic: binomial coefficients modulo a prime.

Hockey stick. $\sum_{j=0}^{k-1}\binom{j}{i} = \binom{k}{i+1}$, from telescoping Pascal's rule. Generalisation: any sum of a polynomial over an interval is again a polynomial of one higher degree, and the binomial basis $\binom{k}{i}$ is the basis in which summation is a shift of index; this is the discrete analogue of integrating x^i.

Cyclic permutations and orders. A permutation of a set of size m whose orbit through one point has size m is an m-cycle and has order m. The affine group $\mathrm{AGL}_n(\mathbb{F}_p) = \mathbb{F}_p^n \rtimes \mathrm{GL}_n(\mathbb{F}_p)$ has order $p^n \prod_{i=0}^{n-1}(p^n - p^i)$.

Circuit Bank

- **WHEN** a map on a finite set is required to have all its iterates distinct **DO** say "the map is a single cycle", deduce that its order equals the size of the set, and start reasoning inside the group it belongs to **BECAUSE** distinctness is not a usable property but transitivity and order are, and order interacts with every group theoretic constraint available
- **WHEN** an element of a matrix group over $\mathbb{F}_p$ has order a power of p **DO** write $M = I + N$ and assert N is nilpotent, then $N^n = 0$ **BECAUSE** Frobenius makes $(M - I)^{p^m} = M^{p^m} - I$, so

p-power order and unipotence are the same statement in characteristic p

- **WHEN** you must iterate an affine map **DO** write $G^{(k)}(x) = M^k x + (I + M + \cdots + M^{k-1})v$ immediately **BECAUSE** it is the matrix version of $u_{k+1} = au_k + b$, and separating linear part from constant term separates the information in any identity between affine maps
- **WHEN** a nilpotent matrix appears inside a power or a sum of powers **DO** expand binomially and truncate at the nilpotency index **BECAUSE** the expansion has only n terms, so a matrix statement becomes a statement about n binomial coefficients modulo p
- **WHEN** you need many binomial coefficients to vanish modulo p simultaneously **DO** evaluate at k equal to a power of p **BECAUSE** $\binom{p^m}{t}$ is divisible by p for all $0 < t < p^m$, which is the cleanest simultaneous vanishing available

Backtesting

1. First useful cue: p^n distinct iterates in a space of size p^n, so G is one cycle and its order is p^n, a power of the characteristic.
2. Method the expression suggested: compare linear parts in $G^{(p^n)} = \mathrm{id}$ to get $M^{p^n} = I$, then Frobenius to get unipotence, then the binomial closed form.
3. Where you would hesitate, and the thought to reinforce: at the step $(M - I)^{p^n} = M^{p^n} - I$. Reinforce it by checking the sign for $p = 2$ separately and by remembering that the identity needs commutativity, which holds because everything lives in $\mathbb{F}_p[M]$.
4. The trap the question sets: believing the answer is only $n = 1$ and missing the sporadic $(2, 2)$, or quoting a construction for $n = 1$ without checking it (the constant map $x \mapsto 1$ fails, the translation $x \mapsto x + 1$ works). An exclusion proof with no verified constructions scores badly.
5. What makes you recognise it faster next time: the pairing of a finite field of characteristic p with an orbit length that is a power of p. That pairing is never an accident and always resolves through unipotence and binomial coefficients.
6. The reusable decision rule: convert full orbit distinctness into an order statement, then let the characteristic act on that order.

Mirror drills

1. Show that an element of $\mathrm{GL}_n(\mathbb{F}_p)$ of order a power of p has 1 as its only eigenvalue. *Answer:* $M^{p^s} = I$ gives $(M - I)^{p^s} = 0$ by Frobenius, so $M - I$ is nilpotent and every eigenvalue of M equals 1.
2. Let p be odd, $N \neq 0$ with $N^2 = 0$, $M = I + N$ and $Nv \neq 0$. Find the least $k > 0$ with $G^{(k)}(0) = 0$ for $G(x) = v + Mx$. *Answer:* $G^{(k)}(0) = kv + \binom{k}{2}Nv$, and $k = p$ makes both coefficients vanish modulo p because p is odd, so the answer is p, not p^2; this is exactly why odd primes fail at $n = 2$.
3. Prove $\sum_{j=0}^{k-1} \binom{j}{i} = \binom{k}{i+1}$ and state its continuous analogue. *Answer:* telescope $\binom{j+1}{i+1} - \binom{j}{i+1} = \binom{j}{i}$; the analogue is $\int_0^x t^i \, dt = x^{i+1}/(i+1)$, with $\binom{k}{i}$ playing the role of $k^i/i!$.

14.6 A6. Unit area rectangles force the function to vanish

Let $f(x,y)$ be a continuous, real-valued function on $\mathbb{R}^2$. Suppose that, for every rectangular region R of area 1, the double integral of $f(x,y)$ over R equals 0. Must $f(x,y)$ be identically 0?

Read it out loud

Read the structure before the content. You are handed a family of linear constraints on one unknown function, indexed by the family of admissible domains, and asked whether the constraints have only the trivial solution. That is a rigidity question, and rigidity questions are always won by differentiating the constraint with respect to the parameters of the family, because a derivative of an averaged statement is a less averaged statement. Count the parameters: two for the position of the rectangle, one for its shape, and one for its orientation, with the area pinned at 1. That pinned area is the entire difficulty and the entire hint.

Say why. The reflex for going from averages to point values is to shrink the domain and divide by its measure: the Lebesgue differentiation theorem, or at this level the mean value theorem for integrals. Here you cannot shrink, because every admissible rectangle has area exactly 1. If the hypothesis had been "for every rectangle of area at most 1", the problem would be a one line exercise, since the average over a shrinking square converges to f at the centre. The examiners removed exactly that move. So the family of shapes, which is the hyperbola (width) times (height) equal to 1, has to be exploited as a curve you travel along, not as a family you shrink.

The fields are real analysis (uniform continuity, the fundamental theorem of calculus, differentiating a domain integral in a parameter), a little functional equation theory at the end, and, if you want to see why the answer is yes before you can prove it, harmonic analysis (the Fourier transform of the indicator of a rectangle factorises into two sine factors, and the common zero set over all shapes is only the pair of coordinate axes).

Within thirty seconds you should have written the constraint in coordinates,

$$\int_x^{x+c}\int_y^{y+1/c} f(u,v)\,dv\,du = 0 \qquad (x,y\in\mathbb{R},\ c>0),$$

and noticed that you now have three free parameters to differentiate: x, y and c. The first two differentiations are translations and give you edge identities; the third is the shape derivative and is where the area constraint pays.

Approaches, ranked worst to best

1. **Shrink the rectangle onto a point.** *Cue:* a vanishing integral over a small region. *Circuit:* Lebesgue differentiation. Dead on arrival, because area 1 is not negotiable: as the width goes to 0 the height goes to infinity. Worth ten seconds precisely so that you know which tool has been confiscated and stop reaching for it.

2. **Fourier transform.** *Cue:* the hypothesis is a convolution of f with indicators, vanishing identically. *Circuit:* a convolution vanishes when the spectrum avoids the zeros of the transform of the kernel. The transform of the indicator of a c by $1/c$ rectangle is a product of two sine over argument factors, so a spectral component at frequency (ξ,η) survives only if $\sin(\pi c\xi)\sin(\pi\eta/c)=0$ for every $c>0$, which forces $\xi=0$ or $\eta=0$. That says the spectrum sits on the axes, so morally $f(x,y)=g(x)+h(y)$, and then one dimensional averages finish it. Superb for conviction and inadmissible as a proof, because a merely continuous f with no growth control need not be a tempered distribution. Name this gap rather than hiding it.

3. **Sign chase on a tiling.** *Cue:* if two area 1 rectangles overlap in a rectangle, the two complementary pieces have equal integrals. *Circuit:* propagate a sign along a chain of rectangles until

it contradicts itself. Assume f is somewhere positive, find a nearby point where it is negative, and build an explicit chain of area 1 rectangles that forces the integral over a region where $f < 0$ to be positive. Entirely elementary, no differentiation, but the bookkeeping of the chain is where marks are lost.

4. **Corner parallelogram law, then rotate.** *Cue:* "every rectangular region", so tilted rectangles are allowed too. *Circuit:* a translation invariant vanishing integral gives a discrete periodicity, then use density of an irrational rotation orbit. Show first that opposite corners of any area 1 rectangle satisfy $f(A) + f(C) = f(B) + f(D)$, then apply this around a circle to see that the sum of f at the ends of a diameter is independent of the diameter, then conclude that f restricted to any line is constant. Beautiful, and it uses the full strength of the hypothesis, including orientations. Two costs: you need the density of $\{n\alpha\}$ modulo 2π, and you must handle circles of small diameter separately.

5. **Differentiate the constraint in all three parameters.** *Cue:* area pinned at 1 means the shape parameter is a curve you can travel along. *Circuit:* a vanishing family of domain integrals differentiates to a boundary identity. Two translation derivatives and one shape derivative produce three identities; combining them shows that the integral of f over a horizontal segment of fixed length does not depend on where the segment starts, and a continuous solution of Jensen's functional equation finishes. Shortest, fully rigorous, and it proves more than was asked, since only axis parallel rectangles are used.

The ideal solution

The answer is yes: f must be identically 0. We prove more, namely that it is enough to assume the double integral vanishes over every rectangle of area 1 whose sides are parallel to the coordinate axes.

Setting up a differentiable object. Define

$$F(x,y) = \int_0^x \int_0^y f(u,v)\,dv\,du.$$

Since f is continuous, the inner integral $u \mapsto \int_0^y f(u,v)\,dv$ is continuous in u (uniform continuity of f on compact sets), so the fundamental theorem of calculus gives $F_x(x,y) = \int_0^y f(x,v)\,dv$; by Fubini's theorem the order of integration may be reversed, and the same argument gives $F_y(x,y) = \int_0^x f(u,y)\,du$. Both partial derivatives are continuous on $\mathbb{R}^2$, so F is of class C^1 and the chain rule may be applied to it freely. The hypothesis states that for all $x, y \in \mathbb{R}$ and all $c > 0$,

$$\Phi(x,y,c) := F(x+c, y+\tfrac{1}{c}) - F(x, y+\tfrac{1}{c}) - F(x+c, y) + F(x,y) = 0. \tag{H}$$

First derivative: translate horizontally. Differentiating (H) with respect to x and grouping the two pairs of terms,

$$0 = \Phi_x = \Big[F_x(x+c, y+\tfrac{1}{c}) - F_x(x+c, y)\Big] - \Big[F_x(x, y+\tfrac{1}{c}) - F_x(x, y)\Big],$$

that is,

$$\int_y^{y+1/c} f(x+c,v)\,dv = \int_y^{y+1/c} f(x,v)\,dv. \tag{P}$$

In words: for an area 1 rectangle, the integrals of f over the left and right edges are equal.

Second derivative: move along the shape curve. Now differentiate (H) with respect to c, which is legitimate because F is C^1 and $c \mapsto (x+c,\ y+1/c)$ is smooth on $c > 0$:

$$0 = \Phi_c = \Big[F_x(x+c, y+\tfrac{1}{c}) - F_x(x+c, y)\Big] - \frac{1}{c^2}\Big[F_y(x+c, y+\tfrac{1}{c}) - F_y(x, y+\tfrac{1}{c})\Big],$$

that is,

$$\int_y^{y+1/c} f(x+c,v)\,dv = \frac{1}{c^2}\int_x^{x+c} f(u,y+\tfrac{1}{c})\,du. \tag{S}$$

In words: the mean of f over the right edge equals the mean of f over the top edge (multiply both sides by c and note the edges have lengths $1/c$ and c).

Combining. Substituting (P) into (S) gives, for all $x, y \in \mathbb{R}$ and $c > 0$,

$$\int_x^{x+c} f(u,y+\tfrac{1}{c})\,du = c^2\int_y^{y+1/c} f(x,v)\,dv. \tag{C}$$

Fix $w \in \mathbb{R}$ and $c > 0$, and apply (C) with $y = w - \frac{1}{c}$, first at x and then at $x + c$:

$$\int_x^{x+c} f(u,w)\,du = c^2\int_{w-1/c}^{w} f(x,v)\,dv, \qquad \int_{x+c}^{x+2c} f(u,w)\,du = c^2\int_{w-1/c}^{w} f(x+c,v)\,dv.$$

By (P) with $y = w - \frac{1}{c}$ the two right hand sides are equal, so

$$\int_x^{x+c} f(u,w)\,du = \int_{x+c}^{x+2c} f(u,w)\,du \qquad \text{for all } x \in \mathbb{R},\ c > 0. \tag{J}$$

From (J) to constancy along horizontal lines. Fix w and set $\Theta(x) = \int_0^x f(u,w)\,du$, a C^1 function with $\Theta' = f(\cdot, w)$. Identity (J) says

$$\Theta(x+c) - \Theta(x) = \Theta(x+2c) - \Theta(x+c), \quad \text{i.e.} \quad \Theta(a) + \Theta(b) = 2\,\Theta\!\left(\frac{a+b}{2}\right)$$

for all real $a < b$ (put $a = x$, $b = x + 2c$), and trivially for $a = b$. This is Jensen's functional equation, and a continuous solution of it is affine: the difference between Θ and the affine function agreeing with it at two chosen points is continuous, satisfies the same equation, and vanishes at both points, hence vanishes at every dyadic subdivision point and so, by continuity, everywhere. Therefore $\Theta(u) = \alpha u + \beta$ and

$$f(u,w) = \Theta'(u) = \alpha \quad \text{for all } u \in \mathbb{R}.$$

So f is constant on every horizontal line: writing $g(y) = f(0,y)$, which is continuous, we have $f(x,y) = g(y)$ for all x, y.

Finishing. Substituting into the hypothesis, for all $y \in \mathbb{R}$ and $c > 0$,

$$0 = \int_x^{x+c}\int_y^{y+1/c} g(v)\,dv\,du = c\int_y^{y+1/c} g(v)\,dv,$$

so $\int_y^{y+1/c} g(v)\,dv = 0$ for every y. Fixing c and differentiating with respect to y (again the fundamental theorem of calculus, g being continuous) gives $g(y+\frac{1}{c}) = g(y)$ for all y. As c runs over $(0,\infty)$ the number $\frac{1}{c}$ runs over all positive reals, so every positive real is a period of g, whence g is constant. Then $0 = \int_y^{y+1/c} g = g(0)/c$ forces $g(0) = 0$, so $g \equiv 0$ and $f \equiv 0$. ■

Two remarks worth keeping. First, only axis parallel rectangles were used, so the theorem proved is strictly stronger than the one asked for; noticing that a hypothesis is more generous than your proof needs is itself worth marks, and it is how you should audit every solution. Second, continuity can be relaxed to local integrability if you replace "identically 0" by "0 almost everywhere": the hypothesis is preserved by convolution with any smooth compactly supported bump φ, because $\iint_R (f*\varphi) = \int \varphi(z)\left(\iint_{R-z} f\right)dz = 0$ by Fubini, and $f*\varphi$ is smooth, so the theorem gives $f*\varphi \equiv 0$ for every mollifier and therefore $f = 0$ almost everywhere.

Layer 1: WHAT

Block 1. Commit to the answer yes, and identify the real task: get a pointwise statement out of averaged data when the averaging scale cannot be shrunk.

Block 2. Write the hypothesis in coordinates with the three live parameters x, y, c, the area constraint being built into the height $1/c$.

Block 3. Package the double integral as a second difference of the C^1 function F, so every differentiation you are about to perform is justified in one sentence rather than four.

Block 4. Differentiate in x: the integrals of f over the left and right edges coincide.

Block 5. Differentiate in c: the mean of f over the right edge equals the mean over the top edge. This is the only step that uses the area constraint, and it is the step the confiscated shrinking argument was supposed to replace.

Block 6. Combine the two into (J): the integral of f over a horizontal segment of length c is unchanged when the segment slides by c.

Block 7. Recognise (J) as Jensen's equation for the antiderivative Θ, conclude Θ is affine and f is constant along horizontal lines.

Block 8. Substitute back: f depends only on y, has every positive period, hence is constant, hence 0.

The nominal route is to fix a point, take rectangles shrinking towards it in one direction while stretching in the other, and hope for a limit; it produces nothing because the stretched direction is where all the mass hides. A more sophisticated nominal route is to subtract two overlapping area 1 rectangles and to chase the resulting strip identities combinatorially. The bifurcation point is precisely Block 5: do you differentiate only in the position parameters, which is what everybody does first, or do you differentiate in the shape parameter along the hyperbola $st = 1$? Position derivatives alone give (P), and (P) alone is not enough, since it is satisfied by $f(x, y) = g(y)$ for any g at all. The shape derivative is the only place the number 1 enters the argument.

Layer 2: HOW

Block 1. **Cue.** The question is "must f be 0", a yes or no with a proof attached. **Why natural.** On the Putnam a rigidity question of this shape is almost always yes, because the constraint family is huge (three real parameters plus rotations against one unknown function). Deciding yes immediately is not guessing; it is a decision about which of the two possible pieces of work to start, and you can afford ninety seconds of Fourier heuristic to confirm it before committing. **Link to the next block.** A yes needs a mechanism that converts averages into point values.

Block 2. **Cue.** "For every rectangular region of area 1" with the area fixed. **Why natural.** Fixed area means the family of shapes is one dimensional, parametrised by the width c with height $1/c$. Writing the constraint with all parameters explicit is how you see that you have three derivatives available. Naming them is not bookkeeping; it is the search space of the proof. **Link to the next block.** Before differentiating anything, make the object you are differentiating obviously differentiable.

Block 3. **Cue.** A double integral over a box with variable corners. **Why natural.** Every such integral is a second difference of the antiderivative $F(x, y) = \int_0^x \int_0^y f$. That reduces all four justifications (differentiability, interchange of limits, chain rule) to one statement, namely that F is C^1 with the two partials computed by the fundamental theorem of calculus and Fubini. Graders take marks for differentiating under an integral sign without a word of justification; this packaging makes the justification a single sentence. **Link to the next block.** Now differentiate in the cheapest parameter first.

Block 4. **Cue.** The constraint holds for every x, so $\Phi_x = 0$. **Why natural.** Differentiating a translation invariant constraint measures the flux through the two moving edges: the right edge sweeps outward and the left edge sweeps inward at the same rate, so their contributions must cancel. This is the first variation of a domain integral, and here it says left edge integral equals right edge integral. **Link to the next block.** This identity is genuinely weaker than the problem, since it holds for every function of y alone, so you know it cannot be the last step and you know what the next step must break: the freedom in y.

Block 5. **Cue.** The height is $1/c$, a function of the width; the shape is a curve, not a free parameter. **Why natural.** Differentiating along that curve moves the right edge outward at speed 1 and the top edge downward at speed $1/c^2$, so the flux identity reads (right edge integral) equals $1/c^2$ times (top edge integral). Because the two edges are perpendicular, this couples the horizontal and vertical behaviour of f for the first time, which is exactly what (P) could not do. The general principle is the transport or first variation formula: the derivative of $\int_{R(\lambda)} f$ is the integral of f times the normal speed of the boundary. **Link to the next block.** You now have a horizontal quantity equal to a vertical quantity, so you can eliminate the vertical one using (P).

Block 6. **Cue.** Two identities sharing the term "integral over a vertical edge". **Why natural.** Eliminating a shared term is the standard way to combine two constraints, and here the elimination is engineered by choosing $y = w - 1/c$ so that the top edge sits on the line you care about. What comes out, (J), is a purely one dimensional statement about the single variable function $u \mapsto f(u, w)$, which is a decisive simplification: the plane has been eliminated. **Link to the next block.** A statement about equal integrals over adjacent intervals is a statement about the antiderivative.

Block 7. **Cue.** Equal increments of Θ over adjacent intervals of equal length, for every length. **Why natural.** This is exactly the midpoint condition $\Theta(a) + \Theta(b) = 2\Theta(\frac{a+b}{2})$. You already know the theory: Cauchy and Jensen functional equations have only the affine continuous solutions, by dyadic subdivision plus continuity. Quoting a named theorem here is not laziness, it is the correct use of an installed circuit; the alternative finish (divide by the length and let it tend to 0, using the mean value theorem for integrals to recover $f(0, w) = f(z, w)$) is equally short if the name escapes you in the arena. **Link to the next block.** Constancy along horizontal lines collapses f to a function of y.

Block 8. **Cue.** A function of one variable with every positive number as a period. **Why natural.** Periods form an additive subgroup; a subgroup of $\mathbb{R}$ containing every positive real is everything, so g is constant, and one more substitution into the hypothesis pins the constant at 0. Do not stop at "f is constant", because the problem asked for the value.

Layer 3: WHY

The mathematical truth underneath is the first variation of a domain functional, sometimes met as the Reynolds transport theorem:

$$\frac{d}{d\lambda} \int_{R(\lambda)} f = \int_{\partial R(\lambda)} f \, (V \cdot n) \, dS,$$

where V is the velocity of the moving boundary and n its outward normal. All three identities in the solution are this one formula applied to three different motions of the same rectangle: translation in x (normal speed $+1$ on the right edge and -1 on the left), translation in y, and the shape motion along the hyperbola (normal speed $+1$ on the right edge and $-1/c^2$ on the top). Once you see the hypothesis as "this functional is constant along a three parameter family of motions", the proof

writes itself: constancy of a functional means the flux through the boundary vanishes for every admissible motion, and vanishing flux for enough motions pins the boundary values, which by continuity pins the function.

The discernment mechanism is a single question: what continuous symmetries does the constraint family have, and which of them changes the shape rather than the position? Translations are free and always give you edge identities, but a family closed only under translations can never determine a function beyond its behaviour in the direction of the translations. The area constraint looks like a restriction, but it is the only source of a shape motion, so it is the asset, not the obstacle. Train this by rewriting standard results in the same language: the mean value property of harmonic functions comes from differentiating the average over a sphere in the radius, Green's identities come from differentiating a domain integral in a deformation parameter, and the classical result that a continuous function whose integral over every interval of a fixed length L vanishes must be L-periodic with zero mean is the one dimensional shadow of (P).

The examiners nudged in three ways. They fixed the area at 1 rather than allowing all small rectangles, which removes the Lebesgue differentiation reflex and forces the shape derivative. They said "rectangular region" rather than "rectangle with sides parallel to the axes", which offers the rotation route as a second, more geometric solution and rewards anyone who notices that the extra generosity is not needed. And they asked "must f be identically 0" rather than "prove that f is 0", to keep you honest about the possibility of a counterexample; a candidate who has not at least tried $f(x,y) = \sin(2\pi x)$ and seen why it fails (its integral over a c by $1/c$ rectangle is not identically 0 in c) has not earned the answer.

Revolving concepts

Differentiating a domain integral. For a C^1 family of domains, the derivative of $\int_{R(\lambda)} f$ is the boundary integral of f against the normal speed. In one dimension this is the fundamental theorem of calculus with variable endpoints, $\frac{d}{d\lambda}\int_{a(\lambda)}^{b(\lambda)} f = f(b)b' - f(a)a'$, and every identity in this problem is a two dimensional instance of it. See Book 1, field: Calculus and integrals, topic: Leibniz rule and variable limits.

Regularity bookkeeping. If f is continuous then $F(x,y) = \int_0^x \int_0^y f$ is C^1 with $F_x = \int_0^y f(x,v)dv$ and $F_y = \int_0^x f(u,y)du$, and $F_{xy} = F_{yx} = f$. The ingredients are uniform continuity on compact sets, the fundamental theorem of calculus and Fubini's theorem for continuous functions on a rectangle. This packaging is the cheapest way to make several differentiations legitimate at once. See Book 1, field: Real analysis, topic: differentiation under the integral sign.

Jensen's functional equation. A continuous $\Theta : \mathbb{R} \to \mathbb{R}$ with $\Theta(a) + \Theta(b) = 2\Theta(\frac{a+b}{2})$ for all a, b is affine. The proof is the standard dyadic recursion: the equation propagates the value at the midpoint, hence at every dyadic rational subdivision point of any interval, and continuity fills the rest. Generalised recursion: the same subdivision argument proves that continuous solutions of the Cauchy equation $h(a+b) = h(a) + h(b)$ are linear, that midpoint convex continuous functions are convex, and in general that a functional equation which determines the midpoint value determines the function on a dense set, so continuity upgrades "determined on the dyadics" to "determined everywhere". See Book 1, field: Functional equations, topic: Cauchy and Jensen.

Periods of a function. The set of periods of a function is an additive subgroup of $\mathbb{R}$; if it is not cyclic it is dense, and if the function is continuous with a dense set of periods it is constant. Here every positive real is a period, so constancy is immediate. See Book 1, field: Real analysis, topic: periodic functions and subgroups of the reals.

Mean value theorem for integrals and Lebesgue differentiation. $\frac{1}{h}\int_a^{a+h} f \to f(a)$ as $h \to 0^+$ for continuous f; the almost everywhere version for locally integrable f is the Lebesgue differentiation theorem. This is the standard route from averaged information to pointwise information, and recognising that the fixed area hypothesis blocks it is the first strategic decision in the problem.

Fourier transform of an indicator, and the honest gap. The transform of the indicator of $[0,c]\times[0,1/c]$ is a product of factors of the form $\sin(\pi c\xi)/(\pi\xi)$ and $\sin(\pi\eta/c)/(\pi\eta)$ up to phases, so the common zero set over all $c>0$ is the union of the two axes. This explains, in ninety seconds, why the answer must be yes and why the extremal near-counterexamples look like functions of one variable. Turning it into a proof requires the theory of tempered distributions and a growth hypothesis on f that the problem does not supply, so it is a knowledge gap to name and not a solution to submit. See Book 1, field: Real analysis, topic: Fourier transforms and convolution.

Density of an irrational rotation orbit. The alternative rotation solution needs the fact that $\{n\alpha \bmod 2\pi\}$ is dense when $\alpha/2\pi$ is irrational, together with continuity to pass from a dense set of angles to all of them. This is Weyl's equidistribution theorem in its weakest form and is beyond A-level; install it. See Book 1, field: Number theory, topic: equidistribution.

Circuit Bank

- **WHEN** a hypothesis says an integral over every member of a parametrised family of domains vanishes **DO** differentiate the identity with respect to each parameter and read the result as a boundary or edge identity **BECAUSE** a constant functional has zero first variation, and the first variation of a domain integral is the flux of the integrand through the moving boundary
- **WHEN** the size of the domains is pinned so you cannot shrink them **DO** stop hunting for a Lebesgue differentiation argument and look instead for a shape parameter to differentiate along **BECAUSE** a pinned measure removes the shrinking family but creates a one dimensional family of shapes, and that family is the only place the pinned constant can enter the proof
- **WHEN** you are about to differentiate a multiple integral with several variable limits **DO** first rewrite it as a second difference of a single antiderivative F and prove once that F is C^1 **BECAUSE** one justification then covers every subsequent differentiation, and unjustified interchanges are the standard way to lose marks on an otherwise correct analysis solution
- **WHEN** you obtain equal increments of a function over adjacent intervals of arbitrary equal length **DO** recognise Jensen's equation and quote that continuous solutions are affine **BECAUSE** the midpoint condition determines the function on a dense dyadic set, so continuity forces affinity, and its derivative is then constant
- **WHEN** a continuous function turns out to have arbitrarily small periods **DO** conclude at once that it is constant **BECAUSE** the periods form an additive subgroup of the reals, and a non-cyclic subgroup is dense

Backtesting

1. First useful cue: the area is pinned at 1, so the usual shrink and divide route is confiscated and the shape parameter c with height $1/c$ must be the engine.
2. Method the expression suggested: write the constraint with parameters x, y, c, package it as a second difference of a C^1 antiderivative, and differentiate in x and in c.
3. Where you would hesitate, and the thought to reinforce: after getting the edge identity (P), because it looks like progress but is satisfied by every function of y alone. Reinforce the habit of asking which functions satisfy the identity you just derived; if the answer is a big family, you have not used the key hypothesis yet.
4. The trap the question sets: the fixed area, which blocks the reflex; and the word "rectangular region", which tempts you into a rotation argument that is longer than necessary.

5. What makes you recognise it faster next time: any hypothesis of the form "the integral over every domain of exactly this size vanishes". The measure being fixed is the signal to differentiate along the shape family.
6. The reusable decision rule: constraints indexed by a family of domains are differentiated in the family parameters, and the parameter that changes shape rather than position is the one that carries the content.

Mirror drills

1. Let h be continuous on $\mathbb{R}$ with $\int_x^{x+1} h = 0$ for every x. What can you say about h? *Answer:* differentiating in x gives $h(x+1) = h(x)$, so h is 1-periodic with mean zero; $h(x) = \sin 2\pi x$ shows it need not vanish, which is exactly why the two dimensional problem needs a shape derivative as well as a translation derivative.
2. Suppose f is continuous on $\mathbb{R}^2$ and its integral over every square of area 1 (sides parallel to the axes) vanishes. Must f be 0? *Answer:* no. Take $f(x,y) = \sin(2\pi x)$, whose integral over any unit square vanishes; the family of squares has no shape parameter, so only translations are available and the argument of this problem collapses.
3. Show that a continuous Θ with $\Theta(a) + \Theta(b) = 2\Theta(\frac{a+b}{2})$ for all a, b is affine. *Answer:* subtract the affine function agreeing with Θ at 0 and 1; the difference satisfies the same equation and vanishes at 0 and 1, hence at every dyadic rational in $[0,1]$ by repeated midpointing, hence identically by continuity, and the same on every interval.

14.7 B1. closure under addition and composition forces multiplication

Let S be a class of functions from $[0,\infty)$ to $[0,\infty)$ that satisfies:

(i) the functions $f_1(x) = e^x - 1$ and $f_2(x) = \ln(x+1)$ are in S;

(ii) if $f(x)$ and $g(x)$ are in S, the functions $f(x) + g(x)$ and $f(g(x))$ are in S;

(iii) if $f(x)$ and $g(x)$ are in S and $f(x) \geq g(x)$ for all $x \geq 0$, then the function $f(x) - g(x)$ is in S.

Prove that if $f(x)$ and $g(x)$ are in S, then the function $f(x)g(x)$ is also in S.

Read it out loud

This is a closure system, not a functional equation to solve. Condition (i) hands you two named generators, condition (ii) closes S under addition and under composition, and condition (iii) closes S under a subtraction that only fires when the result stays nonnegative. Nowhere is multiplication a primitive operation, so the whole task is to manufacture the pointwise product $f(x)g(x)$ out of sums, compositions and one guarded subtraction. That absence is the entire question, and it is what you should read out loud first: you are being asked to reroute a multiplicative target through additive machinery.

Look at the two generators as a pair rather than as two isolated examples. Compute $f_2(f_1(x)) = \ln(e^x) = x$ and $f_1(f_2(x)) = e^{\ln(x+1)} - 1 = x$: they are mutually inverse on $[0,\infty)$, and they are the shifted exponential and shifted logarithm, shifted by exactly enough to fix 0 and to keep both maps inside $[0,\infty)$. That is the cue to write down within thirty seconds: whenever a closure axiom names two functions that compose to the identity, a conversion pair has been handed to you, and here that pair is exp and log, the pair that turns sums into products and back. The circuit that should fire is immediate: product wanted, only sums and compositions available, so route through log and exp.

Approaches, ranked worst to best

1. **treat $f(g(x))$ as a stand in for $f(x)g(x)$.** *Cue:* condition (ii) lists composition and the target is a product, so the two look related. *Circuit:* misreading the axiom rather than applying it. This route confuses two different operations that happen to sit in the same clause. No manipulation of $f \circ g$, $f \circ f$ or $g \circ g$ alone reduces to the pointwise product for generic f, g; the ink spent here buys nothing, since no such identity exists.

2. **structural induction on how f was generated from f_1, f_2.** *Cue:* S is built up by finitely many applications of (ii) and (iii), so induction on that construction feels natural. *Circuit:* induction on the closure system itself. This can be forced to work, proving at each generating step that multiplying by f_1 or f_2 stays in S, but it reintroduces the real content, the log and exp identity below, as a base case while adding a whole induction scaffold on top for no extra rigour. It costs far more ink than the direct route for the same conclusion.

3. **route the product through logs and exponentials.** *Cue:* f_1, f_2 are a shifted exp and log pair, mutually inverse, fixing 0. *Circuit:* convert a product into a sum via ln, close under (ii), convert back via exp, then subtract the surplus with (iii). This uses every axiom for exactly what it is best at and reaches fg in a fixed number of steps regardless of how f and g were themselves built.

The ideal solution

Fix $f, g \in S$.

Since $f_2(x) = \ln(x+1) \in S$ by (i), condition (ii) applied to the pairs (f_2, f) and (f_2, g) gives

$$\ln(f(x)+1) \in S, \qquad \ln(g(x)+1) \in S.$$

Condition (ii) applied to these two functions gives their sum in S:

$$h(x) := \ln(f(x)+1) + \ln(g(x)+1) \in S.$$

Since $f_1(x) = e^x - 1 \in S$ by (i) and $h \in S$, condition (ii) applied to the pair (f_1, h) gives $f_1(h(x)) \in S$, and

$$f_1(h(x)) = e^{\ln(f(x)+1)+\ln(g(x)+1)} - 1 = (f(x)+1)(g(x)+1) - 1 = f(x)g(x) + f(x) + g(x).$$

Write $F(x) := f(x)g(x) + f(x) + g(x)$, so $F \in S$. Condition (ii) applied directly to f, g gives $G(x) := f(x) + g(x) \in S$. Because f, g take values in $[0, \infty)$, $f(x)g(x) \geq 0$ for every $x \geq 0$, so $F(x) \geq G(x)$ for every $x \geq 0$. Condition (iii) applied to the pair (F, G) is therefore licensed, and gives

$$F(x) - G(x) = f(x)g(x) \in S,$$

which is the claim.

Layer 1: WHAT

Block 1: read f_1, f_2 as a shifted, mutually inverse exp/log pair inside S. Block 2: build $\ln(f+1)$ and $\ln(g+1)$ by composing f_2 onto f and onto g. Block 3: add them, which is secretly $\ln((f+1)(g+1))$. Block 4: compose f_1 back onto that sum, exponentiating to reach $(f+1)(g+1) - 1 = fg + f + g$. Block 5: obtain $f + g$ directly from (ii), untouched since Block 1. Block 6: subtract, licensed by $fg \geq 0$, to isolate fg.

The nominal route stares at $f(x) + g(x)$ and $f(g(x))$ in (ii) and tries to force one of them, or some blend of them, straight onto f and g, never returning to f_1, f_2 after reading (i). The bifurcation point is the line where you ask what the two named generators are actually for: the zombie route treats them as a starting example and forgets them; the ideal route treats them as the conversion machine that (ii) and (iii) alone cannot supply.

Layer 2: HOW

Block 1. **Cue.** Condition (i) names exactly two functions, and $f_1 \circ f_2 = f_2 \circ f_1 = \mathrm{id}$ on $[0, \infty)$. **Why natural.** Multiplication is not a primitive of (ii) or (iii); the only way to reach it is to convert it into an operation those axioms do supply, and a mutually inverse pair that is also exp and log, shifted to respect $[0, \infty)$, is exactly that conversion. **Link to the next block.** Apply f_2 to f and to g separately, since ln is what turns the product you do not yet have into a sum you can build with (ii).

Block 2. **Cue.** $f, f_2 \in S$, and (ii) closes S under composition. **Why natural.** Composing f_2 onto f is evaluating ln at $f(x)$; it is the only way inside this axiom system to apply a fixed function pointwise to f. **Link to the next block.** Two functions of the same shape are now in S, ready to be added, which is the other half of (ii).

Block 3. **Cue.** Two elements of S with nothing left to do to them individually. **Why natural.** (ii) closes S under addition, and $\ln(f+1) + \ln(g+1) = \ln((f+1)(g+1))$ already encodes the product, wrapped in a logarithm. **Link to the next block.** The natural next step is to strip the logarithm off by composing with its inverse, f_1.

Block 4. **Cue.** $f_1 \in S$ from (i), and the sum from Block 3 is in S. **Why natural.** Composing f_1 undoes the logarithm because $f_1 = f_2^{-1}$ on $[0, \infty)$, handing you $(f+1)(g+1) - 1 = fg + f + g$. **Link to the next block.** You have reached fg plus a known surplus $f + g$; the remaining task is to obtain that surplus separately and subtract it.

Block 5. **Cue.** $f, g \in S$ directly, with (ii) untouched since Block 1. **Why natural.** Addition of the two original functions is the most primitive move (ii) offers, used here for exactly the surplus Block 4 left behind. **Link to the next block.** $fg + f + g$ and $f + g$ are both in S and differ pointwise by $fg \geq 0$, which is exactly what condition (iii) needs as its licence.

Block 6. **Cue.** $F(x) = fg + f + g \geq G(x) = f + g$ for every x, because f, g map into $[0, \infty)$. **Why natural.** Condition (iii) is a restricted operation, only available once $F \geq G$ pointwise is certified; here that certification is one line, since S's codomain is nonnegative by definition. **Link to the next block.** Subtracting removes exactly the surplus Block 4 introduced, leaving fg alone in S.

Layer 3: WHY

The route works because multiplication is conjugate to addition through any pair of mutually inverse maps carrying a multiplicative identity to an additive one, which is exactly what log and exp are, and condition (i) hands you a copy of that pair already living inside S. Whenever a closure axiom gives you two named generators, compute whether they compose to the identity in either order; if they do, a conversion pair has been handed to you, and the problem is asking you to move between the two structures it connects.

The shift is the second truth worth naming. Plain ln and plain exp cannot survive inside S, since S's codomain is $[0, \infty)$ rather than $\mathbb{R}$ or $(0, \infty)$; shifting by 1 is what keeps f_1, f_2 mapping $[0, \infty)$ to $[0, \infty)$ and fixing 0. That shift is precisely what produces the surplus $f + g$ in Block 4: exponentiating $\ln(f+1) + \ln(g+1)$ gives $(f+1)(g+1)$, not fg, and the extra 1 on each factor has to be paid for. Train yourself to expect this: a generator that is the standard function, shifted to fix a boundary condition, leaves a correction term downstream proportional to the shift, and the remaining axiom in the problem, here (iii), usually exists specifically to pay that correction off.

Revolving concepts

Closure systems generated by axioms. S is the smallest class containing f_1, f_2 and closed under (ii) and (iii); reasoning about membership means exhibiting an explicit finite chain of operations reaching the target, exactly as for a subgroup or subring generated by a set. See Book 1, field: Abstract algebra, topic: generated substructures.

The log/exp conjugacy. For $a, b > 0$, $ab = e^{\ln a + \ln b}$; more generally a strictly monotonic bijection ϕ conjugates addition on its codomain to the operation $x * y := \phi^{-1}(\phi(x) + \phi(y))$ on its domain. Here $\phi(x) = \ln(x+1)$ conjugates multiplication of $(f+1), (g+1)$ to addition of logs, and the shift by 1 keeps ϕ valued on $[0, \infty)$. See Book 1, field: Real analysis, topic: exponential and logarithmic identities.

Restricted subtraction under pointwise order. Condition (iii) is a partial operation, defined only where $f \geq g$ pointwise; every invocation needs its own one line certificate. See Book 1, field: Inequalities, topic: pointwise ordering of functions.

There is no recursive or inductive structure anywhere in this problem, so no generalised recursion entry applies here.

Circuit Bank

- **WHEN** a closure axiom names two specific generator functions **DO** compute whether they compose to the identity, in either order **BECAUSE** because a mutually inverse pair among your generators is almost always the mechanism the problem is built on, most often an exp/log or a reflection pair
- **WHEN** you need a product but your closure rules only supply sums and compositions **DO** route through ln, add, then route back through exp **BECAUSE** because $e^{a+b} = e^a e^b$ is the elementary identity that turns addition into multiplication, and composition with a fixed function is how a closure system lets you apply it pointwise
- **WHEN** a generator is a shift of a standard function to respect a boundary such as $[0, \infty)$ **DO** expect a residual linear correction term after exponentiating or inverting, and expect it to need cleaning up **BECAUSE** because the shift is extra content that the identity you are using does not automatically cancel
- **WHEN** a closure system gives a restricted subtraction, valid only when $f \geq g$ pointwise **DO** certify the inequality explicitly, in one line, at the point you invoke it **BECAUSE** because an uncertified use of a partial operation is a step of rigour lost, however obvious the inequality is

Backtesting

1. First useful cue: f_1 and f_2 compose to the identity on $[0, \infty)$, so they are a shifted exp/log pair, not two unrelated examples.
2. Method the expression suggested: convert the target product into a sum of logs with f_2, close under addition, then invert with f_1 to exponentiate back.
3. Where you would hesitate, and the thought to reinforce: exponentiating lands on $fg + f + g$, not fg alone; the thought to reinforce is that a shifted generator always leaves a correction term, and the unused axiom exists to remove it.
4. The trap the question sets: reading the two clauses of (ii), addition and composition, as competing candidates for the multiplication rule, and forcing one onto f and g without ever returning to f_1, f_2.
5. What makes you recognise it faster next time: any closure axiom naming two specific functions as generators is worth testing for a compositional inverse relationship before anything else.
6. The reusable decision rule: when a closure system gives addition and composition but the goal is multiplication, search the generators for an exp/log style conjugacy pair, use it to convert, then use the remaining axiom to pay for whatever shift the generators carried.

Mirror drills

1. Let T be the smallest class of functions $\mathbb{R} \to \mathbb{R}$ containing $x \mapsto x$ and the constant 1, closed under addition, subtraction and multiplication. Show every polynomial with integer coefficients lies in T. *Answer:* build each power x^k by repeated multiplication, available directly here, then assemble integer coefficients by repeated addition and subtraction of 1.
2. A class S of positive reals contains 2 and 3 and is closed under $a * b = ab$ and, when $a \geq b$, a/b. Show $2^n 3^m \in S$ for all integers n, m with $2^n 3^m \geq 1$. *Answer:* the same generator discipline as B1, except multiplication is already primitive here, so the only work is bookkeeping the sign of each exponent against the restricted division.

3. A class S of functions $[0,\infty) \to [0,\infty)$ is closed under addition and composition and contains only sin and cos, suitably restricted, as generators beyond the identity. Explain why the B1 route to $f(x)g(x)$ has no foothold here. *Answer:* sin, cos do not form a mutually inverse pair conjugating an additive structure to a multiplicative one, so there is no analogue of ln, exp to route the product through.

14.8 B2. a single face bounds the ball count below

Let P be a given (non-degenerate) polyhedron. Prove that there is a constant $c(P) > 0$ with the following property: if a collection of n balls whose volumes sum to V contains the entire surface of P, then $n > c(P)/V^2$.

Read it out loud

Strip away the word polyhedron for a moment and read the shape of the claim: a count n, a volume V, and an inequality relating them of the form $n >$ constant$/V^2$. That exponent of 2 on V is a command term in disguise. Volumes of balls scale as the cube of radius, so V carries a degree 3 flavour in the radii r_i; an inequality with V^2 on the bottom and a bare count n on top is exactly the shape produced by comparing a degree 2 quantity, cubed, against a degree 3 quantity, squared, which is the signature of Hölder's inequality, or equivalently the power mean inequality, applied to two different powers of the same sequence (r_i).

That tells you what the middle of the proof must be. The two ends still need supplying: a degree 2 quantity in the radii, bounded below by something depending only on P, and a degree 3 quantity, which is V itself up to a constant, bounded above. A sum of squared radii is exactly what controls an area, since a disc of radius r has area πr^2; and the problem gives you balls covering the entire surface of a polyhedron, which certainly includes covering any single flat face. The moment you notice you do not need the whole surface of P, only one flat piece of it, the three dimensional covering problem collapses to a two dimensional one: discs covering a fixed flat region. That is the reroute, and it is the whole saving in this problem.

The cue to write down inside thirty seconds: fix one face F of P with area $A > 0$, which exists because P is non-degenerate, note that balls covering P's surface in particular cover F, and prepare to bound $\sum r_i^2$ below by A/π while $\sum r_i^3$ is bounded above by a multiple of V.

Approaches, ranked worst to best

1. **estimate covering efficiency over the whole surface of P at once.** *Cue:* the hypothesis is phrased in terms of the entire surface. *Circuit:* a direct global covering argument. Faces meet at edges and vertices, where a ball can cover parts of two or three faces at once, at angles that are awkward to bound uniformly; controlling the covering efficiency across every edge and corner of P simultaneously costs far more ink than the target inequality actually requires, and it still has to end in essentially the same Hölder step below.

2. **bound n using Cauchy-Schwarz on $\sum r_i^2$ and $\sum r_i^3$ alone.** *Cue:* Cauchy-Schwarz is the default inequality between two sums of powers. *Circuit:* two term Cauchy-Schwarz applied directly to the radii. Cauchy-Schwarz compares two sums of matched degree; here three different objects are in play at once, the bare count $n = \sum 1$, $\sum r_i^2$ and $\sum r_i^3$, at degrees $0, 2, 3$ respectively, and no single application of Cauchy-Schwarz ties all three together. Chaining two applications reaches the same place as one application of Hölder, at strictly more cost for the same conclusion.

3. **restrict to one flat face, then apply Hölder's inequality with exponents $3/2$ and 3 to $\sum r_i^2$.** *Cue:* a face gives a clean two dimensional covering bound, and the target inequality's V^2 signals a degree matching problem across three quantities. *Circuit:* extremal restriction to a sub-structure, followed by a single Hölder step chosen to match degrees $0, 2, 3$. This is the ideal route: it discards everything about P except that it has one flat face of positive area, reduces the whole problem to bounding how much of a plane a disc can cover, and closes the degree gap between count, area and volume in one clean inequality.

The ideal solution

Since P is a non-degenerate polyhedron, fix one of its faces F, a planar polygon of area $A > 0$, lying in a plane Π.

Suppose n balls of radii $r_1, \dots, r_n$, with $\sum_{i=1}^{n} \frac{4}{3}\pi r_i^3 = V$, together contain the entire surface of P; in particular they cover F. The intersection of the i-th ball with the plane Π is either empty or a disc of some radius $\rho_i \le r_i$: if the ball's centre lies at perpendicular distance d_i from Π with $d_i < r_i$, the intersection is a disc of radius $\sqrt{r_i^2 - d_i^2} \le r_i$; if $d_i \ge r_i$, the intersection with Π is empty, covered by taking $\rho_i = 0 \le r_i$. This disc has area at most $\pi\rho_i^2 \le \pi r_i^2$, so the portion of F that the i-th ball can possibly cover has area at most πr_i^2. Since the n balls together cover all of F,

$$\sum_{i=1}^{n} \pi r_i^2 \;\ge\; A, \qquad \text{that is} \qquad \sum_{i=1}^{n} r_i^2 \;\ge\; \frac{A}{\pi}.$$

Next, from the volume constraint,

$$\sum_{i=1}^{n} r_i^3 = \frac{3V}{4\pi}.$$

Now connect n, $\sum r_i^2$ and $\sum r_i^3$ by Hölder's inequality with conjugate exponents $p = 3/2$ and $q = 3$, so that $1/p + 1/q = 2/3 + 1/3 = 1$, applied to the sequences $(r_i^2)_{i=1}^n$ and $(1)_{i=1}^n$:

$$\sum_{i=1}^{n} r_i^2 \cdot 1 \;\le\; \left(\sum_{i=1}^{n} (r_i^2)^{3/2}\right)^{2/3} \left(\sum_{i=1}^{n} 1^3\right)^{1/3} = \left(\sum_{i=1}^{n} r_i^3\right)^{2/3} n^{1/3}.$$

Cubing both sides,

$$\left(\sum_{i=1}^{n} r_i^2\right)^3 \;\le\; \left(\sum_{i=1}^{n} r_i^3\right)^2 n.$$

Substituting the two bounds established above,

$$\left(\frac{A}{\pi}\right)^3 \;\le\; \left(\sum_{i=1}^{n} r_i^2\right)^3 \;\le\; \left(\sum_{i=1}^{n} r_i^3\right)^2 n = \left(\frac{3V}{4\pi}\right)^2 n,$$

so

$$n \;\ge\; \frac{A^3/\pi^3}{9V^2/(16\pi^2)} = \frac{16A^3}{9\pi V^2}.$$

This holds for every collection of balls covering the surface of P, so any constant $c(P)$ with $0 < c(P) < \dfrac{16A^3}{9\pi}$ satisfies $n > c(P)/V^2$ for every such collection: the strict inequality $n > c(P)/V^2$ is safe even though the derived bound on n is only $\ge$, precisely because $c(P)$ is chosen strictly smaller than the constant the argument produces. Taking

$$\boxed{c(P) = \frac{8A^3}{9\pi}}$$

where A is the area of any fixed face of P, completes the proof.

Layer 1: WHAT

Block 1: fix one flat face F of P, area $A > 0$. Block 2: note the covering balls cover F too, and bound the area of F a single ball of radius r_i can cover by πr_i^2, giving $\sum r_i^2 \ge A/\pi$. Block 3: rewrite the volume constraint as $\sum r_i^3 = 3V/(4\pi)$. Block 4: apply Hölder's inequality with exponents $3/2$ and 3 to link n, $\sum r_i^2$ and $\sum r_i^3$ in one line. Block 5: substitute the two bounds and solve for n, then choose $c(P)$ strictly below the resulting constant to convert $\ge$ into $>$.

The nominal route tries to control the covering efficiency across the whole boundary of P at once, tracking how balls straddle edges and vertices between faces. The bifurcation point is the decision to discard almost all of P's structure and keep only one flat face: the zombie route treats "surface of P" as a single indivisible object to be estimated globally; the ideal route notices that a lower bound only needs one witness, and the cheapest witness available is a single flat face reduced to a two dimensional disc covering problem.

Layer 2: HOW

Block 1. **Cue.** P is non-degenerate, so it has at least one genuine flat polygonal face of positive area. **Why natural.** A lower bound on n only needs one sufficient obstruction, not a full account of P's geometry; a flat face is the cheapest piece of P to reason about, since it reduces three dimensional covering to two dimensional covering. **Link to the next block.** Since the given balls cover the whole surface of P, they in particular cover F, so every subsequent bound only has to reason about how balls interact with one fixed plane.

Block 2. **Cue.** A ball intersected with a plane is a disc of radius at most the ball's own radius. **Why natural.** This is the sharpest elementary bound available: the cross section of a sphere by a plane at perpendicular distance d from the centre has radius $\sqrt{r^2 - d^2} \leq r$, with equality only when the plane passes through the centre. **Link to the next block.** Summing the maximum possible area each ball contributes to covering F gives a lower bound on $\sum r_i^2$ purely in terms of A, the one geometric invariant of P that survives into the proof.

Block 3. **Cue.** The hypothesis already gives V as a sum of ball volumes. **Why natural.** Expressing V as $\sum \frac{4}{3}\pi r_i^3$ makes $\sum r_i^3$, the other quantity Hölder needs, immediately available as a constant multiple of V. **Link to the next block.** Two power sums of the radii, at degrees 2 and 3, are now both bounded in terms of the data; the count n is the third and last quantity to bring in.

Block 4. **Cue.** Three quantities of different degrees in the radii, n at degree 0, $\sum r_i^2$ at degree 2, $\sum r_i^3$ at degree 3, need to be tied together in one inequality. **Why natural.** Hölder's inequality is exactly the tool for comparing sums of different powers of the same sequence; choosing exponents $3/2$ and 3 on the pair $((r_i^2), (1))$ is forced by wanting $\sum r_i^2$ cubed to sit against $\sum r_i^3$ squared and n, which is precisely the shape A^3/V^2 demanded by the problem. **Link to the next block.** The resulting inequality $(\sum r_i^2)^3 \leq (\sum r_i^3)^2 n$ only needs the two bounds from Blocks 2 and 3 substituted in.

Block 5. **Cue.** $(\sum r_i^2)^3 \leq (\sum r_i^3)^2 n$, with $\sum r_i^2 \geq A/\pi$ and $\sum r_i^3 = 3V/(4\pi)$. **Why natural.** Substituting both bounds and simplifying the constants gives an explicit lower bound $n \geq 16A^3/(9\pi V^2)$, valid for every covering collection, so any constant strictly below $16A^3/(9\pi)$ turns the problem's required strict inequality $n > c(P)/V^2$ into a true statement for every such collection. **Link to the next block.** The proof is complete once $c(P)$ is named explicitly as a positive number depending only on P, through the area A of a chosen face, exactly as the problem demands.

Layer 3: WHY

The underlying truth is dimensional. A covering-by-balls hypothesis on a k dimensional object, here a flat face with $k = 2$, bounds a sum of r_i^2 terms below, while a volume hypothesis in ambient 3 dimensional space bounds a sum of r_i^3 terms above. Any inequality asked for between a count n and a ratio of a lower dimensional bound to a higher dimensional one is, structurally, a request to interpolate between the degree 0 sum $n = \sum r_i^0$ and two other power sums, and interpolation between power sums of the same sequence at different exponents is exactly what Hölder's inequality measures. The discernment mechanism is to read the exponents straight off the target inequality,

cube the area bound, square the volume bound, and choose Hölder's conjugate exponents to reproduce exactly that shape, rather than guessing an inequality and hoping the algebra matches.

The second truth is about restriction. The problem states a hypothesis about the entire surface of P, but a lower bound on n is an existential statement in disguise: you only need to exhibit one obstruction that already forces n to be large, not to characterise every possible covering. Any time a covering or packing problem gives an entire boundary or domain as the hypothesis but only asks for an inequality, not an exact characterisation, test whether a single convenient sub-piece of that domain, chosen for the invariant it carries, already carries the whole force of the bound. Here that invariant is flatness: a face's cross sections with a fixed plane are cleanly bounded by the ball's own radius, which is not true of a curved or creased piece of the surface. Training this discernment means asking, before attempting any global estimate, what is the cheapest sub-object that already forces the conclusion.

Revolving concepts

Hölder's inequality. For conjugate exponents $p, q > 1$ with $1/p + 1/q = 1$ and nonnegative reals a_i, b_i, $\sum a_i b_i \leq (\sum a_i^p)^{1/p} (\sum b_i^q)^{1/q}$; taking $b_i \equiv 1$ recovers a power mean comparison between $\sum a_i$ and $\sum a_i^p$ weighted by $n^{1/q}$. Choosing p, q to match the degrees appearing in a problem, here $p = 3/2, q = 3$ to compare degree 2 and degree 3 power sums against a bare count, is the whole skill. See Book 1, field: Inequalities, topic: Hölder and power mean inequalities.

Cross sections of a ball. The intersection of a sphere of radius r centred at O with a plane at perpendicular distance d from O is empty if $d > r$, a single point if $d = r$, and a circle of radius $\sqrt{r^2 - d^2}$ if $d < r$; in every case the radius of the cross section is at most r. See Book 1, field: Geometry, topic: spheres and planar sections.

Existential constants depending only on a fixed object. A claim of the form "there is a constant $c(P) > 0$ such that..." asks you to exhibit one explicit formula for $c(P)$ from the data attached to P, here the area of one face, then verify the inequality for every instance of the varying data, n, the radii, and V. See Book 1, field: Inequalities, topic: existential bounding constants.

There is no recursive or inductive structure in this problem; the argument is a single chain of estimates, not a process that repeats on smaller instances of itself, so no generalised recursion entry applies here.

Circuit Bank

- **WHEN** an inequality is demanded between a count n and powers of two other quantities built from the same underlying variables **DO** identify the degree of each quantity in those variables and choose Hölder's, or power mean, exponents to reproduce exactly the powers demanded **BECAUSE** because the target inequality's exponents are not decoration, they specify which conjugate exponents will close the proof in one line
- **WHEN** a hypothesis is stated about an entire surface, boundary or domain but the goal is only an inequality, not a full characterisation **DO** search for the cheapest convenient sub-piece of that domain that already forces the inequality, and discard the rest **BECAUSE** because a lower or upper bound is an existential claim, one witness sub-piece is enough, and a flat or otherwise structured piece is always easier to estimate than the whole
- **WHEN** a ball, sphere or disc is intersected with a fixed plane or line **DO** bound the resulting cross section's size by the object's own radius, with equality exactly when the plane or line passes through the centre **BECAUSE** because this is the sharp elementary bound and almost always as much geometric information as a covering argument needs
- **WHEN** a problem asks you to produce a constant $c(P)$ depending only on a fixed object P **DO** derive the sharpest constant your argument naturally produces, then explicitly weaken

it, for instance by halving it, to convert a nonstrict bound into the strict inequality asked for **BECAUSE** because the problem's strict inequality is a request to certify that some smaller constant works, not to find the exact extremal constant

Backtesting

1. First useful cue: the target inequality has V^2 in the denominator against a bare count n, and V is cubic in the radii, which signals a degree matching problem across three power sums.
2. Method the expression suggested: fix one flat face, bound $\sum r_i^2$ below via disc covering, bound $\sum r_i^3$ above via the volume constraint, and connect all three quantities with Hölder's inequality.
3. Where you would hesitate, and the thought to reinforce: choosing which exponents to feed Hölder's inequality; the thought to reinforce is to read the exponents straight off the target inequality's own powers rather than guessing.
4. The trap the question sets: trying to estimate the covering efficiency across the whole polyhedron at once, including its edges and vertices, when a single flat face already carries the entire argument.
5. What makes you recognise it faster next time: any "balls or discs cover a region, bound a count by a size" problem is, structurally, a Hölder or power mean problem waiting for you to name the exponents.
6. The reusable decision rule: when a count, an area type bound, and a volume type bound must be reconciled in one inequality, restrict to the smallest sub-piece that carries the area bound, then apply Hölder with exponents chosen to reproduce the target inequality's own powers.

Mirror drills

1. n discs of radii $r_1, \ldots, r_n$ cover a fixed segment of length ℓ on a line, with $\sum r_i = L$. Find the correct exponent matching to bound n below in terms of ℓ and L. *Answer:* covering a length ℓ needs $\sum 2r_i \geq \ell$; only two power sums are in play here, degree 1 and degree 0, so plain Cauchy-Schwarz already suffices, and the three way Hölder step of B2 is unnecessary because no degree 3 quantity is present.
2. Balls of total volume V cover the entire surface of a cube of side s. Find an explicit constant c, in terms of s, such that the number of balls exceeds c/V^2. *Answer:* take F to be one face of area $A = s^2$, run the B2 argument verbatim, giving $n \geq 16s^6/(9\pi V^2)$, so any $c < 16s^6/(9\pi)$ works.
3. n balls of volumes summing to V cover a fixed disc of area A in a plane, with no ambient solid to restrict to: the disc is the whole covered object. Does the same Hölder argument bound n in terms of A and V, and how does it compare with B2's bound. *Answer:* yes, identically, since B2's argument never used any property of P beyond the existence of one flat face of area A; the bound $n \geq 16A^3/(9\pi V^2)$ is unchanged, confirming that B2's real content is a statement about discs and balls, with the polyhedron only ever supplying the face.

14.9 B3. Hall's theorem matches days to distinct winners

A round-robin tournament of $2n$ teams lasted for $2n-1$ days, as follows. On each day, every team played one game against another team, with one team winning and one team losing in each of the n games. Over the course of the tournament, each team played every other team exactly once. Can one necessarily choose one winning team from each day without choosing any team more than once?

Read it out loud

Peel off the tournament dressing and read the combinatorial skeleton: there are $2n-1$ days, each day equipped with a set of n winning teams out of $2n$ teams total, since each day's games partition all $2n$ teams into n winners and n losers. The question asks whether you can pick, for every day, one of that day's winners, such that the $2n-1$ picks are all different teams. That is a system of distinct representatives: one representative per day, drawn from that day's set of eligible teams, no repeats. Any time a question reads "choose one X from each Y, all distinct," the underlying object is a bipartite graph between the Y's and the X's, and the tool that decides whether a full matching exists is Hall's marriage theorem. That is the circuit that should fire the instant the phrase registers.

So the real work is not in Hall's theorem itself, which is a black box once invoked, but in verifying its hypothesis: for every collection of m days, the teams that won on at least one of those days must number at least m. That hypothesis has to be proved using the one piece of tournament structure not yet used, that each team plays every other team exactly once across the whole tournament, so no team ever faces the same opponent twice. The cue to write down inside thirty seconds: this is a distinct-representatives question, name the bipartite graph, days against teams, edges are wins, and go hunting for why Hall's condition holds, using the "exactly once" clause as the engine of that argument.

Approaches, ranked worst to best

1. **induct on n by removing one team and one day from the tournament.** *Cue:* a smaller tournament with $2n-2$ teams over $2n-3$ days looks like the same problem one size down. *Circuit:* structural induction on the tournament itself. Removing a team destroys the perfect matching structure of every day that team played in, since each day was a pairing of all $2n$ teams; rebuilding this into a genuine $(2n-2)$-team, $(2n-3)$-day round-robin tournament needs delicate schedule surgery, and even after that work an argument essentially equivalent to Hall's theorem is still needed for the base case or the reattachment step. Far more ink for no extra rigour.

2. **choose a winner for each day greedily, in some fixed order of the days.** *Cue:* the goal is a sequence of choices, so make them one at a time. *Circuit:* a greedy algorithm with no global certificate. Greedy can paint itself into a corner: having freely picked winners for earlier days, a later day's entire set of winners might already be exhausted, and there is no way to know in advance, without something equivalent to Hall's condition, that a different order or different earlier choices were not required. A greedy construction with no proof of correctness is not a solution.

3. **build the bipartite graph of days against teams, edges are wins, and verify Hall's condition via a pigeonhole argument on a hypothetical team that lost every game in a chosen set of days.** *Cue:* the goal is a system of distinct representatives, one per day, and Hall's theorem is the exact tool for that. *Circuit:* Hall's marriage theorem, with the "exactly once" clause supplying the pigeonhole step that proves its hypothesis. This is the ideal route: it isolates the one nontrivial fact to prove, Hall's condition, proves it in a few lines using a fact the problem

hands you for free, no repeated opponents, and then cites a named theorem for the rest, rather than reconstructing a matching argument from scratch.

The ideal solution

The answer is yes.

Build a bipartite graph H with one part D, the $2n-1$ days, and the other part T, the $2n$ teams, joining a day d to a team t by an edge exactly when t won its game on day d. Since every day has exactly n winners, every vertex of D has degree n in H. The problem asks whether H has a matching that saturates D, that is, a set of $2n-1$ edges, no two sharing an endpoint, one incident to each day. By Hall's marriage theorem, such a matching exists if and only if, for every subset $D' \subseteq D$, the neighbourhood $N(D') \subseteq T$, the teams that won on at least one day in D', satisfies $|N(D')| \geq |D'|$.

Lemma 14.9.1. *For every nonempty $D' \subseteq D$ with $|D'| = m$, $|N(D')| \geq m$.*

Proof. Suppose instead $|N(D')| \leq m-1$. Since every team plays exactly one game on every day, every team either wins or loses on each day in D'. Because $|N(D')| \leq m-1 < 2n$, some team L is not in $N(D')$, meaning L lost its game on every day in D'. Over the m days in D', L played m games, one per day, against m opponents; since each pair of teams meets exactly once over the whole tournament, these m opponents are pairwise distinct. On each such day, the opponent that beat L won that day's game, so each of these m distinct opponents belongs to $N(D')$, by way of the specific day it played, and beat, L. This exhibits m distinct teams inside $N(D')$, contradicting $|N(D')| \leq m-1$. □

Hall's condition holds for every subset of D, so H has a matching saturating D: a set of $2n-1$ winning teams, one chosen for each day, all distinct. Hence

Yes: such a choice can always be made.

Layer 1: WHAT

Block 1: translate "choose one winner per day, all distinct" into bipartite matching language, days against teams, edges are wins. Block 2: name Hall's marriage theorem as the tool that decides whether a matching saturating the days exists. Block 3: isolate the one thing left to prove, Hall's condition, as a standalone lemma about an arbitrary subset of days. Block 4: prove that lemma by contradiction, producing a team L that lost on every day in the subset. Block 5: use the "each pair meets exactly once" clause to show L's m opponents on those days are distinct, hence m distinct winners exist, closing the contradiction.

The nominal route tries to construct the choice directly, day by day, either by induction on the tournament's size or by a greedy pass through the days, and gets stuck justifying that the construction cannot fail partway through. The bifurcation point is recognising that "does a valid simultaneous choice exist" is a question about a bipartite graph having a certain matching, which converts an open ended constructive search into a single hypothesis to verify, Hall's condition, with a named theorem supplying the rest for free.

Layer 2: HOW

Block 1. Cue. The question asks to choose one item, a winning team, from each of several sets, each day's winners, with all choices distinct. **Why natural.** This is the definition of a system of distinct representatives, which is exactly a matching in the bipartite graph whose two parts are the collection of sets and the underlying elements, with an edge for membership. **Link to the next block.** Once the bipartite graph is named, "can one always choose" becomes "does this graph always have a matching saturating one side," a question with a standard named answer.

Block 2. **Cue.** A bipartite graph, one side fixed at $2n-1$ vertices, the days, and a matching saturating that side is wanted. **Why natural.** Hall's marriage theorem states exactly that such a matching exists if and only if every subset of that side has a neighbourhood at least as large as itself; it is the correct tool because the graph here has no further special structure, such as regularity on both sides, to exploit more directly. **Link to the next block.** Applying the theorem reduces the entire problem to a single combinatorial claim: verify Hall's condition for this specific graph.

Block 3. **Cue.** Hall's condition must hold for every subset of days, not just special ones. **Why natural.** Stating it as a lemma for an arbitrary subset D' of size m isolates exactly the fact that needs proof and separates it cleanly from the cited theorem, which is good proof hygiene: the lemma can then be checked entirely on its own terms. **Link to the next block.** Proving the lemma by contradiction is the natural move, since a claim that a count is "at least m" is often easiest to secure by assuming fewer than m and deriving an impossibility.

Block 4. **Cue.** Assuming fewer than m distinct teams won across the m days in D', while every team plays every day. **Why natural.** Since fewer than $m < 2n$ teams won on those days, some team is excluded from every winner set for those days, meaning it lost every one of those m games; such a team must exist because the total team count exceeds the winner count. **Link to the next block.** That team, L, played m games during D', one per day, all losses; the next question is what can be said about L's m opponents.

Block 5. **Cue.** Each pair of teams meets exactly once across the entire tournament. **Why natural.** This is the one clause of the problem not yet used, and it is precisely what forces L's m opponents, one per day in D', to be pairwise distinct, since a repeated opponent would mean two games between the same pair. **Link to the next block.** Each of those m distinct opponents won on the specific day it beat L, a day inside D', so all m belong to $N(D')$, contradicting the assumption that $|N(D')| < m$ and closing the lemma, and with it the whole proof once Hall's theorem is invoked.

Layer 3: WHY

The underlying truth is that "choose one representative per set, all distinct" is never solved by inspecting the sets one at a time; it is a global existence question about a bipartite graph, and Hall's theorem converts that global question into a local, checkable condition, namely that no small collection of sets can conspire to share too few elements between them. The discernment mechanism is linguistic: any phrase of the form "choose one from each, without repeats" should immediately trigger the bipartite graph translation, at which point the only remaining mathematical content is verifying Hall's condition, almost always by contradiction, almost always using whichever hypothesis of the problem has not yet been spent.

The second truth is about where the contradiction comes from. In this problem, the clause "each team played every other team exactly once" reads, on a first pass, like background scaffolding explaining the tournament's structure rather than an active ingredient. It only becomes active once you are inside a proof by contradiction and need a source of distinctness, here, the distinctness of a single team's m opponents across m days. Train yourself to notice, when a Hall's condition proof stalls, that the missing distinctness fact is usually sitting in an unused clause of the problem statement, phrased as an "exactly" or "at most once" condition, because that is precisely the shape of fact that rules out repetition.

Revolving concepts

Hall's marriage theorem. For a bipartite graph with parts X, Y, a matching saturating X exists if and only if $|N(S)| \geq |S|$ for every $S \subseteq X$, where $N(S)$ is the set of vertices in Y adjacent to some

vertex of S. This is the standard tool for any system of distinct representatives question. See Book 1, field: Discrete structures and graphs, topic: bipartite matching and Hall's theorem.

Systems of distinct representatives. Given sets $A_1, \ldots, A_k$, a system of distinct representatives is a choice of distinct elements $a_1 \in A_1, \ldots, a_k \in A_k$; such a choice exists exactly when Hall's condition holds for the bipartite graph of memberships. The days-versus-winners structure here is a direct instance, with A_d the set of teams that won on day d. See Book 1, field: Combinatorics, topic: systems of distinct representatives.

Proof by contradiction via an extremal or exceptional witness. To show a count is at least m, assume it is at most $m-1$ and produce an explicit object, here a team that lost on every day of a chosen subset, whose existence is guaranteed by simple counting, fewer winners than teams, and then derive m new distinct witnesses from that one object's behaviour. See Book 1, field: Combinatorics, topic: pigeonhole and counting contradictions.

Hall's marriage theorem is itself proved by an induction that is worth generalising once: induct on $|X|$; if every nonempty proper subset S of X has $|N(S)| > |S|$ strictly, match any single vertex of X to any neighbour and delete both, leaving Hall's condition intact on the smaller graph by induction; otherwise some proper nonempty S is tight, $|N(S)| = |S|$, so match S to $N(S)$ by induction on the smaller graph, then check Hall's condition survives on $(X \setminus S)$ against $(Y \setminus N(S))$, which holds because any shortfall there would combine with S to violate Hall's condition on the original graph. Beyond this problem, that same two case induction, either shrink by one matched pair or split at a tight set, is the template for proving any theorem about matchings by peeling off a manageable piece and recursing on what remains, and it is the same idea that underlies deficiency versions of Hall's theorem and König's theorem on bipartite graphs.

Circuit Bank

- **WHEN** a question asks to choose one item from each of several sets, or one representative per day, round or group, with no repeats allowed **DO** translate immediately into a bipartite graph, sets or days on one side, items or teams on the other, edges for membership or eligibility, and invoke Hall's marriage theorem **BECAUSE** because this phrasing is the definition of a system of distinct representatives, and Hall's theorem is the standard tool that decides existence without a constructive search

- **WHEN** Hall's theorem has been invoked and only its hypothesis remains to prove **DO** state the hypothesis as a standalone lemma for an arbitrary subset, then prove it by contradiction, assuming the neighbourhood is too small **BECAUSE** because isolating the lemma keeps the citation of the theorem clean and turns an open ended existence question into one bounded counting argument

- **WHEN** a proof by contradiction produces an exceptional witness object that must fail some property on every element of a chosen subset **DO** look for an unused clause of the problem statement to supply distinctness or another needed property among that witness's interactions with the subset **BECAUSE** because problem statements rarely include a clause that plays no role, and a stalled Hall's condition proof is almost always one unused hypothesis away from closing

- **WHEN** a combinatorial structure has an "exactly once" or "at most once" clause you have not yet used **DO** treat it as the source of a pigeonhole style distinctness argument the moment a proof needs m pairwise distinct objects **BECAUSE** because "exactly once" is precisely the hypothesis that rules out silent repetition, which is what a pigeonhole count needs to go through

Backtesting

1. First useful cue: "choose one from each day, all distinct" is the definition of a system of distinct representatives, so the graph to draw is days against teams.
2. Method the expression suggested: Hall's marriage theorem, applied to the bipartite graph of days and their winners, reduces the whole problem to one counting lemma.
3. Where you would hesitate, and the thought to reinforce: proving Hall's condition itself; the thought to reinforce is to assume the condition fails, extract a specific team that lost throughout the chosen days, and look for what forces that team's opponents to be distinct.
4. The trap the question sets: trying to build the required choice directly, by induction on the tournament size or by a greedy pass over the days, without a certificate that the construction cannot fail.
5. What makes you recognise it faster next time: the phrase "choose one X from each Y without repeats" should fire the Hall's theorem circuit before any attempt at direct construction.
6. The reusable decision rule: whenever a distinct-choice question appears, build the bipartite membership graph, cite Hall's theorem, then spend all remaining effort proving Hall's condition by contradiction using whichever hypothesis has not yet been used.

Mirror drills

1. $2n$ students each need a distinct exam slot, and every group of k students is jointly acceptable to at least k of the $2n-1$ available slots. Show every student can be assigned a slot, all different. *Answer:* this is Hall's theorem verbatim, students and slots as the two parts of the bipartite graph; no further argument beyond checking the hypothesis is needed.
2. In a different event with $2n-1$ rounds among $2n$ competitors, every pair meets at most once, but each round crowns a single overall winner rather than n separate winners. Explain why "choose a different winner per round" becomes a much weaker question here. *Answer:* with one winner per round, the rounds-to-winners graph has every round of degree 1, so Hall's condition reduces to asking whether any two rounds ever share their unique winner, a condition about repeat champions rather than a genuine matching argument.
3. A bipartite graph has parts X, Y with $|X| = |Y| = k$ and every vertex of X of degree at least $k/2$. Give a short argument, in the spirit of B3's lemma, for why any $S \subseteq X$ with $|S| > k/2$ has $N(S) = Y$. *Answer:* if some $y \in Y$ were missed by $N(S)$, every vertex of S would avoid y, so y's degree would be at most $k - |S| < k/2$, contradicting the degree hypothesis; the same "assume a witness is missed, count its forced exclusions" pattern as B3's lemma closes the argument.

14.10 B4. A recursion that shadows the natural logarithm

Suppose that $a_0 = 1$ and that $a_{n+1} = a_n + e^{-a_n}$ for $n = 0, 1, 2, \ldots$. Does $a_n - \log n$ have a finite limit as $n \to \infty$? (Here $\log n = \log_e n = \ln n$.)

Read it out loud

Say it in words. This is a first order autonomous recursion, meaning a_{n+1} depends on a_n alone through a fixed function, and the increment e^{-a_n} is positive and decreasing in a_n. So the sequence increases, and the faster it climbs the slower it climbs next. You are then handed a specific growth scale, $\log n$, and asked whether the error $a_n - \log n$ settles down. Notice the shape of the command term: not *find* the limit, but *does it have a finite limit*. That word finite, paired with an explicit comparison function, is the guiding light of the whole problem. You are being asked for an existence statement, and the cheapest existence theorem in analysis is monotone convergence: a decreasing sequence bounded below converges. No candidate value required.

Where did $\log n$ come from, and why should you trust it? Run the shadow ODE. Replace the discrete increment $a_{n+1} - a_n$ by a derivative $y'(x)$ and the value a_n by $y(x)$. The recursion becomes $y' = e^{-y}$, which separates: $e^y y' = 1$, so $e^y = x + C$ and $y = \log(x + C)$. That is a scouting report, not a proof, but it tells you two things at once. First, $\log n$ is the right scale. Second, the free constant C sits *inside* the logarithm, and $\log(n + C) - \log n = \log(1 + C/n) \to 0$, so the shadow model predicts not merely that the error converges but that it converges to 0. Keep that prediction; you will collect it at the end as the plus alpha.

The fields wired into this question are real analysis (limits of monotone sequences), sequences and recurrences (comparison of a recursion against a closed form), and the elementary inequality theory of the exponential. There is no clever machinery anywhere in this problem. There is one inequality, $e^x \geq 1 + x$, used twice, in opposite directions, and one induction.

Within thirty seconds of reading, you should have written three lines on the page: $d_n = a_n - \log n$; the difference $d_{n+1} - d_n = e^{-a_n} - \log\left(1 + \frac{1}{n}\right)$; and the note *need a lower bound on a_n, of logarithmic type, to control e^{-a_n}*. Everything after that is bookkeeping. The reason those three lines are forced is that you have decided to prove monotonicity of d_n, monotonicity is a statement about $d_{n+1} - d_n$, and the recursion computes that difference for you exactly.

Approaches, ranked worst to best

1. **Numerical scouting.** *Cue:* the phrase "does it have a finite limit" invites a quick experiment. *Circuit:* compute terms and look at them. Compute a thousand terms, watch $a_n - \log n$ shrink, and assert convergence. This scores zero. Worse, it is actively misleading: the sequence $\log\log n$ also looks like it is settling down over any finite range you can compute, and it diverges. Use numerics to check a claim you have proved, never to replace the proof.

2. **Closed form hunt.** *Cue:* a clean elementary increment suggests the recursion might be solvable. *Circuit:* invert the recursion into an explicit formula. The substitution $b_n = e^{a_n}$ turns the recursion into $b_{n+1} = b_n e^{1/b_n}$, which is no more solvable than what you started with. There is no closed form, and hunting for one costs the whole exam session. The tell is that e^{-a_n} mixes the sequence with a transcendental function of itself; only linear or fractional linear recursions surrender closed forms for free.

3. **Two sided sandwich.** *Cue:* $\log n$ is given, so bound a_n on both sides. *Circuit:* squeeze. Prove $\log(n + 1) < a_n < \log n + C$ and conclude that d_n is bounded. This is more work than the winning route and it does not finish: boundedness is not convergence. You would still have to prove monotonicity or a Cauchy estimate on top. More ink for a weaker conclusion.

4. **Exponential substitution and increment averaging.** *Cue:* the increment is e^{-a_n}, so the increment of e^{a_n} is nearly constant. *Circuit:* change variable so the increment linearises. Put $b_n = e^{a_n}$. Then $b_{n+1} = b_n e^{1/b_n} = b_n + 1 + R_n$ with $0 \le R_n \le 1/b_n$, so b_n is n plus a controlled error, and $b_n/n \to 1$. This is a genuine route and it proves strictly more than the question asks, namely that the limit is 0. It costs a Taylor remainder bound and a harmonic sum estimate. Rank it second and keep it in your pocket for the plus alpha.
5. **Monotone convergence on the error term.** *Cue:* "finite limit" plus an explicit comparison function. *Circuit:* decreasing and bounded below implies convergent. Show $a_n > \log(n+1)$ by induction, deduce that $d_n = a_n - \log n$ is strictly decreasing and strictly positive, and quote monotone convergence. One induction, one inequality used twice. This is the shortest complete proof available and it is what a full ten points looks like.

The ideal solution

The answer is $\boxed{\text{Yes}}$: the limit exists.

Lemma 14.10.1 (tangent line)**.** *For every real $x > -1$ we have $\log(1+x) \le x$, with equality only at $x = 0$.*

Proof. The function $\varphi(x) = x - \log(1+x)$ satisfies $\varphi'(x) = 1 - \frac{1}{1+x} = \frac{x}{1+x}$ on $(-1, \infty)$, which is negative for $x \in (-1, 0)$ and positive for $x > 0$. So φ attains its minimum at $x = 0$, where $\varphi(0) = 0$, and is strictly positive elsewhere. □

Lemma 14.10.2 (lower bound)**.** *For every integer $n \ge 0$ we have $a_n > \log(n+1)$.*

Proof. Induction on n. The increments e^{-a_n} are positive, so (a_n) is increasing and $a_n \ge a_0 = 1 > 0$ for all n. For $n = 0$, $a_0 = 1 > 0 = \log 1$.

Assume $a_n > \log(n+1) \ge 0$. Let $f(x) = x + e^{-x}$, so that $a_{n+1} = f(a_n)$. Since $f'(x) = 1 - e^{-x} > 0$ for $x > 0$, f is strictly increasing on $[0, \infty)$, and both a_n and $\log(n+1)$ lie there. Hence

$$a_{n+1} = f(a_n) > f\big(\log(n+1)\big) = \log(n+1) + \frac{1}{n+1}.$$

By the tangent line lemma with $x = \frac{1}{n+1}$ we have $\frac{1}{n+1} \ge \log\left(1 + \frac{1}{n+1}\right)$, so the right hand side is at least $\log(n+1) + \log\frac{n+2}{n+1} = \log(n+2)$. Therefore $a_{n+1} > \log(n+2)$, closing the induction. □

Now set $d_n = a_n - \log n$ for $n \ge 1$.

Step 1: $(d_n)_{n \ge 1}$ is strictly decreasing. Using the recursion,

$$d_{n+1} - d_n = e^{-a_n} + \log\frac{n}{n+1}.$$

By the lower bound lemma, $a_n > \log(n+1)$, so $e^{-a_n} < e^{-\log(n+1)} = \frac{1}{n+1}$. By the tangent line lemma with $x = -\frac{1}{n+1} \in (-1, 0)$,

$$\log\frac{n}{n+1} = \log\left(1 - \frac{1}{n+1}\right) \le -\frac{1}{n+1}.$$

Adding the two estimates, $d_{n+1} - d_n < \frac{1}{n+1} - \frac{1}{n+1} = 0$.

Step 2: $(d_n)_{n \ge 1}$ is bounded below. By the lower bound lemma again, $d_n = a_n - \log n > \log(n+1) - \log n > 0$.

A strictly decreasing sequence of real numbers bounded below converges, by the monotone convergence theorem. Hence $a_n - \log n$ has a finite limit, and that limit is at least 0. This answers the question.

Remark 14.10.3. The limit is in fact exactly 0, which the question does not ask for but which you should be able to produce on demand. Put $b_n = e^{a_n}$, so $b_0 = e$ and $b_{n+1} = b_n e^{1/b_n}$. Since (a_n) increases from 1, we have $b_n \geq e$ and $u_n = 1/b_n \in (0, 1/e]$. For $0 < u \leq 1$,

$$1 + u \leq e^u = 1 + u + \sum_{k\geq 2} \frac{u^k}{k!} \leq 1 + u + u^2 \sum_{k\geq 2} \frac{1}{k!} \leq 1 + u + u^2.$$

Multiplying through by b_n with $u = 1/b_n$ gives $b_n + 1 \leq b_{n+1} \leq b_n + 1 + \frac{1}{b_n}$. The left inequality and induction give $b_n \geq n + e$. Feeding that back into the right inequality and summing,

$$b_n \leq e + n + \sum_{i=0}^{n-1} \frac{1}{b_i} \leq e + n + \sum_{i=0}^{n-1} \frac{1}{i+2} = e + n + H_{n+1} - 1 \leq n + e + \log(n+1),$$

where $H_m \leq 1 + \log m$ by comparing $\sum_{k=2}^{m} \frac{1}{k}$ with $\int_1^m \frac{dt}{t}$. Hence for $n \geq 1$,

$$1 < \frac{b_n}{n} \leq 1 + \frac{e + \log(n+1)}{n} \longrightarrow 1,$$

and therefore $a_n - \log n = \log \frac{b_n}{n} \to 0$. A finer expansion of the same estimates gives $a_n = \log n + \frac{\log n}{2n} + O\left(\frac{1}{n}\right)$, which is why the numerical convergence is so slow: at $n = 10^6$ the error is still about 9×10^{-6}.

Layer 1: WHAT

Block 1. Read the command term "finite limit" as an existence request and commit to monotone convergence as the mechanism.

Block 2. Recover the scale $\log n$ from the shadow ODE $y' = e^{-y}$, so that you know the object of study is the error $d_n = a_n - \log n$ and not a_n itself.

Block 3. Difference the error: $d_{n+1} - d_n = e^{-a_n} + \log \frac{n}{n+1}$. This exposes that everything hinges on comparing e^{-a_n} with $\frac{1}{n+1}$.

Block 4. Therefore you need a lower bound $a_n > \log(n+1)$. Prove it by induction, using that $x \mapsto x + e^{-x}$ is increasing.

Block 5. Feed the bound into the difference; the two halves cancel exactly, courtesy of $\log(1 + x) \leq x$ applied at $x = \pm\frac{1}{n+1}$. Monotonicity done.

Block 6. The same bound gives positivity. Monotone convergence closes the problem.

The nominal route, and the juxtaposition. The zombie writes $a_n \approx \log n + c + o(1)$, starts expanding $e^{-a_n} \approx \frac{1}{n}\left(1 - \frac{c}{1}\right)$ and similar unjustified asymptotics, and produces three quarters of a page of O symbols with no theorem underneath any of them. Alternatively the zombie tries to solve the recursion. The bifurcation point is precise: it is the line where you either write $a_n \approx \log n + c$ (an asymptotic ansatz you cannot yet justify) or write $d_{n+1} - d_n$ (an exact identity you can). Exact identity beats ansatz. The difference costs one line to write and it converts an asymptotics problem into an inequality problem, and inequalities are decidable at A-level with one derivative.

Layer 2: HOW

Block 1. Cue. The question asks whether a limit is finite, not what it is, and it supplies the comparison function itself. **Why natural.** When you are not asked for a value, do not chase one. The theorems that give existence without a value are monotone convergence, Cauchy, and compactness. Of these, monotone convergence needs the least: one inequality about consecutive terms and one crude bound. So the plan is fixed before any calculation happens. **Link to the next block.** A monotone convergence plan needs a specific sequence to be monotone. Which sequence? Not a_n, which diverges. The question names the candidate: $a_n - \log n$.

Block 2. **Cue.** The comparison function $\log n$ appears from nowhere in the problem statement, so it must be forced by the recursion. **Why natural.** Increments become derivatives when the steps are small, and $a_{n+1} - a_n = e^{-a_n}$ is a difference equation whose continuous shadow $y' = e^{-y}$ separates instantly. The solution $y = \log(x + C)$ explains the $\log n$ and predicts the limit is 0. Doing this in the arena costs fifteen seconds and buys certainty about what you are aiming for. **Link to the next block.** Knowing that the true behaviour is $\log(n + C)$ tells you the error is $\log(1 + C/n)$, of size $1/n$. So all the estimates you will need are of size $1/n$. That calibrates every inequality that follows.

Block 3. **Cue.** Monotonicity of d_n is a statement about $d_{n+1} - d_n$, and the recursion gives $a_{n+1} - a_n$ exactly. **Why natural.** Never estimate a sequence when you can compute its difference exactly. Here $d_{n+1} - d_n = e^{-a_n} - \log \frac{n+1}{n}$, in which the first term is unknown and the second is explicit. That single line converts the whole problem into: is e^{-a_n} smaller than $\log \frac{n+1}{n}$? **Link to the next block.** $\log \frac{n+1}{n}$ is slightly larger than $\frac{1}{n+1}$ (the tangent line lemma again), so it suffices to prove $e^{-a_n} \le \frac{1}{n+1}$, that is, $a_n \ge \log(n + 1)$. The required lemma has now been generated by the calculation rather than guessed.

Block 4. **Cue.** You need $a_n > \log(n + 1)$, a statement about all n with a recursive definition available: induction is forced. **Why natural.** The step needs $a_{n+1} = f(a_n)$ with $f(x) = x + e^{-x}$ increasing on $[0, \infty)$, which is one derivative. Then $a_n > \log(n + 1)$ gives $a_{n+1} > f(\log(n + 1)) = \log(n + 1) + \frac{1}{n+1}$, and you need that to exceed $\log(n + 2)$, which is exactly $\log(1 + x) \le x$ at $x = \frac{1}{n+1}$. Note the strengthening: had you tried to prove the weaker $a_n > \log n$, the step would demand $\log n + \frac{1}{n} \ge \log(n + 1)$, which is true, so that also closes; but the shifted version is what you need in Block 5, so prove the shifted version once and use it twice. **Link to the next block.** The same lemma now feeds both remaining tasks: it bounds e^{-a_n} from above (monotonicity) and d_n from below (boundedness).

Block 5. **Cue.** Two estimates of the same size $\frac{1}{n+1}$, pointing in opposite directions. **Why natural.** This is the signature of a problem that has been designed to close exactly. The inequality $\log(1 + x) \le x$ is applied at $x = +\frac{1}{n+1}$ inside the induction and at $x = -\frac{1}{n+1}$ inside the monotonicity step; the two applications are the same fact seen from either side of 0. When a problem uses one inequality twice with opposite signs, you are on the intended route. **Link to the next block.** Strict decrease is established. Only a lower bound remains.

Block 6. **Cue.** A decreasing sequence needs any lower bound whatsoever, not a sharp one. **Why natural.** $d_n > \log(n + 1) - \log n > 0$ is free from the lemma you already have. Be lazy: do not compute a sharp bound when 0 works. **Link to the next block.** Monotone convergence applies and the answer is yes.

Layer 3: WHY

The mathematical truth beneath this route is that convergence of a real sequence is a statement about the completeness of $\mathbb{R}$, and completeness is delivered in three standard packages: monotone convergence, Cauchy, and Bolzano and Weierstrass. Monotone convergence is the cheapest because its hypotheses are checkable term by term. Any Putnam question phrased as "does the limit exist" or "show that the limit exists" is a request to select one of those three packages, and your first move should always be to test whether the quantity in question is eventually monotone. Here it is, and the check costs two lines.

The second truth is the discrete to continuous correspondence. A recursion $a_{n+1} = a_n + g(a_n)$ with $g > 0$ small is a Riemann sum in disguise. Set G with $G' = 1/g$; then $G(a_{n+1}) - G(a_n) = \int_{a_n}^{a_{n+1}} \frac{dt}{g(t)} \approx \frac{a_{n+1} - a_n}{g(a_n)} = 1$, so $G(a_n) \approx n$ and $a_n \approx G^{-1}(n)$. That single formula predicts the answer

to an entire family of Putnam problems. With $g(y) = e^{-y}$ you get $G(y) = e^{y}$ and $a_n \approx \log n$. With $g(y) = 1/y$ you get $G(y) = y^2/2$ and $a_n \approx \sqrt{2n}$. With $g(y) = -y^3/6$, which is what sin does near 0, you get $a_n \approx \sqrt{3/n}$. Install the formula, not the individual answers.

The third truth is about which inequality to reach for. Every comparison between a discrete increment of size $\frac{1}{n}$ and a logarithmic difference is mediated by the tangent line inequality $e^x \geq 1 + x$, equivalently $\log(1 + x) \leq x$, equivalently $\frac{x}{1+x} \leq \log(1 + x)$ after substitution. There is no second inequality in this circle of problems. If you find yourself wanting a Taylor expansion of $\log(1 + 1/n)$ to two terms with an unjustified remainder, stop: the exact one sided inequality is available, is sharper than you need, and is provable in one line.

How to train the discernment. Take any recursion problem and ask three questions in order. What does the shadow ODE say the growth is? What quantity does the question compare it to? Is that difference monotone? If the answers are available, you have a proof; if the difference is not monotone, only then reach for Cauchy or for a summable bound on $|d_{n+1} - d_n|$. The examiners nudged you toward exactly this by putting $\log n$ in the question rather than asking for the asymptotics of a_n: they handed you the comparison function so that the monotone route was available without any guessing at all.

Revolving concepts

Monotone convergence theorem. A nonincreasing sequence bounded below converges to its infimum; a nondecreasing sequence bounded above converges to its supremum. It is equivalent to the least upper bound axiom for $\mathbb{R}$, so it is not a technique but the definition of the real numbers in usable form. See Book 1, field: Real analysis, topic: completeness and monotone sequences.

The tangent line inequality for the exponential. $e^x \geq 1 + x$ for all real x, with equality only at $x = 0$, since e^x is convex and $1 + x$ is its tangent at 0. Equivalent forms: $\log(1 + x) \leq x$ for $x > -1$; $\frac{x}{1+x} \leq \log(1 + x)$ for $x > -1$ (substitute $x \mapsto \frac{x}{1+x}$ in the first form and rearrange); $\left(1 + \frac{1}{n}\right)^n < e < \left(1 + \frac{1}{n}\right)^{n+1}$. See Book 1, field: Inequalities, topic: convexity and tangent lines.

First order autonomous recursions, generalised. Let $a_{n+1} = a_n + g(a_n)$ with $g > 0$ and g decreasing, and let G be any antiderivative of $1/g$. Then $G(a_n) = n + O(\text{error})$ and $a_n = G^{-1}(n)(1 + o(1))$. The rigorous version of the heuristic is the pair of bounds obtained by substituting the exact recursion into G and using monotonicity of $1/g$ on the interval $[a_n, a_{n+1}]$. The special cases worth memorising: $g(y) = y^{-\alpha}$ gives $a_n \sim ((\alpha + 1)n)^{1/(\alpha+1)}$; $g(y) = e^{-y}$ gives $a_n = \log n + o(1)$; a decreasing sequence with $a_{n+1} = a_n - ca_n^k$ gives $a_n \sim (c(k - 1)n)^{-1/(k-1)}$. See Book 1, field: Sequences and recurrences, topic: asymptotics of first order recursions.

Harmonic sums. $H_n = \sum_{k=1}^{n} \frac{1}{k}$ satisfies $\log n < H_n \leq 1 + \log n$ by comparison with $\int \frac{dt}{t}$, and $H_n = \log n + \gamma + \frac{1}{2n} + O(n^{-2})$ where γ is the Euler and Mascheroni constant. Any time a recursion accumulates increments of size c/i, the total is $c \log n$ plus a constant, which is why $b_n = n + \frac{1}{2}\log n + O(1)$ in the remark. See Book 1, field: Real analysis, topic: harmonic sums and Euler and Mascheroni constant.

Shadow ODE heuristic. Replacing $a_{n+1} - a_n$ by y' converts a difference equation into a differential equation whose solution predicts the leading behaviour. It is never a proof, because it silently assumes the increments are small and smooth, but it is the correct scouting tool and it is free. See Book 1, field: Sequences and recurrences, topic: difference equations and their continuous shadows.

Circuit Bank

- **WHEN** a question asks whether a limit exists, and does not ask for its value **DO** test whether the quantity is eventually monotone and bounded, and quote monotone convergence **BECAUSE** existence in $\mathbb{R}$ is completeness, and the monotone package has the cheapest hypotheses

- **WHEN** a recursion of the form $a_{n+1} = a_n + g(a_n)$ appears **DO** write the shadow ODE $y' = g(y)$ and solve it by separation before anything else **BECAUSE** it costs fifteen seconds and it tells you the growth scale, hence what to compare against
- **WHEN** you must decide whether a sequence d_n is monotone **DO** compute $d_{n+1} - d_n$ exactly from the recursion instead of estimating d_n **BECAUSE** an exact identity converts an asymptotics problem into an inequality problem, and inequalities are decidable
- **WHEN** a comparison between an increment of size $1/n$ and a difference of logarithms is needed **DO** apply $\log(1+x) \le x$ at $x = \pm\frac{1}{n+1}$ **BECAUSE** that one tangent line inequality mediates every such comparison, so no Taylor remainder is ever required
- **WHEN** an induction step almost closes but leaves a gap of one index **DO** strengthen the statement by shifting the index inside the comparison function, for example prove $a_n > \log(n+1)$ rather than $a_n > \log n$ **BECAUSE** a stronger hypothesis gives a stronger step, and the shifted form is usually the one downstream needs

Backtesting

1. First useful cue: the words "finite limit" next to an explicitly supplied comparison function $\log n$. That combination means existence by monotone convergence, applied to the difference.
2. Method the expression suggested: the increment e^{-a_n} separates in the shadow ODE $y' = e^{-y}$, giving $y = \log(x + C)$ and hence the target $d_n = a_n - \log n$.
3. Where you would hesitate, and the thought to reinforce: at the induction, wondering whether to prove $a_n > \log n$ or $a_n > \log(n+1)$. Reinforce this: prove the version that the later step needs, and if unsure, prove the stronger one, because the induction step gets a stronger hypothesis too and usually closes more easily.
4. The trap the question sets: the answer is 0 and the convergence is extremely slow (about $\frac{\log n}{2n}$), so numerical work is inconclusive and any attempt to guess the limit first and prove convergence second wastes the session. The question is deliberately phrased so that you never need the value.
5. What makes you recognise it faster next time: any recursion whose increment is a positive decreasing function of the current term is a Riemann sum for $\int \frac{dt}{g(t)}$, and any question comparing such a recursion to a named function is asking about the error, not the growth.
6. The reusable decision rule: existence question plus explicit comparison function equals monotone convergence on the difference, with the lower bound supplied by the same induction that controls the increment.

Mirror drills

1. Let $a_0 = 1$ and $a_{n+1} = a_n + \frac{1}{a_n}$. Show that $a_n/\sqrt{2n} \to 1$. *Answer:* square the recursion, $a_{n+1}^2 = a_n^2 + 2 + a_n^{-2}$, so $a_n^2 \ge 2n + 1$ and $a_n^2 \le 1 + 2n + \sum_{i<n} \frac{1}{2i+1} = 2n + O(\log n)$; divide by $2n$.
2. Let $a_0 \in (0, \pi)$ and $a_{n+1} = \sin a_n$. Show that $\sqrt{n}\, a_n \to \sqrt{3}$. *Answer:* $a_n \downarrow 0$, and $a_{n+1}^{-2} - a_n^{-2} \to \frac{1}{3}$ from $\sin x = x - \frac{x^3}{6} + O(x^5)$; by the Cesaro mean, $n a_n^2 \to 3$.
3. Let $a_0 = 1$ and $a_{n+1} = a_n + e^{-2a_n}$. Does $a_n - \frac{1}{2}\log n$ converge? *Answer:* yes; the shadow ODE gives $e^{2y} = 2x + C$, so run the identical proof with the induction hypothesis $a_n > \frac{1}{2}\log(2n+1)$ and the same tangent line inequality.

14.11 B5. Convex envelopes turn a supremum into lines

Prove that, for any two bounded functions $g_1, g_2 : \mathbb{R} \to [1, \infty)$, there exist functions $h_1, h_2 : \mathbb{R} \to \mathbb{R}$ such that, for every $x \in \mathbb{R}$,

$$\sup_{s \in \mathbb{R}} \left(g_1(s)^x g_2(s) \right) = \max_{t \in \mathbb{R}} \left(x h_1(t) + h_2(t) \right).$$

Read it out loud

Read the two sides separately and the problem dissolves. On the left, x is the variable and s is the index; $g_1(s)^x g_2(s)$ is an exponential in x, because $g_1(s)^x = e^{x \log g_1(s)}$. So the left hand side is the pointwise supremum of a family of functions $x \mapsto e^{\alpha x + \beta}$, indexed by s, where $\alpha = \log g_1(s)$ and $\beta = \log g_2(s)$. On the right, $x h_1(t) + h_2(t)$ is an *affine function of* x for each fixed t, so the right hand side is the upper envelope of a family of straight lines, indexed by t. The problem therefore says: this particular function of x is the upper envelope of a family of lines, and moreover the envelope is attained at every x.

Which functions are upper envelopes of families of lines? Exactly the convex ones. That is not a piece of folklore, it is the content of the supporting line theorem: a finite convex function on $\mathbb{R}$ lies above a line through each point of its graph, so it is the supremum of its affine minorants, and the supremum is attained at each point by the supporting line at that very point. So the entire problem reduces to one word: prove the left hand side is a convex function of x. Everything else is bookkeeping about which line to name $h_1(t), h_2(t)$.

Now read the hypotheses as a checklist rather than as decoration. The codomain $[1, \infty)$ says $\log g_i(s) \geq 0$. Bounded says $\log g_i(s) \leq \log M$ for some $M \geq 1$. Together: the pair (α, β) ranges over a bounded box $[0, \log M]^2$. That is exactly what is needed to make the supremum on the left finite for every real x, negative x included. If g_1 were allowed to take values in $(0, 1)$ and to approach 0, the supremum would be $+\infty$ for $x < 0$ and there would be nothing to prove and no theorem to prove it with. The hypotheses are not scenery; they are the finiteness of F. Note also what is *not* assumed: no continuity, no measurability, nothing at all about g_1, g_2 as functions of s. The index set may as well be an arbitrary nonempty set. That tells you no analysis of s can be involved.

The fields are convexity and real analysis, with a Legendre transform lurking underneath. What you should have on the page inside thirty seconds: $g_1(s)^x g_2(s) = e^{\alpha_s x + \beta_s}$; $F(x) := \sup_s$ of that; the observation that each member of the family is convex in x; and the target sentence, "a finite convex function on $\mathbb{R}$ has a supporting line at every point".

Approaches, ranked worst to best

1. **Compute the supremum explicitly.** *Cue:* the left side is a supremum you would like to evaluate. *Circuit:* find a formula, then differentiate it. You cannot. The functions g_1, g_2 are arbitrary, the supremum need not be attained at any s, and F need not be differentiable. Any route that begins "let $s(x)$ be the maximising s" is dead on arrival, since no such s need exist.

2. **Define h_1, h_2 by differentiating F.** *Cue:* the supporting line at t is the tangent when F is smooth. *Circuit:* tangent line construction. This is the right picture but the wrong hypothesis. F can have corners, for instance when the family is two lines crossing. You would have to prove one sided differentiability first, which is the same work as the supporting line lemma but presented as if it were about calculus rather than about convexity.

3. **Fenchel biconjugation.** *Cue:* the shape $\max_t (x h_1(t) + h_2(t))$ is a Legendre transform. *Circuit:* quote $F^{**} = F$ for closed proper convex F. Set $h_1(t) = t$ and $h_2(t) = -F^*(t)$ where $F^*(y) = \sup_x (xy - F(x))$. This is correct and it is the conceptual home of the problem, but as a written solution it imports a heavier theorem than you need, and it still owes the grader an argument

that the supremum over t is attained. Attainment is precisely the supporting line statement again, so nothing was saved.

4. **Take logarithms first, then exponentiate.** *Cue:* $g_1(s)^x$ begs for logarithms. *Circuit:* sup of affine functions is convex, then compose with exp. Define $f(x) = \sup_s(x \log g_1(s) + \log g_2(s))$, note f is convex as a supremum of affine functions, then argue e^f is convex because exp is convex and increasing. Entirely correct, and it is the official route. It has one redundant stage: you build a convex function f only to compose it back. Ink is time.

5. **Sup of convex is convex, directly.** *Cue:* each member of the family is $e^{\alpha x+\beta}$, already convex in x. *Circuit:* pointwise supremum of convex functions is convex, then supporting lines. Skip f altogether. Each $x \mapsto g_1(s)^x g_2(s)$ has second derivative $\alpha_s^2 e^{\alpha_s x+\beta_s} \geq 0$, so it is convex; a pointwise supremum of convex functions is convex by a two line check; finiteness comes from the bounded box; supporting lines finish. This is the shortest complete proof and it is the one to write.

The ideal solution

Since g_1, g_2 are bounded and take values in $[1, \infty)$, fix $M \geq 1$ with $1 \leq g_i(s) \leq M$ for all $s \in \mathbb{R}$ and $i = 1, 2$. Put

$$\alpha_s = \log g_1(s) \in [0, \log M], \qquad \beta_s = \log g_2(s) \in [0, \log M],$$

so that $g_1(s)^x g_2(s) = e^{\alpha_s x+\beta_s} =: \Phi_s(x)$, and define

$$F(x) = \sup_{s\in\mathbb{R}} \Phi_s(x).$$

Step 1: F is real valued. For every s and every x we have $\alpha_s x + \beta_s \leq |x| \log M + \log M$, so $0 < \Phi_s(x) \leq M^{|x|+1}$. Taking the supremum over the nonempty index set, $0 < F(x) \leq M^{|x|+1} < \infty$ for every $x \in \mathbb{R}$. This is the only place the two hypotheses are used, and they are exactly enough.

Step 2: F is convex. Each Φ_s is convex on $\mathbb{R}$, since $\Phi_s''(x) = \alpha_s^2 \Phi_s(x) \geq 0$. Let $x_1, x_2 \in \mathbb{R}$ and $\lambda \in [0, 1]$. For every s,

$$\Phi_s\big(\lambda x_1 + (1-\lambda)x_2\big) \leq \lambda \Phi_s(x_1) + (1-\lambda)\Phi_s(x_2) \leq \lambda F(x_1) + (1-\lambda)F(x_2).$$

The right hand side is independent of s, so taking the supremum over s on the left gives $F(\lambda x_1 + (1-\lambda)x_2) \leq \lambda F(x_1) + (1-\lambda)F(x_2)$. Thus F is a finite convex function on $\mathbb{R}$.

Lemma 14.11.1 (chords)**.** *Let $F : \mathbb{R} \to \mathbb{R}$ be convex and let $u < v < w$. Then*

$$\frac{F(v) - F(u)}{v - u} \leq \frac{F(w) - F(u)}{w - u} \leq \frac{F(w) - F(v)}{w - v}.$$

Proof. Write $v = \lambda u + (1-\lambda)w$ with $\lambda = \frac{w-v}{w-u} \in (0, 1)$, so that $v - u = (1-\lambda)(w-u)$ and $w - v = \lambda(w-u)$. Convexity gives $F(v) \leq \lambda F(u) + (1-\lambda)F(w)$. Subtracting $F(u)$ from both sides gives $F(v) - F(u) \leq (1-\lambda)\big(F(w) - F(u)\big)$; dividing by $v - u = (1-\lambda)(w-u) > 0$ gives the first inequality. Subtracting both sides from $F(w)$ gives $F(w) - F(v) \geq \lambda\big(F(w) - F(u)\big)$; dividing by $w - v = \lambda(w-u) > 0$ gives the second. □

Lemma 14.11.2 (supporting line)**.** *Let $F : \mathbb{R} \to \mathbb{R}$ be convex and let $t \in \mathbb{R}$. Then there exists $m \in \mathbb{R}$ such that $F(x) \geq F(t) + m(x - t)$ for all $x \in \mathbb{R}$, with equality at $x = t$.*

Proof. Put $m = \inf_{x>t} \frac{F(x)-F(t)}{x-t}$. By the chord lemma applied to $u < t < w$, every difference quotient to the left of t is at most every difference quotient to the right of t; in particular the set over which the infimum is taken is bounded below by $F(t) - F(t-1)$, so m is a real number. For $x > t$, the definition of the infimum gives $\frac{F(x)-F(t)}{x-t} \geq m$, that is $F(x) \geq F(t) + m(x-t)$. For $x < t$, the same chord comparison gives $\frac{F(t)-F(x)}{t-x} \leq m$, that is $F(t) - F(x) \leq m(t-x)$, which rearranges to $F(x) \geq F(t) + m(x-t)$. At $x = t$ both sides agree. □

Step 3: construct h_1 and h_2. For each $t \in \mathbb{R}$ let m_t be a real number supplied by the supporting line lemma for F at t, and define

$$h_1(t) = m_t, \qquad h_2(t) = F(t) - t\,m_t.$$

These are functions $\mathbb{R} \to \mathbb{R}$, as required. For all $x, t \in \mathbb{R}$,

$$xh_1(t) + h_2(t) = F(t) + m_t(x - t) \le F(x),$$

by the supporting line inequality, and taking $t = x$ turns this into an equality. Hence the set $\{xh_1(t) + h_2(t) : t \in \mathbb{R}\}$ has largest element $F(x)$, attained at $t = x$: the maximum on the right hand side exists and

$$\max_{t\in\mathbb{R}} \big(xh_1(t) + h_2(t)\big) = F(x) = \sup_{s\in\mathbb{R}} \big(g_1(s)^x g_2(s)\big)$$

for every $x \in \mathbb{R}$. This is the required identity.

Remark 14.11.3. If you would rather not quote or prove any convexity theory at all, the bounded box gives a second, fully explicit construction. Fix t and choose s_k with $\Phi_{s_k}(t) \to F(t)$. The pairs $(\alpha_{s_k}, \beta_{s_k})$ lie in the compact box $[0, \log M]^2$, so by Bolzano and Weierstrass a subsequence converges to some (α, β) in that box. Writing $E(x) = e^{\alpha x+\beta}$, continuity gives $E(t) = F(t)$, while $E(x) = \lim_k \Phi_{s_k}(x) \le F(x)$ for every x. The tangent to E at t, namely $\ell(x) = E(t)\big(1 + \alpha(x - t)\big)$, satisfies $\ell \le E \le F$ because $e^u \ge 1 + u$, and $\ell(t) = F(t)$. So one may take

$$h_1(t) = \alpha\,F(t), \qquad h_2(t) = (1 - \alpha t)F(t),$$

with α produced as above. This route also shows that the supremum on the left is attained by the limiting exponent pair, even though it need not be attained by any single s.

Layer 1: WHAT

Block 1. Put x where it belongs: $g_1(s)^x g_2(s) = e^{\alpha_s x+\beta_s}$, an exponential of a line in x.

Block 2. Read the right hand side as an upper envelope of lines in x, and recall that the functions expressible that way are exactly the finite convex ones. The target becomes: prove F is convex.

Block 3. Cash in the hypotheses for finiteness: $[1, \infty)$ and bounded together say the exponent pair lives in a bounded box, so $F(x)$ is finite for every real x.

Block 4. Each family member is convex in x; a pointwise supremum of convex functions is convex. Two lines, done.

Block 5. Supporting line at each t gives the slope m_t; name $h_1(t) = m_t$ and $h_2(t) = F(t) - tm_t$.

Block 6. Equality at $t = x$ upgrades the supremum on the right to a maximum, which the problem demanded by writing max rather than sup.

The nominal route, and the juxtaposition. The zombie tries to evaluate $\sup_s g_1(s)^x g_2(s)$, or invents a maximising $s(x)$ and differentiates through it, or reaches for Holder and weighted means because the expression looks like an interpolation. All of these are attempts to compute an object that is not computable, since g_1, g_2 are arbitrary. The bifurcation point is the sentence where you stop asking *what is* the left hand side and start asking *what property of the left hand side* the right hand side is claiming. The problem is testing whether you can read $\max_t(xh_1(t) + h_2(t))$ as the phrase "is convex", and nothing else.

Layer 2: HOW

Block 1. Cue. A base raised to the power x, where the base is indexed by a parameter and x is the variable. **Why natural.** Powers with a variable exponent are exponentials; put the variable in the exponent and the parameter in the coefficients. Then x appears affinely, and affine dependence is the only kind convexity arguments can see. This move costs one line and it is forced by which symbol is the variable. **Link to the next block.** Once every family member is $e^{\alpha x+\beta}$, the family is closed under everything convexity cares about, and the right hand side can be compared with it.

Block 2. **Cue.** The right hand side is a maximum over t of something affine in x. **Why natural.** An affine function of x is a line; a maximum of lines is an upper envelope; upper envelopes of lines are convex, and conversely every finite convex function on $\mathbb{R}$ is the upper envelope of its supporting lines. So the identity to be proved is not an identity to be computed, it is the statement that F is convex, plus the attainment claim. Recognising this is the whole problem. **Link to the next block.** Convexity has to be proved for F, so you need to know F is real valued first: an extended real valued supremum equal to $+\infty$ somewhere cannot have a supporting line there.

Block 3. **Cue.** Two hypotheses that look like decoration: codomain $[1, \infty)$, and bounded. **Why natural.** Every hypothesis in a Putnam problem is load bearing; find the load. Here $[1, \infty)$ gives $\alpha_s, \beta_s \geq 0$ and boundedness gives $\alpha_s, \beta_s \leq \log M$. The bound $\Phi_s(x) \leq M^{|x|+1}$ then holds uniformly in s, for negative x as well as positive. Without the lower bound 1 the supremum blows up for negative x; without the upper bound it blows up for positive x. **Link to the next block.** With F finite you may now speak of convexity of F as an ordinary function $\mathbb{R} \to \mathbb{R}$.

Block 4. **Cue.** A supremum of a family, and a property (convexity) defined by an inequality. **Why natural.** Any property defined by an inequality of the form "value at a convex combination is at most the convex combination of values" passes to pointwise suprema, because you bound each family member by the supremum on the right, then take the supremum on the left. Write that argument out once and you own it forever. Note it fails for infima, and it fails for concavity. **Link to the next block.** Convexity in hand, the only remaining question is which line to attach to which t.

Block 5. **Cue.** You must produce h_1, h_2 as *functions of* t, so you need one line per real number t. **Why natural.** The supporting line at t is the canonical choice, and it is the only one that will make the envelope tight at $x = t$. Its existence for a finite convex function is the standard lemma; prove it from chord slope monotonicity so that nothing is quoted. The slope is the infimum of right hand difference quotients, which is finite because left hand quotients bound it below. **Link to the next block.** Having equality at $x = t$ is exactly what turns a supremum into a maximum.

Block 6. **Cue.** The problem wrote max, not sup. **Why natural.** Examiners choose that symbol deliberately. It adds the requirement that the envelope be attained, and it also hints at how: whatever family you build must contain, for each x, a member that touches. The supporting line at $t = x$ touches at x by construction, so the hint has been followed. **Link to the next block.** The two sides now agree for every x, which is the statement to be proved.

Layer 3: WHY

The truth beneath the route is the elementary case of the Fenchel and Moreau theorem: on $\mathbb{R}$, a finite convex function equals the pointwise supremum of its affine minorants, and at each point the supremum is attained by a supporting line. Read backwards, that says the class of functions representable as $\max_t(xh_1(t) + h_2(t))$ is exactly the class of finite convex functions on $\mathbb{R}$. So this problem is a one directional instance of a characterisation, and the problem setter chose a family, exponentials with bounded exponents, whose supremum is guaranteed convex and finite. Once you see the classification, there is nothing left to invent.

The second truth is a stability principle worth stating in general: convexity is preserved by pointwise suprema, by nonnegative combinations, by composition with an increasing convex function on the outside, and by affine substitution on the inside. Those four closure rules cover almost every convexity claim you will need. In this problem two of them apply and either finishes: sup of the convex functions Φ_s, or exp of the convex function $f = \log F$, which is the official route.

Knowing the closure list is what lets you choose the cheaper one instantly rather than reconstructing an argument.

The third truth is the duality bookkeeping. With $h_1(t) = m_t$ a supporting slope at t and $h_2(t) = F(t) - tm_t$, the pair satisfies $h_2(t) = -F^*(h_1(t))$, where $F^*(y) = \sup_x(xy - F(x))$ is the Legendre transform: indeed the supremum defining $F^*(m_t)$ is attained at $x = t$, precisely because the line supports there. So the h_1, h_2 you produced are a parametrisation of the graph of $-F^*$, and the identity you proved is $F = F^{**}$ in disguise. That is worth knowing because the Legendre transform is the standard tool for turning a supremum over one variable into a maximum over a dual variable, which is what the question asked you to do.

The discernment mechanism, stated as a rule you can apply under pressure: when an identity has an unevaluable supremum on one side and a family of affine functions on the other, the content is a structural property (convexity, subadditivity, monotonicity), never a computation. Train it by rewriting each side in the same variables and asking which side is a definition and which side is a theorem. Here the left side is a definition of a function of x, and the right side is a theorem about that function. The examiners nudged you three times: with $[1, \infty)$ (finiteness for $x < 0$), with bounded (finiteness for $x > 0$), and with max (attainment). Three hypotheses, three jobs, no spare parts.

Revolving concepts

Convex function. F is convex on an interval if $F(\lambda x_1 + (1-\lambda)x_2) \le \lambda F(x_1) + (1-\lambda)F(x_2)$ for all $\lambda \in [0,1]$. For twice differentiable F this is $F'' \ge 0$, but the inequality definition is the one that survives corners and suprema. See Book 1, field: Inequalities, topic: convexity.

Closure rules. If each F_i is convex then so is $\sup_i F_i$ (pointwise, wherever finite), $\sum c_i F_i$ for $c_i \ge 0$, $x \mapsto F(ax+b)$, and $\varphi \circ F$ for φ convex and nondecreasing. Infima do not preserve convexity: $\min(x, -x) = -|x|$ is concave. See Book 1, field: Inequalities, topic: operations preserving convexity.

Chord slope monotonicity and one sided derivatives. For convex F the difference quotient $\frac{F(v)-F(u)}{v-u}$ is nondecreasing in each argument. It follows that $F'_-(t)$ and $F'_+(t)$ exist with $F'_-(t) \le F'_+(t)$, that F is continuous, that F is differentiable off a countable set, and that every $m \in [F'_-(t), F'_+(t)]$ supports at t. The set of such m is the subdifferential $\partial F(t)$. See Book 1, field: Real analysis, topic: convex functions on the line.

Supporting line and supporting hyperplane. Every finite convex function on $\mathbb{R}$ has a supporting line at every point; in $\mathbb{R}^n$ the same statement holds with hyperplanes, proved by separating the point from the epigraph. Consequently a finite convex function is the supremum of its affine minorants. See Book 1, field: Real analysis, topic: supporting hyperplane theorem.

Legendre and Fenchel transform. $F^*(y) = \sup_x(xy - F(x))$; then $F^{**} = F$ for closed proper convex F, and $y \in \partial F(x)$ if and only if $xy = F(x) + F^*(y)$. Worked example: $F(x) = e^x$ gives $F^*(y) = y\log y - y$ for $y > 0$. See Book 1, field: Real analysis, topic: convex duality.

Log convexity, for contrast. A positive function is log convex if $\log F$ is convex; log convex implies convex, but not conversely. The official solution builds $\log F$ first, which shows the left hand side here is log convex, a strictly stronger fact than needed. Log convexity is the natural home of the Bohr and Mollerup characterisation of the gamma function and of Holder type interpolation bounds. See Book 1, field: Inequalities, topic: log convexity.

Recursion content. There is none in this problem: no sequence, no induction, no self similar structure. Do not manufacture one.

Circuit Bank

- **WHEN** a variable sits in an exponent as a^x with the base carrying a parameter **DO** rewrite as $e^{\alpha x+\beta}$ with α, β the parameters **BECAUSE** the variable then enters affinely, which is the only dependence convexity arguments can exploit

- **WHEN** one side of a required identity is a supremum or maximum of affine functions of the variable **DO** stop computing and prove that the other side is convex and finite, then produce supporting lines **BECAUSE** finite convex functions on $\mathbb{R}$ are exactly the upper envelopes of families of lines, and the envelope is attained at each point
- **WHEN** a problem hands you boundedness together with a codomain bounded away from the bad end **DO** convert both into the statement that the parameters lie in a compact box **BECAUSE** that box is what makes the supremum finite for every sign of the variable, and it also gives compactness for free
- **WHEN** a statement writes max where sup would have been safer **DO** treat attainment as a separate deliverable and design the construction so that the index $t = x$ touches **BECAUSE** examiners choose max deliberately, and an unattained supremum makes the statement false as written
- **WHEN** a property defined by an inequality must be proved for a supremum of a family **DO** bound each family member by the supremum on the far side, then take the supremum on the near side **BECAUSE** the far side no longer depends on the index, so the supremum passes through the inequality

Backtesting

1. First useful cue: x appears as an exponent on the left and as a coefficient of $h_1(t)$ on the right, so both sides are envelopes of one parameter families and the right hand family is affine in x.
2. Method the expression suggested: rewrite the family as $e^{\alpha_s x+\beta_s}$, observe each member is convex in x, and take the supremum, since supremum preserves convexity.
3. Where you would hesitate, and the thought to reinforce: at the existence of the supporting line, wondering whether F is differentiable. Reinforce this: convexity, not differentiability, is what produces supporting lines, and the proof is four lines from chord slope monotonicity. Never make smoothness assumptions you were not given.
4. The trap the question sets: trying to evaluate the supremum, or assuming it is attained at some s. It need not be attained at any s, and no formula for it exists.
5. What makes you recognise it faster next time: the shape $\max_t(xh_1(t) + h_2(t))$ is a Legendre transform written without its name. Any time you meet it, translate it immediately as "the left hand side is convex".
6. The reusable decision rule: an identity between an unevaluable supremum and a family of lines is a convexity statement, and the supporting line at $t = x$ is the construction that makes the maximum attained.

Mirror drills

1. Let $g : \mathbb{R} \to [1, M]$ be surjective. Compute $F(x) = \sup_s g(s)^x$ explicitly and write down h_1, h_2. *Answer:* $F(x) = \max(1, M^x)$; take $h_1(t) = 0, h_2(t) = 1$ for $t \leq 0$ and $h_1(t) = M^t \log M$, $h_2(t) = M^t(1 - t \log M)$ for $t > 0$, which is the supporting line at each t and shows the corner at $x = 0$ causes no trouble.
2. Prove that a pointwise supremum of convex functions is convex, and give an example showing that a pointwise infimum need not be. *Answer:* the two line argument of Step 2; for the counterexample, $\min(x, -x) = -|x|$, an infimum of two lines that is concave.

3. For $F(x) = e^x$, compute the Legendre transform $F^*(y) = \sup_x(xy - F(x))$ and verify $h_2(t) = -F^*(h_1(t))$ for the tangent line data $h_1(t) = e^t$, $h_2(t) = e^t(1-t)$. *Answer:* $F^*(y) = y\log y - y$ for $y > 0$, so $-F^*(e^t) = -\left(te^t - e^t\right) = e^t(1-t) = h_2(t)$.

14.12 B6. The sign of cubing modulo a prime

Let p be an odd prime number such that $p \equiv 2 \pmod 3$. Define a permutation π of the residue classes modulo p by $\pi(x) \equiv x^3 \pmod p$. Show that π is an even permutation if and only if $p \equiv 3 \pmod 4$.

Read it out loud

A permutation of a finite set is defined by an arithmetic operation, and you are asked for its sign. The answer is a congruence condition modulo 4. Three separate structures are therefore in play: the group S_p and its sign homomorphism, the arithmetic of $\mathbb{Z}/p$, and an ordering of the residues, because sign has nothing to do with order but inversions do. Any solution must decide how those three interact.

Start with the hypothesis $p \equiv 2 \pmod 3$, because a hypothesis stated first is usually the one that makes the object exist. It says $p - 1 \equiv 1 \pmod 3$, so $3 \nmid p - 1$, so $\gcd(3, p - 1) = 1$. That is exactly the condition for $x \mapsto x^3$ to be a bijection on the nonzero classes: choose u with $3u \equiv 1 \pmod{p-1}$ and $x \mapsto x^u$ inverts it by Fermat. Without that hypothesis π is three to one on a subgroup and the word permutation in the statement would be false. So hypothesis one buys existence.

Now the structural reading. Cubing is a nonlinear operation on $\mathbb{Z}/p$ under addition, but it is the cube map on a group under multiplication, and $(\mathbb{Z}/p)^\times$ is cyclic of order $p - 1$. Fix a primitive root g and write every nonzero class as g^i. Then cubing sends g^i to g^{3i}: in the exponent coordinate the map is multiplication by 3 on $\mathbb{Z}/(p-1)$, which is linear. Since sign depends only on cycle type, and cycle type is unchanged by relabelling the set, you may compute the sign of the relabelled map instead. That single move, the discrete logarithm, is the whole problem. Everything on the far side is a finite count.

The fields are group theory (sign homomorphism, conjugation invariance, cyclic groups), elementary number theory (primitive roots, orders, and behind the scenes quadratic reciprocity), and combinatorics (inversions, floor functions, lattice point counts). Within thirty seconds you should have on the page: $q = p - 1$; $\sigma(i) = 3i \bmod q$; the note that $\operatorname{sgn}\pi = \operatorname{sgn}\sigma$; the note that $\operatorname{sgn}\sigma = (-1)^N$ with N the number of inversions; and $q = 3m + 1$ where $m = (p-2)/3$, because $3 \nmid q$ means the multiples of 3 wrap the interval $[0, q)$ in a pattern controlled by m.

A knowledge gap declared up front, because pretending otherwise wastes your time. This problem is unapproachable without three facts that are not in A-level or Further Mathematics: the sign homomorphism on S_n and its invariance under relabelling; the existence of primitive roots modulo p, that is, the cyclicity of $(\mathbb{Z}/p)^\times$; and the identity $\operatorname{sgn} = (-1)^{\text{number of inversions}}$ relative to any fixed linear order. Install those three and the problem is a page of counting. Miss any one and no amount of ingenuity substitutes.

Approaches, ranked worst to best

1. **Tabulate small primes and match the pattern.** *Cue:* the conclusion is a congruence, so small cases look tempting. *Circuit:* compute and guess. With $p = 5, 11, 17, 23$ you will see the pattern immediately and prove nothing. Use this to check the final parity formula, never as the argument.

2. **Attack the cycle structure of $x \mapsto x^3$ directly on $\mathbb{Z}/p$.** *Cue:* sign is determined by cycle type. *Circuit:* decompose into cycles. Without the discrete logarithm you are asking for the orbit lengths of cubing on residues, which are the multiplicative orders of 3 modulo the various orders of elements. You will discover this only after fighting the multiplicative structure by hand, which is the same structure the discrete logarithm hands you for free.

3. **Product formula and quadratic reciprocity, after Elkies.** *Cue:* sign equals $\prod_{x<y} \frac{\pi(x)-\pi(y)}{x-y}$.

Circuit: evaluate a symmetric product modulo p. Since $\frac{x^3-y^3}{x-y} = x^2+xy+y^2$, the sign is congruent modulo p to $\prod_{0\le x<y<p}(x^2+xy+y^2)$, and you then need how often the binary quadratic form x^2+xy+y^2 takes each nonzero value. Its discriminant is -3, and -3 is a quadratic residue modulo p if and only if $p \equiv 1 \pmod 3$, so for our primes the form is anisotropic and each nonzero value is taken exactly $p+1$ times over ordered pairs. This route is beautiful, it is essentially a proof of a reciprocity law, and it costs the most: conics over $\mathbb{F}_p$, the quadratic character of -3, and a primitive root computation at the end. A caveat worth clarifying: the published write up states $\chi(-3) = +1$ for these primes, which is the wrong value, and it also drops a sign in the count of unordered pairs; the two slips cancel and its final exponent $(p+\chi(-c))/2$ is correct. If you reconstruct this route, recompute both steps and test against $p = 5$.

4. **Count cycles of even length, after van Hoeij.** *Cue:* sign is $(-1)^{\text{number of even cycles}}$. *Circuit:* orbit counting by element order. Elements of multiplicative order d are permuted among themselves and fall into $\varphi(d)/f(d)$ orbits, where $f(d) = \operatorname{ord}_d(3)$. The summands for odd d and for $2d$ coincide and cancel in pairs, so only d divisible by 4 matter; if $p \equiv 3 \pmod 4$ there are none. Short and conceptual, but the case $4 \mid d$ needs the structure of $(\mathbb{Z}/2^a)^\times$ as $\mathbb{Z}/2 \times \mathbb{Z}/2^{a-2}$ and a two adic valuation comparison. Rank it above the reciprocity route for economy, below the inversion count for prerequisites.

5. **Discrete logarithm, then count inversions.** *Cue:* cubing is multiplicative and the group is cyclic. *Circuit:* conjugate the permutation into a linear map on $\mathbb{Z}/(p-1)$, then count inversions. Cubing becomes $i \mapsto 3i \bmod (p-1)$, whose inversions split into three blocks according to $\lfloor 3i/(p-1) \rfloor \in \{0,1,2\}$, and each cross block pair contributes the same triangular number. Self contained, one page, no reciprocity, no two adic analysis. This is the solution to write.

The ideal solution

Throughout, write $q = p-1$ and $m = \frac{p-2}{3}$, which is an integer because $p \equiv 2 \pmod 3$. Then $q = 3m+1$, and m is odd because $p = 3m+2$ is odd.

Step 0: π is a permutation. Since $p \equiv 2 \pmod 3$ we have $q \equiv 1 \pmod 3$, so $\gcd(3,q) = 1$. Choose u with $3u \equiv 1 \pmod q$. For $x \not\equiv 0$, $(x^3)^u = x^{3u} = x$ by Fermat's little theorem, and $\pi(0) = 0$, so π is a bijection of $\mathbb{Z}/p$.

Step 1: transport the problem to the exponents. Let g be a primitive root modulo p and let $\phi : \{0,1,\dots,q-1\} \to (\mathbb{Z}/p)^\times$ be the bijection $\phi(i) = g^i$. Define σ on $\{0,1,\dots,q-1\}$ by $\sigma(i) = 3i \bmod q$; this is a permutation because $\gcd(3,q) = 1$. For every i,

$$\pi(\phi(i)) = (g^i)^3 = g^{3i} = g^{3i \bmod q} = \phi(\sigma(i)),$$

using $g^q = 1$. Hence the restriction of π to $(\mathbb{Z}/p)^\times$ equals $\phi \circ \sigma \circ \phi^{-1}$, so it has the same cycle type as σ, hence the same sign. Adjoining the fixed point 0 adds a cycle of length 1 and changes nothing. Therefore $\operatorname{sgn}\pi = \operatorname{sgn}\sigma$.

Step 2: sign as an inversion count. For a permutation σ of a finite totally ordered set, $\operatorname{sgn}\sigma = (-1)^N$ where

$$N = \#\{(i,j) : i < j,\ \sigma(i) > \sigma(j)\}$$

is the number of inversions. (Reason: swapping two adjacent positions changes N by exactly 1 and changes the sign, and the identity has $N = 0$.) We order $\{0,1,\dots,q-1\}$ in the usual way and count.

Step 3: split the index range into three blocks. For $0 \le i \le q-1$ we have $0 \le 3i < 3q$, so $\lfloor 3i/q \rfloor \in \{0,1,2\}$, and $3i$ is never equal to q or to $2q$ because $3 \nmid q$. Using $q = 3m+1$:

$$3i < q \iff i \le m, \qquad q < 3i < 2q \iff m+1 \le i \le 2m, \qquad 2q < 3i < 3q \iff 2m+1 \le i \le 3m.$$

So put $B_0 = \{0, \dots, m\}$, $B_1 = \{m+1, \dots, 2m\}$, $B_2 = \{2m+1, \dots, 3m\}$, of sizes $m+1, m, m$, and note $\sigma(i) = 3i - rq$ for $i \in B_r$.

If $i < j$ lie in the same block then $\sigma(i) - \sigma(j) = 3(i-j) < 0$, so there is no inversion inside a block. The blocks are increasing intervals, so every inversion has $i \in B_r$ and $j \in B_{r'}$ with $r < r'$, and the condition $\sigma(i) > \sigma(j)$ reads

$$3i - rq > 3j - r'q \iff 3(j-i) < (r'-r)q.$$

Step 4: the three counts. Since $q = 3m+1$, the condition $3(j-i) < q$ is equivalent to $j - i \le m$, and $3(j-i) < 2q$ is equivalent to $j - i \le 2m$.

- $r = 0, r' = 1$: for $i \in B_0$ we need $j \in B_1$ with $j \le i + m$. As $i \le m$ we have $i + m \le 2m$, so j runs over $\{m+1, \dots, i+m\}$, giving exactly i values. Total $\sum_{i=0}^{m} i = \frac{m(m+1)}{2}$.
- $r = 1, r' = 2$: for $i \in B_1$ we need $j \in B_2$ with $j \le i + m$. As $m+1 \le i \le 2m$ we have $2m+1 \le i+m \le 3m$, so j runs over $\{2m+1, \dots, i+m\}$, giving exactly $i - m$ values. Total $\sum_{i=m+1}^{2m} (i-m) = \frac{m(m+1)}{2}$.
- $r = 0, r' = 2$: for $i \in B_0$ we need $j \in B_2$ with $j \le i + 2m$. As $i \le m$ we have $i + 2m \le 3m$, so j runs over $\{2m+1, \dots, i+2m\}$, giving exactly i values. Total $\frac{m(m+1)}{2}$.

Hence

$$N = 3 \cdot \frac{m(m+1)}{2}, \qquad m = \frac{p-2}{3},$$

equivalently $N = \frac{(p-2)(p+1)}{6}$. (Check with $p = 11$: $m = 3$, $N = 18$, and the sequence $\sigma(0), \dots, \sigma(9) = 0, 3, 6, 9, 2, 5, 8, 1, 4, 7$ does have 18 inversions.)

Step 5: parity. Since p is odd and $p = 3m + 2$, the integer m is odd, so $m + 1$ is even and

$$N = 3m \cdot \frac{m+1}{2} \equiv \frac{m+1}{2} \pmod 2,$$

because $3m$ is odd. Thus N is even if and only if $4 \mid m+1$, that is $m \equiv 3 \pmod 4$. Finally $p = 3m + 2$, so $m \equiv 3 \pmod 4$ gives $p \equiv 11 \equiv 3 \pmod 4$, and $m \equiv 1 \pmod 4$ (the only other option, as m is odd) gives $p \equiv 5 \equiv 1 \pmod 4$. Therefore

$$\operatorname{sgn} \pi = \operatorname{sgn} \sigma = (-1)^N = +1 \iff p \equiv 3 \pmod 4,$$

which is the assertion.

Layer 1: WHAT

Block 1. Convert the hypothesis $p \equiv 2 \pmod 3$ into $\gcd(3, p-1) = 1$ and observe that this is what makes π a permutation at all.

Block 2. Recall that sign depends only on cycle type, hence is unchanged by relabelling the underlying set; so you are free to choose coordinates.

Block 3. Choose the coordinates in which cubing is linear: the discrete logarithm turns $x \mapsto x^3$ on $(\mathbb{Z}/p)^\times$ into $i \mapsto 3i$ on $\mathbb{Z}/(p-1)$.

Block 4. Sign needs an order, so fix representatives $0, 1, \dots, q-1$ and use $\operatorname{sgn} = (-1)^{\text{inversions}}$.

Block 5. Inversions can only occur when the wrap count $\lfloor 3i/q \rfloor$ jumps, so split the range into three blocks and count cross block pairs.

Block 6. In each of the three cross block pairs the inversion condition is a single inequality $j - i \le$ constant, so each count is the same triangular number $\frac{m(m+1)}{2}$.

Block 7. Carry only the parity: m is odd, $N \equiv \frac{m+1}{2}$, and translating back through $p = 3m + 2$ gives the congruence modulo 4.

The nominal route, and the juxtaposition. The zombie starts from the definition of even permutation, writes x^3 as a permutation of $0, 1, \dots, p-1$, and tries to find its cycles. That is not stupid; it

is simply the wrong coordinate system, and it leads into orders of 3 modulo divisors of $p-1$, which is a harder count than the one the good coordinates give. The bifurcation point is the line where you write $x = g^i$. Before that line the map is nonlinear and the set has no useful structure; after it the map is multiplication by 3 on a cyclic group and the entire remaining problem is counting lattice points below a line. If you take one sentence from this problem, take this: cubing is nonlinear on residues and linear on exponents, so work with exponents.

Layer 2: HOW

Block 1. **Cue.** A congruence hypothesis stated before the object is even defined. **Why natural.** The statement calls π a permutation, so something must guarantee injectivity of cubing. On a cyclic group of order q the map $x \mapsto x^k$ is bijective exactly when $\gcd(k, q) = 1$, and $p \equiv 2 \pmod 3$ is precisely $\gcd(3, p-1) = 1$. Reading a hypothesis as the reason an object exists is a habit that pays on every algebra problem. **Link to the next block.** Knowing $\gcd(3, q) = 1$, you also know that 3 is invertible modulo q, which is what makes the exponent side of the problem a permutation too.

Block 2. **Cue.** The question asks for a sign, not a formula, and sign is a group homomorphism $S_n \to \{\pm 1\}$. **Why natural.** Sign is constant on conjugacy classes, and conjugation is exactly relabelling of the underlying set. So you are permitted to replace the set $\mathbb{Z}/p$ by any set in bijection with it, provided you transport the map. That freedom is enormous and it is free; the only question left is which relabelling makes the map simple. **Link to the next block.** The map is multiplicative, so the relabelling should be one that linearises multiplication. That is the definition of a logarithm.

Block 3. **Cue.** The multiplicative group modulo a prime is cyclic, and π is a power map. **Why natural.** Under $x = g^i$ the operation $x \mapsto x^3$ becomes $i \mapsto 3i$ modulo $p-1$. Nothing was invented here: this is what cyclicity is for. The fixed point 0 is discarded first because a fixed point contributes a cycle of length one and cannot affect the sign. **Link to the next block.** You now hold a linear map on $\mathbb{Z}/q$. Sign, however, is not defined for a map on an unordered set of residues; you must fix representatives, and that is where the counting starts.

Block 4. **Cue.** Sign of a concrete permutation of a totally ordered finite set. **Why natural.** The inversion formula is the only computational handle that turns sign into a count you can do by hand, and it is available the moment you choose the representatives $0, 1, \ldots, q-1$. Note that the number of inversions depends on the chosen order but its parity does not, because the sign does not. **Link to the next block.** The permutation is $i \mapsto 3i - q\lfloor 3i/q \rfloor$, so its failure to be increasing is entirely caused by the subtracted multiple of q. That tells you where to look for inversions.

Block 5. **Cue.** A floor function taking only three values on the index range. **Why natural.** If $\lfloor 3i/q \rfloor = \lfloor 3j/q \rfloor$ then σ is increasing on that pair, so no inversion. Partitioning the index range into the level sets of the floor function is therefore not a convenience, it is the exact decomposition the problem has: inversions live only between different levels. This is the standard move for any map of the form $i \mapsto ai \bmod q$. **Link to the next block.** Within a fixed pair of levels the inversion condition simplifies to a linear inequality in $j - i$, because the wrap contribution is a constant.

Block 6. **Cue.** Three cross block pairs, each with a condition of the form $j - i \le$ constant. **Why natural.** Counting integer pairs in an interval subject to $j \le i + c$ is a triangular number; there is nothing to be clever about. Doing all three and finding the same answer $\frac{m(m+1)}{2}$ three times is a useful internal check that no block boundary was mishandled. **Link to the next block.** You now have $N = 3\binom{m+1}{2}$, an exact count, which is more than you need; only its parity matters.

Block 7. Cue. The target is a congruence modulo 4 and you hold a product of an odd number with $\frac{m+1}{2}$. **Why natural.** Carry parity, not size. $3m$ is odd, so $N \equiv \frac{m+1}{2} \pmod 2$, and the condition $4 \mid m+1$ translates through $p = 3m+2$ into $p \equiv 3 \pmod 4$. The factor 3 multiplying the triangular number is invisible to parity, which is why the three block counts did not need to be added carefully, only counted. **Link to the next block.** The chain $\operatorname{sgn}\pi = \operatorname{sgn}\sigma = (-1)^N$ closes the argument.

Layer 3: WHY

The first truth is that sign is a class function: conjugate permutations have equal sign, because conjugation preserves cycle type. On a finite set with no distinguished structure, that means the sign of a map is an invariant of the map up to relabelling, so you may always compute it in the most convenient coordinates. Almost every Putnam problem about the parity of a permutation is really a problem about choosing coordinates, and the coordinates are chosen by asking which group operation the map respects. Cubing respects multiplication; therefore work in the group where multiplication is addition.

The second truth is the cyclicity of $(\mathbb{Z}/p)^\times$, that is, the existence of primitive roots. It is what makes the discrete logarithm exist. It is not obvious and it is not cheap to prove (the standard argument counts elements of each order using $\sum_{d\mid n}\varphi(d) = n$ and the fact that $x^d = 1$ has at most d roots in a field), but you may quote it. Once quoted it converts multiplicative questions modulo p into additive questions modulo $p-1$, and the difficulty of a problem often collapses when you cross that bridge. Notice that this is why the answer involves $p-1$ arithmetic (a condition modulo 4) rather than p arithmetic.

The third truth is that counting inversions of $i \mapsto ai \bmod q$ is a lattice point count under a line, which is precisely the shape that appears in the Eisenstein and Gauss lemma proofs of quadratic reciprocity. That is not a coincidence. Zolotarev's lemma states that for odd n and $\gcd(a,n) = 1$, the sign of multiplication by a on $\mathbb{Z}/n$ is the Jacobi symbol $\left(\frac{a}{n}\right)$, and reciprocity itself can be derived from that identity. Here the modulus $q = p-1$ is even, so Zolotarev does not apply and the count has to be done by hand, but the family resemblance explains why the answer is a congruence modulo 4: parities of triangular numbers are governed by residues modulo 4, and that is the same mechanism that puts $\frac{p-1}{2}$ and $\frac{p^2-1}{8}$ into the supplementary laws.

The discernment mechanism, stated so you can apply it cold: when a permutation is defined by an arithmetic formula, ask which group operation it is a homomorphism for, and relabel so that it becomes linear. Then ask whether you need cycle structure (use orbit lengths and orders) or a sign (use inversions). Cycle counting is better when the orbit lengths are uniform; inversion counting is better when the map is linear on an ordered range, because then inversions are governed by a floor function with few values. Train it by taking the maps $x \mapsto x^{-1}$, $x \mapsto ax$, $x \mapsto x^k$ modulo p and computing all three signs by both methods, once.

Revolving concepts

Sign of a permutation. The unique homomorphism $\operatorname{sgn} : S_n \to \{\pm 1\}$ with $\operatorname{sgn}(\text{transposition}) = -1$. A cycle of length ℓ has sign $(-1)^{\ell-1}$, so $\operatorname{sgn}\sigma = (-1)^{n-c(\sigma)}$ where c counts cycles including fixed points, and equivalently $\operatorname{sgn}\sigma = (-1)^{\#\{\text{even length cycles}\}}$. Sign is constant on conjugacy classes. See Book 1, field: Abstract algebra, topic: symmetric group and sign.

Inversions. For a permutation of a totally ordered set, $\operatorname{sgn}\sigma = (-1)^N$ with N the number of pairs out of order. The count N itself depends on the order chosen; the parity does not. See Book 1, field: Combinatorics, topic: inversions and the sign statistic.

Primitive roots and the discrete logarithm. $(\mathbb{Z}/p)^\times$ is cyclic of order $p-1$; a generator is a primitive root g, and $x \mapsto \log_g x$ is an isomorphism onto $\mathbb{Z}/(p-1)$. Consequences used here: $x \mapsto x^k$ is bijective if and only if $\gcd(k, p-1) = 1$; the number of elements of order d dividing $p-1$ is $\varphi(d)$;

the k-th powers form the subgroup of index $\gcd(k, p-1)$. See Book 1, field: Number theory, topic: primitive roots.

Order of an element and orbit lengths. If x has multiplicative order d and $\gcd(3, d) = 1$, the orbit of x under cubing has length $\mathrm{ord}_d(3)$, the multiplicative order of 3 modulo d. This is the engine of the cycle counting route, and it generalises the recursion $x_{n+1} = x_n^3$: the orbit closes after k steps exactly when $3^k \equiv 1 \pmod{d}$, so every question about iterating a power map reduces to a multiplicative order in the exponent group. See Book 1, field: Number theory, topic: orders and iterated power maps.

Zolotarev's lemma and quadratic reciprocity. For odd n and $\gcd(a, n) = 1$, the sign of $x \mapsto ax$ on $\mathbb{Z}/n$ equals the Jacobi symbol $\left(\frac{a}{n}\right)$; quadratic reciprocity follows. For p prime and $a = g^k$, the direct proof is short: multiplication by a splits $(\mathbb{Z}/p)^\times$ into $(p-1)/d$ cycles of length $d = \mathrm{ord}_p(a)$, so the sign is $(-1)^{(d-1)(p-1)/d}$, which equals $+1$ exactly when a is a quadratic residue. See Book 1, field: Number theory, topic: quadratic reciprocity and Zolotarev.

Quadratic character of small numbers. $\left(\frac{-1}{p}\right) = 1$ iff $p \equiv 1 \pmod{4}$; $\left(\frac{2}{p}\right) = 1$ iff $p \equiv \pm 1 \pmod{8}$; $\left(\frac{-3}{p}\right) = 1$ iff $p \equiv 1 \pmod{3}$. The last one is what governs the binary form $x^2 + xy + y^2$, whose discriminant is -3, and hence the reciprocity route to this problem. See Book 1, field: Number theory, topic: Legendre symbol and supplementary laws.

Structure of $(\mathbb{Z}/2^a)^\times$. For $a \geq 3$ it is $\mathbb{Z}/2 \times \mathbb{Z}/2^{a-2}$, generated by -1 and 3, with $\mathrm{ord}_{2^a}(3) = 2^{a-2}$. This is what the cycle counting route needs in order to show $\varphi(d)/\mathrm{ord}_d(3)$ is even for every d divisible by 4 other than $d = 4$ itself. See Book 1, field: Number theory, topic: units modulo prime powers.

Circuit Bank

- **WHEN** a permutation of a finite set is defined by an arithmetic formula and you are asked for its sign or parity **DO** ask which group operation the formula is linear for, and relabel the set by an isomorphism so that it becomes linear **BECAUSE** sign depends only on cycle type, so relabelling is free, and a linear map is countable by hand
- **WHEN** a power map $x \mapsto x^k$ appears modulo a prime **DO** take discrete logarithms and study $i \mapsto ki$ modulo $p-1$ **BECAUSE** $(\mathbb{Z}/p)^\times$ is cyclic, so power maps are multiplication maps in the exponent, and bijectivity is just $\gcd(k, p-1) = 1$
- **WHEN** you must compute the sign of a concrete permutation of $\{0, 1, \ldots, n\}$ **DO** count inversions, and expect the count to be a sum of triangular numbers **BECAUSE** $\mathrm{sgn} = (-1)^{\text{inversions}}$ turns an algebraic invariant into a finite count you can organise by cases
- **WHEN** a map has the form $i \mapsto ai \bmod q$ and you must count inversions **DO** partition the range by the value of $\lfloor ai/q \rfloor$ and count only cross block pairs **BECAUSE** within a level set the map is increasing, so all disorder comes from the jumps of the floor function
- **WHEN** an exact count has been obtained but the question asks only for a parity **DO** strip odd factors immediately and reduce the expression modulo 2 before substituting back **BECAUSE** carrying the full expression invites arithmetic slips, and parity questions are usually decided by one factor

Backtesting

1. First useful cue: the hypothesis $p \equiv 2 \pmod{3}$, read as $\gcd(3, p-1) = 1$, which both makes π a permutation and makes multiplication by 3 invertible modulo $p-1$.
2. Method the expression suggested: cubing is multiplicative, so pass to exponents through a primitive root; the map becomes $i \mapsto 3i$ on $\mathbb{Z}/(p-1)$, and sign becomes an inversion count.

3. Where you would hesitate, and the thought to reinforce: at the claim that the sign is unchanged when you relabel the set. Reinforce this: conjugate permutations have the same cycle type, and cycle type determines sign; a fixed point contributes a cycle of length one and never affects sign.
4. The trap the question sets: doing the count on $\mathbb{Z}/p$ instead of on $\mathbb{Z}/(p-1)$, or forgetting that 3 divides neither q nor $2q$, which is what guarantees the three blocks have no boundary cases and that $m = (p-2)/3$ is an integer.
5. What makes you recognise it faster next time: any sign or parity question about $x \mapsto x^k$ modulo p is a question about $i \mapsto ki$ modulo $p-1$, and any question about $i \mapsto ki$ modulo q is a lattice count under a line, which is the Zolotarev and Eisenstein circle of ideas.
6. The reusable decision rule: relabel until the map is linear, order the labels, count inversions by level sets of the floor function, then carry only the parity.

Mirror drills

1. Compute the sign of the inversion map $x \mapsto x^{-1}$ on the nonzero classes modulo an odd prime p. *Answer:* it fixes ± 1 and pairs the other $p-3$ classes into $\frac{p-3}{2}$ transpositions, so the sign is $(-1)^{(p-3)/2}$, which is even exactly when $p \equiv 3 \pmod 4$.
2. Prove Zolotarev's lemma for a prime modulus: for $\gcd(a,p) = 1$, the sign of $x \mapsto ax$ on $\mathbb{Z}/p$ is $\left(\frac{a}{p}\right)$. *Answer:* the map splits the nonzero classes into $(p-1)/d$ cycles of length $d = \operatorname{ord}_p(a)$, so the sign is $(-1)^{(d-1)(p-1)/d}$; if d is odd this is $+1$ and a is a square, and if d is even it is $(-1)^{(p-1)/d}$, which matches the residue character of $a = g^k$ since $\gcd(k, p-1) = (p-1)/d$.
3. For odd q, count the inversions of $i \mapsto 2i \bmod q$ on $\{0, 1, \ldots, q-1\}$ and deduce the supplementary law for 2. *Answer:* with $q = 2m+1$ there are two blocks and the count is $\frac{m(m+1)}{2}$, which is even exactly when $q \equiv \pm 1 \pmod 8$; by Zolotarev this is $\left(\frac{2}{q}\right)$.

Chapter 15

Putnam 2011

15.1 A1. counting lattice points a spiral cannot reach

Define a growing spiral *in the plane to be a sequence of points with integer coordinates* $P_0 = (0,0), P_1, \ldots, P_n$ *such that* $n \geq 2$ *and:*

- *the directed line segments* $P_0P_1, P_1P_2, \ldots, P_{n-1}P_n$ *are in the successive coordinate directions east (for* P_0P_1*), north, west, south, east, etc.;*
- *the lengths of these line segments are positive and strictly increasing.*

How many of the points (x, y) *with integer coordinates* $0 \leq x \leq 2011$, $0 \leq y \leq 2011$ *cannot be the last point* P_n *of any growing spiral?*

Read it out loud

Read the definition aloud and you hear two separate instructions stacked on top of each other: a direction rule (east, north, west, south, repeating) and a length rule (positive, strictly increasing). The question then asks for a count of failures, how many lattice points in a fixed square cannot be produced as the endpoint of such a path. That framing, "how many cannot," is itself a cue: it is the complement of an existence question, and existence questions built from a strictly increasing sequence are almost always easy once you find the one constraint that is doing the real work.

The field wired into this problem is combinatorics, specifically lattice point counting over a bounded box, but the tool that cracks it open belongs to sequences: an alternating sum of a monotone sequence. The moment you read "lengths... strictly increasing" sitting next to a direction cycle that reverses sign every other step in each axis, you should recognise that each coordinate of P_n is nothing more than a signed sum of a strictly increasing sequence of positive integers. That recognition is the first useful cue, and within thirty seconds of reading the problem you should have written down that P_n's x and y coordinates are each an alternating sum over one of two interleaved increasing subsequences of the step lengths.

The other thing worth saying aloud is what the question is NOT asking. It is not asking you to characterise every growing spiral, and it is not asking for a formula in n. It is asking for a count inside a 2012×2012 grid, which tells you the answer will almost certainly be a small, structured region near the origin (where a short prefix of the increasing sequence still has room to constrain things) and nothing at all once the coordinates are far enough from the axes for the sequence to have exhausted its constraining power.

Approaches, ranked worst to best

1. **brute enumeration of small spirals, hunting visually for a pattern.** *Cue:* a bounded grid 0 to 2011 invites listing. *Circuit:* no circuit, pure search. (why it ranks here: reaching coordinate 2011 forces spirals with dozens of segments and many different splits of the increasing sequence between the two axes; even after staring at pictures you have proved nothing, and the work does not shrink as the bound grows, it grows with it)
2. **hunt for a parity or modular invariant on the coordinates.** *Cue:* integer lattice, direction cycle of period four. *Circuit:* look for an invariant mod a small number. (why it ranks here: a period four rotation tempts you toward an invariant mod 2 or mod 4, but the true excluded set is a boundary strip near small coordinates, not a congruence class; you burn the first minutes of the arena on a dead end before falling back to the real structure)
3. **write each coordinate as an alternating sum of the two length subsequences, pin its sign and size by pairing consecutive terms, then construct explicitly once the pairing bound is exhausted.** *Cue:* "strictly increasing" lengths riding a sign-reversing direction cycle. *Circuit:* alternating sum of a monotone sequence, pairing lemma. (why it ranks first: it turns a

geometric definition into two one line algebraic objects with no case search, the same pairing device both rules points out and tells you exactly how many terms you need before every remaining point becomes free, and it converts the "prove impossible" half and the "prove possible" half into two short, independent arguments)

The ideal solution

Write the segment lengths of a growing spiral in order as $\ell_1 < \ell_2 < \cdots < \ell_n$, positive integers (integers because every P_i has integer coordinates and every step moves an integer amount along one axis). Split them by parity of index into the horizontal lengths $x_1 = \ell_1 < x_2 = \ell_3 < x_3 = \ell_5 < \cdots$ and the vertical lengths $y_1 = \ell_2 < y_2 = \ell_4 < y_3 = \ell_6 < \cdots$. Since these are alternate terms of one increasing sequence, they satisfy the finer interleaved chain

$$x_1 < y_1 < x_2 < y_2 < x_3 < y_3 < \cdots .$$

If P_n uses $a = \lceil n/2 \rceil$ horizontal steps and $b = \lfloor n/2 \rfloor$ vertical steps, tracking east, north, west, south directly gives

$$P_n = \Big(x_1 - x_2 + x_3 - \cdots + (-1)^{a-1} x_a,\ \ y_1 - y_2 + y_3 - \cdots + (-1)^{b-1} y_b \Big).$$

Lemma 15.1.1 (Sign of an alternating monotone sum)**.** *For a strictly increasing sequence of positive integers $c_1 < c_2 < \cdots < c_k$, put $A_k = c_1 - c_2 + c_3 - \cdots + (-1)^{k-1} c_k$. Then $A_k \geq c_1 > 0$ if k is odd, and $A_k < 0$ if k is even.*

Proof. If k is odd, pair every term after the first: $A_k = c_1 + (c_3 - c_2) + (c_5 - c_4) + \cdots + (c_k - c_{k-1})$; every bracket is positive since the sequence increases, so $A_k \geq c_1$. If k is even, pair from the start: $A_k = (c_1 - c_2) + (c_3 - c_4) + \cdots + (c_{k-1} - c_k)$; every bracket is negative, so $A_k < 0$. □

Apply the lemma to the x subsequence: the x coordinate of any P_n with $n \geq 1$ is positive when a is odd and negative when a is even. In particular it is never 0, so no point with $x = 0$ beyond P_0 is ever produced. This rules out $(0, y)$ for every $0 \leq y \leq 2011$: 2012 points.

Now apply the lemma to the y subsequence for $n \geq 4$, where $b = \lfloor n/2 \rfloor \geq 2$. If b is even, $y < 0$, outside our square entirely. If b is odd, then since $b \geq 2$ we in fact have $b \geq 3$, and we can sharpen the lemma using the interleaving: write $y = y_1 + (y_3 - y_2) + (y_5 - y_4) + \cdots + (y_b - y_{b-1})$. Each bracket $y_{2j+1} - y_{2j}$ is not merely positive, because the interleaved chain places x_{2j+1} strictly between y_{2j} and y_{2j+1}, forcing $y_{2j+1} \geq y_{2j} + 2$. So every bracket is at least 2, and $y_1 \geq 2$ since $y_1 > x_1 \geq 1$, giving $y \geq y_1 + 2 \geq 4$. So every point beyond P_3 has y-coordinate either negative or at least 4.

That leaves only P_2 and P_3 (where $b = 1$, $y = y_1$) to examine inside the square. $P_3 = (x_1 - x_2, y_1)$ has negative x-coordinate by the first part of the lemma ($a = 2$), so it never lies in our region. $P_2 = (x_1, y_1)$ has $x = x_1$ and $y = y_1$ with $1 \leq x_1 < y_1$ free, and conversely every pair $1 \leq x < y \leq 2011$ is achieved this way by the two step spiral with lengths $x < y$. So the entire region $0 < x < y$, for y from 1 to 2011, is achievable, and no point with $y \in \{0, 1\}$ is ever achievable (there is no valid x_1 with $1 \leq x_1 < 1$, and $y = 0$ never occurs for $n \geq 2$ at all), and for $y \in \{2, 3\}$ only $x < y$ is achievable, i.e. $x \geq 2$ is excluded when $y = 2$ and $x \geq 3$ is excluded when $y = 3$.

It remains to show every point with $x \geq y \geq 4$ is achievable, since $0 < x < y$ is already covered by P_2 and $y \leq 3$ is already fully classified. Take the six step spiral with lengths

$$1,\ 2,\ 3,\ x + 1,\ x + 2,\ x + y - 1$$

(strictly increasing exactly when $x \geq y \geq 4$, since then $x + 1 > 3$ and $x + y - 1 > x + 2$). Its endpoint has

$$X = 1 - 3 + (x + 2) = x, \qquad Y = 2 - (x + 1) + (x + y - 1) = y,$$

so $P_6 = (x, y)$ exactly. Hence every (x, y) with $x \geq y \geq 4$ is achievable, and combined with $0 < x < y$ for all y, every point with $y \geq 4$ is achievable.

Collecting the excluded points, which fall into five disjoint groups:

$$\underbrace{(0,y),\ 0 \leq y \leq 2011}_{2012} \cup \underbrace{(x,0),\ 1 \leq x \leq 2011}_{2011} \cup \underbrace{(x,1),\ 1 \leq x \leq 2011}_{2011} \cup \underbrace{(x,2),\ 2 \leq x \leq 2011}_{2010} \cup \underbrace{(x,3),\ 3 \leq x \leq 201}_{2009}$$

the total is

$$2012 + 2011 + 2011 + 2010 + 2009 = \boxed{10053}.$$

Layer 1: WHAT

Block 1: translate the geometric spiral into algebra, writing each coordinate of P_n as an alternating sum over one of the two interleaved increasing subsequences of lengths. Block 2: prove the x coordinate is never 0, by pairing the alternating sum from the front and reading off a strict sign from the leftover first term. Block 3: push the same pairing device on the y coordinate, then sharpen it using the interleaving with the x subsequence to get a gap of 2 per bracket instead of 1, which is what forces $y \geq 4$ rather than merely $y \geq 3$ once three or more vertical steps are involved. Block 4: dispose of the two small cases P_2 and P_3 directly, since they have too few vertical steps for the sharpened bound to apply, and note P_2 alone already supplies the entire region $0 < x < y$. Block 5: construct an explicit six step spiral with two free parameters to cover the remaining region $x \geq y \geq 4$, then assemble the five disjoint excluded strips and add.

The nominal route stares at pictures of spirals for $n = 2, \ldots, 10$ hoping the excluded pattern announces itself, then tries to extend that pattern to a square of side 2011 by analogy, with no proof that nothing new goes wrong further out. The ideal route never draws a single spiral for $n > 6$; it proves a sign lemma once and reuses it for every n simultaneously. The bifurcation point sits right after you read "lengths...strictly increasing": the zombie route reaches for graph paper, the ideal route reaches for the alternating sum.

Layer 2: HOW

Block 1. **Cue.** The direction cycle east, north, west, south repeats with period four, so odd numbered steps are horizontal and even numbered steps are vertical. **Why natural.** Splitting a single increasing sequence by parity of index is the natural move whenever a path alternates between two independent axes, because it converts one geometric object into two purely one dimensional ones that only interact through their shared global ordering. **Link to the next block.** Once each coordinate is a signed sum, the question of which points are reachable becomes a question about the sign and size of an alternating sum, which is a solved shape of problem the moment you name it.

Block 2. **Cue.** The lengths are strictly increasing, so consecutive terms in each subsequence differ, not merely stay level. **Why natural.** Strict monotonicity is precisely the hypothesis that lets you pair terms and get a bracket of a definite sign, which is why the lemma needs strictness and would fail for a merely nondecreasing sequence. **Link to the next block.** The sign of x never being zero rules out the entire $x = 0$ column in one stroke, without touching y at all, so the two coordinates can be analysed independently from here.

Block 3. **Cue.** Once $b \geq 3$, the vertical subsequence has enough terms for the pairing to leave brackets rather than a single leftover term. **Why natural.** Each bracket in the y sum is not just an increasing gap in the vertical subsequence, it is a gap that has a horizontal length interleaved strictly inside it, so the gap must be at least 2, not 1; this is where the specific interleaving of this problem, not a generic fact about increasing sequences, earns you the sharper constant. **Link to the next**

block. A gap of at least 4 for every $b \geq 3$ means the only y values still in doubt near the bottom of the square are 2 and 3, both of which come only from $b = 1$.

Block 4. **Cue.** $b = 1$ has no bracket to pair at all, since there is only one vertical term. **Why natural.** With one term, $y = y_1$ is simply the length itself, unconstrained except by $y_1 > x_1 \geq 1$, and $x = x_1$ is likewise free below y_1; two free positive integers with one strict inequality between them is exactly the region $0 < x < y$. **Link to the next block.** This single two step spiral already answers "which (x, y) with $y \leq 3$ are achievable," leaving only the region $x \geq y \geq 4$ needing a separate construction.

Block 5. **Cue.** The remaining region has two free coordinates and no sign constraint left to exploit, so this half of the proof is a construction, not an inequality. **Why natural.** Pinning every length except the last two to its smallest legal value maximises the room left for those two lengths, which is exactly enough freedom to solve for two target coordinates simultaneously; six steps is the smallest step count with two vertical and two horizontal degrees of freedom left after the three smallest lengths are spent on satisfying strict increase. **Link to the next block.** With every point in the square now classified as excluded or explicitly constructed, the final count is a straightforward disjoint sum over five strips, no inclusion or exclusion correction needed since the strips share no points.

Layer 3: WHY

The mathematical truth underneath this problem is that an alternating sum of a strictly increasing sequence behaves like a restoring force: each new term overcorrects the previous partial sum, and pairing consecutive terms from whichever end leaves the parity of the count correct turns that overcorrection into a definite, computable sign. This is the same mechanism behind the alternating series remainder bound in analysis, transplanted here into a finite, integer setting where "overcorrects by at least the gap between consecutive terms" becomes a genuine quantitative bound rather than a limiting statement.

The discernment mechanism for choosing this route over brute search is structural: any time a path or sequence definition hands you strict monotonicity plus a sign reversal, that combination is a standing invitation to pair terms rather than track a running total numerically. You train that discernment by refusing to touch a specific numerical case (spiral of length 2011, say) until you have exhausted what the general shape of the definition tells you for free.

The examiners nudged you toward this by choosing a 2012×2012 square far larger than the actual obstruction, which only reaches coordinate 3. That gap between the stated bound and the true size of the phenomenon is itself a signal: if the excluded set genuinely grew with the box size in a complicated way, the problem would have to tell you a formula in terms of the bound, but it only asks for a number, which means the interesting part of the answer is a small fixed pattern near the origin, and the rest of the square is a bookkeeping exercise in one flexible construction. Recognising that shape, a small structured core plus a fully free remainder, before you have proven a single line, is what separates a twenty minute solve from a search that never terminates.

Revolving concepts

Alternating sum of a monotone sequence. For strictly increasing positive $c_1 < \cdots < c_k$, the sum $c_1 - c_2 + \cdots \pm c_k$ has sign $(-1)^{k-1}$, proved by pairing from whichever end leaves a single leftover term; the same pairing gives a quantitative lower bound on the magnitude whenever you know a lower bound on consecutive gaps. See Book 1, field: Sequences and recurrences, topic: alternating sums and telescoping sign arguments. Wherever a further sequence is interleaved between the terms of the one you are summing, as here, the effective gap bound doubles or otherwise inflates, and that inflation is exactly what should be checked first whenever a pairing bound looks one short of what a problem needs.

Lattice point counting by complement. Rather than counting an achievable set directly, count its complement in a bounded box when the complement is structurally simpler; this is most powerful exactly when the achievable set is "everything past a small threshold," so the complement collapses to finitely many boundary rows and columns. See Book 1, field: Combinatorics, topic: complement counting in bounded lattices.

Existence via degree of freedom matching. To prove a whole two parameter family of targets is reachable by a constrained increasing sequence construction, spend the minimum number of terms pinning down the ordering constraint at its smallest legal values, then solve the two remaining free terms directly against the two target values. See Book 1, field: Combinatorics, topic: explicit constructions with matched degrees of freedom; this generalises to any problem asking you to realise an m parameter target using a sequence with more than m degrees of freedom once the ordering constraints are satisfied by the cheapest possible prefix.

Circuit Bank

- **WHEN** a quantity is defined as an alternating sum of a strictly increasing sequence **DO** pair terms from whichever end leaves the correct number of leftover terms for the sign you expect **BECAUSE** each pair then has a definite sign and the leftover term fixes the overall sign, handing you both existence of sign and a lower bound on magnitude for free
- **WHEN** a competition asks how many lattice points in a bounded box cannot arise from a construction **DO** look for the complement to split into a small boundary strip near the axes plus everything past a fixed threshold reached by one flexible construction **BECAUSE** monotonicity or increasing parameters in a construction typically exhaust their constraining power quickly, so obstructions rarely persist far from the origin
- **WHEN** you need to hit an arbitrary target pair with a constrained increasing sequence construction **DO** fix every length except the last two at its smallest legal value, then solve the remaining two lengths explicitly against the two target coordinates **BECAUSE** this isolates exactly two degrees of freedom mapping onto the two targets, and the smallest legal prefix leaves the largest possible range for those two free lengths
- **WHEN** a path alternates direction with period four across two axes **DO** split the step index set into odds and evens immediately and treat the two resulting subsequences as independent **BECAUSE** the two coordinates never interact except through the one shared global ordering constraint on all the lengths

Backtesting

1. First useful cue: the phrase "strictly increasing" attached to lengths whose directions cycle through all four compass points.
2. Method the expression suggested: split the increasing sequence by parity into two interleaved increasing subsequences, and read each coordinate off as an alternating sum of one of them.
3. Where you would hesitate, and the thought to reinforce: after proving x is never 0 with a bound of $x \geq x_1$, you might assume the analogous bound $y \geq y_1$ settles every case; reinforce that once a third term enters, the interleaving with the other axis's subsequence forces a strictly larger gap than the generic pairing bound gives.
4. The trap the question sets: the box is stated up to 2011, far past where the obstruction actually lives, tempting you to believe the excluded set depends on the bound in a complicated way and to search case by case instead of proving a fixed, small pattern.

5. What makes you recognise it faster next time: any "how many lattice points cannot occur" question built from a strictly increasing construction is a strong signal that the achievable set is everything past a bounded threshold, because monotonicity is exactly the hypothesis that eventually removes every obstruction.
6. The reusable decision rule: whenever directions cycle and lengths strictly increase, write both coordinates as alternating sums of the two parity classes before doing anything else, and look for a pairing that pins sign first, size second.

Mirror drills

1. Using the same square, which points on the diagonal $x = y$, for $1 \leq x \leq 2011$, cannot be the last point of a growing spiral? *Answer:* exactly $x = y \in \{1, 2, 3\}$ are excluded (row exclusions at $y = 1, 2, 3$ each still bite on the diagonal), and every $x = y \geq 4$ is achievable by the six step construction with equal target coordinates.
2. What is the smallest n for which P_n can have y-coordinate exactly 4? *Answer:* $n = 2$, using lengths $1 < 4$; the bound $y \geq 4$ proved for $b \geq 3$ is a lower bound for that family only, it does not mean every occurrence of $y = 4$ needs three vertical steps.
3. Generalise the square to $0 \leq x \leq N$, $0 \leq y \leq N$ for $N \geq 4$ and find the excluded count as a function of N; check it reduces to 10053 at $N = 2011$. *Answer:* the same five strips give $(N+1) + N + N + (N-1) + (N-2) = 5N - 2$, and $5(2011) - 2 = 10053$.

15.2 A2. telescoping an underspecified bounded recurrence

Let $a_1, a_2, \ldots$ and $b_1, b_2, \ldots$ be sequences of positive real numbers such that $a_1 = b_1 = 1$ and $b_n = b_{n-1}a_n - 2$ for $n = 2, 3, \ldots$. Assume that the sequence (b_j) is bounded. Prove that

$$S = \sum_{n=1}^{\infty} \frac{1}{a_1 a_2 \cdots a_n}$$

converges, and evaluate S.

Read it out loud

Read the recurrence aloud: b_n is built from b_{n-1} and a_n multiplicatively, then shifted down by 2. The initial values pin $a_1 = b_1 = 1$. The one hypothesis separating a solvable problem from an impossible one is that (b_j) stays bounded, nothing more specific than that. You are not told (b_j) converges, or is monotone, or takes any particular values, only that it never runs off to infinity. And yet the problem asks you to evaluate S exactly, down to a single number. That mismatch, an almost entirely unspecified sequence paired with a fully specified answer, is the loudest cue in the problem: whatever route you take, it has to produce a value for S that cannot depend on which bounded (b_j) you were handed, which means the mechanism underneath has to be an identity holding termwise, not a limiting property special to one sequence.

The field wired into this problem is sequences and recurrences, riding underneath a real analysis convergence argument, and the object under the hood, $\sum 1/(a_1 \cdots a_n)$, is a sum of reciprocal partial products, exactly the shape that appears whenever a recursively defined sequence is compared against a shifted or rescaled copy of itself. The first useful cue, readable within the first thirty seconds, is the single word "evaluate": had the problem only asked you to prove convergence, a boundedness or comparison argument alone would suffice, but asking you to evaluate S tells you a closed form exists and is waiting to be extracted, which should point you at once toward telescoping rather than estimation.

Solve $b_n = b_{n-1}a_n - 2$ for a_n and you get $a_n = (b_n + 2)/b_{n-1}$, a ratio of consecutive shifted terms of the same sequence b. A product of such ratios, $a_1 a_2 \cdots a_n$, is a textbook telescoping shape in disguise: whenever a term is a ratio of consecutive values of one underlying sequence, a product of many such terms collapses. Here the numerator runs over $b_k + 2$ and the denominator over b_{k-1}, two different partial products built from the same raw data, and that separation is exactly the auxiliary object you should build by hand before touching a single specific value of b_j.

Approaches, ranked worst to best

1. **hunt for an explicit closed form for b_n or a_n as a function of n.** *Cue:* a recursively defined sequence invites solving the recursion outright. *Circuit:* no circuit, treats an underspecified family as if it were one fixed sequence. (why it ranks here: (b_j) is only assumed positive and bounded, not pinned to any particular values, so no single closed form for b_n exists to find; every bounded positive sequence with $b_1 = 1$ is a valid instance of the problem, and chasing a formula that cannot exist wastes the opening minutes)

2. **compute S_1, S_2, S_3 explicitly in terms of $b_2, b_3, \ldots$, simplify, and guess the closed form for S_m before proving it by induction.** *Cue:* "evaluate S" implies a closed form exists to be found. *Circuit:* pattern spotting from small cases, then induction. (why it ranks here: it does work, and it is a legitimate fallback, but it spends ink grinding two or three partial sums by hand before the telescoping shape becomes visible, and a candidate who does not simplify aggressively can misidentify the constant or miss the pattern entirely)

3. **solve the recursion for a_n as a ratio of shifted b values, recognise $a_1 \cdots a_n$ as a ratio of two partial products, and build the auxiliary sequence $T_n = \prod b_i/(b_i+2)$ directly from that structure.** *Cue:* $a_n = (b_n+2)/b_{n-1}$ is a ratio of consecutive terms of one underlying sequence. *Circuit:* telescoping via a constructed partial product auxiliary sequence. (why it ranks first: it manufactures the closed form from the shape of the recursion itself rather than guessing it, the identity $1/(a_1 \cdots a_n) = \frac{3}{2}(T_{n-1} - T_n)$ falls out by direct algebra, and it makes transparent exactly where the specific value $b_1 = 1$ enters, namely only in fixing the constant $3/2$)

The ideal solution

For $n \geq 2$, since $b_{n-1} > 0$, solve the recurrence for a_n:

$$a_n = \frac{b_n + 2}{b_{n-1}}.$$

Since $a_1 = 1$, telescoping the product of these ratios gives, for every $n \geq 1$ (reading an empty product as 1),

$$a_1 a_2 \cdots a_n = \prod_{k=2}^{n} \frac{b_k + 2}{b_{k-1}} = \frac{(b_2+2)(b_3+2)\cdots(b_n+2)}{b_1 b_2 \cdots b_{n-1}},$$

so

$$\frac{1}{a_1 \cdots a_n} = \frac{b_1 b_2 \cdots b_{n-1}}{(b_2+2)(b_3+2)\cdots(b_n+2)}.$$

Define $T_0 = 1$ and $T_n = \prod_{i=1}^{n} \frac{b_i}{b_i+2}$ for $n \geq 1$. Then

$$T_{n-1} - T_n = \frac{b_1 \cdots b_{n-1}}{(b_1+2)\cdots(b_{n-1}+2)} \left(1 - \frac{b_n}{b_n+2}\right) = \frac{2\, b_1 \cdots b_{n-1}}{(b_1+2)\cdots(b_{n-1}+2)(b_n+2)}.$$

Since $b_1 = 1$, the factor $b_1 + 2 = 3$ is the only one that fails to cancel against $\frac{1}{a_1 \cdots a_n}$ above, leaving

$$T_{n-1} - T_n = \frac{2}{3} \cdot \frac{b_1 \cdots b_{n-1}}{(b_2+2)\cdots(b_n+2)} = \frac{2}{3} \cdot \frac{1}{a_1 \cdots a_n},$$

that is,

$$\frac{1}{a_1 \cdots a_n} = \frac{3}{2}\big(T_{n-1} - T_n\big) \qquad \text{for every } n \geq 1.$$

(At $n = 1$ both sides are checked directly against the empty product convention: the left side is $1/a_1 = 1$, and the right side is $\frac{3}{2}(T_0 - T_1) = \frac{3}{2}(1 - \frac{1}{3}) = 1$, since $T_1 = b_1/(b_1+2) = 1/3$.)

Summing this telescoping identity from $n = 1$ to m,

$$S_m := \sum_{n=1}^{m} \frac{1}{a_1 \cdots a_n} = \frac{3}{2}(T_0 - T_m) = \frac{3}{2}(1 - T_m).$$

Every factor $b_i/(b_i+2)$ lies strictly between 0 and 1 since $b_i > 0$, so $0 < T_m < 1$ for every m, which already shows (S_m) is increasing (each summand is positive) and bounded above by $3/2$. That alone proves (S_m) converges. To pin the exact limit, use the boundedness hypothesis: say $0 < b_j \leq B$ for every j. The function $t \mapsto t/(t+2)$ is strictly increasing on $(0, \infty)$, since its derivative is $2/(t+2)^2 > 0$, so

$$\frac{b_j}{b_j+2} \leq \frac{B}{B+2} =: r < 1 \qquad \text{for every } j.$$

Hence

$$0 < T_m = \prod_{i=1}^{m} \frac{b_i}{b_i+2} \leq r^m \longrightarrow 0 \qquad \text{as } m \to \infty,$$

and so

$$S = \lim_{m \to \infty} S_m = \frac{3}{2}(1 - 0) = \boxed{\frac{3}{2}}.$$

Layer 1: WHAT

Block 1: solve the recurrence for a_n, exposing it as a ratio of shifted terms of the single sequence b. Block 2: recognise that the product $a_1 \cdots a_n$ telescopes into a ratio of two separate partial products, one running over b_k and one over $b_k + 2$. Block 3: construct the auxiliary sequence T_n directly from that same ratio shape, and identify the reciprocal product as a constant multiple of $T_{n-1} - T_n$, using $b_1 = 1$ only to fix the constant. Block 4: sum the telescoping identity to get S_m in closed form, and read off both monotonicity and boundedness directly from that formula, settling convergence with no separate analytic estimate. Block 5: use the boundedness hypothesis on (b_j) to force $T_m \to 0$ via a uniform ratio bound, and conclude the exact value of S.

The nominal route computes S_1, S_2, S_3 by hand, stares at the growing mess of $b_2, b_3, \ldots$ terms, and tries to guess the pattern in S_m before proving it by induction. The ideal route never computes a single numerical partial sum; it builds the auxiliary sequence T_n straight from the algebraic shape of a_n and verifies the telescoping identity symbolically. The bifurcation point sits right after solving $a_n = (b_n + 2)/b_{n-1}$: the zombie route starts substituting $n = 1, 2, 3$ into the definition of S_m, the ideal route asks what partial product, built from the same raw pieces, has a consecutive difference equal to this term.

Layer 2: HOW

Block 1. **Cue.** The recurrence gives b_n in terms of b_{n-1} and a_n multiplicatively, with b_{n-1} guaranteed nonzero since every b_j is positive. **Why natural.** Isolating a_n is the only way to access the product $a_1 \cdots a_n$ that the sum actually needs, since the sum is stated directly in terms of that product, not in terms of b_n alone. **Link to the next block.** Once a_n is a ratio of shifted b terms, the product $a_1 \cdots a_n$ becomes a product of such ratios, precisely the shape that collapses under telescoping.

Block 2. **Cue.** The product $\prod_{k=2}^{n} (b_k + 2)/b_{k-1}$ is a product of fractions with no shared factors forced together. **Why natural.** Multiplying fractions never entangles numerator and denominator sequences; they regroup independently into one partial product over $b_k + 2$ and a separate partial product over b_{k-1}, which is just b reindexed down by one. **Link to the next block.** Seeing two cleanly separated partial products is what tells you to compare ONE auxiliary sequence built from each piece, rather than trying to simplify the whole fraction as a single object.

Block 3. **Cue.** The two partial products differ only by a shift of $+2$ inside one factor at a time. **Why natural.** A ratio of "the same partial product with one factor shifted" is the canonical telescoping shape, because $T_{n-1} - T_n$ automatically isolates the difference contributed by a single factor, $1 - b_n/(b_n + 2)$, while every other factor cancels unchanged between the two partial products. **Link to the next block.** This identity holds symbolically before you ever plug in $b_1 = 1$, so summing it from here on is purely mechanical, not exploratory.

Block 4. **Cue.** A telescoping identity summed over n collapses to its two endpoint terms. **Why natural.** That collapse is the defining property of telescoping, and it hands you S_m as an explicit two term closed form with no residual sum left over. **Link to the next block.** An explicit closed form for S_m answers both halves of the problem at once: convergence becomes a one line check of monotonicity and boundedness on that formula, not a separate estimate built from scratch.

Block 5. **Cue.** The only hypothesis left unused at this point is that (b_j) is bounded. **Why natural.** Boundedness is exactly what lets a product of many factors, each strictly less than 1, be pushed all the way down to 0, via one uniform ratio bound $r < 1$; without a uniform bound the factors could creep toward 1 too fast for the product to vanish. **Link to the next block.** Once $T_m \to 0$ is

established, the closed form from Block 4 hands you the numerical value of S immediately, with no further work.

Layer 3: WHY

The mathematical truth underneath this problem is that a product of ratios of consecutive or shifted terms of one sequence always telescopes: more generally, any sum whose n-th term can be written as $c(U_{n-1} - U_n)$ for some auxiliary sequence U_n built from the same raw data collapses to $c(U_0 - \lim U_n)$. The discernment mechanism for reaching for this rather than estimation is structural: whenever a sum's general term is a product of ratios built from one recursively defined sequence, actively construct the matching partial product auxiliary sequence before attempting any inequality. You train that discernment by asking, of any product-of-ratios term, what is the simplest partial product built from the same pieces whose consecutive difference reproduces this term up to a constant.

It is also worth naming explicitly why the underspecification of (b_j) is not an obstacle but a guarantee. A Putnam problem that asks you to "evaluate" a quantity built from a sequence constrained only by an inequality, here boundedness, is telling you the answer cannot depend on information you were never given; the only way that can happen is if the quantity is pinned by an identity true for every admissible sequence, not by an estimate tailored to one. Once you internalise that, "evaluate" versus "estimate" or "bound" becomes a reliable signal for which kind of argument the examiners are asking for, before you have written a single line of algebra.

Finally, it is worth checking exactly how much of the given numerical data is load bearing, because that tells you what generalises. Replace the shift -2 in the recurrence by a general $-k$ for $k > 0$, keeping $a_1 = b_1 = 1$, and the identical derivation with $T_n = \prod b_i/(b_i + k)$ gives $1/(a_1 \cdots a_n) = \frac{1+k}{k}(T_{n-1} - T_n)$, so under the same boundedness hypothesis $S = 1 + \frac{1}{k}$; setting $k = 2$ recovers $3/2$. The constant $3/2$ is not a numerological accident of this specific problem, it is $(b_1 + 2)/2$ read off directly from the algebra, and the examiners' choice of shift 2 together with $b_1 = 1$ is simply one convenient instance of a whole family of identical arguments.

Revolving concepts

Telescoping via a constructed auxiliary partial product. When a term is a product of ratios of shifted values of one recursively defined sequence (b_i), build $T_n = \prod_{i \le n} b_i/(b_i + k)$ from the same raw pieces; the term equals a constant multiple of $T_{n-1} - T_n$, and the sum collapses to that constant times $T_0 - \lim T_n$. See Book 1, field: Sequences and recurrences, topic: telescoping sums from recursively defined data. The recursion generalises cleanly: replacing the shift 2 by any $k > 0$ (with $a_1 = b_1 = 1$ fixed) gives $S = 1 + 1/k$ under the same boundedness hypothesis, so the specific numbers in this problem are one instance of a one parameter family of identical telescoping arguments.

Monotone bounded sequences and geometric domination. A sum of positive terms whose partial sums have an explicit closed form $c(1 - T_m)$ with $0 < T_m < 1$ is automatically increasing and bounded above by c, proving convergence without a separate estimate; if in addition every factor of T_m is bounded above by a fixed ratio $r < 1$, then $T_m \le r^m \to 0$ geometrically, pinning the exact limit. See Book 1, field: Real analysis, topic: monotone convergence and geometric bounds on products.

Monotone functions converting a bound on one quantity into a bound on a related one. For fixed $k > 0$, $t \mapsto t/(t + k)$ is strictly increasing on $(0, \infty)$, so a uniform upper bound $b_j \le B$ converts directly into a uniform upper bound $b_j/(b_j + k) \le B/(B + k)$ on the ratio. See Book 1, field: Real analysis, topic: monotone transformations of bounds.

Circuit Bank

- **WHEN** a series' general term is a product of ratios built from one recursively defined sequence **DO** construct the auxiliary partial product sequence from the same raw pieces and look for a constant multiple of its consecutive difference **BECAUSE** such products are exactly the telescoping shape, and the constant is fixed by the boundary values already given in the problem
- **WHEN** a problem asks you to "evaluate" a sum or limit built from an underspecified sequence, constrained only by an inequality such as boundedness **DO** expect an identity valid for every admissible sequence rather than an estimate tailored to one particular choice **BECAUSE** the answer cannot depend on information you were never given, so only a universal identity can pin it down
- **WHEN** you have reduced a partial sum to an explicit two term closed form **DO** read off monotonicity and boundedness directly from that formula for the convergence half, and the limiting behaviour of the remaining term for the value half **BECAUSE** a single closed form answers both halves of a "prove convergence and evaluate" question at once, with no separate analytic estimate needed
- **WHEN** a boundedness hypothesis on a positive sequence is the only unused hypothesis left **DO** convert it into a uniform ratio bound $r < 1$ via a monotone function of that sequence, then bound a product of such factors by r^m **BECAUSE** a single uniform ratio bound is exactly what lets a product of many sub unit factors be squeezed to zero geometrically

Backtesting

1. First useful cue: the instruction "prove S converges, and evaluate S" attached to a sum whose sequence (b_j) is left almost entirely unspecified.
2. Method the expression suggested: solve the recurrence for a_n as a ratio of shifted b terms, exposing $a_1 \cdots a_n$ as a ratio of two partial products built from the same data.
3. Where you would hesitate, and the thought to reinforce: after seeing the product-of-ratios shape, you might reach for estimation, bounding S_m above and below, rather than an exact identity; reinforce that "evaluate" is a guarantee an exact telescoping identity exists, since the answer cannot depend on an unspecified sequence.
4. The trap the question sets: it hands you a boundedness hypothesis that looks built for a comparison or estimation argument, tempting you to skip the exact algebraic identity and reach straight for inequalities.
5. What makes you recognise it faster next time: any sum whose n-th term is a product of ratios built from one recursively defined sequence is a telescoping candidate; build the matching partial product auxiliary sequence before attempting any estimate.
6. The reusable decision rule: when a series' term is a product of ratios of one sequence's shifted values, construct the auxiliary partial product from the same pieces, express the term as a constant multiple of its consecutive difference, sum it in closed form, and only then bring in any remaining hypothesis, such as boundedness, to evaluate the limit of the auxiliary sequence.

Mirror drills

1. Same recurrence but with shift -4 in place of -2 (still $a_1 = b_1 = 1$, (b_j) positive and bounded). Find $S = \sum 1/(a_1 \cdots a_n)$. *Answer:* by the same telescoping with $k = 4$, $S = 1 + \frac{1}{4} = \frac{5}{4}$.

2. Suppose (b_j) were only known to be bounded in absolute value, not necessarily positive, with a_n still defined by the same recurrence and assumed nonzero. Does the proof that $T_m \to 0$ still go through? *Answer:* no; $t/(t+2)$ need not lie in $(0,1)$, or can even be undefined near $t=-2$, once t is allowed negative, so the uniform bound $r<1$ collapses entirely, showing positivity of every b_j is load bearing, not a convenience.

3. If (b_j) is bounded above by $B=4$ specifically, give an explicit rate at which S_m approaches $\frac{3}{2}$. *Answer:* $r = B/(B+2) = 2/3$, so $S - S_m = \frac{3}{2}T_m \le \frac{3}{2}(2/3)^m$, a geometric rate.

15.3 A3. the endpoint order forced by a vanishing cosine

Find a real number c and a positive number L for which

$$\lim_{r\to\infty} \frac{r^c \displaystyle\int_0^{\pi/2} x^r \sin x\, dx}{\displaystyle\int_0^{\pi/2} x^r \cos x\, dx} = L.$$

Read it out loud

Read the limit aloud: a ratio of two integrals that look almost identical, same power weight x^r, same interval $[0, \pi/2]$, one dressed with $\sin x$ and one with $\cos x$, multiplied by a power of r whose exponent you must supply, with the goal of landing on a finite positive constant. The command "find a real number c" rather than "evaluate the limit" for a given c is itself informative: it tells you upfront that the raw ratio, with no correction at all, does not converge to a finite nonzero number. It must diverge, to 0 or to infinity, at some definite polynomial rate in r, and your entire job is to name that rate and the constant that survives it.

The field wired in is calculus and integrals, specifically the asymptotics of a family of integrals as a parameter $r \to \infty$, riding on a real analysis squeeze argument. Both integrands share the factor x^r, an increasingly sharply peaked weight concentrating near the right endpoint $x = \pi/2$ as r grows, since for any fixed $x_0 < \pi/2$ the ratio $(x_0/(\pi/2))^r$ decays geometrically, making the contribution to either integral from $x \le x_0$ negligible in comparison. The first useful cue, readable within thirty seconds, is a single pair of values: $\cos(\pi/2) = 0$ while $\sin(\pi/2) = 1$. Because the weight x^r concentrates exactly where cosine vanishes and sine does not, the cosine integral must be systematically smaller than the sine integral, by a factor that grows with r. This single fact, before any calculation, already tells you c must be negative to counteract a ratio that blows up, and it tells you $\cos(\pi/2) = 0$ is not incidental, it is the entire reason a nontrivial exponent is needed at all.

Naming that boundary value also suggests which direction integration by parts should run. If you integrate the cosine integral by parts differentiating $\cos x$, the factor that vanishes at $\pi/2$, rather than integrating it, whatever boundary term appears dies automatically instead of needing to be estimated. That observation converts the cosine integral into a sine integral one step higher in r, divided by $r + 1$, an exact identity rather than an approximation, and it means you only ever need one asymptotic estimate, of the sine integral, applied at two consecutive values of r.

Approaches, ranked worst to best

1. **repeated integration by parts to a closed finite form, then extract asymptotics termwise.** *Cue:* a definite integral of x^r against a trigonometric factor invites reducing the power of x step by step. *Circuit:* no circuit, brute reduction. (why it ranks here: this only terminates cleanly for integer r, produces an r term alternating sum of boundary values, and still forces you to argue separately that the boundary term at $\pi/2$ dominates the rest, doing strictly more work to reach the same leading order that a direct estimate gives at once)

2. **argue heuristically that x^r concentrates at $\pi/2$, substitute $x = \pi/2 - t/r$, and approximate $\sin x \approx 1$, $\cos x \approx 2t/\pi$ near the endpoint.** *Cue:* x^r is an increasingly sharp bump at the right endpoint as $r \to \infty$. *Circuit:* Laplace's method intuition, unproven. (why it ranks here: this correctly identifies why the answer is what it is, and is exactly the right picture to carry in your head, but turning it into a fully rigorous bound needs more machinery, controlled error terms on the substitution, than a direct elementary inequality supplies, so a full write up this way costs more ink for no extra rigor)

3. **integrate the cosine integral by parts in the direction that kills the boundary term exactly, reducing it to a shifted sine integral, then pin the sine integral's exact leading asymptotic with a two sided elementary inequality.** *Cue:* $\cos(\pi/2) = 0$ paired with an interval ending exactly at $\pi/2$. *Circuit:* exact self reduction via integration by parts, plus Jordan's inequality squeeze. (why it ranks first: it is fully rigorous with a single elementary inequality, it needs only one asymptotic estimate rather than two, and it exploits the vanishing boundary value exactly rather than approximately)

The ideal solution

Write $f(r) = \displaystyle\int_0^{\pi/2} x^r \sin x\,dx$ and $g(r) = \displaystyle\int_0^{\pi/2} x^r \cos x\,dx$.

Step 1: an exact reduction of g to f. Integrate $g(r)$ by parts with $u = \cos x$, $dv = x^r\,dx$, so $du = -\sin x\,dx$ and $v = x^{r+1}/(r+1)$:

$$g(r) = \left[\cos x \cdot \frac{x^{r+1}}{r+1}\right]_0^{\pi/2} + \int_0^{\pi/2} \frac{x^{r+1}}{r+1} \sin x\,dx.$$

At $x = \pi/2$ the bracket vanishes because $\cos(\pi/2) = 0$; at $x = 0$ it vanishes because $x^{r+1} \to 0$ (for $r > -1$). So the boundary term is exactly 0, and

$$g(r) = \frac{f(r+1)}{r+1}.$$

Step 2: the exact leading asymptotic of f. On $[0, \pi/2]$, $\sin x < 1$, so

$$f(r) < \int_0^{\pi/2} x^r\,dx = \frac{(\pi/2)^{r+1}}{r+1}.$$

Since sin is concave on $[0, \pi/2]$, its graph lies above the chord from $(0,0)$ to $(\pi/2, 1)$, which is Jordan's inequality $\sin x \geq \frac{2x}{\pi}$, so

$$f(r) > \int_0^{\pi/2} \frac{2x^{r+1}}{\pi}\,dx = \frac{2}{\pi} \cdot \frac{(\pi/2)^{r+2}}{r+2} = \frac{(\pi/2)^{r+1}}{r+2}.$$

Multiplying the two bounds by $r(2/\pi)^{r+1} > 0$,

$$\frac{r}{r+2} < r\left(\frac{2}{\pi}\right)^{r+1} f(r) < \frac{r}{r+1},$$

and both outer expressions tend to 1 as $r \to \infty$, so by the squeeze theorem,

$$\lim_{r\to\infty} r\left(\frac{2}{\pi}\right)^{r+1} f(r) = 1. \qquad (\star)$$

Step 3: the ratio of consecutive values. Write $u_r = r(2/\pi)^{r+1} f(r)$, so $u_r \to 1$ by $(\star)$, and $f(r) = u_r \cdot \frac{1}{r}(\pi/2)^{r+1}$. Then

$$\frac{f(r)}{f(r+1)} = \frac{u_r}{u_{r+1}} \cdot \frac{r+1}{r} \cdot \frac{(\pi/2)^{r+1}}{(\pi/2)^{r+2}} = \frac{u_r}{u_{r+1}} \cdot \frac{r+1}{r} \cdot \frac{2}{\pi} \longrightarrow 1 \cdot 1 \cdot \frac{2}{\pi} = \frac{2}{\pi} \qquad (r \to \infty).$$

Step 4: assembling the answer. Using Step 1,

$$\frac{r^c f(r)}{g(r)} = \frac{r^c f(r)(r+1)}{f(r+1)} = r^c(r+1) \cdot \frac{f(r)}{f(r+1)}.$$

By Step 3 the factor $\frac{f(r)}{f(r+1)}$ tends to $\frac{2}{\pi}$, and $(r+1) \sim r$, so the whole expression is asymptotically $\frac{2}{\pi} r^{c+1}$. This tends to a finite nonzero limit exactly when $c+1=0$: if $c > -1$ the expression diverges to infinity, and if $c < -1$ it tends to 0. So $c = -1$ is forced, and with $c = -1$,

$$\lim_{r\to\infty} \frac{r+1}{r} \cdot \frac{f(r)}{f(r+1)} = 1 \cdot \frac{2}{\pi} = \frac{2}{\pi}.$$

Hence

$$\boxed{(c, L) = \left(-1, \frac{2}{\pi}\right)},$$

and this is the only such pair, since no other value of c produces a finite positive limit.

Layer 1: WHAT

Block 1: reduce the cosine integral exactly to a sine integral one step higher in r, using integration by parts in the direction that kills the boundary term because $\cos(\pi/2) = 0$. Block 2: bound the sine integral above and below by the same power integral, using $\sin x < 1$ and Jordan's inequality, giving matching leading order. Block 3: squeeze the two bounds to pin the exact constant in front of $(\pi/2)^{r+1}/r$, not merely its order. Block 4: reuse that single estimate at two consecutive parameters, r and $r+1$, to get the ratio $f(r)/f(r+1) \to 2/\pi$. Block 5: combine the exact reduction and the ratio limit, and match polynomial orders in r to force $c = -1$ and read off L.

The nominal route tries to compute $f(r)$ and $g(r)$ in closed form by repeated integration by parts, or else waves at "the mass concentrates near $\pi/2$" without ever pinning a constant. The ideal route reduces g to f exactly in one step, then estimates f exactly once with an elementary two sided inequality, and reuses that one estimate twice. The bifurcation point sits right after noticing $\cos(\pi/2) = 0$: the zombie route starts grinding a closed form for $g(r)$ term by term, the ideal route asks which integration by parts direction makes the boundary term vanish for free and stops after a single step.

Layer 2: HOW

Block 1. **Cue.** $\cos(\pi/2) = 0$ sits exactly at the endpoint where x^r concentrates as $r \to \infty$. **Why natural.** Integration by parts has two directions; choosing to differentiate the factor that already vanishes at the endpoint of interest, rather than integrate it, forces the resulting boundary term to be exactly 0 rather than merely small, turning an approximate comparison into an exact algebraic identity. **Link to the next block.** Once $g(r) = f(r+1)/(r+1)$ is exact, the entire problem reduces to understanding one function, f, at two nearby arguments, so no second, structurally different estimate is ever needed.

Block 2. **Cue.** $\sin x$ is trivially below 1 and, being concave on $[0, \pi/2]$, lies above the straight chord joining its endpoint values. **Why natural.** A concave function always lies above any chord connecting two of its points, and the chord from $(0,0)$ to $(\pi/2, 1)$ is the specific linear function $2x/\pi$, giving a lower bound of exactly the same shape, a constant times x, as the trivial upper bound of 1. **Link to the next block.** Two bounds of matching shape, both constants times the pure power x^r or x^{r+1}, integrate to two expressions that differ only in their constant, setting up a squeeze rather than a one sided estimate.

Block 3. **Cue.** Both bounds, once integrated and multiplied by $r(2/\pi)^{r+1}$, become rational functions of r that visibly tend to the same limit, 1. **Why natural.** The squeeze theorem applies whenever two bounds sandwiching a quantity converge to a common limit; here that common limit is forced by $r/(r+1) \to 1$ and $r/(r+2) \to 1$, both elementary. **Link to the next block.** Pinning the

exact constant, not just the exact power of r, is what makes the final answer L computable rather than merely "some positive number."

Block 4. Cue. The same asymptotic statement $(\star)$ is true for every large r, in particular for r and for $r+1$ simultaneously. **Why natural.** Dividing the statement at r by the statement at $r+1$ cancels the unknown behaviour $u_r, u_{r+1} \to 1$ in the limit and leaves only the explicit, elementary factors $r/(r+1)$ and a fixed power of $2/\pi$, so the ratio of consecutive values of f is extracted with no new estimation. **Link to the next block.** This ratio is exactly the quantity the exact reduction from Block 1 needs, since $g(r)$ was expressed in terms of $f(r+1)$, not $f(r)$ directly.

Block 5. Cue. The assembled expression $r^c(r+1)f(r)/f(r+1)$ is asymptotically a nonzero constant times r^{c+1}. **Why natural.** A power of r has a finite nonzero limit as $r \to \infty$ if and only if that power is exactly r^0, which pins the exponent $c+1=0$ uniquely; any other value of c sends the expression to 0 or to infinity. **Link to the next block.** This is the last step precisely because it needs both earlier results at once, the exact reduction for the extra factor of $r+1$, and the ratio limit for the constant $2/\pi$, so it could not have been done first.

Layer 3: WHY

The mathematical truth underneath this problem is a special case of endpoint concentration asymptotics: for a continuous weight w on $[a,b]$ with $w(b) > 0$, $\int_a^b x^r w(x)\,dx \sim \dfrac{b^{r+1}w(b)}{r}$ as $r \to \infty$, because the mass of x^r concentrates geometrically fast at the right endpoint, so only the value of w right at b, not its behaviour elsewhere on $[a,b]$, controls the leading order. When $w(b)=0$, as for cosine at $\pi/2$, this leading order estimate gives 0, telling you the true order is one power of r smaller, exactly matching the extra factor of $1/(r+1)$ that fell out of the exact integration by parts identity. The elementary squeeze used here, sandwiching $\sin x$ between 1 and $2x/\pi$, is nothing but a fully rigorous, self contained substitute for that general endpoint lemma in this one case, avoiding any need to import or prove the general statement.

The discernment mechanism for reaching for this rather than brute computation is to check, before any calculus, the value of each weight function exactly at the endpoint where the dominant power concentrates. Whenever one weight is nonzero there and another vanishes there, expect the two integrals to differ by a definite power of r, and expect that power to be extractable by comparing orders alone, with the constant supplied afterward by whichever elementary bound is tightest at the endpoint. You train this discernment by refusing to compute a single integral numerically until you have named both endpoint values and predicted, in words, which integral should be smaller and by roughly what order.

The examiners nudged you here by choosing sine and cosine on precisely $[0, \pi/2]$, an interval whose right endpoint makes one function equal to 1 and the other equal to 0, manufacturing exactly a one order of r gap between two otherwise structurally identical integrals. The phrasing "find c and L", rather than stating c outright, forces you to notice that gap and name it yourself, rather than being handed the answer's shape in advance.

Revolving concepts

Endpoint concentration asymptotics for power weighted integrals. For w continuous on $[a,b]$ with $w(b) > 0$, $\int_a^b x^r w(x)\,dx \sim b^{r+1}w(b)/r$ as $r \to \infty$; the weight's value elsewhere on the interval is asymptotically irrelevant. See Book 1, field: Calculus and integrals, topic: Laplace type endpoint asymptotics. This generalises as a two term recursion in r: consecutive members of such a family satisfy $f(r)/f(r+1) \to 1/b$, and a weight vanishing at b drops the order of the integral by exactly

one full power of r relative to a weight that does not vanish there, which is the entire mechanism behind this problem's answer.

Jordan's inequality. $\sin x \geq 2x/\pi$ for $x \in [0, \pi/2]$, an immediate consequence of the concavity of sin on that interval, since a concave function lies above every chord joining two of its points, here the chord from $(0, 0)$ to $(\pi/2, 1)$. See Book 1, field: Inequalities, topic: concavity based bounds on trigonometric functions.

Integration by parts direction chosen to annihilate a boundary term. When one factor of a product vanishes at an endpoint of interest, differentiate that factor rather than integrate it, so the resulting boundary term is exactly zero, converting a comparison that would otherwise need an error estimate into an exact identity. See Book 1, field: Calculus and integrals, topic: strategic choice of u and dv in integration by parts.

Circuit Bank

- **WHEN** two integrals $\int_a^b x^r w_1(x)\,dx$ and $\int_a^b x^r w_2(x)\,dx$ share the same power weight concentrating as $r \to \infty$ **DO** compare the two weights' values at the concentration endpoint b before doing any calculation **BECAUSE** whichever weight vanishes there produces a systematically smaller integral, off by one full power of r, and that comparison is free, needing no integration at all
- **WHEN** integrating a product by parts, and one factor vanishes at the endpoint of interest **DO** differentiate that vanishing factor rather than integrate it, so the boundary term dies automatically **BECAUSE** this turns an approximate comparison into an exact identity with no error term left to control
- **WHEN** you need the leading asymptotic order of $\int_a^b x^r w(x)\,dx$ and w is continuous and positive at b **DO** squeeze between two matching bounds, such as $w(b) - \epsilon$ and $w(b) + \epsilon$, or a specific two sided inequality like Jordan's, each integrated against the same pure power **BECAUSE** the mass concentrates at b geometrically fast, so only the value of w right at b controls the leading order, and a matched two sided bound pins the exact constant, not merely the order
- **WHEN** a limit is phrased as "find c and L" rather than "evaluate this limit" **DO** expect the raw, uncorrected expression to provably diverge to 0 or infinity at a definite polynomial rate, and treat your job as naming that rate and the surviving constant **BECAUSE** the phrasing itself is telling you the uncorrected limit does not exist finitely, so guessing $c = 0$ wastes time that direct order matching would save

Backtesting

1. First useful cue: $\cos(\pi/2) = 0$ while $\sin(\pi/2) = 1$, sitting exactly at the endpoint where x^r concentrates as $r \to \infty$.
2. Method the expression suggested: integrate the cosine integral by parts differentiating $\cos x$, the factor that vanishes at $\pi/2$, so the boundary term dies exactly, reducing it to a shifted sine integral.
3. Where you would hesitate, and the thought to reinforce: after obtaining the exact reduction $g(r) = f(r+1)/(r+1)$, it is easy to think the hard work is done; reinforce that an exact identity between g and f at a shifted argument is not itself a number, you still need to know how f scales with r to extract a ratio.
4. The trap the question sets: it asks you to "find" c and L rather than stating c, tempting a candidate to guess a round number like $c = 0$ or $c = 1$ by pattern matching rather than deriving the rate from the vanishing of cosine at the endpoint.

5. What makes you recognise it faster next time: whenever a limit compares two integrals of x^r against different weights as $r \to \infty$, check the weights' values exactly at the endpoint where x^r concentrates before doing any calculus; a vanishing weight there is a guaranteed source of one extra power of r in the final answer.
6. The reusable decision rule: reduce a comparison between two differently weighted power integrals to a comparison of the SAME weighted integral at two nearby parameters whenever an exact identity, typically via integration by parts choosing the boundary killing direction, makes that possible, so only one asymptotic estimate is ever needed, applied twice.

Mirror drills

1. Find real c and positive L with $\lim_{r\to\infty} \dfrac{r^c \int_0^{\pi/2} x^r\,dx}{\int_0^{\pi/2} x^r \cos x\,dx} = L$. *Answer:* the numerator is exactly $(\pi/2)^{r+1}/(r+1)$, of the same nonvanishing order as $f(r)$ since the constant weight 1 agrees with $\sin(\pi/2) = 1$ at the endpoint; the same reduction for $g(r)$ applies unchanged, giving $(c, L) = (-1, 2/\pi)$ again.
2. Swap the roles: find real c and positive L with $\lim_{r\to\infty} \dfrac{r^c \int_0^{\pi/2} x^r \cos x\,dx}{\int_0^{\pi/2} x^r \sin x\,dx} = L$. *Answer:* this is the reciprocal of the original ratio, so $c = 1$ (not -1) and $L = \pi/2$, the reciprocal of $2/\pi$.
3. Without fully computing it, give the leading asymptotic order of $\int_0^{\pi/2} x^r \sin^2 x\,dx$ as $r \to \infty$. *Answer:* since $\sin^2(\pi/2) = 1 \neq 0$, the endpoint concentration lemma gives order $(\pi/2)^{r+1}/r$ with constant 1, the same leading behaviour as $f(r)$ itself.

15.4 A4. Prescribed parities of dot products

For which positive integers n is there an $n \times n$ matrix with integer entries such that every dot product of a row with itself is even, while every dot product of two different rows is odd?

Read it out loud

Say the structure out loud before you write anything. You are given n vectors and every one of the n^2 pairwise dot products is prescribed, but only modulo 2. The entries are integers, yet no size information is used anywhere in the hypotheses. The only thing about the matrix that the question can see is the reduction of each entry modulo 2. So the honest restatement is: for which n is there a matrix $\overline{M}$ over the two element field $\mathbb{F}_2$ whose Gram matrix $\overline{M}\,\overline{M}^T$ has zero diagonal and ones off the diagonal?

The fields wired into the problem are linear algebra over a finite field, Gram matrices and rank, and the behaviour of squaring in characteristic 2. What is really being tested is whether you convert a parity statement into a statement over $\mathbb{F}_2$ immediately, instead of dragging integers around, and whether you package n^2 separate conditions into one matrix equation instead of manipulating them one at a time.

Within thirty seconds you should have two lines on the page. Line one: $\overline{M}\,\overline{M}^T = J - I$ over $\mathbb{F}_2$, where J is the all ones matrix and I the identity. Line two: over $\mathbb{F}_2$ we have $x^2 = x$, so the dot product of a row with itself equals the sum of the entries of that row; hence the diagonal condition says exactly $\overline{M}v = 0$, where v is the all ones column vector. Two conditions, both linear algebraic, one about the Gram matrix and one about the kernel. The whole problem is whether those two can coexist.

Approaches, ranked worst to best

1. **Hand search for examples.** *Cue:* the phrase "for which positive integers n" invites experiment. *Circuit:* small case testing. You build examples for $n = 1, 2, 3, 4$ and guess that odd n works. That is worth sixty seconds as reconnaissance, but it can never close the problem: for even n you would have to exhaust an infinite family of integer matrices. It ranks last because it produces a conjecture and no obstruction, and on the Putnam an answer with no obstruction scores near zero.

2. **Global parity count over $\mathbb{Z}$.** *Cue:* sums of dot products look like an expanded square. *Circuit:* expand $\|\sum_i r_i\|^2$. Write $s = r_1 + \cdots + r_n$ and expand: $\|s\|^2 = \sum_i \|r_i\|^2 + \sum_{i \neq j} r_i \cdot r_j \equiv 0 + n(n-1) \pmod 2$. For n even this gives $\|s\|^2$ even, which is true and useless. It ranks above pure search because it is a real attempt at an obstruction, but it stalls for a nameable reason: it extracts one scalar from an $n \times n$ table of information, and one scalar cannot see the difference between odd and even n.

3. **Linear dependence of the rows over $\mathbb{F}_2$.** *Cue:* a singular matrix has a dependency among its rows. *Circuit:* turn singularity into a nonempty index set S with $\sum_{i \in S} r_i = 0$. From $\overline{M}v = 0$ the matrix is singular, so some nonempty S has $\sum_{i \in S} r_i = 0$. Dotting with r_j for $j \in S$ gives $|S| - 1 \equiv 0$, so $|S|$ is odd; dotting with r_j for $j \notin S$ gives $|S| \equiv 0$, so no such j exists, so S is everything and $n = |S|$ is odd. This is complete and elementary. It ranks here only because it splits into two cases that the next two routes avoid.

4. **Gram determinant over $\mathbb{F}_2$.** *Cue:* all pairwise dot products means Gram matrix, and nonexistence means determinant. *Circuit:* $\det(\overline{M})^2 = \det(J - I)$. Over $\mathbb{Z}$, $\det(J - I) = (-1)^{n-1}(n-1)$, which is odd when n is even, so $\det \overline{M} \neq 0$ and $\overline{M}$ is invertible, contradicting $\overline{M}v = 0$. Short and correct. It ranks below the winner only because it imports the spectrum of J and a determinant computation that you do not need.

5. **Gram rank over $\mathbb{F}_2$, via the involution.** *Cue:* $J - I$ squares to something recognisable. *Circuit:* $\operatorname{rank}(AB) \le \operatorname{rank}(A)$. Over $\mathbb{F}_2$, $(J-I)^2 = J^2 - 2J + I = nJ + I$. For n even this is I, so $J - I$ is its own inverse, hence of rank n, hence $\overline{M}$ has rank n and trivial kernel, contradicting $\overline{M}v = 0$. For n odd the same computation gives $(J-I)^2 = J - I$, an idempotent of rank $n-1$ with kernel spanned by v, which is not a contradiction but a blueprint for the construction. Best route: three lines, no determinants, and it hands you the example.

The ideal solution

The answer is $\boxed{n \text{ odd}}$.

Throughout, J is the $n \times n$ all ones matrix, I the identity, and v the all ones column vector. Write $r_1, \dots, r_n$ for the rows.

Odd n works. Take $M = J - I$ over $\mathbb{Z}$, that is, zeros on the diagonal and ones elsewhere. The dot product of a row with itself counts the $n-1$ entries equal to 1, giving $n-1$, which is even. For $i \neq j$, rows i and j have a 1 in a common position exactly at the $n-2$ positions different from both i and j, giving $n-2$, which is odd. Both conditions hold precisely because n is odd.

Even n fails. Suppose n is even and M is an integer matrix with the stated properties. Every hypothesis is a congruence modulo 2, and reduction modulo 2 is a ring homomorphism, so we may pass to $\overline{M} \in M_n(\mathbb{F}_2)$ and work there.

First, the hypotheses in one equation. The (i,j) entry of $\overline{M}\,\overline{M}^T$ is $r_i \cdot r_j$ read in $\mathbb{F}_2$, which is 0 when $i = j$ and 1 when $i \neq j$. Thus

$$\overline{M}\,\overline{M}^T = J - I \qquad \text{over } \mathbb{F}_2.$$

Second, the diagonal condition again, read differently. In $\mathbb{F}_2$ every element satisfies $x^2 = x$, so

$$r_i \cdot r_i = \sum_{t=1}^{n} \overline{M}_{it}^2 = \sum_{t=1}^{n} \overline{M}_{it},$$

the sum of the entries of row i. The hypothesis that this is even says that every row of $\overline{M}$ sums to 0, that is,

$$\overline{M}v = 0, \qquad v = (1, 1, \dots, 1)^T \neq 0.$$

So $\overline{M}$ has nontrivial kernel and $\operatorname{rank}(\overline{M}) \le n-1$.

Third, the Gram matrix is invertible. Since $J^2 = nJ$ and $2 = 0$ in $\mathbb{F}_2$,

$$(J-I)^2 = J^2 - 2J + I = nJ + I = I,$$

the last step because n is even. Hence $J - I$ is its own inverse and $\operatorname{rank}(J-I) = n$. But

$$n = \operatorname{rank}(J-I) = \operatorname{rank}\big(\overline{M}\,\overline{M}^T\big) \le \operatorname{rank}(\overline{M}) \le n-1,$$

a contradiction. Therefore no such M exists for even n, and the answer is exactly the odd n. ∎

Layer 1: WHAT

Block 1. Notice that no hypothesis sees the size of an entry, only its parity, and move the whole problem into $M_n(\mathbb{F}_2)$.

Block 2. Package all n^2 hypotheses as the single matrix equation $\overline{M}\,\overline{M}^T = J - I$.

Block 3. Read the diagonal hypothesis a second time using $x^2 = x$ in $\mathbb{F}_2$: it also says $\overline{M}v = 0$, so $\overline{M}$ is singular.

Block 4. Compute $(J-I)^2 = nJ + I$ in $\mathbb{F}_2$ and split on the parity of n: for even n the Gram matrix is invertible, which is incompatible with Block 3.

Block 5. For odd n the same computation makes $J - I$ idempotent with kernel exactly the span of v, which is what Block 3 demands, so the idempotent itself is the construction.

Juxtaposition. The nominal route is: build examples by hand, notice that $n = 2$ and $n = 4$ resist, then try to prove impossibility by expanding sums of dot products over $\mathbb{Z}$, get $n(n-1) \equiv 0$, and grind. The bifurcation point is Block 1 into Block 2: the instant you decide that the object of study is not the integer matrix but its image in $M_n(\mathbb{F}_2)$, and that the object to write down is the Gram matrix rather than the individual dot products. Everything after that bifurcation is four lines; everything before it is unbounded.

Layer 2: HOW

Block 1. **Cue.** The words "even" and "odd" attached to every hypothesis, and the total absence of any inequality or magnitude condition. **Why natural.** Parity is the ring homomorphism $\mathbb{Z} \to \mathbb{F}_2$. If every hypothesis and the desired conclusion are statements about images under a homomorphism, then working upstairs carries information you will never use and cannot exploit. Downstairs you gain something concrete: $\mathbb{F}_2$ is a field, so the entire apparatus of linear algebra (rank, kernel, invertibility) is available and behaves exactly as over $\mathbb{R}$. **Link to the next block.** Once you are in a vector space, you should be looking for a linear algebraic object that carries all the hypotheses at once.

Block 2. **Cue.** "Every dot product of a row with itself" and "every dot product of two different rows" is a complete table of pairwise inner products. **Why natural.** A complete table of pairwise dot products of the rows of M is by definition the matrix MM^T. This is not a manoeuvre, it is a translation: the table and the product are the same object. Writing $\overline{M}\,\overline{M}^T = J - I$ replaces n^2 scalar conditions by one equation, and equations between matrices can be attacked by rank, determinant, trace and by squaring. **Link to the next block.** You now have one equation and one unknown matrix. To get a contradiction you need a second, independent consequence of the hypotheses, ideally one that is linear rather than quadratic in $\overline{M}$.

Block 3. **Cue.** A self dot product in characteristic 2. **Why natural.** In $\mathbb{F}_2$ the Frobenius map $x \mapsto x^2$ is the identity, so the quadratic form $r \mapsto r \cdot r$ collapses to the linear form $r \mapsto \sum_t r_t$. A condition you thought was quadratic is secretly linear, and a linear condition on all rows at once is a kernel statement. This is the single most important structural fact in the problem, and it is a fact about the field, not about the matrix. **Link to the next block.** You have singularity from the diagonal and a Gram identity from the off diagonal. The only remaining question is whether the right hand side $J - I$ is singular too.

Block 4. **Cue.** J satisfies $J^2 = nJ$, so any polynomial in J collapses; and you want the rank of $J - I$. **Why natural.** Do not reach for a determinant expansion. J generates a two dimensional algebra spanned by J and I, so squaring $J - I$ must return to that algebra, and the computation is one line. The parity of n enters through the coefficient n of J, which is exactly where the answer to the problem lives. **Link to the next block.** For n even you have your contradiction and the impossibility half is finished. For n odd the obstruction evaporates, which is the signal to stop trying to prove impossibility and start constructing.

Block 5. **Cue.** For odd n the necessary conditions say: find $\overline{M}$ with $\overline{M}\,\overline{M}^T = J - I$ and $v \in \ker \overline{M}$, where $J - I$ is idempotent of rank $n - 1$ with kernel spanned by v. **Why natural.** An idempotent symmetric matrix E satisfies $EE^T = E$ automatically. So the shortest candidate for $\overline{M}$ is $J - I$ itself, and its kernel is already the span of v. The construction is not a lucky guess; it is the unique obvious matrix satisfying both constraints, and lifting it back to $\mathbb{Z}$ costs nothing because $J - I$ has integer

entries. **Link to the next block.** Verify over $\mathbb{Z}$ directly (dot products $n-1$ and $n-2$) so the write up never has to justify a lift.

Layer 3: WHY

The mathematical truth beneath this route is that a prescribed table of pairwise inner products is a Gram matrix, and the only invariant of a matrix that survives the passage $M \mapsto MM^T$ in general is rank. Over $\mathbb{R}$ you also get positive semidefiniteness, which is why real Gram problems are usually solved by eigenvalue positivity; over $\mathbb{F}_2$ there is no order, so positivity is gone and rank is all you have. Recognising which invariants survive in which field is the discernment being tested.

The second truth is the collapse of squaring in characteristic 2. Over any field of characteristic 2 the map $x \mapsto x^2$ is a ring homomorphism, and over the prime field $\mathbb{F}_2$ it is the identity. Consequently the quadratic form $q(r) = r \cdot r$ is a linear functional there, and the symmetric bilinear form $r \cdot s$ is alternating precisely when that functional vanishes. In this problem the diagonal hypothesis says the form restricted to the row space is alternating, and alternating forms have even rank; that is the same obstruction in different clothing, and it is why even n is the doomed case.

The discernment mechanism, stated as a decision, is this. When every hypothesis is a congruence modulo a prime p, reduce and use linear algebra over $\mathbb{F}_p$; when the hypotheses are about all pairwise dot products of a list of vectors, write the Gram matrix; when you must prove that no object exists, look for two consequences of the hypotheses that force incompatible ranks. Nonexistence proofs on the Putnam almost always come from an invariant, and rank is the cheapest invariant in linear algebra.

How to train it: take any problem phrased with the words even and odd and rewrite it, before doing anything else, as a sentence over $\mathbb{F}_2$. Then ask which standard $\mathbb{F}_2$ object is being described: a Gram matrix, a parity check matrix of a binary code, an incidence matrix of a set system. All three occur constantly, and the family of problems around this one (Oddtown, Eventown, Fisher type inequalities) is the same circuit fired at different targets.

What the examiners nudged. They wrote "integer entries" rather than "entries in $\{0,1\}$" to see whether you would notice that only the parity matters. They chose the diagonal to be even and the off diagonal odd, rather than the other way round, because the reverse condition gives $\overline{M}\,\overline{M}^T = I$, which is possible for every n (take $M = I$) and would have made the problem empty. The asymmetry between the two parities is the content.

Revolving concepts

The field $\mathbb{F}_p$. Definition: the integers modulo a prime p form a field, so every theorem of linear algebra proved over an abstract field applies verbatim: rank, kernel, rank plus nullity, invertibility equivalent to nonzero determinant, and the fact that a square matrix with a right inverse is invertible. What does not transfer is anything relying on order or positivity. See Book 1, field: Linear algebra, topic: linear algebra over finite fields.

Frobenius in characteristic p. The map $x \mapsto x^p$ is a ring homomorphism on any commutative ring of characteristic p, so $(a+b)^p = a^p + b^p$. On $\mathbb{F}_p$ itself Fermat gives $x^p = x$, so for $p = 2$ the squaring map is the identity and quadratic conditions degenerate to linear ones. See Book 1, field: Abstract algebra, topic: Frobenius endomorphism.

Gram matrices and rank. For any M over any field, MM^T records all pairwise dot products of the rows, and $\operatorname{rank}(AB) \le \min\{\operatorname{rank}(A), \operatorname{rank}(B)\}$, so $\operatorname{rank}(MM^T) \le \operatorname{rank}(M)$. Over $\mathbb{R}$ equality always holds; over other fields it can fail, and the failure is exactly what problems like this one exploit. See Book 1, field: Matrices and determinants, topic: Gram matrix.

The all ones matrix. J has $J^2 = nJ$, so $\{aJ + bI\}$ is closed under multiplication with $(aJ + bI)(cJ + dI) = (acn + ad + bc)J + bdI$. Over $\mathbb{R}$ its eigenvalues are n once and 0 with multiplicity $n-1$, so $\det(aJ + bI) = b^{n-1}(an + b)$ and in particular $\det(J - I) = (-1)^{n-1}(n-1)$. The trivial

recursion here is $J^m = n^{m-1}J$ for $m \geq 1$, which generalises to any rank one matrix uw^T with $w^Tu = c$, giving $(uw^T)^m = c^{m-1}uw^T$, and then to the rank one update formula $\det(A + uw^T) = \det(A)(1 + w^TA^{-1}u)$. See Book 1, field: Matrices and determinants, topic: rank one matrices.

Set systems with prescribed intersection parities. Characteristic vectors turn $|A_i \cap A_j|$ into a dot product over $\mathbb{F}_2$. Oddtown: if all $|A_i|$ are odd and all $|A_i \cap A_j|$ $(i \neq j)$ are even, then the Gram matrix is I, the vectors are independent, and there are at most n sets. Eventown is the complementary statement with bound $2^{\lfloor n/2 \rfloor}$. This problem is the third member of that family: even diagonal, odd off diagonal. See Book 1, field: Combinatorics, topic: linear algebra method in combinatorics.

Circuit Bank

- **WHEN** the words even and odd attached to algebraic quantities **DO** rewrite the entire problem over $\mathbb{F}_2$ on line one **BECAUSE** parity is a ring homomorphism, so every hypothesis survives the reduction while all unusable magnitude information is discarded, and $\mathbb{F}_2$ is a field where linear algebra is available
- **WHEN** a complete table of pairwise dot products of a list of vectors **DO** write the Gram matrix MM^T **BECAUSE** one matrix equation carries all n^2 hypotheses, and matrix equations expose rank, determinant and trace
- **WHEN** a self dot product in characteristic 2 **DO** replace it by the sum of the coordinates **BECAUSE** $x^2 = x$ in $\mathbb{F}_2$, so a quadratic condition on a vector is a linear one, and linear conditions on every row are a kernel statement
- **WHEN** you must prove that no object exists **DO** hunt for two consequences of the hypotheses that force incompatible ranks **BECAUSE** rank is the cheapest invariant that both sides of an identity are forced to reveal, and nonexistence needs an invariant, not a search
- **WHEN** a construction is needed and the constraints are invariant under permuting rows **DO** try the family $aJ + bI$ first **BECAUSE** if the constraint set is symmetric, a symmetric solution usually exists, and that family is two dimensional and closed under multiplication

Backtesting

1. First useful cue: every hypothesis is a parity statement about a dot product, so the ambient object is a matrix over $\mathbb{F}_2$ and the ambient equation is a Gram identity.
2. Method the expression suggested: package the table of dot products as $\overline{M}\,\overline{M}^T = J - I$, then extract a second, linear consequence from the diagonal using $x^2 = x$.
3. Where you would hesitate, and the thought to reinforce: at the moment the Gram equation is written and no contradiction is visible. The thought to reinforce is that one quadratic equation is never enough, you need a linear one, and the only untouched hypothesis is the diagonal, so read it again in the new field.
4. The trap the question sets: staying over $\mathbb{Z}$ and expanding $\|r_1 + \cdots + r_n\|^2$. That computation is correct and yields nothing, and it can absorb twenty minutes.
5. What makes you recognise it faster next time: the phrase "dot product of two different rows" is the signature of a Gram matrix problem, and the word "even" is the signature of $\mathbb{F}_2$. Together they name the whole solution before you have written a line.
6. The reusable decision rule: parity hypotheses become $\mathbb{F}_2$ statements, pairwise inner products become Gram matrices, and impossibility comes from rank.

Mirror drills

1. (Oddtown) Let $A_1, \dots, A_m$ be subsets of $\{1, \dots, n\}$ with every $|A_i|$ odd and every $|A_i \cap A_j|$ even for $i \neq j$. Prove $m \leq n$. *Answer:* the characteristic vectors in $\mathbb{F}_2^n$ have Gram matrix I, which is invertible, so they are linearly independent, so $m \leq n$.

2. Compute $\det(J - I)$ over $\mathbb{Z}$ and decide for which n the matrix $J - I$ is invertible over $\mathbb{F}_2$ and over $\mathbb{F}_3$. *Answer:* $\det(J - I) = (-1)^{n-1}(n-1)$, so it is invertible over $\mathbb{F}_p$ exactly when $p \nmid n - 1$; over $\mathbb{F}_2$ that means n even, over $\mathbb{F}_3$ it means $n \not\equiv 1 \pmod 3$.

3. Let n be odd and let $\overline{M}$ over $\mathbb{F}_2$ satisfy $\overline{M}\,\overline{M}^{T} = J - I$. Show $\operatorname{rank}(\overline{M}) = n - 1$ and that $\ker \overline{M}$ is spanned by the all ones vector. *Answer:* for odd n, $(J - I)x = (\sum_t x_t)\mathbf{1} + x$ vanishes only for $x = 0$ and $x = \mathbf{1}$, so $\operatorname{rank}(J - I) = n - 1$, forcing $\operatorname{rank}(\overline{M}) \geq n - 1$; the diagonal condition puts $\mathbf{1}$ in the kernel, giving equality.

15.5 A5. Rectify the field, then pigeonhole

Let $F : \mathbb{R}^2 \to \mathbb{R}$ and $g : \mathbb{R} \to \mathbb{R}$ be twice continuously differentiable functions with the following properties:

- *$F(u,u) = 0$ for every $u \in \mathbb{R}$;*
- *for every $x \in \mathbb{R}$, $g(x) > 0$ and $x^2 g(x) \leq 1$;*
- *for every $(u,v) \in \mathbb{R}^2$, the vector $\nabla F(u,v)$ is either $\mathbf{0}$ or parallel to the vector $\langle g(u), -g(v)\rangle$.*

Prove that there exists a constant C such that for every $n \geq 2$ and any $x_1, \ldots, x_{n+1} \in \mathbb{R}$, we have

$$\min_{i \neq j} |F(x_i, x_j)| \leq \frac{C}{n}.$$

Read it out loud

Out loud: you are given a two variable function whose gradient direction is prescribed everywhere, which vanishes on the diagonal, and you must show that among any $n+1$ real numbers two of them are sent by F to something of size at most C/n. The arithmetic of the conclusion is the loudest cue on the page. Objects numbered $n+1$, a bound with n in the denominator: that is the signature of pigeonhole into n boxes, and the boxes must be intervals of length proportional to $1/n$ inside a set of finite total length. So before touching the analysis you already know the shape of the endgame: find a map $\Phi : \mathbb{R} \to (A, B)$ into a bounded interval such that $|F(u,v)|$ is controlled by $|\Phi(u) - \Phi(v)|$.

Now read the hypotheses as a list of jobs. The gradient condition is a first order linear partial differential equation in disguise, because "parallel to a given nonzero vector" is the same as "orthogonal to the perpendicular of that vector". The condition $F(u,u) = 0$ is a normalisation, and it will be what turns a Lipschitz estimate into a bound on $|F|$ itself. The condition $g > 0$ makes the antiderivative of g strictly increasing, hence invertible. The condition $x^2 g(x) \leq 1$ is an integrability bound: $g(x) \leq 1/x^2$, and $1/x^2$ is integrable at both ends of the line, so the antiderivative of g is bounded. That is exactly the bounded interval the pigeonhole needs. Every hypothesis has one job and there are no spare parts.

The fields wired in: multivariable calculus (gradients and the chain rule), first order linear partial differential equations and the method of characteristics, one variable analysis (inverse function theorem, mean value theorem, improper integrals), and the pigeonhole principle.

Within thirty seconds you should have written the gradient hypothesis as the scalar equation

$$\frac{1}{g(u)} \frac{\partial F}{\partial u}(u,v) + \frac{1}{g(v)} \frac{\partial F}{\partial v}(u,v) = 0,$$

and next to it the substitution $x = G(u)$ with $G' = g$, which is the only change of variable that turns the operator on the left into $\partial_x + \partial_y$.

Approaches, ranked worst to best

1. **Contradiction with no structure.** *Cue:* the conclusion is a minimum over pairs, so assume every pair is large. *Circuit:* negate and count. You suppose $|F(x_i, x_j)| > C/n$ for all $i \neq j$ and look for an incompatibility. It ranks last because you have no quantitative information about F at all yet, so there is nothing for the assumption to collide with. Negating a conclusion before you have understood the hypotheses is zombie maths.
2. **Bound F by integrating its gradient.** *Cue:* $F(u,u) = 0$ and you want $|F(u,v)|$ small, so integrate from the diagonal. *Circuit:* fundamental theorem of calculus along a segment. Write

$F(u,v) = \int_0^1 \frac{d}{dt} F(u, u + t(v-u))\, dt$ and try to bound the integrand. This stalls immediately and for a nameable reason: the hypothesis specifies the direction of ∇F and says nothing whatsoever about its magnitude, since the scalar of proportionality is an arbitrary function of (u,v). No size estimate can ever be extracted from a direction. Recognising this early is worth more than any computation.

3. **Instantiate the hypotheses on the simplest example.** *Cue:* $x^2 g(x) \le 1$ with $g > 0$ smells like $g(x) = 1/(1+x^2)$. *Circuit:* test the hypotheses on the cheapest object that satisfies them. With $g(x) = 1/(1+x^2)$ you get $G = \arctan$ and $F(u,v) = \arctan u - \arctan v$ satisfies all three hypotheses, and the conclusion holds with $C = \pi$ by pigeonholing $n+1$ arctangents in an interval of length π. This is not a proof and must never be submitted as one, but it ranks above the previous two because it tells you the true theorem: F is a function of $G(u) - G(v)$ where G is bounded. In the arena this is the two minutes that buy you the whole solution.

4. **Method of characteristics in the original variables.** *Cue:* a first order linear PDE means F is constant along the integral curves of a vector field. *Circuit:* solve $\dot{u} = 1/g(u)$, $\dot{v} = 1/g(v)$ and find the invariant. Along such a curve $\frac{d}{dt} G(u) = g(u)\dot{u} = 1$ and likewise for v, so $G(u) - G(v)$ is constant, and F is a function of it. Correct and conceptually the same as the winner, but you must argue that the flow exists for all relevant times and that the orbits cover the plane, which is more analytic bookkeeping than a change of variable needs.

5. **Change coordinates once, globally.** *Cue:* the operator has separated coefficients $\frac{1}{g(u)}\partial_u + \frac{1}{g(v)}\partial_v$. *Circuit:* rectification: choose new coordinates in which the field is constant. Substituting $x = G(u)$, $y = G(v)$ turns the equation into $\Phi_x + \Phi_y = 0$ on a bounded square, whose solutions are exactly the functions of $x - y$. The same substitution simultaneously compresses $\mathbb{R}$ into a bounded interval, which is what pigeonhole needs. Best route: one substitution does two jobs, and the rest is the mean value theorem and counting boxes.

The ideal solution

Step 1: the gradient hypothesis is a linear PDE. Since $g > 0$ everywhere, the vector $\langle g(u), -g(v)\rangle$ is never zero, so the hypothesis says precisely that for every (u,v) there is a real number $\lambda(u,v)$, possibly zero, with

$$\frac{\partial F}{\partial u}(u,v) = \lambda(u,v)\, g(u), \qquad \frac{\partial F}{\partial v}(u,v) = -\lambda(u,v)\, g(v).$$

Dividing and adding eliminates λ:

$$\frac{1}{g(u)}\frac{\partial F}{\partial u}(u,v) + \frac{1}{g(v)}\frac{\partial F}{\partial v}(u,v) = 0 \qquad \text{for all } (u,v) \in \mathbb{R}^2. \tag{1}$$

Step 2: the antiderivative of g is a bounded diffeomorphism. Put $G(x) = \int_0^x g(t)\, dt$. Then $G' = g > 0$, so G is strictly increasing, and G is three times continuously differentiable because g is twice continuously differentiable. For $x \ge 1$ the hypothesis $g(t) \le 1/t^2$ gives

$$0 < G(x) - G(1) = \int_1^x g(t)\, dt \le \int_1^x \frac{dt}{t^2} = 1 - \frac{1}{x} < 1,$$

and for $x \le -1$ the same estimate gives $0 < G(-1) - G(x) < 1$. Hence G is bounded above by $G(1) + 1$ and below by $G(-1) - 1$. A continuous strictly increasing function with finite limits at both ends maps $\mathbb{R}$ bijectively onto the open interval

$$(A, B), \qquad A = \lim_{x \to -\infty} G(x), \quad B = \lim_{x \to +\infty} G(x),$$

and $\ell := B - A \le G(1) - G(-1) + 2 < \infty$. Since $G' = g$ never vanishes, the inverse function theorem gives that $G^{-1} : (A, B) \to \mathbb{R}$ is continuously differentiable with

$$(G^{-1})'(x) = \frac{1}{g(G^{-1}(x))}.$$

Step 3: rectify. Define Φ on the open square $S = (A, B) \times (A, B)$ by

$$\Phi(x, y) = F\big(G^{-1}(x), G^{-1}(y)\big).$$

It is continuously differentiable as a composition of such maps, and by the chain rule together with (1),

$$\frac{\partial \Phi}{\partial x} + \frac{\partial \Phi}{\partial y} = \frac{F_u\big(G^{-1}(x), G^{-1}(y)\big)}{g(G^{-1}(x))} + \frac{F_v\big(G^{-1}(x), G^{-1}(y)\big)}{g(G^{-1}(y))} = 0.$$

Step 4: solve the rectified equation. Let $m = (A + B)/2$ and define

$$h(s) = \Phi\Big(m + \tfrac{s}{2},\, m - \tfrac{s}{2}\Big), \qquad |s| < \ell,$$

which is legitimate because $|s| < \ell$ forces $m \pm s/2 \in (A, B)$. The function h is continuously differentiable on $(-\ell, \ell)$.

Fix $(x, y) \in S$ and set $s = x - y$, so $|s| < \ell$. The points (x, y) and $(m + s/2,\, m - s/2)$ both lie in S and differ by a multiple of $(1, 1)$; since S is convex, the segment joining them lies in S. Along that segment, parametrised as $\gamma(t) = (m + s/2 + t,\, m - s/2 + t)$, we have

$$\frac{d}{dt}\Phi(\gamma(t)) = \Phi_x + \Phi_y = 0,$$

so Φ is constant on it. Hence $\Phi(x, y) = h(x - y)$ for every $(x, y) \in S$, and therefore

$$F(u, v) = h\big(G(u) - G(v)\big) \qquad \text{for all } u, v \in \mathbb{R}. \tag{2}$$

Step 5: normalise and get a Lipschitz constant. Taking $u = v$ in (2) and using $F(u, u) = 0$ gives $h(0) = 0$. Since h' is continuous on $(-\ell, \ell)$, it is bounded on the compact subinterval $[-\ell/2, \ell/2]$; put

$$M = \max_{|s| \le \ell/2} |h'(s)| < \infty.$$

Step 6: pigeonhole and finish. Let $C = M\ell$, a constant depending only on F and g. Let $n \ge 2$ and let $x_1, \dots, x_{n+1}$ be any reals. The $n + 1$ numbers $G(x_1), \dots, G(x_{n+1})$ all lie in (A, B), which is contained in the union of the n intervals

$$I_t = \Big[A + \tfrac{(t-1)\ell}{n},\, A + \tfrac{t\ell}{n}\Big), \qquad t = 1, \dots, n.$$

By the pigeonhole principle two of them, say $G(x_i)$ and $G(x_j)$ with $i \ne j$, lie in the same I_t, so

$$|G(x_i) - G(x_j)| < \frac{\ell}{n} \le \frac{\ell}{2},$$

the last step because $n \ge 2$. The number $G(x_i) - G(x_j)$ therefore lies in $[-\ell/2, \ell/2]$, and by (2), $h(0) = 0$ and the mean value theorem applied to h on the interval between 0 and $G(x_i) - G(x_j)$,

$$|F(x_i, x_j)| = \big|h(G(x_i) - G(x_j)) - h(0)\big| \le M|G(x_i) - G(x_j)| \le \frac{M\ell}{n} = \frac{C}{n}.$$

Hence $\min_{i \ne j} |F(x_i, x_j)| \le C/n$, with C independent of n and of the points. ∎

Layer 1: WHAT

Block 1. Convert "∇F is zero or parallel to $\langle g(u), -g(v)\rangle$" into the single scalar equation $F_u/g(u) + F_v/g(v) = 0$, valid at every point including where the gradient vanishes.

Block 2. Recognise the left side as the derivative of F along the vector field $(1/g(u), 1/g(v))$, and ask for the coordinate in which that field is constant.

Block 3. Answer: $x = G(u)$ with $G' = g$. Check separately that this same G is bounded, because $g(t) \le 1/t^2$ is integrable at both ends. One substitution, two jobs.

Block 4. In the new coordinates the equation is $\Phi_x + \Phi_y = 0$ on a bounded convex square, so Φ depends only on $x - y$; translating back, $F(u, v) = h(G(u) - G(v))$.

Block 5. Use the diagonal hypothesis: $h(0) = 0$. Use continuity of h' on a compact subinterval: h is Lipschitz there with constant M.

Block 6. Put $n + 1$ values of G into n boxes of length ℓ/n, take the colliding pair, apply the mean value theorem.

Juxtaposition. The nominal route reads "twice continuously differentiable" and reaches for Taylor expansions and second derivative bounds; then it reads "$F(u, u) = 0$" and tries the mean value inequality $|F(u, v)| \le \sup|\nabla F| \cdot |u - v|$; then it discovers that no supremum is available and starts inventing hypotheses. The bifurcation point is inside Block 1 to Block 2: the decision to stop estimating F and to solve for F exactly. The hypothesis is not an estimate, it is an equation, and equations should be solved, not bounded. Once (2) is on the page the problem is a one variable exercise.

Layer 2: HOW

Block 1. Cue. "Parallel to" a specified nonzero vector, with the specified vector built from g evaluated separately at u and at v. **Why natural.** Parallelism is a statement with one degree of freedom, the scalar λ, and λ is exactly the information the problem refuses to give you. So eliminate it. Dotting ∇F with the perpendicular vector $\langle g(v), g(u)\rangle$ kills λ in one line and leaves a genuine equation. The case $\nabla F = \mathbf{0}$ is not a separate case at all, it is $\lambda = 0$, which is why the equation holds at every point without exception. **Link to the next block.** You now have a first order linear homogeneous PDE. The standard theory says its solutions are exactly the functions that are constant along the integral curves of the associated vector field, so the next question is what those curves are.

Block 2. Cue. The coefficients are separated: the coefficient of ∂_u depends only on u, and the coefficient of ∂_v only on v. **Why natural.** A vector field with separated coefficients is rectified by a product substitution, one variable at a time. If $x = \phi(u)$ then $\partial_u = \phi'(u)\partial_x$, so the coefficient of ∂_x becomes $\phi'(u)/g(u)$; choosing $\phi' = g$ makes it 1. Nothing was guessed: the requirement $\phi' = g$ is forced by the demand that the coefficient equal 1. **Link to the next block.** The substitution is $x = G(u)$ where G is an antiderivative of g. Before using it, find out what its range is, because the range is the new domain and the whole endgame depends on the new domain being bounded.

Block 3. Cue. The hypothesis $x^2 g(x) \le 1$ has not been used, and it is an inequality about the size of g at infinity. **Why natural.** An unused hypothesis is a misread problem. This one rearranges to $g(x) \le 1/x^2$ for $x \neq 0$, and the only standard fact about $1/x^2$ is that it is integrable away from the origin. So the hypothesis exists to make $\int^{\infty} g$ and $\int_{-\infty} g$ finite, which is to say to make G bounded. The problem author put it there to compactify the line, and no other reading of it is possible. **Link to the next block.** G is a strictly increasing bounded smooth bijection onto a bounded interval (A, B), with a differentiable inverse by the inverse function theorem. Now push F through it.

Block 4. Cue. The transport equation $\Phi_x + \Phi_y = 0$. **Why natural.** The derivative of Φ in the direction $(1, 1)$ is zero, so Φ is constant along every line of slope 1 that stays in the domain. To

turn "constant on each line" into "a function of $x - y$" you need the lines to meet the domain in connected sets, which is exactly convexity of the square. That is why the write up says the word convex: it is not decoration, it is the reason a local statement becomes a global one. **Link to the next block.** $F(u, v) = h(G(u) - G(v))$ is now an exact formula for F. Every remaining hypothesis should be applied to h.

Block 5. **Cue.** $F(u, u) = 0$, the one hypothesis not yet used. **Why natural.** In the new description, $u = v$ means argument zero, so the diagonal hypothesis is precisely $h(0) = 0$. That is what converts a derivative bound into a value bound: $|h(s)| = |h(s) - h(0)| \leq M|s|$. Without $h(0) = 0$ the conclusion would be false, since $F \equiv 1$ satisfies every other hypothesis. **Link to the next block.** The bound is now linear in $|G(x_i) - G(x_j)|$, so all that remains is to make some pair of the $G(x_i)$ close.

Block 6. **Cue.** $n + 1$ points, a bound with n in the denominator, and a domain of finite length ℓ. **Why natural.** This is the pigeonhole principle in its continuous form. Cutting an interval of length ℓ into n pieces of length ℓ/n and dropping $n + 1$ points into it forces a collision, and the collision distance is at most ℓ/n. The restriction $n \geq 2$ in the problem statement exists so that $\ell/n \leq \ell/2$, keeping the difference inside the compact interval on which h' was bounded. **Link to the next block.** Mean value theorem, and the constant $C = M\ell$ depends only on F and g, which is what "there exists a constant C" demands.

Layer 3: WHY

The mathematical truth beneath this route is the classification of solutions of a first order linear homogeneous PDE. The equation $a(u, v)F_u + b(u, v)F_v = 0$ says that ∇F is orthogonal to the field (a, b), that is, that F is constant along the integral curves of that field. Where the field is nonvanishing, the rectification theorem says there are local coordinates in which the field is ∂_x, and then the general solution is an arbitrary function of the remaining coordinate. Here the field is nonvanishing everywhere and the rectifying coordinates are global, so the general solution is global: F is an arbitrary function of the single invariant $G(u) - G(v)$. This is the reason the problem is solvable at all. You are not being asked to estimate an unknown function; you are being handed a complete description of it with one unknown one variable function inside, and the remaining hypotheses pin down what you need about that function.

The second truth is that compactness has to come from somewhere. The conclusion is uniform over all real inputs, so some hypothesis must confine the line. The only candidate is the growth bound on g, and it confines the line not directly but through the invariant: the set of achievable values of $G(u) - G(v)$ is the bounded interval $(-\ell, \ell)$. Once the invariant lives in a bounded set, the pigeonhole principle is available and the $1/n$ appears for free.

The discernment mechanism is a habit rather than a theorem: before writing anything, list the hypotheses and assign each one a job. Here there are four items (diagonal vanishing, positivity of g, decay of g, gradient direction) and four jobs (normalise h, make G invertible, make the range bounded, produce the PDE). If your proof does not use every hypothesis, either your proof is wrong or you have proved something stronger, and on the Putnam the first is far more likely. Notice also which hypothesis is stronger than needed: continuous differentiability of F and g is enough for everything above, and the second derivatives are never used. Saying so out loud is part of understanding the problem, not a criticism of it.

How to train it: whenever a hypothesis constrains a derivative's direction rather than its size, immediately write the associated PDE and try to solve it exactly. Whenever a hypothesis is an integrability or decay bound, immediately name the integral it makes finite. Whenever a conclusion has the shape "among $N + 1$ objects two are within c/N", write "pigeonhole into N boxes" at the bottom of your page before you start, and then work towards it.

What the examiners nudged. They wrote the direction vector as $\langle g(u), -g(v)\rangle$ with the arguments separated, so the characteristic system decouples and the change of variable is a product of one variable maps. They wrote $x^2 g(x) \le 1$ rather than the cleaner "g is integrable", so that you would have to perform the comparison yourself and thereby notice why it is there. They wrote $n+1$ points and C/n rather than n points and C/n, which is the pigeonhole telling you its box count out loud.

Revolving concepts

First order linear homogeneous PDE and characteristics. For $a(u,v)F_u + b(u,v)F_v = 0$ with (a,b) continuous and nonvanishing, the characteristic curves solve $\dot u = a$, $\dot v = b$, and F is a solution exactly when it is constant on each characteristic. If $w(u,v)$ is a function with nonvanishing gradient that is constant on characteristics (a first integral), the general solution is $F = h \circ w$ for an arbitrary function h of the same regularity. Here $w = G(u) - G(v)$. See Book 1, field: Calculus and integrals, topic: first order PDE and characteristics.

The transport equation. $\Phi_x + \Phi_y = 0$ on a convex domain implies $\Phi(x,y) = h(x-y)$; the proof is the one line $\frac{d}{dt}\Phi(x+t, y+t) = \Phi_x + \Phi_y = 0$. The same one line proves that $\Phi_x - \Phi_y = 0$ gives a function of $x+y$, and that $u\Phi_u + v\Phi_v = 0$ gives a function of v/u (homogeneity of degree zero, Euler's relation). See Book 1, field: Calculus and integrals, topic: transport equation.

Rectification of a vector field. A vector field that does not vanish at a point admits local coordinates in which it is the constant field ∂_x. In one dimension the rectifying coordinate is $\int dx/(\text{field})$; with separated coefficients, apply the one dimensional recipe in each variable. See Book 1, field: Real analysis, topic: flows and rectification.

Inverse function theorem in one variable. If G is C^k on an interval with $G' \neq 0$ throughout, then G is a C^k diffeomorphism onto its image and $(G^{-1})' = 1/(G' \circ G^{-1})$. Combined with monotone convergence of G at the ends, this identifies the image as an open interval. See Book 1, field: Real analysis, topic: inverse function theorem.

Comparison test for improper integrals. If $0 \le f \le F$ on $[a,\infty)$ and $\int_a^\infty F$ converges then so does $\int_a^\infty f$; the model case is $\int_1^\infty t^{-p}\,dt$, convergent exactly for $p > 1$. See Book 1, field: Calculus and integrals, topic: improper integrals.

Pigeonhole in continuous form, and Dirichlet approximation. $N+1$ points in an interval of length ℓ force two within ℓ/N. The same circuit, applied to the fractional parts $\{0\}, \{\alpha\}, \ldots, \{N\alpha\}$ in $[0,1)$, gives integers $1 \le q \le N$ with $|q\alpha - p| < 1/N$, which is Dirichlet's approximation theorem. See Book 1, field: Number theory, topic: Dirichlet approximation.

Iteration, flows and Abel's equation (the recursion behind the rectification). There is no discrete recursion in this problem, but the underlying structure is the continuous version of one. A vector field generates a one parameter family of maps Ψ_t with $\Psi_{s+t} = \Psi_s \circ \Psi_t$, the flow, which is the continuous analogue of iterating a single map, $f^{m+1} = f \circ f^m$. The generalised statement is: to understand an iteration or a flow, find coordinates in which it becomes a translation. For a flow that means solving $\phi'(u) \cdot (\text{field}) = 1$, that is, $\phi = \int du/(\text{field})$, which is precisely the G used here; for an iteration f it means solving Abel's functional equation $\phi(f(u)) = \phi(u) + 1$, after which $\phi(f^m(u)) = \phi(u) + m$ and every asymptotic question about f^m becomes arithmetic. The same generalisation solves the recurrence $a_{m+1} = f(a_m)$ for slowly converging f (for example $a_{m+1} = \sin a_m$, where $\phi = 1/a^2$ up to constants gives $a_m \sim \sqrt{3/m}$). See Book 1, field: Sequences and recurrences, topic: conjugating an iteration to a translation.

Circuit Bank

- **WHEN** a hypothesis fixes the direction of a gradient but not its size **DO** eliminate the scalar of proportionality, write the resulting first order PDE, and solve it exactly **BECAUSE** direction

carries no magnitude information, so no estimate can ever be extracted from it, whereas the level sets are completely determined

- **WHEN** a first order operator with separated coefficients $\frac{1}{g(u)}\partial_u + \frac{1}{g(v)}\partial_v$ **DO** substitute $x = G(u)$, $y = G(v)$ where $G' = g$ **BECAUSE** that is the unique substitution making both coefficients equal to 1, and a constant field has an obvious invariant, namely $x - y$
- **WHEN** a hypothesis is a decay or integrability bound such as $x^2 g(x) \leq 1$ **DO** name the integral it makes finite, and expect that finiteness to be the compactness the proof needs **BECAUSE** a Putnam hypothesis is never decoration, and a bound on a function at infinity can only act through an integral or a limit
- **WHEN** the conclusion reads "among $N + 1$ objects, two within c/N" **DO** write "pigeonhole into N boxes" at the foot of the page before starting **BECAUSE** the arithmetic of the statement is the examiner naming the last step, and knowing the last step fixes what the earlier steps must produce
- **WHEN** you must bound $|f|$ and you know f vanishes somewhere **DO** write $f(s) = f(s) - f(s_0)$ and apply the mean value theorem with a derivative bound on a compact set **BECAUSE** a value bound needs one value plus a derivative bound, and continuity of the derivative supplies the bound on any compact set by the extreme value theorem

Backtesting

1. First useful cue: $n + 1$ objects with a C/n conclusion, which fixes pigeonhole as the endgame and therefore demands a bounded interval to pigeonhole into.
2. Method the expression suggested: eliminate the proportionality scalar to get $F_u/g(u) + F_v/g(v) = 0$, then rectify with $x = G(u)$, $G' = g$, which is also the map whose boundedness the hypothesis $x^2 g(x) \leq 1$ was placed there to guarantee.
3. Where you would hesitate, and the thought to reinforce: at the words "twice continuously differentiable", which pull you towards Taylor expansions and second derivative estimates. Reinforce this instead: the smoothness is there only to make the substitution and the mean value theorem legal, and the real content is that the hypothesis is an equation to be solved.
4. The trap the question sets: attempting to bound $|F(u, v)|$ by $\sup|\nabla F| \cdot |u - v|$. There is no such supremum, because the proportionality scalar is unconstrained, and the conclusion is uniform in the points, so no bound depending on $|u - v|$ could ever finish.
5. What makes you recognise it faster next time: two signatures. "Gradient parallel to a given field" equals "solve the PDE". "$N + 1$ objects, bound c/N" equals "pigeonhole into N boxes".
6. The reusable decision rule: solve exact constraints rather than estimating them, and get compactness from the one hypothesis that bounds growth.

Mirror drills

1. Let $\Phi : \mathbb{R}^2 \to \mathbb{R}$ be continuously differentiable with $\Phi_x + \Phi_y = 0$. Show $\Phi(x, y) = h(x - y)$ for some continuously differentiable h. *Answer:* put $h(s) = \Phi(s, 0)$; then $\frac{d}{dt}\Phi(x + t, y + t) = 0$, so $\Phi(x, y) = \Phi(x - y, 0) = h(x - y)$.
2. Let g be continuous and positive on $\mathbb{R}$ with $\int_{-\infty}^{\infty} g < \infty$. Show that $G(x) = \int_0^x g$ is a bijection from $\mathbb{R}$ onto a bounded open interval whose inverse is continuously differentiable. *Answer:* G is strictly increasing with finite limits at both ends, so it is a continuous bijection onto (A, B) with $B - A = \int_{-\infty}^{\infty} g$; the inverse is C^1 by the inverse function theorem since $G' = g > 0$.

3. Prove Dirichlet's theorem: for every real α and integer $N \geq 1$ there are integers p and $1 \leq q \leq N$ with $|q\alpha - p| < 1/N$. *Answer:* the $N+1$ fractional parts $\{k\alpha\}$ for $k = 0, \ldots, N$ lie in $[0, 1)$, which is the union of N intervals of length $1/N$; two share an interval, and subtracting gives the claim with q the difference of the two indices.

15.6 A6. Fourier analysis of a group random walk

Let G be an abelian group with n elements, and let

$$\{g_1 = e, g_2, \ldots, g_k\} \subsetneq G$$

be a (not necessarily minimal) set of distinct generators of G. A special die, which randomly selects one of the elements $g_1, g_2, \ldots, g_k$ with equal probability, is rolled m times and the selected elements are multiplied to produce an element $g \in G$. Prove that there exists a real number $b \in (0,1)$ such that

$$\lim_{m\to\infty} \frac{1}{b^{2m}} \sum_{x\in G} \left(\text{Prob}(g = x) - \frac{1}{n}\right)^2$$

is positive and finite.

Read it out loud

Out loud: this is a random walk on a finite abelian group, started at the identity, whose step distribution is uniform on a generating set that contains the identity. You are asked to show that the walk approaches the uniform distribution, that the approach is exponential in m, and that it has an exact exponential rate, not merely an upper and a lower bound with different constants.

Three phrases fix the whole solution. First, "rolled m times and the selected elements are multiplied": the distribution of the product is the m fold convolution of the step distribution, so the object is a power of one operator and the question is a question about eigenvalues. Second, "$\sum_{x\in G}(\cdot)^2$": that is a squared L^2 norm on the space of functions on G, and squared norms are computed in whichever orthogonal basis diagonalises the operator, never coordinate by coordinate. Third, "abelian": for a finite abelian group the diagonalising basis exists, is explicit, and is the group of characters. A finite abelian group has exactly n characters and they form an orthogonal basis of the n dimensional space of complex functions on G. That is not a convenience, it is the reason the problem is stated for abelian G.

The fields wired in: representation theory of finite abelian groups (characters and orthogonality), Fourier analysis on a finite group (convolution theorem, inversion, Parseval), Markov chains and the spectral gap, and linear algebra (diagonalisation, and the fact that the asymptotics of $T^m v$ are governed by the second largest eigenvalue modulus).

Within thirty seconds you should have written $P_m = P_{m-1} * P$, then $\widehat{P_m} = \widehat{P}^m$, then $\lambda_\chi = \frac{1}{k}\sum_{i=1}^{k} \chi(g_i)$, and finally $b = \max\{|\lambda_\chi| : \chi \neq \widehat{e}\}$. The remaining work is entirely to show $0 < b < 1$, and the two hypotheses you have not yet used, that $g_1 = e$ and that the generating set is a *proper* subset, are exactly the two facts that deliver those two inequalities.

Approaches, ranked worst to best

1. **Direct combinatorial counting.** *Cue:* "rolled m times and multiplied" is a word count. *Circuit:* count words of length m in the generators equal to x. You would write $\text{Prob}(g = x) = k^{-m}\#\{(i_1, \ldots, i_m) : g_{i_1} \cdots g_{i_m} = x\}$ and try to evaluate. Ranks last: the count is a sum of multinomial coefficients over solutions of a linear system in the exponents, it has no closed form in general, and even in the cyclic case the closed form is a character sum, so you arrive at the right answer by the longest possible road. Worse, the quantity asked for is an exponential rate, and rates are invisible in a formula built out of binomial coefficients.

2. **Markov chain plus Perron Frobenius.** *Cue:* a repeated random step with a fixed rule is a Markov chain. *Circuit:* transition matrix, stationary distribution, convergence theorem. Set $M_{x,y} = \frac{1}{k}[\, x^{-1}y \in \{g_i\}\,]$; the chain is doubly stochastic, irreducible because the g_i generate, and aperiodic because $g_1 = e$ gives every state a self loop, so the distribution converges

to uniform. Ranks above counting because it is the right picture and it explains where the hypotheses act. It ranks below the winner because Perron Frobenius gives only that the subdominant eigenvalue governs the decay; without diagonalisability, which here needs the character basis anyway, you get a lim sup and not a limit, and the problem explicitly asks for a limit.

3. **Characters of cyclic groups plus the structure theorem.** *Cue:* the smallest abelian groups are $\mathbb{Z}/n$, where characters are powers of a root of unity. *Circuit:* reduce to cyclic factors and take products. Every finite abelian group is a direct product of cyclic groups, and the characters of a product are products of characters, so you can build everything from roots of unity. Correct and self contained if you already know the structure theorem, but the bookkeeping over the factors is pure ink: nothing in the argument ever needs the decomposition.

4. **Fourier analysis on G directly.** *Cue:* "abelian" and "$\sum_x(\cdot)^2$" name the character basis and Parseval. *Circuit:* convolution becomes multiplication, then read off the m th power. Take the Fourier transform on G, note that the m fold convolution has transform λ_χ^m, apply Parseval to convert the sum over $x \in G$ into a sum over characters, and let b be the largest $|\lambda_\chi|$ over nontrivial χ. The limit then falls out because every term with $|\lambda_\chi| < b$ dies and every term with $|\lambda_\chi| = b$ contributes exactly 1. Best route: no case analysis, no structure theorem, and the two remaining hypotheses are consumed by two short claims.

The ideal solution

Notation and standard facts. Let $\widehat{G}$ denote the group of characters of G, that is, the homomorphisms $\chi : G \to \mathbb{C}^\times$, and let $\widehat{e}$ denote the trivial character. Since $x^n = e$ for every $x \in G$, every value $\chi(x)$ is an n th root of unity, so $|\chi(x)| = 1$ and $\overline{\chi(x)} = \chi(x^{-1})$. We use the standard facts that $|\widehat{G}| = n$ and that the two orthogonality relations hold:

$$\sum_{x\in G} \chi(x)\overline{\psi(x)} = n\,[\chi = \psi], \qquad \sum_{\chi\in\widehat{G}} \chi(x) = n\,[x = e].$$

For functions $\varphi, \psi : G \to \mathbb{C}$ write

$$\widehat{\varphi}(\chi) = \sum_{x\in G} \varphi(x)\chi(x), \qquad (\varphi * \psi)(x) = \sum_{y\in G} \varphi(y)\psi(y^{-1}x).$$

Three consequences are immediate from the orthogonality relations and will be used without further comment: inversion, $\varphi(x) = \frac{1}{n}\sum_\chi \widehat{\varphi}(\chi)\overline{\chi(x)}$; Parseval, $\sum_x|\varphi(x)|^2 = \frac{1}{n}\sum_\chi|\widehat{\varphi}(\chi)|^2$; and the convolution theorem, $\widehat{\varphi * \psi} = \widehat{\varphi}\,\widehat{\psi}$, the last because $\chi(x) = \chi(y)\chi(y^{-1}x)$.

Step 1: the walk is a convolution power. Let $S = \{g_1, \dots, g_k\}$ and let $P(x) = \frac{1}{k}$ for $x \in S$ and $P(x) = 0$ otherwise, the distribution of one roll. Let $P_m(x) = \mathrm{Prob}(g = x)$ after m rolls. Conditioning on the product of the first $m - 1$ rolls and using independence,

$$P_m(x) = \sum_{y\in G} P_{m-1}(y)\,P(y^{-1}x), \qquad \text{that is} \qquad P_m = P_{m-1} * P,$$

with $P_1 = P$. Taking Fourier transforms and iterating, $\widehat{P_m} = \widehat{P}^m$, so with

$$\lambda_\chi := \widehat{P}(\chi) = \frac{1}{k}\sum_{i=1}^{k} \chi(g_i)$$

we have $\widehat{P_m}(\chi) = \lambda_\chi^m$ for every $\chi \in \widehat{G}$ and every $m \geq 1$.

Step 2: Parseval turns the sum over G into a sum over characters. Let $U(x) = 1/n$ be the uniform distribution. Its transform is $\widehat{U}(\chi) = \frac{1}{n}\sum_x \chi(x) = [\chi = \widehat{e}]$, and $\lambda_{\widehat{e}} = \frac{1}{k}\cdot k = 1$, so $\widehat{P_m - U}(\chi) = 0$ for

$\chi = \widehat{e}$ and $= \lambda_\chi^m$ otherwise. Parseval therefore gives, for every $m \geq 1$,

$$D_m := \sum_{x \in G} \Big(P_m(x) - \frac{1}{n} \Big)^2 = \frac{1}{n} \sum_{\chi \neq \widehat{e}} |\lambda_\chi|^{2m}. \tag{1}$$

Step 3: the eigenvalues have modulus at most 1*, with equality only for* $\widehat{e}$*.* Each $\chi(g_i)$ is a complex number of modulus 1, so the triangle inequality gives $|\lambda_\chi| \leq 1$. Suppose $|\lambda_\chi| = 1$ for some χ, that is, $|z_1 + \cdots + z_k| = k$ with $z_i = \chi(g_i)$ of modulus 1. Equality in the triangle inequality forces all the z_i to be nonnegative real multiples of one common unit vector, and since each has modulus 1 they are all equal. Because $g_1 = e$ we have $z_1 = \chi(e) = 1$, hence $\chi(g_i) = 1$ for every i. The set of elements on which χ takes the value 1 is the kernel of χ, a subgroup; it contains all the g_i, which generate G, so it is all of G and $\chi = \widehat{e}$. Therefore

$$|\lambda_\chi| < 1 \quad \text{for every } \chi \neq \widehat{e}. \tag{2}$$

Step 4: the group has a nontrivial character and b *is positive.* Since $S \subsetneq G$ we have $k < n$, and $k \geq 1$, so $n \geq 2$ and $\widehat{G} \neq \{\widehat{e}\}$ because $|\widehat{G}| = n \geq 2$. Define

$$b := \max\{|\lambda_\chi| : \chi \in \widehat{G},\ \chi \neq \widehat{e}\},$$

a maximum over a finite nonempty set, so it is attained. By (2), $b < 1$. Suppose $b = 0$, so that $\lambda_\chi = 0$ for every $\chi \neq \widehat{e}$. Then inversion gives, for every $x \in G$,

$$P(x) = \frac{1}{n} \sum_{\chi} \lambda_\chi \overline{\chi(x)} = \frac{1}{n} \lambda_{\widehat{e}} = \frac{1}{n},$$

so P is the uniform distribution. But S is a proper subset of G, so there is an $x \notin S$, and for it $P(x) = 0 \neq 1/n$. This contradiction gives $b > 0$, so

$$b \in (0, 1).$$

Step 5: the limit. Let

$$N = \#\{\chi \in \widehat{G} : |\lambda_\chi| = b\},$$

noting that $\widehat{e}$ is not counted since $|\lambda_{\widehat{e}}| = 1 > b$, and that $1 \leq N \leq n - 1$ by the definition of b. Divide (1) by b^{2m}:

$$\frac{D_m}{b^{2m}} = \frac{1}{n} \sum_{\chi \neq \widehat{e}} \left(\frac{|\lambda_\chi|}{b} \right)^{2m}.$$

The sum is finite, with N terms equal to 1 for every m and every other term of the form r^{2m} with $0 \leq r < 1$, which tends to 0. Hence

$$\lim_{m \to \infty} \frac{1}{b^{2m}} \sum_{x \in G} \Big(\mathrm{Prob}(g = x) - \frac{1}{n} \Big)^2 = \frac{N}{n},$$

which is positive and finite, and $b \in (0, 1)$. ■

Two remarks worth carrying. First, the value of the limit is N/n and not N; the published solution loses the factor $1/n$ because it computes a squared norm of a vector of characters, each of which has squared norm n. The problem only asks for positive and finite, so nothing is lost, but you should compute constants correctly out of habit. Second, if you set the problem up with a transition matrix rather than with convolution, check which way your matrix multiplies: the natural matrix $M_{ij} = \frac{1}{k}[\, h_j = h_i g \text{ for some } g \in S\,]$ applied to the indicator of e produces the distribution of g^{-1}, not of g. Inversion is a bijection of G, so the sum in the problem is unchanged, but a grader is entitled to ask.

Layer 1: WHAT

Block 1. Name the object: $\text{Prob}(g = x)$ is the m fold convolution power of the uniform measure on the generating set, so the recursion is $P_m = P_{m-1} * P$, a linear recursion in the n dimensional space of functions on G.

Block 2. Diagonalise that linear recursion. Because G is abelian, the diagonalising basis is the character basis, and the eigenvalues are $\lambda_\chi = \frac{1}{k}\sum_i \chi(g_i)$.

Block 3. Convert the requested sum of squares into the character basis by Parseval, obtaining $D_m = \frac{1}{n}\sum_{\chi \neq \widehat{e}} |\lambda_\chi|^{2m}$, with the trivial character removed by the subtraction of $1/n$.

Block 4. Identify b as the largest nontrivial eigenvalue modulus, and spend the two unused hypotheses: $g_1 = e$ with generation gives $b < 1$, and properness of the generating set gives $b > 0$.

Block 5. Divide by b^{2m} and observe that a finite sum of terms r^{2m} with $r \leq 1$ converges to the number of terms with $r = 1$.

Juxtaposition. The nominal route tries to compute $\text{Prob}(g = x)$ for a general x, or fixes a small group such as $\mathbb{Z}/3$ to see a pattern, or quotes a convergence theorem for Markov chains and then tries to upgrade it to an exact rate by hand. The bifurcation point is Block 1 into Block 2: the decision that you will never compute a single value of $\text{Prob}(g = x)$, because the question does not ask for one. It asks for a squared norm, and squared norms are computed in the eigenbasis. Everything after that decision is forced.

Layer 2: HOW

Block 1. Cue. "Rolled m times and the selected elements are multiplied." **Why natural.** The distribution of a product of independent group valued random variables is the convolution of their distributions; that is the definition of convolution, not a theorem. Writing $P_m = P_{m-1} * P$ turns the words of the problem into a linear recursion $v_m = Tv_{m-1}$ on the n dimensional space of functions on G, and every linear recursion is solved by understanding T. **Link to the next block.** A power of a linear operator asks for an eigenbasis, so the next question is which basis diagonalises convolution by a fixed function.

Block 2. Cue. The word "abelian", which is doing no probabilistic work whatsoever, and therefore must be doing algebraic work. **Why natural.** Convolution operators commute with all translations of G. A commuting family of diagonalisable operators has a simultaneous eigenbasis, and for the translations of a finite abelian group that basis is exactly the set of characters: $\chi(y^{-1}x) = \overline{\chi(y)}\chi(x)$ says precisely that translating χ rescales it. So characters are eigenvectors of every convolution operator at once, with eigenvalue the Fourier coefficient of the convolving function. This is the single fact the problem is built on, and it fails for non abelian G because then the irreducible representations have dimension greater than one. **Link to the next block.** With eigenvalues named, the quantity to compute should be re expressed in the eigenbasis rather than in the standard basis.

Block 3. Cue. $\sum_{x \in G}(\cdot)^2$, a squared Euclidean norm over the whole group. **Why natural.** A squared norm is basis independent up to the normalisation of the basis, so compute it in whichever orthogonal basis makes the vector simple. Parseval is that statement for the character basis, and here it is decisive: in the standard basis $P_m - U$ has n complicated entries, while in the character basis it has entries λ_χ^m, which are literally m th powers. Subtracting the uniform distribution corresponds exactly to deleting the trivial character, because the uniform distribution is the projection of anything onto the constant function. **Link to the next block.** The expression $\frac{1}{n}\sum_{\chi \neq \widehat{e}} |\lambda_\chi|^{2m}$ is a finite sum of exponentials, and the asymptotics of such a sum are controlled by the largest base. That defines b with no choice involved.

Block 4. **Cue.** Two hypotheses remain unused: $g_1 = e$, and the generating set is a proper subset. The conclusion demands $b \in (0,1)$, two strict inequalities. **Why natural.** Count the unused hypotheses against the unproved claims: two and two. That pairing tells you what to prove with what before you have any idea how. The identity element in the list is what makes the walk aperiodic, and aperiodicity is exactly the statement that no nontrivial character has $|\lambda_\chi| = 1$; the properness is what stops the walk from being uniform already after one step, which is exactly $b > 0$. If the identity were absent the walk could be periodic: on $\mathbb{Z}/2$ with the single generator 1, the walk alternates deterministically, the nontrivial eigenvalue is -1, and the sum in the problem is constantly $1/2$, so no $b \in (0,1)$ can work. **Link to the next block.** The mechanism for $b < 1$ is equality in the triangle inequality; the mechanism for $b > 0$ is Fourier inversion, which turns "all nontrivial coefficients vanish" into "the function is constant".

Block 5. **Cue.** A finite sum of terms r_χ^{2m} divided by b^{2m} where $b = \max r_\chi$. **Why natural.** Each ratio $(r_\chi/b)^{2m}$ is either identically 1 or a geometric sequence with ratio less than 1. A finite sum of sequences may be passed to the limit term by term, so the limit is the count of maximisers. This is also the reason the problem is stated with b^{2m} rather than with a general normalisation: the ties must be allowed to count themselves, and they contribute an integer. **Link to the next block.** Positivity of the limit is the existence of at least one maximiser, which holds because $\widehat{G} \setminus \{\widehat{e}\}$ is nonempty; finiteness is the finiteness of $\widehat{G}$.

Layer 3: WHY

The truth beneath this route is that a convolution operator on a finite abelian group is a G circulant matrix, and every G circulant matrix is diagonalised by the same explicit unitary, the Fourier matrix. Concretely, if $M_{x,y} = f(x^{-1}y)$ then M commutes with every translation, the translations generate a commutative family of unitary operators, and a commuting family of unitary operators on a complex inner product space has a common orthogonal eigenbasis. For $G = \mathbb{Z}/n$ this is the familiar statement that circulant matrices are diagonalised by the discrete Fourier transform with eigenvalues $\sum_j c_j \omega^{jt}$. The Putnam problem is that statement dressed as a die roll.

The second truth is the general principle governing $T^m v$ for a diagonalisable T. Write $v = \sum_j c_j w_j$ in an eigenbasis with eigenvalues μ_j; then $T^m v = \sum_j c_j \mu_j^m w_j$. If exactly one eigenvalue has maximal modulus and it equals 1 with eigenvector w_0, then $T^m v \to c_0 w_0$ and the error is $\sum_{j \neq 0} c_j \mu_j^m w_j$, whose norm behaves like b^m where b is the second largest modulus. Dividing by b^{2m} and letting m grow keeps exactly the tying terms. This is the spectral gap principle: the gap $1 - b$ controls the mixing time of the walk, roughly $\log(n)/(1-b)$ steps to reach near uniformity, and the problem is asking you to prove that the gap is genuinely strictly between 0 and 1.

The discernment mechanism has three parts, and all three are transferable. First, when the quantity asked for is a norm rather than a value, change basis instead of computing entries. Second, when an operator is built from a group action, look for the representation theory that diagonalises it before doing anything analytic. Third, and most important here, audit your unused hypotheses against your unproved claims. This problem hands you two decorative looking conditions, $g_1 = e$ and $\subsetneq$, and asks for two strict inequalities, $b > 0$ and $b < 1$. That pairing is not a coincidence and it is visible before you know a single technique. Both conditions are sharp: dropping $g_1 = e$ allows periodic walks, and dropping properness allows the walk to be exactly uniform from the first step, which makes D_m identically zero and the limit 0 rather than positive.

How to train it: install the character theory of finite abelian groups as a reflex, then practise saying "convolution becomes multiplication" whenever you meet a sum over pairs whose indices multiply to something fixed. Then practise the audit: on every problem, list the hypotheses and cross them off as you spend them, and never submit a proof with an unspent hypothesis.

What the examiners nudged. They said abelian, which names the tool. They put e into the generating set explicitly, with the notation $g_1 = e$ rather than in prose, because it is the aperiodicity

hypothesis and they wanted it impossible to miss. They wrote $\subsetneqq$ with the strict symbol rather than saying "a set of generators", because properness is a hypothesis and not a description. And they asked for b^{2m} with a square, matching the square in the sum, so that b is the eigenvalue modulus rather than its square, which is the natural normalisation from the eigenvector expansion.

Revolving concepts

Characters and the dual group. A character of a finite abelian group G is a homomorphism $\chi : G \to \mathbb{C}^\times$; the values are n th roots of unity, the characters form a group $\widehat{G}$ under pointwise multiplication, and $\widehat{G} \cong G$ (not canonically). Explicitly, by the structure theorem $G \cong \mathbb{Z}/d_1 \times \cdots \times \mathbb{Z}/d_r$, and the characters are $\chi_{\mathbf{a}}(\mathbf{x}) = \prod_t \exp(2\pi i a_t x_t / d_t)$. See Book 1, field: Abstract algebra, topic: characters of finite abelian groups.

The two orthogonality relations. $\sum_x \chi(x)\overline{\psi(x)} = n[\chi = \psi]$ and $\sum_\chi \chi(x)\overline{\chi(y)} = n[x = y]$. The first says the characters are an orthogonal basis of the space of functions on G; the second is the first applied to $\widehat{G}$. Both follow from the one line argument: if $\chi \neq \widehat{e}$ pick y with $\chi(y) \neq 1$; then $\chi(y)\sum_x \chi(x) = \sum_x \chi(yx) = \sum_x \chi(x)$, so the sum vanishes. See Book 1, field: Abstract algebra, topic: orthogonality relations.

Fourier analysis on a finite abelian group. Transform $\widehat{\varphi}(\chi) = \sum_x \varphi(x)\chi(x)$; inversion $\varphi(x) = \frac{1}{n}\sum_\chi \widehat{\varphi}(\chi)\overline{\chi(x)}$; Parseval $\sum_x |\varphi|^2 = \frac{1}{n}\sum_\chi |\widehat{\varphi}|^2$; convolution theorem $\widehat{\varphi * \psi} = \widehat{\varphi}\widehat{\psi}$. For $G = \mathbb{Z}/n$ this is the discrete Fourier transform and the corresponding matrices are the circulants. See Book 1, field: Discrete structures and graphs, topic: Fourier analysis on finite groups.

Kernels and generation. The set where a homomorphism takes the value 1 is a subgroup, so a character trivial on a generating set is trivial. Conversely, for any proper subgroup $H \subsetneqq G$ there is a nontrivial character trivial on H, obtained by lifting a nontrivial character of G/H. See Book 1, field: Abstract algebra, topic: quotients and lifting characters.

Equality in the triangle inequality. If $z_1, \ldots, z_k \in \mathbb{C}$ and $|z_1 + \cdots + z_k| = |z_1| + \cdots + |z_k|$, then all nonzero z_i have the same argument; if in addition all have modulus 1, they are all equal. This is the standard mechanism for turning an extremal modulus into a rigidity statement, and it is how aperiodicity is proved in every finite Markov chain argument. See Book 1, field: Complex numbers, topic: equality cases.

Markov chains, stationarity and the spectral gap. A doubly stochastic transition matrix has the uniform distribution as stationary; irreducibility follows from the steps generating the group, aperiodicity from a self loop. The convergence theorem then gives $P_m \to U$, and the rate is governed by the second largest eigenvalue modulus b, with mixing time of order $\log(n)/(1-b)$. See Book 1, field: Probability, topic: Markov chains and mixing.

The generalised recursion (this is the load bearing generalisation). The problem contains the recursion $P_m = P_{m-1} * P$, that is, $v_m = Tv_{m-1}$ with T linear on a finite dimensional space. The general solution: if T is diagonalisable with eigenpairs (μ_j, w_j) and $v_0 = \sum_j c_j w_j$, then $v_m = \sum_j c_j \mu_j^m w_j$. If $\mu_0 = 1$ is the unique eigenvalue of maximal modulus, then $v_m \to c_0 w_0$ and, with b the next largest modulus and J the set of indices attaining it,

$$\lim_{m\to\infty} \frac{\|v_m - c_0 w_0\|^2}{b^{2m}} = \sum_{j\in J} |c_j|^2 \|w_j\|^2$$

whenever the eigenvectors attaining b are pairwise orthogonal and the μ_j with $|\mu_j| = b$ are equal in modulus but the cross terms vanish, which holds here because distinct characters are orthogonal and the sum was expanded in an orthogonal basis. The same template solves every linear recurrence $a_m = c_1 a_{m-1} + \cdots + c_d a_{m-d}$ through its companion matrix, where the dominant root controls growth and ties among roots of equal modulus produce oscillating or period doubling corrections. See Book 1, field: Sequences and recurrences, topic: linear recursions by diagonalisation.

Circuit Bank

- **WHEN** a probability is described as the product of m independent group valued choices **DO** write the m fold convolution and take the Fourier transform on the group **BECAUSE** convolution becomes multiplication, so the m th step becomes an m th power, and powers have visible asymptotics while convolutions do not
- **WHEN** the quantity to evaluate is $\sum(\cdot)^2$ over a finite index set **DO** apply Parseval in the basis that diagonalises the relevant operator, never compute entries **BECAUSE** a squared norm is basis independent, so you may choose the basis in which the vector has explicit coordinates
- **WHEN** a hypothesis looks decorative, such as "the identity is among the generators" or a strict subset symbol **DO** find the exact inequality it produces before writing anything, and count unused hypotheses against unproved claims **BECAUSE** Putnam hypotheses are never decoration, and the pairing between spare hypotheses and missing inequalities tells you the shape of the proof in advance
- **WHEN** a sum of complex numbers of modulus one attains the maximal possible modulus **DO** conclude that all of them are equal **BECAUSE** equality in the triangle inequality forces a common argument, and equal modulus then forces equality; this is the standard aperiodicity mechanism
- **WHEN** you must produce an exact exponential rate rather than bounds **DO** expand in eigenvectors, take b to be the largest subdominant eigenvalue modulus, and divide **BECAUSE** terms below the maximum vanish geometrically and tied terms contribute exactly one each, so the limit exists and counts the maximisers

Backtesting

1. First useful cue: the product of m independent uniform picks is a convolution power, and the requested quantity is a squared L^2 distance to uniform, so the problem lives in the character basis.
2. Method the expression suggested: Fourier transform on G, convolution theorem to get λ_χ^m, Parseval to get $D_m = \frac{1}{n}\sum_{\chi\neq\hat{e}}|\lambda_\chi|^{2m}$, then b as the maximum.
3. Where you would hesitate, and the thought to reinforce: at proving $0 < b < 1$, which looks like the technical part but is the part the hypotheses were designed for. Reinforce the audit habit: two unused hypotheses, two strict inequalities, match them before you look for techniques.
4. The trap the question sets: computing $\text{Prob}(g = x)$ for individual x, or working with a transition matrix without checking whether it produces the distribution of g or of g^{-1}. The first burns the clock, the second is a real correctness issue that a careful grader can penalise.
5. What makes you recognise it faster next time: the word abelian in a probability problem is never about probability, it is naming the diagonalising basis; and any request for an exact exponential rate is a request for a subdominant eigenvalue.
6. The reusable decision rule: convolution plus abelian group equals characters; sum of squares equals Parseval; exact rate equals second eigenvalue.

Mirror drills

1. Let $G = \mathbb{Z}/n$ with $n \geq 3$ and $S = \{0, 1\}$. Compute the eigenvalues λ_t, the value of b, and the number of maximisers. *Answer:* $\lambda_t = (1 + \omega^t)/2$ with $\omega = e^{2\pi i/n}$, so $|\lambda_t| = |\cos(\pi t/n)|$; the maximum over $t \neq 0$ is $\cos(\pi/n)$, attained at $t = 1$ and $t = n - 1$, so $b = \cos(\pi/n)$ and $N = 2$.

2. Show that a character χ of a finite abelian group that is trivial on a generating set is trivial. *Answer:* the set $\{x : \chi(x) = 1\}$ is the kernel of χ, hence a subgroup; it contains a generating set, so it is the whole group.

3. Show that $|\lambda_\chi| = 1$ for some nontrivial χ if and only if the step set S is contained in a single coset of a proper subgroup of G. *Answer:* if $|\lambda_\chi| = 1$ then all $\chi(g_i)$ share a value z, so $S \subseteq \chi^{-1}(z)$, a coset of $\ker\chi \subsetneq G$; conversely if $S \subseteq cH$ with H proper, lift a nontrivial character of G/H to χ, and then χ is constant on S, so $|\lambda_\chi| = 1$. This is exactly aperiodicity, and $e \in S$ rules it out.

15.7 B1. squeezing a scaled root difference into a window

Let h and k be positive integers. Prove that for every $\epsilon > 0$, there are positive integers m and n such that

$$\epsilon < |h\sqrt{m} - k\sqrt{n}| < 2\epsilon.$$

Read it out loud

This is a request to hit a moving target. You are not asked to solve an equation, you are asked to land a specific real quantity, $h\sqrt{m} - k\sqrt{n}$, inside a window of width exactly ϵ, for every window you are handed. The variables you control are m and n, both positive integers; h, k, ϵ are fixed before you choose. So the fields wired into this are elementary number theory (density of the rationals, existence of integer solutions) and real analysis (an approximation argument), dressed up as an algebra problem about square roots.

The first cue: two objects of the same shape, $h\sqrt{m}$ and $k\sqrt{n}$, subtracted. Whenever you see a difference of two square roots, or scaled square roots, the reaction circuit is to rationalise: multiply and divide by the conjugate sum, because $\sqrt{X} - \sqrt{Y} = \frac{X-Y}{\sqrt{X}+\sqrt{Y}}$ converts an uncontrolled difference of two irrational, wildly varying quantities into a controlled quotient, where the numerator is an integer you choose directly. That is the guiding light: this problem is not really about square roots, it is about engineering a quotient to land in an interval, and the factor of exactly 2 in "ϵ to 2ϵ" is the examiners' signal that some slack is available, you do not have to hit ϵ on the nose.

The second cue, "for every $\epsilon > 0$", is the density-of-the-rationals cue: whenever a problem asks you to get arbitrarily close to a target using only integers, and the target can be phrased as a ratio, density of $\mathbb{Q}$ in $\mathbb{R}$ is the tool, once the problem has been reduced to choosing that ratio.

Within thirty seconds you should have written down "difference of scaled roots, rationalise", and be hunting for a way to force $h\sqrt{m}$ and $k\sqrt{n}$ to both be a fixed multiple of a square root, so that h and k cancel and you are left with a bare, one-parameter-family difference of roots.

Approaches, ranked worst to best

1. **Guess and check small cases.** *Cue:* a specific small ϵ, an existence claim. *Circuit:* unstructured pattern matching. (why it ranks here: it might locate a working pair (m, n) for one numerical ϵ, but it proves nothing about "every $\epsilon > 0$" and gives no method to run in the arena when ϵ is handed to you as a free variable; it costs unbounded ink and never terminates into a proof.)

2. **Fix n, sweep m, and use a vanishing step size.** *Cue:* an unbounded increasing sequence must eventually cross any fixed window. *Circuit:* a discrete intermediate value argument. (why it ranks here: fix any n and look at $x_m = h\sqrt{m} - k\sqrt{n}$ for $m = 1, 2, 3, \ldots$. This is strictly increasing to infinity, and consecutive gaps $x_{m+1} - x_m = h/(\sqrt{m+1} + \sqrt{m})$ shrink to 0, so eventually the sequence cannot jump over an interval of length ϵ without landing inside it. This works, but it forces you to separately guarantee that the gap is already below ϵ before the sequence reaches the window, which means choosing n large first, purely to buy room; that is a layer of bookkeeping the algebraic route below does not need, and it hands you an existence claim with no explicit formula for m, n.)

3. **Rationalise, force a common scale, then invoke density of the rationals.** *Cue:* difference of two scaled square roots, "for every ϵ". *Circuit:* conjugate rationalisation plus a scaling trick to patch a side condition. (the ideal solution: fully explicit formulas for m and n in terms of two auxiliary positive integers, no asymptotic bookkeeping, minimal casework.)

The ideal solution

Fix positive integers h, k and $\epsilon > 0$, and write $N = hk$.

Choose positive integers a, b with

$$\frac{3\epsilon}{N} < \frac{b}{a} < \frac{4\epsilon}{N}.$$

Such a, b exist because $\left(\frac{3\epsilon}{N}, \frac{4\epsilon}{N}\right)$ is a nonempty open interval of positive reals, of length $\epsilon/N > 0$, and the rationals are dense in the reals, so some rational lies in it; write that rational as b/a with a, b positive integers.

If $3a^2 \leq b$, replace the pair (a, b) by (ta, tb) for a positive integer t large enough that $3(ta)^2 > tb$, that is, $t > b/(3a^2)$. This replacement leaves the ratio b/a unchanged, so the displayed inequality above still holds, and now $3a^2 > b$. Relabel so a, b satisfy both the ratio bound and $3a^2 > b$.

Set

$$m = k^2(a^2 + b), \qquad n = h^2a^2.$$

Both are positive integers, and

$$h\sqrt{m} = hk\sqrt{a^2 + b}, \qquad k\sqrt{n} = hka = Na,$$

so

$$h\sqrt{m} - k\sqrt{n} = N\left(\sqrt{a^2 + b} - a\right).$$

Since $b > 0$, $\sqrt{a^2 + b} > a$, so this quantity is positive and equals its own absolute value. It remains to show

$$\frac{\epsilon}{N} < \sqrt{a^2 + b} - a < \frac{2\epsilon}{N}.$$

Rationalise:

$$\sqrt{a^2 + b} - a = \frac{b}{\sqrt{a^2 + b} + a}.$$

Bound the denominator on both sides using the two facts already in hand. Since $b > 0$, $\sqrt{a^2 + b} > a$, so $\sqrt{a^2 + b} + a > 2a$. Since $3a^2 > b$, $a^2 + b < 4a^2$, so $\sqrt{a^2 + b} < 2a$, giving $\sqrt{a^2 + b} + a < 3a$. Hence

$$\frac{b}{3a} < \sqrt{a^2 + b} - a < \frac{b}{2a}.$$

Substitute the ratio bound $\frac{3\epsilon}{N} < \frac{b}{a} < \frac{4\epsilon}{N}$:

$$\sqrt{a^2 + b} - a > \frac{b}{3a} > \frac{1}{3} \cdot \frac{3\epsilon}{N} = \frac{\epsilon}{N}, \qquad \sqrt{a^2 + b} - a < \frac{b}{2a} < \frac{1}{2} \cdot \frac{4\epsilon}{N} = \frac{2\epsilon}{N}.$$

This is exactly the sandwich required, so

$$\epsilon < N(\sqrt{a^2 + b} - a) = h\sqrt{m} - k\sqrt{n} = |h\sqrt{m} - k\sqrt{n}| < 2\epsilon,$$

with $m = k^2(a^2 + b)$ and $n = h^2a^2$ positive integers, as required.

Layer 1: WHAT

Block 1: see a difference of two scaled square roots and reach for the rationalising identity before touching anything else. Block 2: force $h\sqrt{m}$ and $k\sqrt{n}$ to both carry the factor $N = hk$, by choosing m, n as k^2 and h^2 times other squares, so the coefficients cancel and only a bare $\sqrt{a^2 + b} - a$ survives. Block 3: rationalise that bare difference into $b/(\sqrt{a^2 + b} + a)$, converting a two-variable geometric statement into a one-variable ratio statement. Block 4: sandwich the denominator using $\sqrt{a^2 + b} > a$ (always true) and $\sqrt{a^2 + b} < 2a$ (true once $3a^2 > b$), turning the quotient into the range $b/(3a)$ to $b/(2a)$. Block 5: choose a, b by density of the rationals so b/a lands in the interval $(3\epsilon/N, 4\epsilon/N)$, an interval reverse engineered from Block 4 to land exactly on $(\epsilon/N, 2\epsilon/N)$. Block 6: patch the side condition $3a^2 > b$ by scaling $(a, b) \mapsto (ta, tb)$, which preserves the ratio.

Juxtapose the nominal route against this one. The nominal route (Approach 2) never leaves the two-variable world: it tracks an increasing sequence x_m and argues by vanishing step size that it cannot leap over a window, which then needs a separate argument for why the sequence is already fine grained by the time it reaches the window. The bifurcation point is Block 2: the ideal route immediately asks whether the coefficients h, k can be forced to cancel by a clever choice of m, n, collapsing the problem to one free ratio b/a; the nominal route leaves h, k inside the sequence and pays for it later in bookkeeping about how large n must be chosen first.

Layer 2: HOW

Block 1. **Cue.** A difference of two scaled square roots, $h\sqrt{m} - k\sqrt{n}$. **Why natural.** Whenever an expression is a difference of square roots, its size is not controlled directly, since $\sqrt{}$ is not linear; the standard reroute is to multiply by the conjugate sum, which turns the difference into a quotient whose numerator is a plain integer difference and is therefore fully under your control. **Link to the next block.** Before rationalising anything, first ask whether the coefficients h, k can be absorbed so the two radicands share a common structure.

Block 2. **Cue.** The presence of two independent coefficients h, k sitting outside two independent radicands m, n. **Why natural.** Carrying four independent symbols through an inequality argument costs more ink than carrying two. Setting $n = h^2a^2$ makes $k\sqrt{n} = kha$ a bare product with no leftover root, and setting $m = k^2(a^2 + b)$ makes $h\sqrt{m} = hk\sqrt{a^2 + b}$; both choices are forced by wanting the coefficient in front of the remaining square root to be the same combined constant $N = hk$ in both terms. **Link to the next block.** With h, k cancelled, the entire remaining problem is a statement about the single expression $\sqrt{a^2 + b} - a$, which is exactly the shape rationalising was built for.

Block 3. **Cue.** The now-bare expression $\sqrt{a^2 + b} - a$. **Why natural.** This is precisely $\sqrt{X} - \sqrt{Y}$ with $X = a^2 + b$, $Y = a^2$, so the identity $\sqrt{X} - \sqrt{Y} = (X - Y)/(\sqrt{X} + \sqrt{Y})$ applies directly, and $X - Y = b$, an integer you chose yourself. **Link to the next block.** A quotient is easy to bound once its denominator is sandwiched, so the next task is pinning $\sqrt{a^2 + b} + a$ between two explicit multiples of a.

Block 4. **Cue.** The denominator $\sqrt{a^2 + b} + a$, needing bounds on both sides. **Why natural.** The lower bound $\sqrt{a^2 + b} > a$ is automatic since $b > 0$; the upper bound $\sqrt{a^2 + b} < 2a$ is not automatic, it needs $a^2 + b < 4a^2$, that is $3a^2 > b$, an inequality you have not yet arranged. Naming this missing inequality explicitly, rather than assuming it, is what lets you fix it cleanly one step later. **Link to the next block.** Once the denominator is trapped between $2a$ and $3a$, the whole quotient $b/(\sqrt{a^2 + b} + a)$ is trapped between $b/(3a)$ and $b/(2a)$, so the target sandwich on $\sqrt{a^2 + b} - a$ has become a target sandwich on the single ratio b/a.

Block 5. **Cue.** A target interval for the ratio b/a, and the freedom to choose positive integers a, b. **Why natural.** Density of the rationals in the reals says every nonempty open interval of reals contains a rational; reading the target sandwich backwards through the factors of 3 and 2 from Block 4 pins down exactly which interval, $(3\epsilon/N, 4\epsilon/N)$, the ratio b/a must land in. **Link to the next block.** The only loose end is the side condition $3a^2 > b$ assumed in Block 4, which the chosen a, b may not yet satisfy, and which scaling both by the same integer t repairs without disturbing the ratio.

Layer 3: WHY

The truth underneath is that density of $\mathbb{Q}$ in $\mathbb{R}$ is the only tool that turns "for every $\epsilon > 0$" into "there exist integers"; everything else in the solution exists to reduce the original two-variable geometric statement down to a bare statement about a ratio of two free integers, because only a ratio can be placed anywhere by density. Rationalising is the reroute that makes this reduction possible: a raw difference of roots is not proportional to anything simple, but a quotient of the form $b/(\text{something near } 2a)$ is a bounded multiple of b/a, and bounded multiples of a freely placeable ratio are themselves freely placeable, up to the known bound.

The discernment mechanism: whenever a problem says "for every ϵ, find integers with an expression trapped in a window", check first whether the expression can be rewritten as a bounded multiple of a ratio of two free integer parameters. If it can, density of the rationals finishes the problem in two lines; if it cannot, you are in a genuinely harder approximation problem and may need the fuller machinery of Diophantine approximation. Train this discernment by practising the rationalising reroute on sight, every time a difference of roots appears, rather than reaching for it only after other routes stall.

The examiners nudged you toward this exact construction by choosing the window as ϵ to 2ϵ, a clean factor of two, not ϵ to $\epsilon + \delta$ for some fussy separate δ. That factor of two is exactly the slack a $2a$-to-$3a$ denominator sandwich produces on its own, which is a strong hint, once you notice it, that the intended route runs through exactly this kind of rationalise-then-sandwich construction rather than through a finer analytic estimate.

Revolving concepts

Density of $\mathbb{Q}$ in $\mathbb{R}$: between any two reals $x < y$ there is a rational, provable via the Archimedean property by choosing a denominator q with $1/q < y - x$ and rounding qx up to the nearest integer. See Book 1, field: Real analysis, topic: density and the Archimedean property.

Rationalising conjugates: $\sqrt{X} - \sqrt{Y} = (X - Y)/(\sqrt{X} + \sqrt{Y})$, the standard device for converting an additive difference of roots into a quotient; the same identity is the reason $\sqrt{n+1} - \sqrt{n} \to 0$ and underlies estimates like the one this year's A3 needs for x^r against x^{r+1} type ratios. See Book 1, field: Sequences and recurrences, topic: rationalising estimates.

Scaling a pair of integers defined only through their ratio, to force a side inequality: whenever a, b are only constrained by their ratio b/a, you may replace (a, b) by (ta, tb) for any positive integer t without disturbing the ratio; this is a free extra degree of freedom, useful whenever one side of a needed inequality grows faster in t than the other (here $3a^2$, quadratic in t, eventually beats b, linear in t). See Book 1, field: Number theory, topic: scaling arguments in Diophantine constructions.

Diophantine approximation, as context beyond what this problem needs: the general theory of approximating reals by rationals, including Dirichlet's approximation theorem, that for any real α and positive integer N there exist integers p, q with $1 \le q \le N$ and $|\alpha - p/q| < 1/(qN)$. This problem is a miniature instance of the same theme, trapping a constructed real inside a window using integer choices, though it needs only bare density, not the sharper Dirichlet bound.

Circuit Bank

- **WHEN** you see a difference of two scaled or shifted square roots **DO** multiply and divide by the conjugate sum immediately, before attempting any other manipulation **BECAUSE** it converts an unruly additive difference of irrationals into a quotient whose numerator is an integer you can choose directly
- **WHEN** a problem hands you free integer parameters that only interact through their ratio **DO** look for a target interval for that ratio and invoke density of the rationals to place it there

BECAUSE density of $\mathbb{Q}$ in $\mathbb{R}$ is the tool that turns "arbitrarily close" into "there exist integers", at no cost

- **WHEN** a construction needs a side inequality between two auxiliary integers that were only defined by a ratio **DO** scale both by a large positive integer t **BECAUSE** scaling preserves the ratio exactly while a term linear in t is eventually beaten by a term quadratic in t, in whichever direction the side inequality needs

Backtesting

1. First useful cue: a difference of two scaled square roots, $h\sqrt{m} - k\sqrt{n}$, reading as "rationalise me".
2. Method the expression suggested: force the coefficients h, k to cancel by choosing m, n as perfect-square multiples, then rationalise the remaining bare difference of square roots into a quotient.
3. Where you would hesitate, and the thought to reinforce: after rationalising you have $b/(\sqrt{a^2+b}+a)$, and it is tempting to approximate the denominator as roughly $2a$ without proving a clean two-sided bound; the thought to reinforce is that "roughly" is not a proof, so pin the denominator explicitly between $2a$ and $3a$ using a side condition you can always arrange by scaling.
4. The trap the question sets: the window is exactly ϵ to 2ϵ, a factor of two, which is precisely the slack the $2a$-to-$3a$ sandwich produces; a tighter window would force a finer sandwich, so read the given factor of two as a gift rather than an arbitrary choice of constants.
5. What makes you recognise it faster next time: any "for every ϵ, trap this integer-built expression in a window" problem is a density-of-the-rationals problem in disguise, the moment the expression can be rewritten as a bounded multiple of a free ratio.
6. The reusable decision rule: difference of roots, rationalise first; a free ratio of integer parameters, invoke density; a side inequality on auxiliary integers, fix it by scaling.

Mirror drills

1. Prove that for every $\epsilon > 0$ there are positive integers m, n with $0 < \sqrt{m} - \sqrt{n} < \epsilon$. *Answer:* take $n = a^2$, $m = a^2 + 1$ for a large enough that $1/(2a) < \epsilon$, then rationalise exactly as above.
2. Prove that $\{\sqrt{m} - \sqrt{n} : m, n \in \mathbb{Z}^+\}$ is dense in $[0, \infty)$. *Answer:* the same rationalise-and-scale construction, run against a target interval $(\epsilon, 2\epsilon)$ centred anywhere, not only near 0.
3. For fixed positive integers h, k, show there are infinitely many pairs (m, n) with $0 < h\sqrt{m} - k\sqrt{n} < 1$. *Answer:* run this problem's construction with $\epsilon = 1/2^j$ for $j = 1, 2, 3, \ldots$, producing infinitely many distinct pairs since the trapped value shrinks each time.

15.8 B2. why two and five dominate this set

Let S be the set of all ordered triples (p, q, r) of prime numbers for which at least one rational number x satisfies $px^2 + qx + r = 0$. Which primes appear in seven or more elements of S?

Read it out loud

This is a counting question wearing a quadratic-equation costume. The surface object is "a quadratic with prime coefficients has a rational root"; the real question, once that is stripped away, is "how many ordered triples of primes can a single prime sit inside". The phrase "at least one rational number x satisfies" is the command term: for a quadratic with integer coefficients, a rational root exists exactly when the discriminant is a perfect square, since an integer discriminant has a rational square root only when it has an integer one. That translation is the first thing you write down, inside thirty seconds: $(p, q, r) \in S \iff q^2 - 4pr = a^2$ for some integer $a \geq 0$.

The fields wired in are elementary number theory, parity and divisibility arguments, residues modulo a small number, a difference-of-squares factorisation, dressed as algebra. The word "prime" attached to all three of p, q, r, rather than "integer", is a second guiding light: it means every divisor-pair argument you run terminates in at most two or three cases, because a prime has almost no divisors beyond 1 and itself. Whenever a Putnam problem hands you primes as the objects being counted or factored, the reaction circuit is to force a small, explicit, finite structure, a factor pair or a residue class, out of that primality, rather than treating p, q, r as generic integers and hoping an inequality closes the argument.

The third cue is the number "seven": the question is not asking for the maximum, it is asking you to find a threshold and show only certain primes clear it. "Seven or more" signals a two-part proof: an explicit construction reaching the threshold for the primes that do, and a uniform upper bound ruling out every other prime, with the argument built so the bound is tight, not generous.

Approaches, ranked worst to best

1. **Numerically search small primes.** *Cue:* "which primes appear", a finite-looking search. *Circuit:* unstructured trial. (why it ranks here: you can stumble onto the fact that $(2, 5, 2)$ and similar triples work, but primes are unbounded, so a search alone cannot prove that no other prime ever reaches seven appearances; it produces examples, never a bound, and the problem needs both.)

2. **Treat p, q, r as generic integers and search the discriminant condition directly.** *Cue:* $q^2 - 4pr = a^2$. *Circuit:* brute algebraic manipulation that ignores primality. (why it ranks here: without forcing the primality of p, q, r into a small divisor structure early, you re-solve a two-variable Diophantine condition from scratch for every candidate prime, at high cost in ink, and you will not notice the forced appearance of 2 without the residue argument that primality of q and oddness of the discriminant actually supplies.)

3. **Force primality into a factor pair via a modulus-8 parity argument and a difference-of-squares factorisation, then pigeonhole modulo 3 to cap large primes.** *Cue:* primes, discriminant a perfect square, "seven or more". *Circuit:* parity forcing, a finite divisor-pair count, and a residue argument to bound every other candidate. (the ideal solution: converts an apparently unbounded counting problem into checking the primality of four explicit small numbers per candidate prime.)

The ideal solution

Reformulating membership in S. For primes p, q, r, the quadratic $px^2 + qx + r = 0$ has a rational root exactly when its discriminant $q^2 - 4pr$ is the square of a rational number. Since p, q, r are

integers, $q^2 - 4pr$ is an integer, and an integer is the square of a rational number exactly when it is the square of an integer. So

$$(p,q,r) \in S \iff q^2 - 4pr = a^2 \text{ for some integer } a \geq 0.$$

Forcing $q \geq 5$ and a 2 into every triple. Since $p, r \geq 2$, $4pr \geq 16$, so $q^2 = a^2 + 4pr \geq 16$, giving $q \geq 4$; as q is prime, $q \geq 5$, so q is odd. Then $a^2 = q^2 - 4pr$ is odd, so a is odd too. Odd squares are 1 modulo 8, so

$$4pr = q^2 - a^2 \equiv 1 - 1 = 0 \pmod 8,$$

hence $pr \equiv 0 \pmod 2$, so $2 \in \{p, r\}$.

Solving for q when $r = 2$. Suppose $r = 2$. Then $8p = q^2 - a^2 = (q-a)(q+a)$. Both factors are positive, since $a < q$ (as $a^2 = q^2 - 4pr < q^2$), and both are even, since q, a are both odd. Write $q - a = 2s$, $q + a = 2t$ with $0 < s < t$ integers, strict since $a \geq 1$; then $st = 2p$ and $s + t = q$. Since p is prime, the factor pairs of $2p$ with $s < t$ are $(1, 2p)$ and, when $p > 2$, $(2, p)$, giving

$$q = 1 + 2p = 2p + 1 \quad \text{or} \quad q = 2 + p = p + 2.$$

(When $p = 2$, only $(1, 4)$ has $s < t$, giving $q = 5$, consistent with $q = 2p + 1$.)

By the identical argument with the roles of p, r exchanged, since $q^2 - 4pr$ is symmetric in p, r, if instead $p = 2$, then

$$q = 2r + 1 \quad \text{or} \quad q = r + 2.$$

Counting appearances of a fixed prime $\ell \neq 2$. Every triple in S has $2 \in \{p, r\}$. Fix a prime $\ell \neq 2$ and count triples in which ℓ appears as p, q, or r.

If $r = 2$: ℓ can appear as p, giving $(\ell, q, 2)$ with $q \in \{2\ell + 1, \ell + 2\}$, contributing one triple for each of $2\ell + 1, \ell + 2$ that is prime; or ℓ can appear as q, giving $q = \ell \in \{2p + 1, p + 2\}$, so $p = (\ell - 1)/2$ or $p = \ell - 2$, contributing one triple for each of $(\ell - 1)/2, \ell - 2$ that is prime, since $q = \ell$ is automatically prime. So the branch $r = 2$ contributes at most one triple for each prime appearing in the list

$$2\ell + 1, \quad \ell + 2, \quad \frac{\ell - 1}{2}, \quad \ell - 2.$$

By the mirrored argument, the branch $p = 2$ contributes at most one further triple for each prime in this same list. So ℓ appears in at most twice as many elements of S as there are primes in this four-number list, hence at most 8, and strictly fewer whenever one of the four fails to be prime.

Ruling out every prime $\ell \geq 7$. The three numbers $\ell - 2, \ell, \ell + 2$ form an arithmetic progression with common difference 2, and $\gcd(2, 3) = 1$, so they occupy all three residues modulo 3 between them; exactly one is divisible by 3. Since $\ell \geq 7$ is prime, ℓ itself is not divisible by 3, so the multiple of 3 among the three is $\ell - 2$ or $\ell + 2$; both exceed 3 (as $\ell \geq 7$), so that one is composite. Hence at least one of the four list entries above, $\ell - 2$ or $\ell + 2$, is not prime, so at most three of the four are prime, and ℓ appears in at most $2 \times 3 = 6 < 7$ elements of S.

Ruling out $\ell = 3$. The list becomes $2(3) + 1 = 7$ (prime), $3 + 2 = 5$ (prime), $(3 - 1)/2 = 1$ (not prime), $3 - 2 = 1$ (not prime). Only two of the four are prime, so 3 appears in at most $2 \times 2 = 4 < 7$ elements of S.

The prime 2. The bound just proved assumed $\ell \neq 2$; it does not apply to $\ell = 2$, since 2 is the distinguished prime that appears in every branch of the case split, $r = 2$ or $p = 2$, not merely as an occasional coincidence produced by the four-number list. So 2 is not restricted by this counting argument at all, and its appearances, like 5's, must be confirmed directly by example.

Explicit examples reaching seven. The seven triples

$$(2,5,2),\ (2,5,3),\ (2,7,5),\ (2,11,5),\ (3,5,2),\ (5,7,2),\ (5,11,2)$$

each lie in S: their discriminants $q^2 - 4pr$ are $25 - 16 = 9$, $25 - 24 = 1$, $49 - 40 = 9$, $121 - 40 = 81$, $25 - 24 = 1$, $49 - 40 = 9$, $121 - 40 = 81$ respectively, all perfect squares $(3^2, 1^2, 3^2, 9^2, 1^2, 3^2, 9^2)$.

Every one of these seven triples contains a 2, so 2 appears in at least seven elements of S; every one of these seven triples also contains a 5, as p, q, or r, so 5 appears in at least seven elements of S too.

Answer. The primes appearing in seven or more elements of S are exactly 2 and 5.

Layer 1: WHAT

Block 1: translate "has a rational root" into "discriminant is a perfect square", the standard quadratic-formula cue. Block 2: use $p, r \geq 2$ to force $q \geq 5$, hence q odd, hence, via odd squares modulo 8, force a factor of 2 into $\{p, r\}$. Block 3: turn $8p = (q-a)(q+a)$, or its mirror $8r = (q-a)(q+a)$, into an explicit finite list of possible q, using the divisor pairs of $2p$ and that p is prime to limit those pairs to at most two. Block 4: reframe "how many triples contain ℓ" as "how many of four explicit numbers built from ℓ are prime", collapsing an unbounded search into a four-item checklist. Block 5: use a pigeonhole on residues modulo 3 over $\ell-2, \ell, \ell+2$ to show every prime $\ell \geq 7$ fails at least one item on the checklist, capping its count at six. Block 6: check $\ell = 3$ directly, which fails two of the four items, and exhibit seven explicit triples confirming both $\ell = 2$ and $\ell = 5$ reach seven.

Juxtapose the nominal route against this one. Approach 2 treats p, q, r as generic integers and tries to bound $q^2 - 4pr = a^2$ by raw size or search; it never notices that primality of p (or r) collapses an infinite divisor search into at most two cases, because it never runs the difference-of-squares factorisation with a prime specifically in view. The bifurcation point is Block 3: the moment you factor $8p = (q-a)(q+a)$ and ask how many ways a prime can split as a product of two positive integers, the whole problem becomes finite and explicit; skip that step, and you are searching an unbounded two-variable lattice by hand.

Layer 2: HOW

Block 1. **Cue.** "At least one rational number x satisfies $px^2 + qx + r = 0$." **Why natural.** A rational root of an integer quadratic exists exactly when the quadratic formula's square root term is rational, and the square root of an integer is rational only when that integer is a perfect square; this is a fixed reaction, not a case-by-case check. **Link to the next block.** With S rewritten as a perfect-square condition on $q^2 - 4pr$, the next move is to exploit that p, q, r are specifically primes, not arbitrary integers.

Block 2. **Cue.** Three prime variables p, q, r, and the size bound $4pr \geq 16$ forced by $p, r \geq 2$. **Why natural.** A crude size bound on q is the cheapest way to extract a first piece of structure before any casework; it costs one line and immediately restricts q to being odd, since the only even prime, 2, is now excluded. **Link to the next block.** Oddness of q, and hence of a, invites a residue argument, since odd squares have a single fixed residue modulo 8.

Block 3. **Cue.** Two odd squares q^2, a^2, both congruent to 1 modulo 8. **Why natural.** Subtracting forces $4pr \equiv 0 \pmod 8$, that is pr even, and since p, r are primes this can only mean one of them literally equals 2; a residue argument on a difference of two odd squares is the standard route to planting a factor of 2 into a product of primes. **Link to the next block.** With $2 \in \{p, r\}$ established, the discriminant equation for the case $r = 2$ becomes $8p = q^2 - a^2$, a genuine difference-of-squares factorisation with the prime p isolated on one side.

Block 4. **Cue.** $8p = (q-a)(q+a)$, with p prime. **Why natural.** A prime has only two positive factor pairs, $1 \times p$ and (when $p \neq 2$) $2 \times (p/1)$ scaled appropriately; running that through $st = 2p$ with $s < t$ produces exactly two candidate values of q, no more, because there is nothing else $2p$ can factor as. **Link to the next block.** Because the same relation holds with p, r exchanged by symmetry, every triple in S has its q pinned down by one of two linear formulas in whichever of p, r equals the non-2 prime, and this is the moment to ask how a fixed prime ℓ can appear across all such triples.

Block 5. **Cue.** A fixed prime ℓ, appearing either as the non-2 prime or as q itself, across two mirrored branches. **Why natural.** Each role ℓ could play corresponds to one of the four numbers $2\ell+1, \ell+2, (\ell-1)/2, \ell-2$ being prime, and the two branches ($p=2$ versus $r=2$) mirror each other exactly, so the count of triples containing ℓ is bounded by twice the count of primes in this fixed, explicit list. **Link to the next block.** Turning an unbounded question about appearances of ℓ into a bounded question about primality of four fixed numbers is what makes a pigeonhole argument on ℓ itself possible.

Block 6. **Cue.** Three numbers $\ell-2, \ell, \ell+2$ in arithmetic progression with common difference 2. **Why natural.** Since $\gcd(2,3)=1$, these three numbers hit every residue class modulo 3 exactly once, so one of them is always a multiple of 3; once $\ell \geq 7$, that multiple cannot be ℓ itself (which is prime), so it must be $\ell-2$ or $\ell+2$, and being a multiple of 3 exceeding 3, it is composite. **Link to the next block.** This composite entry knocks the four-item checklist down to at most three primes for every $\ell \geq 7$, capping appearances at six and leaving only $\ell=2$ and $\ell=5$ as candidates once $\ell=3$ is checked by hand and $\ell=2$ is set aside as exempt from the bound.

Layer 3: WHY

The mathematical truth underneath is that a prime has exactly the divisor pairs $1\times p$ and (for $p\neq 2$) its own analogue, nothing else; every restriction in this solution that looks engineered is really this one fact, applied three times over, once to bound q by raw size, once to force a 2 by parity, once to force the finite q-list by the divisor pairs of $2p$.

The discernment mechanism: whenever a Putnam problem asks you to count or classify objects built from primes, check first whether a size bound, a parity or residue bound, or a divisor-pair count can shrink an infinite search to a finite, explicit one, before you touch any inequality chasing. Train that discernment by asking, every time "prime" appears as a hypothesis rather than merely "integer", exactly what divisor structure that hypothesis is silently handing you.

The examiners nudged you toward the residue step by setting the threshold at "seven", a number just above the maximum any prime other than 2 and 5 can reach, namely six, which only becomes visible once you have built the four-item checklist. That is the signature of a bound meant to be tight rather than generous, and tight bounds of this kind are almost always proved by a residue or parity argument, not by a size estimate; a size estimate alone would leave you unable to explain why the threshold is exactly seven rather than, say, ten.

Revolving concepts

Rational root of an integer quadratic if and only if its discriminant is a perfect square: a direct reading of the quadratic formula, combined with the fact that the square root of an integer is rational only when it is already an integer. See Book 1, field: Polynomials, topic: discriminants and rational roots.

Quadratic residues modulo 8: every odd square is congruent to 1 modulo 8; a standard, load-bearing fact whenever a difference or sum of two odd squares appears. See Book 1, field: Number theory, topic: quadratic residues modulo small moduli.

Difference-of-squares factorisation combined with divisor-pair counting against a known, largely prime, factorisation: the standard route from "a number is a difference of two squares" to an explicit finite list of representations. See Book 1, field: Number theory, topic: difference of squares and divisor pairs.

Pigeonhole on residues modulo a small prime, applied to an arithmetic progression whose common difference is coprime to that modulus: guarantees exactly one term of any full residue cycle is divisible by it. See Book 1, field: Number theory, topic: pigeonhole on residue classes.

Circuit Bank

- **WHEN** a quadratic with integer coefficients is said to have a rational root **DO** rewrite the condition immediately as "the discriminant is a perfect square" **BECAUSE** the quadratic formula makes the root rational exactly when the square root term is rational, and an integer's square root is rational only when it is already an integer
- **WHEN** a Diophantine condition involves primes specifically, not generic integers **DO** look for a size bound, a parity or residue bound, or a divisor-pair count that the primality forces **BECAUSE** a prime's only factor pairs are 1 and itself, which collapses an unbounded search the instant you factor something as a product involving that prime
- **WHEN** you must bound how often a parameter ℓ can appear across a family of constructions built from small linear translates of ℓ **DO** reduce to checking primality of a short explicit list of numbers derived from ℓ, then use a residue pigeonhole on a small modulus to show large ℓ cannot pass every check **BECAUSE** an arithmetic progression with common difference coprime to a small modulus always hits every residue class, forcing one entry to become composite once ℓ is large enough

Backtesting

1. First useful cue: "at least one rational number x satisfies", translating directly to a perfect-square discriminant.
2. Method the expression suggested: force primality of p, q, r into a small explicit structure, first via a size bound on q, then via a parity argument planting a 2 into every triple.
3. Where you would hesitate, and the thought to reinforce: after the modulus-8 step gives $2 \in \{p, r\}$, it is tempting to handle $p = 2$ and $r = 2$ as one case by symmetry without saying so; the thought to reinforce is to name the symmetry explicitly, since $q^2 - 4pr$ is unchanged under swapping p, r, so the grader sees both branches were actually handled rather than assumed.
4. The trap the question sets: the threshold "seven" looks arbitrary until you have derived that every prime other than 2 and 5 tops out at six; do not hunt for the threshold's significance before finishing the general bound, or you will chase a pattern that only appears at the very end.
5. What makes you recognise it faster next time: any Putnam question counting or bounding "how many configurations of primes satisfy an algebraic condition" very likely routes through a perfect-square-discriminant translation plus a small-modulus parity or pigeonhole argument, because primes carry almost no divisor structure beyond size and residues.
6. The reusable decision rule: a quadratic with a rational root, translate to a perfect-square discriminant; primes doing the structural work, hunt for a forced parity or divisor-pair collapse; a numeric threshold in the question, prove the general bound first and let the threshold's significance fall out at the end.

Mirror drills

1. For which primes p does $x^2 + px + 1 = 0$ have a rational root? *Answer:* never; the discriminant $p^2 - 4$ would need to equal a^2 with $(p - a)(p + a) = 4$, forcing $p \leq 3$, and direct checking rules out $p = 2, 3$.
2. How many ordered pairs of primes (p, q) satisfy $q^2 - 4p$ a perfect square and $q < 20$? *Answer:* run the same $8p = (q - a)(q + a)$-style factorisation and list the finitely many candidates by hand for $q < 20$.

3. Fix a prime $\ell \geq 7$; show directly, without reference to S, that $\ell - 2, \ell, \ell + 2$ cannot all be prime. *Answer:* the same modulus-3 pigeonhole: one of the three is always a multiple of 3 exceeding 3.

15.9 B3. recovering differentiability from a product and quotient

Let f and g be real valued functions defined on an open interval containing 0, with g nonzero and continuous at 0. If fg and f/g are differentiable at 0, must f be differentiable at 0?

Read it out loud

This is a "must it follow" question, which folds two questions into one sentence: does the implication hold, and why does it need both hypotheses, fg differentiable and f/g differentiable, rather than either alone. Notice immediately that g carries the weakest hypothesis on the list, merely continuous and nonzero at 0, not differentiable; that is the guiding light. If g were assumed differentiable, this would be a non-problem: the product rule and quotient rule would hand you $f = (fg)/g$ differentiable instantly. The entire content of the problem is squeezed into the gap between "continuous" and "differentiable" for g.

The second cue is structural: you are handed a product fg and a quotient f/g of the same two functions. Whenever two given quantities are a product and a quotient of the same underlying objects, the reaction circuit is to multiply them, because the "other" factor cancels: $(fg) \cdot (f/g) = f^2$. That single line eliminates g entirely and produces a new differentiable function, f^2, built only from the two hypotheses you were actually given.

The third thing to notice, before writing a single further line, is that recovering f from f^2 is not free. $f = \pm\sqrt{f^2}$, and the square root function is differentiable at a value only when that value is positive, not at 0. So the moment you write $f^2 = (fg)(f/g)$, you have located the fault line of the whole problem: what happens when $f(0) = 0$, exactly where the square root loses differentiability. A full-marks solution announces this fault line before working through the cases, rather than discovering it midway and patching around it.

A caveat worth clarifying up front: g nonzero on the interval, together with continuity at 0, is what lets f/g even be defined near 0 in the first place; that housekeeping fact is used silently everywhere below and is worth stating once, at the start, rather than re-justifying each time.

Approaches, ranked worst to best

1. **Formally differentiate $f = \sqrt{(fg)(f/g)}$ without checking $f(0)$.** *Cue:* f looks like the square root of a differentiable product. *Circuit:* blind chain-rule differentiation of a composite. (why it ranks here: the resulting formula $f'(0) = \dfrac{(fg)'(0)(f/g)(0) + (fg)(0)(f/g)'(0)}{2f(0)}$ divides by $2f(0)$, silently assuming $f(0) \neq 0$; it produces a correct-looking answer that is simply undefined exactly in the case the problem is really testing, and a grader withholds full marks for the unhandled case even when the rest of the algebra is right.)

2. **Chase the definitions of $(fg)'(0)$ and $(f/g)'(0)$ directly, with no algebraic reduction first.** *Cue:* differentiability is ultimately a limit statement. *Circuit:* unstructured limit manipulation. (why it ranks here: it can be made to work, since it re-derives essentially the same combination of difference quotients as the ideal route, but without first isolating $f^2 = (fg)(f/g)$ or an equivalent identity, you are juggling three unknown limits at once with no organising principle, at real risk of an unjustified step or a missed case.)

3. **Multiply to get $f^2 = (fg)(f/g)$, split on whether $f(0)$ is zero, and handle each case on its own terms.** *Cue:* a product and a quotient of the same two functions. *Circuit:* cancellation, followed by a case split forced exactly where the square root loses differentiability. (the ideal solution: motivated at every step, and the case split is not a defect, it is the precise location the problem is testing.)

The ideal solution

Let I be the open interval containing 0 on which f, g are defined. Since g is continuous and nonzero at 0, there is a subinterval around 0 on which $g \neq 0$, so f/g is defined there; shrink I to that subinterval.

Step 1: f is continuous at 0. Since f/g is differentiable at 0, it is continuous at 0; since g is continuous at 0 by hypothesis, the product

$$f = (f/g) \cdot g$$

is continuous at 0, being a product of two functions continuous at 0.

Step 2: eliminate g. Since fg and f/g are both differentiable at 0, and a product of two functions differentiable at 0 is differentiable at 0,

$$(fg) \cdot (f/g) = f^2$$

is differentiable at 0.

Step 3: the case $f(0) \neq 0$. By Step 1, f is continuous at 0, so f has constant sign on some neighbourhood of 0, matching the sign of $f(0)$. Let $\delta = 1$ if $f(0) > 0$ and $\delta = -1$ if $f(0) < 0$; then $\delta f > 0$ on that neighbourhood, and there

$$\delta f = \sqrt{f^2},$$

since δf is the positive square root of $(\delta f)^2 = f^2$ whenever $\delta f > 0$. The function f^2 is differentiable at 0 (Step 2) and takes the positive value $f(0)^2$ there, and $\sqrt{\ }$ is differentiable at every positive real. By the chain rule, $\sqrt{f^2}$ is differentiable at 0, hence so is δf, hence so is $f = \delta \cdot (\delta f)$.

Step 4: the case $f(0) = 0$. Then $(f/g)(0) = f(0)/g(0) = 0$, so

$$(f/g)'(0) = \lim_{x \to 0} \frac{(f/g)(x) - (f/g)(0)}{x} = \lim_{x \to 0} \frac{f(x)}{xg(x)}.$$

This limit exists and is finite, since f/g is differentiable at 0. Since g is continuous at 0 with $g(0) \neq 0$, $g(x) \to g(0) \neq 0$ as $x \to 0$, so

$$\frac{f(x)}{x} = \frac{f(x)}{xg(x)} \cdot g(x) \longrightarrow (f/g)'(0) \cdot g(0)$$

as $x \to 0$, a finite limit, being a product of two convergent limits. Since $f(0) = 0$,

$$f'(0) = \lim_{x \to 0} \frac{f(x) - f(0)}{x} = \lim_{x \to 0} \frac{f(x)}{x}$$

exists, so f is differentiable at 0.

Assembling the two cases. In both cases f is differentiable at 0, so yes, f must be differentiable at 0.

Remark, a route avoiding the case split. The split in Steps 3 and 4 can be avoided by eliminating $g(x) - g(0)$ directly, rather than eliminating g multiplicatively. Write $u = fg$, $v = f/g$, and expand

$$u(x) - u(0) = (f(x) - f(0))g(x) + f(0)(g(x) - g(0)),$$

$$g(0)g(x)\big(v(x) - v(0)\big) = (f(x) - f(0))g(0) - f(0)(g(x) - g(0)).$$

Adding cancels the two $f(0)(g(x) - g(0))$ terms, since they enter with opposite sign, leaving

$$(f(x) - f(0))\big(g(x) + g(0)\big) = \big(u(x) - u(0)\big) + g(0)g(x)\big(v(x) - v(0)\big).$$

Since $g(x) + g(0) \to 2g(0) \neq 0$, dividing by $x\big(g(x) + g(0)\big)$ and letting $x \to 0$ shows $\lim_{x\to 0} \frac{f(x)-f(0)}{x}$ exists directly, in one line, with no reference to the sign of $f(0)$, because every term on the right converges: u, v are differentiable at 0, and g is continuous at 0. This is the same fact, proved uniformly instead of by cases. It is worth keeping as a second circuit, because eliminating an uncontrolled quantity by an additive combination, rather than working around it case by case, is a reroute worth having whenever one hypothesis in a problem is deliberately weaker than the others.

Layer 1: WHAT

Block 1: notice the product-and-quotient-of-the-same-two-functions structure, and multiply them to cancel g, producing $f^2 = (fg)(f/g)$, differentiable at 0 since it is a product of two functions differentiable at 0. Block 2: recognise that recovering f from f^2 needs the square root, which is differentiable only away from 0; this locates the fault line, $f(0) = 0$ against $f(0) \neq 0$, before doing any further work. Block 3: handle $f(0) \neq 0$ via the chain rule on $\sqrt{f^2}$ composed with the sign of $f(0)$. Block 4: handle $f(0) = 0$ separately by returning to f/g's difference quotient directly, dividing by $g(x)$ and then multiplying back by $g(x)$, which tends to the nonzero $g(0)$, to recover the difference quotient of f itself.

Juxtapose the nominal route against this one. Approach 1 writes down exactly the same identity, $f = \sqrt{(fg)(f/g)}$, but differentiates it with a ready-made composite-function formula, dividing by $2f(0)$ without asking whether that division is legal. The bifurcation point is Block 2: the ideal route stops the instant f^2 is known differentiable and asks whether the square root is safe to apply here, before writing any derivative formula; the nominal route writes the formula first and never asks.

Layer 2: HOW

Block 1. **Cue.** A product fg and a quotient f/g of the same two functions, both differentiable at 0. **Why natural.** Whenever a product and a quotient of the same pair of objects are both available, multiplying them cancels the shared factor; here $(fg)(f/g) = f^2$ removes g, the one function whose regularity is deliberately weak, and leaves a function built entirely from the two strong hypotheses. **Link to the next block.** f^2 being differentiable at 0 is not the same statement as f being differentiable at 0; recovering f from f^2 is a further, separate step, with its own restrictions.

Block 2. **Cue.** The recovery step $f = \pm\sqrt{f^2}$. **Why natural.** The square root function is differentiable at every positive real but not at 0, where its difference quotient diverges; since $f^2(0) = f(0)^2$, whether this recovery step is safe depends entirely on whether $f(0)$ itself is zero, which is a fact about f, not about f^2. **Link to the next block.** This is the exact fault line to split on, so the two remaining blocks are simply the two sides of that line, handled by whichever tool fits each side.

Block 3. **Cue.** $f(0) \neq 0$, so f is continuous (Step 1) and hence of constant sign near 0. **Why natural.** Once $\delta f > 0$ near 0 for the correct sign δ, δf literally equals $\sqrt{f^2}$ there, a genuine composition of the differentiable function f^2 with the square root, evaluated at the positive point $f(0)^2$; the chain rule applies without qualification. **Link to the next block.** Nothing here helps when $f(0) = 0$, since then $f^2(0) = 0$ sits exactly at the square root's point of non-differentiability, so a different tool is needed for that case, not a limiting version of this one.

Block 4. **Cue.** $f(0) = 0$, so f^2 vanishes at 0 and the square-root route is unavailable. **Why natural.** Returning to the raw definition of $(f/g)'(0)$ and multiplying its difference quotient back by $g(x)$, which is bounded away from 0 near 0 by continuity, reconstructs $f(x)/x$ as a product of two convergent quantities; this sidesteps the square root entirely rather than trying to rescue it.

Link to the next block. Both cases now closed, the two halves reassemble into one statement, f is differentiable at 0 regardless of the value of $f(0)$, closing the "must" in the question.

Layer 3: WHY

The truth underneath is that differentiability is a local, pointwise property, and the square root's derivative is unbounded at 0; any argument that recovers f from f^2 by taking a square root is only as strong as the square root is differentiable at the specific value $f(0)^2$ actually takes, which depends entirely on whether $f(0)$ is zero. The case split is not an artefact of a clumsy proof, it is a direct restatement of exactly where $\sqrt{}$ fails to be differentiable.

The discernment mechanism: whenever a solution manufactures a new function via a product, a quotient, or a composition of given data, check the domain on which the recovery step, a square root, a reciprocal, a logarithm, whatever undoes the manufacturing, is actually smooth, before trusting the manufactured function to hand back the original. That check is exactly the case split, or the reason there is none. Train this by asking, every time you multiply or compose two hypotheses together to build a new object, "what operation would I need to undo this, and where does that operation stop being differentiable".

The examiners chose "g nonzero and continuous, not differentiable" as the one deliberately weak hypothesis; the whole problem is an exercise in noticing that a weak hypothesis on one function, combined with strong hypotheses on two functions built from it and another, can still force a strong conclusion on the first, provided you find the algebraic combination, the product here, or the additive elimination in the remark, that never needs the weak hypothesis upgraded to anything stronger than it already is.

Revolving concepts

Differentiability implies continuity at the same point: used twice, once directly for f/g giving f continuous at 0 in Step 1, once implicitly wherever a differentiable function is treated as continuous. See Book 1, field: Real analysis, topic: differentiability implies continuity.

The chain rule, together with the domain of differentiability of the outer function: a composition $k(h(x))$ is differentiable at c once h is differentiable at c and k is differentiable at $h(c)$ specifically, not merely defined there; the square root is the standing example of a function differentiable everywhere on its domain except at one boundary point, 0. See Book 1, field: Calculus and integrals, topic: chain rule and domain restrictions on standard functions.

Algebra of limits: sums and products of functions with existing finite limits at a point have limits equal to the sum and product of those limits; used repeatedly to combine separately known limits of difference quotients into the limit defining $f'(0)$. See Book 1, field: Real analysis, topic: algebra of limits.

Eliminating an uncontrolled quantity between two expanded identities by an additive combination chosen so its coefficients cancel: a linear-elimination move, here applied to difference quotients rather than to a literal system of linear equations, useful whenever one function in a problem is assumed only continuous while others built from it are assumed differentiable. See Book 1, field: Real analysis, topic: elimination as a proof technique.

Circuit Bank

- **WHEN** you are given a product and a quotient of the same two functions, each with known but unequal regularity **DO** multiply them to cancel the shared, weaker factor, immediately **BECAUSE** $(fg)(f/g) = f^2$ eliminates the less regular function entirely and produces a new object built only from what was actually given

- **WHEN** you recover a function from its square, or from its image under any function that is not smooth everywhere **DO** check first whether the recovery step is smooth exactly where you need it, and split into cases at the boundary of where it fails **BECAUSE** the square root is differentiable only at positive values, and silently applying a derivative formula across that boundary produces an undefined or wrong answer disguised as a finished proof
- **WHEN** one hypothesis in a problem is visibly weaker than the others, continuous rather than differentiable **DO** look for an algebraic combination of the given data, a product, or an additive elimination of the weak function's difference quotient, that never requires upgrading the weak hypothesis **BECAUSE** the problem is testing exactly how far the weak hypothesis can be leveraged, not whether a stronger one can be smuggled in

Backtesting

1. First useful cue: fg and f/g given differentiable, g only continuous, a product and a quotient of the same pair of functions.
2. Method the expression suggested: multiply the product and the quotient to cancel g and produce f^2, differentiable at 0.
3. Where you would hesitate, and the thought to reinforce: right after writing $f = \pm\sqrt{f^2}$, the thought to reinforce is "is the square root differentiable at the value f^2 actually takes at 0", not "is f^2 differentiable", since those are different questions and only the first can fail.
4. The trap the question sets: applying the composite-derivative formula for $\sqrt{f^2}$ without checking $f(0) \neq 0$ first, dividing by $2f(0)$ when it may be 0.
5. What makes you recognise it faster next time: whenever you are asked to recover a function from a product or combination of two other functions, name the "undoing" operation explicitly, a square root, a division, and ask where it stops being smooth before applying it.
6. The reusable decision rule: a product and a quotient of the same functions, multiply to cancel; recovering the original via an operation that is not smooth everywhere, split exactly at its failure point; a hypothesis visibly weaker than the rest, hunt for a combination that never needs it upgraded.

Mirror drills

1. Let f, g be differentiable at 0 with $g(0) \neq 0$. Show, without quoting the quotient rule, that f/g is differentiable at 0. *Answer:* this is the base case of the same technique: write out the difference quotient of f/g, multiply and divide by $g(x)g(0)$, and use that $g(x) \to g(0) \neq 0$.
2. If f^2 and f^3 are both differentiable at 0, must f be differentiable at 0? *Answer:* yes; $f^3/f^2 = f$ wherever $f \neq 0$, and where $f(0) = 0$ a limit argument analogous to Step 4 applies, mirroring the product and quotient circuit with f^2, f^3 in place of $fg, f/g$.
3. Suppose only fg is known differentiable at 0, not f/g, but g itself is differentiable and nonzero at 0. Is f necessarily differentiable at 0? *Answer:* yes, trivially, since $f = (fg)/g$ with g differentiable and nonzero lets the ordinary quotient rule finish it in one line; contrast this against the original problem to see exactly why weakening g from differentiable to merely continuous is what creates the real difficulty.

15.10 B4. Tournament standings as a Gaussian Gram matrix

In a tournament, 2011 players meet 2011 times to play a multiplayer game. Every game is played by all 2011 players together and ends with each of the players either winning or losing. The standings are kept in two 2011×2011 *matrices,* $T = (T_{hk})$ *and* $W = (W_{hk})$*. Initially,* $T = W = 0$*. After every game, for every* (h, k) *(including for* $h = k$*), if players* h *and* k *tied (that is, both won or both lost), the entry* T_{hk} *is increased by 1, while if player* h *won and player* k *lost, the entry* W_{hk} *is increased by 1 and* W_{kh} *is decreased by 1.*

Prove that at the end of the tournament, $\det(T + iW)$ *is a non-negative integer divisible by* 2^{2010}*.*

Read it out loud

Say the structure to yourself before you touch a single entry. There are 2011 objects (the players). There is a table whose (h, k) entry is built only out of a comparison between object h and object k, accumulated over 2011 independent trials. You are asked for the determinant of that table. That is the whole shape of the question, and it is the shape of an inner product table. A matrix whose (h, k) entry is a pairwise comparison is a Gram matrix until proven otherwise; your first move is not to compute anything, it is to hunt for the vectors.

The second reading is about symmetry type. T is symmetric, because tying is a symmetric relation. W is antisymmetric with $W_{hh} = 0$, because the update rule explicitly writes $+1$ into W_{hk} and -1 into W_{kh}. Therefore $(T + iW)^* = T^{\mathsf{T}} - iW^{\mathsf{T}} = T + iW$: the matrix is Hermitian. The question chose the letter i for you, and it chose it precisely so that the symmetric part and the antisymmetric part could be carried in one Hermitian object. A Hermitian matrix has real determinant, so the word "integer" in the conclusion is already half explained. The word "non-negative" is the loud one. Over $\mathbb{C}$ there is essentially one structural reason a determinant is non-negative: the matrix is A^*A for some A. So you should have written $T + iW = A^*A$ on your paper inside thirty seconds, with A unknown, and then spent your effort finding the alphabet that A is written in.

The third reading is the exponent. Not 2^{2011}, not 2^{1005}, but 2^{2010}, and $2010 = 2011 - 1$, one less than the size of the matrix. Whenever a divisor of an $n \times n$ determinant carries the exponent $n - 1$, the mechanism is almost always the same: every entry of the matrix lies in a single residue class modulo something, so you may subtract one row from the other $n - 1$ rows and pull a common factor out of each of them. The exponent counts the rows you can clear. So the problem has told you, in the exponent, that the matrix you should be reducing has 2011 rows all congruent to each other, and told you, in the word non-negative, that this matrix is not $T + iW$ itself but its square root A.

The fields wired into this problem are linear algebra (Gram matrices, Hermitian forms, multiplicativity of the determinant), complex numbers used as an alphabet rather than as numbers, and a thin slice of the arithmetic of the Gaussian integers $\mathbb{Z}[i]$. That last slice is where the problem hides its real content, and it is where the candidates who get 2^{1005} instead of 2^{2010} lose four points.

Approaches, ranked worst to best

1. **Leibniz expansion.** *Cue:* "prove det of this explicit integer matrix is divisible by ...". *Circuit:* brute expansion of a determinant. You have 2011! signed products and no control over any of them. There is no parity bookkeeping that survives. This is zombie maths: it costs infinite ink and never terminates.

2. **Peel off** W **as a low rank perturbation.** *Cue:* W_{hk} counts wins against losses, which looks like a difference of two one-player statistics. *Circuit:* rank one and rank two determinant updates. This one is worth doing in your head because it produces a true and pretty fact. Write $\varepsilon_{jk} = +1$ if player k won game j and -1 if he lost. Then in game j the quantity $(\varepsilon_{jh} - \varepsilon_{jk})/2$ is $+1$ when h won and k lost, -1 when h lost and k won, and 0 on a tie, which is exactly the update to

W_{hk}. Summing over games, $W_{hk} = (s_h - s_k)/2$ where s_h is player h's wins minus losses. So $W = \frac{1}{2}(s\mathbf{1}^\mathsf{T} - \mathbf{1}s^\mathsf{T})$ has rank at most 2. Now you are tempted to compute $\det(T + iW)$ as a rank two update of $\det T$. It stalls: the update formula drags in T^{-1}, which you have no reason to believe exists, and the factors of $\frac{1}{2}$ generate a 2^{-2011} that you then have to fight. You have learnt something structural and gained nothing computational.

3. **Reduce $T + iW$ itself modulo the Gaussian prime $1 + i$.** *Cue:* the exponent $2010 = 2011 - 1$ shouts "congruent rows". *Circuit:* if all entries of an $n \times n$ matrix are congruent modulo π then π^{n-1} divides the determinant. This is the right circuit fired at the wrong matrix, and it is instructive to see exactly how much it gives. For $h \neq k$, each game either has the pair tying or splitting, so $T_{hk} + (\text{number of splits}) = 2011$, and W_{hk} is a signed count of those splits, hence $W_{hk} \equiv 2011 - T_{hk} \equiv 1 + T_{hk} \pmod 2$. Therefore $T_{hk} + iW_{hk} \equiv T_{hk}(1 + i) + i \equiv i \pmod{1 + i}$, while on the diagonal $T_{hh} + iW_{hh} = 2011 \equiv 1 \pmod{1 + i}$; and $1 \equiv i$ modulo $1 + i$ since $i - 1 = i(1 + i)$. So every entry of $T + iW$ is congruent to every other modulo $1 + i$, and $(1 + i)^{2010}$ divides $\det(T + iW)$ in $\mathbb{Z}[i]$. But $2 = -i(1 + i)^2$, so the prime 2 ramifies, and $(1 + i)^{2010}$ dividing a rational integer only gives 2^{1005}. You are exactly a square root short. That deficiency is the signal to go looking for a square root of the matrix.

4. **Encode results as ± 1 and use the real correlation matrix.** *Cue:* results are binary, so use signs. *Circuit:* Gram matrix over $\mathbb{R}$. With $E = (\varepsilon_{jk})$ you get $(E^\mathsf{T}E)_{hk} = T_{hk} - (2011 - T_{hk})$, so $E^\mathsf{T}E = 2T - 2011J$ where J is all ones. Clean, but fatal: $E^\mathsf{T}E$ is symmetric and therefore cannot see W at all. Real correlations are blind to the order of a pair. That failure is not a dead end, it is a specification: you need an alphabet whose Hermitian products are non-real.

5. **Gram factorisation over $\mathbb{Z}[i]$ with alphabet $\{1, i\}$, then congruent rows on A.** *Cue:* "non-negative" plus "2^{n-1}". *Circuit:* $\det(A^*A) = |\det A|^2$, combined with row clearing modulo $1 - i$. Everything the problem says is used, nothing is computed, and the norm at the end doubles the $(1 - i)$ valuation into exactly the required power of 2. This is the ideal route.

The ideal solution

Write $n = 2011$ throughout, and index games by j and players by h, k.

Choosing the alphabet. Suppose we look for complex numbers u (recording a win) and v (recording a loss) and set $a_{jk} = u$ if player k won game j and $a_{jk} = v$ if he lost. Consider the Hermitian product $\overline{a_{jh}}\, a_{jk}$ contributed by game j to the (h, k) position. We need it to equal 1 when h and k tied, $+i$ when h won and k lost, and $-i$ in the reverse case, because those are precisely the increments to $T_{hk} + iW_{hk}$. The tie conditions force $|u| = |v| = 1$; the split condition forces $\bar{u}v = i$, that is $v = iu$. Taking $u = 1$ gives $v = i$, and every other admissible choice is this one multiplied by a fixed number of modulus one, which changes nothing. The alphabet is not invented, it is solved for.

The factorisation. Let $A = (a_{jk})$ be the $n \times n$ matrix over $\mathbb{Z}[i]$ with $a_{jk} = 1$ if player k won game j and $a_{jk} = i$ if player k lost game j, and let $A^* = \overline{A}^\mathsf{T}$.

Claim 15.10.1. *$T + iW = A^*A$ at the end of the tournament.*

Proof. $(A^*A)_{hk} = \sum_{j=1}^{n} \overline{a_{jh}}\, a_{jk}$. Fix a game j and a pair (h, k), allowing $h = k$. If both players won, the term is $\bar{1} \cdot 1 = 1$; if both lost, it is $\bar{i} \cdot i = |i|^2 = 1$; in both cases the pair tied and game j added exactly 1 to T_{hk} and 0 to W_{hk}. If h won and k lost, the term is $\bar{1} \cdot i = i$, and game j added 1 to W_{hk}, contributing $i \cdot 1$. If h lost and k won, the term is $\bar{i} \cdot 1 = -i$, and game j added -1 to W_{hk}, contributing $i \cdot (-1) = -i$. The case $h = k$ is covered by the first two lines, since a player always ties with himself and the term is $|a_{jh}|^2 = 1$. So in every game and every position the term of A^*A equals the increment of $T + iW$. Both matrices start at 0 and both accumulate these increments over $j = 1, \ldots, n$, so they are equal at the end. □

Lemma 15.10.2 (congruent rows)**.** *Let R be a commutative ring, $\pi \in R$, and M an $n \times n$ matrix over R all of whose entries are congruent to one another modulo π. Then π^{n-1} divides* $\det M$ *in R.*

Proof. Subtract row 1 of M from each of rows $2, \dots, n$. This is a determinant preserving operation. Every entry of every new row is a difference of two entries of M, hence divisible by π; write each such row as π times a row over R. Multilinearity of the determinant in the rows lets us extract one factor of π from each of those $n-1$ rows, so $\det M = \pi^{n-1} \det M'$ for a matrix M' over R. □

Applying the lemma to A. Every entry of A lies in $\{1, i\}$, and $i - 1 = -(1 - i)$, so all entries of A are congruent modulo $\pi = 1 - i$ in $R = \mathbb{Z}[i]$. By the lemma there is $z \in \mathbb{Z}[i]$, say $z = a + bi$ with $a, b \in \mathbb{Z}$, such that

$$\det A = (1-i)^{n-1} z = (1-i)^{2010}(a + bi).$$

Taking the norm. The determinant is multiplicative, and $\det(A^*) = \det(\overline{A}^{\mathsf{T}}) = \det(\overline{A}) = \overline{\det A}$, since transposing preserves the determinant and conjugation commutes with the polynomial expression defining it. Hence

$$\det(T + iW) = \det(A^* A) = \overline{\det A}\, \det A = |\det A|^2 = |1-i|^{2 \cdot 2010}\, |z|^2.$$

Now $|1 - i|^2 = 2$, so $|1 - i|^{2 \cdot 2010} = 2^{2010}$, and $|z|^2 = a^2 + b^2$. Therefore

$$\boxed{\det(T + iW) = 2^{2010}\,(a^2 + b^2)}$$

for some integers a, b. Since $a^2 + b^2$ is a non-negative integer, $\det(T + iW)$ is a non-negative integer divisible by 2^{2010}, as required.

Remark on the hypothesis 2011 games. The equality of the number of games and the number of players is used only to make A square. It is not needed for the conclusion. If there are g games and n players, A is $g \times n$, and the Cauchy and Binet formula gives $\det(A^* A) = \sum_S |\det A_S|^2$, the sum over all n-element subsets S of the games, where A_S is the corresponding $n \times n$ submatrix. Each $\det A_S$ is divisible by $(1 - i)^{n-1}$ by the same lemma, so each $|\det A_S|^2$ is divisible by 2^{n-1}, and the conclusion holds for any number of games (vacuously with determinant 0 when $g < n$).

Layer 1: WHAT

Block 1. Read the (h, k) entry as a pairwise comparison and hypothesise $T + iW = A^* A$.

Block 2. Confirm the hypothesis is legal: T symmetric, W antisymmetric, so $T + iW$ is Hermitian, which is a necessary condition for being $A^* A$.

Block 3. Solve for the alphabet: unimodular u, v with $\bar{u} v = i$, giving $u = 1$, $v = i$.

Block 4. Verify $T + iW = A^* A$ by checking the four outcome cases of one game in one position, including $h = k$.

Block 5. Convert the conclusion: $\det(A^* A) = |\det A|^2$ delivers "real" and "non-negative" at once, and reduces the problem to a divisibility statement about $\det A$.

Block 6. Read 2010 as $n - 1$; all entries of A are congruent modulo $1 - i$; clear $n - 1$ rows and extract $(1 - i)^{2010}$.

Block 7. Take absolute squares. The norm doubles the valuation, turning $(1 - i)^{2010}$ into 2^{2010}, and $|z|^2 = a^2 + b^2$ is a non-negative integer.

Juxtaposition. The nominal route works on $T + iW$ directly: compute its entries in terms of win counts, notice they are all congruent modulo $1 + i$, apply the congruent rows lemma, and land on 2^{1005} with no idea why the problem asked for twice as much. The ideal route never computes an entry of T or W at all; it factorises first and applies the same lemma one level down, to A.

Bifurcation point. The split is at the word *non-negative*. If you treat it as decoration, you work on $T + iW$ and lose half the exponent. If you treat it as a command term meaning "this matrix is $A^* A$, go and find A", every later step is forced. There is no second decision in this problem.

Layer 2: HOW

Block 1. **Cue.** The phrase "for every (h, k)" together with an update rule that depends only on the results of players h and k. Nothing in the rule mentions any third player. **Why natural.** A table indexed by pairs of objects, whose entry depends only on those two objects and accumulates over independent trials, is by definition a sum over trials of a two-object function. If that function factorises as $\overline{\phi(h)}\phi(k)$ for each trial, the table is a Gram matrix. Since determinants of general matrices are unmanageable and determinants of Gram matrices are trivial, you always test this factorisation first. **Link to the next block.** Before hunting for ϕ, check that the target has the symmetry a Gram matrix must have.

Block 2. **Cue.** W_{kh} is decreased by exactly the amount W_{hk} is increased; T treats h and k alike. **Why natural.** A^*A is always Hermitian, so if $T + iW$ were not Hermitian the search would be over. It is: $T^\mathsf{T} = T$ and $W^\mathsf{T} = -W$ give $(T + iW)^* = T^\mathsf{T} - i\,W^\mathsf{T} = T + iW$. The problem manufactured a Hermitian matrix on purpose, which is why i appears in the statement at all. **Link to the next block.** A Hermitian target tells you the products $\overline{a_{jh}}a_{jk}$ must reproduce the increments. Write down what those increments are and solve.

Block 3. **Cue.** The increment is 1 on a tie and $\pm i$ on a split, with the sign flipping when the pair is swapped. **Why natural.** $\overline{a_{jh}}a_{jk}$ swaps to its conjugate when h and k swap. So a symmetric increment needs a real product and an antisymmetric increment needs a purely imaginary one. Setting u, v for win and loss, the tie conditions $\bar{u}u = \bar{v}v = 1$ force modulus one and the split condition $\bar{u}v = i$ forces $v = iu$. This is a two line calculation with a unique answer up to an irrelevant global phase, so nothing here is invention. It also explains, retroactively, why the ± 1 encoding of results cannot work: with real entries the product is symmetric and W is invisible. **Link to the next block.** You now have a candidate A; verify it exhaustively, because a Gram claim is only worth what its case check is worth.

Block 4. **Cue.** "for every (h, k) (including for $h = k$)". The examiners inserted that parenthesis; treat it as an instruction to check the diagonal separately. **Why natural.** There are exactly four outcomes for a pair in a game and the verification is four one line computations. The diagonal case is the one candidates drop, and it is the one that makes $T_{hh} = 2011$ match $\sum_j |a_{jh}|^2 = 2011$. **Link to the next block.** With the factorisation proved, restate the goal in terms of $\det A$.

Block 5. **Cue.** det is multiplicative and $\det(A^*) = \overline{\det A}$. **Why natural.** $\det(A^*A) = |\det A|^2$ converts a statement about a complicated Hermitian matrix into a statement about one Gaussian integer. Reality and non-negativity are now automatic rather than things to prove, which is the whole payoff of factorising. **Link to the next block.** The surviving problem is: why is $|\det A|^2$ divisible by 2^{2010}? Equivalently, why is $\det A$ divisible by something of modulus squared 2, raised to the power 2010?

Block 6. **Cue.** $2010 = 2011 - 1 = n - 1$, and A has only two distinct entries. **Why natural.** The exponent $n - 1$ is the fingerprint of row clearing: you can always make $n - 1$ rows divisible by π if all entries are congruent modulo π, and no more than that, because the first row survives. Here the two entries are 1 and i, whose difference is $-(1 - i)$, so $\pi = 1 - i$ is not chosen, it is read off. The element $1 - i$ has $|1 - i|^2 = 2$, which is exactly the modulus squared you needed in Block 5. **Link to the next block.** Extract $(1 - i)^{2010}$ and take the norm.

Block 7. **Cue.** A Gaussian integer written as $a + bi$, and $|z|^2 = a^2 + b^2$. **Why natural.** Taking $|\cdot|^2$ of $\det A = (1 - i)^{2010}z$ multiplies valuations by two: $|1 - i|^2 = 2$ gives 2^{2010}, and $|z|^2$ is a non-negative

rational integer with no further structure needed. This is where the ramification of 2 in $\mathbb{Z}[i]$ pays you back the factor the nominal route lost. **Link to the next block.** Nothing remains; state the boxed identity and read the conclusion off it.

Layer 3: WHY

Three mathematical truths carry this problem, and each is worth installing on its own.

First, the Gram principle. An $n \times n$ complex matrix M is of the form A^*A if and only if M is Hermitian positive semidefinite, and in that case $\det M \geq 0$ because the determinant is the product of the (real, non-negative) eigenvalues. So "prove this determinant is a non-negative integer" is never a computational instruction; it is a request for a factorisation. In the arena, the reflex must be: non-negative determinant, find the vectors. Here the vectors are the columns of A, one per player, living in $\mathbb{C}^{2011}$ with one coordinate per game, and the entry $T_{hk} + iW_{hk}$ is literally the Hermitian inner product of player h's history with player k's history.

Second, the arithmetic of $\mathbb{Z}[i]$, which is the part of the problem that is not linear algebra at all. In $\mathbb{Z}[i]$ the rational prime 2 is not prime: $2 = -i(1+i)^2$, and $1+i$ (equivalently $1-i$, its associate) is the unique ramified prime, of norm 2. The consequence you must feel in your fingers is that the $(1+i)$ valuation of a rational integer is twice its 2 valuation. That single fact explains the entire design of the problem. The naive route obtains $(1+i)^{2010} \mid \det(T+iW)$ and therefore only 2^{1005}; the ideal route obtains $(1-i)^{2010} \mid \det A$ and then squares the norm, gaining a factor of two in the exponent and reaching 2^{2010}. The examiners chose the exponent 2010 rather than 1005 exactly to detect which of the two you did. The discernment mechanism is therefore quantitative: compute what the lazy route gives, compare it with what is demanded, and if you are short by a factor of two in an exponent over $\mathbb{Z}[i]$, you are being asked for a square root of the matrix.

Third, the congruent rows lemma, which is the general reason exponents of the form $n-1$ appear in divisibility statements about determinants. It holds over any commutative ring and its proof is one row operation. Learn it as a rule about *alphabets*: if the entries of a matrix are drawn from a set that lies in a single residue class modulo π, you get π^{n-1} free. Sets like $\{1,-1\}$ modulo 2, $\{0,1\}$ modulo 1 (useless), $\{1,i,-1,-i\}$ modulo $1-i$, and $\{1,\omega,\omega^2\}$ modulo $1-\omega$ for a primitive cube root of unity, are all worth knowing by sight.

How to train this. Take every past problem in which the answer is a determinant and ask, before computing, which of the three standard structures it has: Gram (A^*A), circulant or group-invariant (diagonalised by characters), or combinatorial incidence (Cauchy and Binet or matrix tree). Then ask what the entries are made of. Two of those three questions solve this problem, and neither of them is a calculation.

Revolving concepts

Gram matrix. For vectors $v_1, \ldots, v_n$ in $\mathbb{C}^m$ the Gram matrix is $G_{hk} = \langle v_h, v_k \rangle = v_h^* v_k$, and $G = A^*A$ where A has the v_k as columns. Orbiting facts: G is Hermitian positive semidefinite; $\operatorname{rank} G = \operatorname{rank} A$; $\det G \geq 0$ with equality if and only if the vectors are linearly dependent; over $\mathbb{R}$, $\det G$ is the square of the volume of the parallelotope they span. Conversely every Hermitian positive semidefinite matrix is a Gram matrix, via its positive semidefinite square root. See Book 1, field: Linear algebra, topic: Gram matrices and positive semidefinite forms.

Hermitian matrices. $M^* = M$; eigenvalues real; $\det M$ and $\operatorname{tr} M$ real; the decomposition of a real matrix pair into $T + iW$ with T symmetric and W antisymmetric is the standard way to package a symmetric and an antisymmetric statistic into one object. See Book 1, field: Linear algebra, topic: Hermitian and unitary matrices.

Cauchy and Binet. For A of size $g \times n$ with $g \geq n$, $\det(A^*A) = \sum_{|S|=n} |\det A_S|^2$, the sum over n-subsets of rows. This is the tool that removes the artificial coincidence between the number of

games and the number of players, and it is the correct generalisation of $\det(A^*A) = |\det A|^2$. See Book 1, field: Matrices and determinants, topic: Cauchy and Binet formula.

Gaussian integers. $\mathbb{Z}[i]$ is a Euclidean domain with norm $N(a+bi) = a^2+b^2$, units $\{\pm 1, \pm i\}$. A rational prime q splits if $q \equiv 1 \pmod 4$, stays prime if $q \equiv 3 \pmod 4$, and ramifies for $q = 2$, where $2 = -i(1+i)^2$. Hence $v_{1+i}(m) = 2v_2(m)$ for $m \in \mathbb{Z}$. Sums of two squares, Fermat's theorem on $q \equiv 1 \pmod 4$, and the classification of Gaussian primes all orbit this. See Book 1, field: Number theory, topic: Gaussian integers and sums of two squares.

The congruent rows lemma. Stated and proved above. Generalise it to: if the rows of M fall into r classes modulo π then π^{n-r} divides $\det M$, by clearing all but one row in each class. The version with $r = 1$ is the one you will use ninety per cent of the time.

The recursion in the problem, generalised. The tournament is a recursion: if M_g denotes $T + iW$ after g games and $v_g \in \mathbb{C}^n$ is the row of game g (entry 1 for a winner, i for a loser), then the update rule says exactly

$$M_g = M_{g-1} + v_g^* v_g, \qquad M_0 = 0.$$

That is the general rank one accumulation recursion, and its general solution is $M_g = \sum_{j \le g} v_j^* v_j = A^*A$ with A the matrix of rows $v_1, \dots, v_g$. Two standard consequences travel with it. The matrix determinant lemma, $\det(M + uu^*) = \det(M)\,(1 + u^* M^{-1} u)$ for invertible M, is the step by step evaluation and is usually the wrong tool. The closed form via Cauchy and Binet is the right one. Whenever a problem describes a process that adds a rank one Hermitian correction at each step, do not iterate: recognise the sum as a Gram matrix and jump straight to the closed form. See Book 1, field: Matrices and determinants, topic: rank one updates and accumulation.

Roots of unity as an alphabet. Encoding an m-valued outcome by m-th roots of unity so that $\overline{u}v$ records the difference of two outcomes is the same device that underlies Gauss sums, characters of finite abelian groups, and the discrete Fourier transform. The specialisation here is the smallest interesting case: two outcomes, ratio i, so that $\overline{u}v$ is purely imaginary and can carry an antisymmetric statistic. See Book 1, field: Complex numbers, topic: roots of unity as encodings.

Circuit Bank

- **WHEN** a matrix whose (h,k) entry is a comparison between object h and object k **DO** look for vectors v_h with entry $= v_h^* v_k$ and write $M = A^*A$ **BECAUSE** a pairwise table is an inner product table until proven otherwise, and factorising hands you reality, non-negativity, rank and multiplicativity of the determinant in one line

- **WHEN** the words "prove that this determinant is non-negative" **DO** stop computing and find the Gram factorisation **BECAUSE** over $\mathbb{C}$ the only structural source of a non-negative determinant is A^*A, so the phrase is a command term, not decoration

- **WHEN** a claim that π^{n-1} divides an $n \times n$ determinant **DO** check whether all entries lie in one residue class modulo π, then subtract row 1 from the other rows **BECAUSE** the exponent $n-1$ counts exactly the rows you can clear by one determinant preserving operation

- **WHEN** a two valued outcome that must be recorded antisymmetrically **DO** encode the outcomes as unimodular numbers whose ratio is i, never as ± 1 **BECAUSE** $\overline{u}v$ is symmetric under swapping the pair precisely when u/v is real, so real alphabets are blind to orientation

- **WHEN** a process defined by "after every step, add one to these entries" **DO** write the step as $M_g = M_{g-1} + v_g^* v_g$ and pass to the closed form A^*A **BECAUSE** rank one accumulation always has a Gram closed form, and iterating the determinant instead is strictly more expensive

Backtesting

1. First useful cue: the pair of words "non-negative integer" attached to a determinant, reinforced by the exponent $2010 = 2011 - 1$.
2. Method the expression suggested: factorise $T + iW$ as A^*A where A records each player's history in an alphabet of two unimodular symbols, then apply the congruent rows lemma to A over $\mathbb{Z}[i]$ and take norms.
3. Where you would hesitate, and the thought to reinforce: choosing i for a loss rather than -1. Reinforce by solving rather than guessing: write u, v unimodular, demand $\bar{u}v = i$, conclude $v = iu$. Any encoding question of this type is a two line linear demand, never a flash of insight.
4. The trap the question sets: applying the congruent rows lemma to $T + iW$ itself. It is a legal argument, it is quick, and it yields only 2^{1005} because 2 ramifies in $\mathbb{Z}[i]$. A candidate who does this and writes "hence divisible by 2^{2010}" has made a valuation error worth the whole problem.
5. What makes you recognise it faster next time: the exponent audit. Compute the exponent the cheap route produces, compare it with the exponent demanded, and let the ratio tell you what structure you are missing. A shortfall by a factor of exactly two in an exponent over a quadratic ring means a norm is waiting to be taken.
6. The reusable decision rule: pairwise table plus non-negativity equals Gram factorisation; exponent $n - 1$ equals congruent rows; and when both appear, apply the second to the factor, not to the product.

Mirror drills

1. Let M be an $n \times n$ integer matrix all of whose entries are odd. Show $2^{n-1} \mid \det M$. *Answer:* all entries are congruent modulo 2; subtract row 1 from the other $n - 1$ rows and pull out a factor of 2 from each.
2. Let M be an $n \times n$ matrix each of whose entries is a fourth root of unity. Show that 2^{n-1} divides $|\det M|^2$. *Answer:* all of $1, i, -1, -i$ are congruent modulo $1 - i$ (their pairwise differences are $1 - i$ times a Gaussian integer), so $(1 - i)^{n-1} \mid \det M$; now take $|\cdot|^2$ and use $|1 - i|^2 = 2$.
3. Same tournament, but with g games instead of 2011, where g may be less than or greater than the number n of players. Show $\det(T + iW)$ is still a non-negative integer divisible by 2^{n-1}. *Answer:* $T + iW = A^*A$ with A of size $g \times n$; apply the Cauchy and Binet formula and the congruent rows lemma to each $n \times n$ minor, noting that the determinant is 0 when $g < n$.

15.11 B5. Overlap integrals of Lorentzians against reciprocal sums

Let $a_1, a_2, \ldots$ be real numbers. Suppose that there is a constant A such that for all n,

$$\int_{-\infty}^{\infty} \left(\sum_{i=1}^{n} \frac{1}{1+(x-a_i)^2} \right)^2 dx \le An.$$

Prove there is a constant $B > 0$ such that for all n,

$$\sum_{i,j=1}^{n} (1+(a_i-a_j)^2) \ge Bn^3.$$

Read it out loud

The hypothesis is the L^2 norm squared of a sum of n identical bumps, each bump a Lorentzian $1/(1+x^2)$ translated to sit over a_i. The conclusion is a double sum over the same index pairs of $1+(a_i-a_j)^2$. Say that out loud and the shape of the problem is already fixed: the hypothesis is a quadratic form in the pairs (i,j), and the conclusion is a quadratic form in the same pairs. The only work is to make the hypothesis look like the conclusion.

The mechanism that turns an L^2 norm of a sum of translates into a pairwise quadratic form is expansion. Never bound a squared sum before you expand it: the cross terms are the entire content, and any pointwise estimate on the sum destroys them. Expanding gives $\sum_{i,j} \int \phi(x-a_i)\phi(x-a_j)\,dx$ with $\phi(x) = 1/(1+x^2)$, and the inner integral is translation invariant, so it depends only on $a_i - a_j$. Call it $f(a_i-a_j)$. Within thirty seconds you should have written

$$\sum_{i,j=1}^{n} f(a_i-a_j) \le An, \qquad f(y) = \int_{-\infty}^{\infty} \frac{dx}{(1+x^2)\big(1+(x+y)^2\big)}.$$

Now look at the two sums side by side. The hypothesis, once you know the size of f, controls $\sum_{i,j} 1/(1+(a_i-a_j)^2)$ from above. The conclusion asks for $\sum_{i,j}(1+(a_i-a_j)^2)$ from below. Those two summands are exact reciprocals. There is no pointwise inequality connecting a positive number to its reciprocal in the direction you need, so the bridge has to be a global one, and the global statement about a family and its reciprocals is the Cauchy and Schwarz inequality: $\big(\sum 1\big)^2 \le \big(\sum u\big)\big(\sum 1/u\big)$. That is not a choice, it is the only mechanism available.

Check the exponent before you commit. There are n^2 index pairs, so $\sum_{i,j} 1 = n^2$, and Cauchy and Schwarz will deliver $\sum u \ge (n^2)^2 / \sum(1/u) \ge n^4/(\text{constant}\cdot n) = \text{constant}\cdot n^3$. The exponent 3 in the problem is the fingerprint of exactly this manoeuvre. The fields here are real analysis (convergence and term by term integration), the calculus of the Cauchy kernel, and inequalities. Note also the phrase "there is a constant $B > 0$": that licenses you to throw away every fixed numerical factor from the first line onward, so you must never compute f exactly.

Approaches, ranked worst to best

1. **Bound the integrand pointwise before expanding.** *Cue:* a squared sum inside an integral looks like something to estimate. *Circuit:* crude majorisation. Any bound of the form $\sum_i \phi(x-a_i) \le$ something, or an application of the Cauchy and Schwarz inequality inside the integral, collapses the pairwise structure into a single number. The conclusion is a statement about pairs (i,j), so information about pairs is exactly what you must not throw away. This is the pure zombie route: it produces true inequalities that are useless.

2. **Contrapositive by clustering.** *Cue:* "prove there is a constant $B > 0$" invites a proof by contradiction. *Circuit:* if the pairwise distances are small then the bumps pile up and the L^2

norm is large. Assume $\sum_{i,j}(1+(a_i-a_j)^2)$ is small, deduce that the a_i crowd into few unit intervals, deduce that $\int \Phi_n^2$ is large. This works, and pushed properly it gives the stronger bound Bn^4: count q_z points in $[z, z+1)$, note $\sum_z q_z^2 \lesssim An$, apply Jensen to see the points occupy at least cn distinct unit intervals, then sum $(|i-j|-1)^2$ over those. It ranks low here because it is strictly more machinery than the question asks for, and because it still requires the expansion step first.

3. **Plancherel.** *Cue:* an L^2 norm of a sum of translates. *Circuit:* energy in frequency space. Since $\widehat{\phi}(\xi) = \pi e^{-|\xi|}$ for $\phi(x) = 1/(1+x^2)$, Plancherel turns the hypothesis into $\frac{\pi}{2}\int e^{-2|\xi|}\left|\sum_j e^{i\xi a_j}\right|^2 d\xi \le An$. This is the conceptual explanation of why the hypothesis is a positive definite quadratic form in the differences, and it is worth knowing. As a write-up it is longer, it needs the Fourier transform of the Cauchy kernel as a quoted fact, and it lands you back at the same pairwise sum.

4. **Evaluate f exactly.** *Cue:* a convolution of two Lorentzians. *Circuit:* partial fractions, contour integration, or the Poisson kernel semigroup. You can compute $f(y) = 2\pi/(4+y^2)$, most cheaply by noting that $P_t(x) = \frac{t}{\pi(t^2+x^2)}$ is the Poisson kernel with $P_s * P_t = P_{s+t}$, so $f = \pi^2(P_1 * P_1) = \pi^2 P_2$. Correct, and the exact constant is pleasant, but the problem only ever needs the order of magnitude of f, and a contour integral with two double poles in the arena is ink you do not have to spend.

5. **Expand, bound f below on one fixed interval, then Cauchy and Schwarz.** *Cue:* reciprocal summands on the two sides, and the exponent $n^3 = (n^2)^2/n$. *Circuit:* a lower bound for an integral of a positive function is obtained by shrinking the domain, and reciprocal families are compared by Cauchy and Schwarz. Three moves, no computation, every constant deliberately crude. This is the ideal route.

The ideal solution

Write $\phi(x) = 1/(1+x^2)$ and $\Phi_n(x) = \sum_{i=1}^n \phi(x-a_i)$, and put

$$f(y) = \int_{-\infty}^{\infty} \frac{dx}{(1+x^2)\big(1+(x+y)^2\big)}.$$

Step 0: the constant A is positive. Apply the hypothesis with $n=1$: $\int_{-\infty}^{\infty} dx/(1+(x-a_1)^2)^2 \le A$. The left side equals $\int_{-\infty}^{\infty} dx/(1+x^2)^2 = \pi/2$, so $A \ge \pi/2 > 0$. This is needed only so that the constant B produced below is positive.

Step 1: expand the hypothesis into a pairwise sum. For each pair (i,j) the integrand $\phi(x-a_i)\phi(x-a_j)$ is continuous and non-negative, and

$$\int_{-\infty}^{\infty} \phi(x-a_i)\phi(x-a_j)\,dx \le \int_{-\infty}^{\infty} \phi(x-a_i)\,dx = \pi < \infty,$$

because $\phi \le 1$ everywhere. So each of the n^2 integrals converges, and since the sum is finite we may expand the square and integrate term by term:

$$\int_{-\infty}^{\infty} \Phi_n(x)^2\,dx = \sum_{i,j=1}^{n} \int_{-\infty}^{\infty} \phi(x-a_i)\phi(x-a_j)\,dx.$$

Substituting $x = t + a_i$ in the (i,j) term turns it into $\int \phi(t)\phi(t+a_i-a_j)\,dt = f(a_i-a_j)$. Hence the hypothesis reads

$$\sum_{i,j=1}^{n} f(a_i-a_j) \le An \qquad \text{for all } n. \tag{1}$$

Step 2: a crude lower bound for f. First, f is even: replacing x by $-x$ in the defining integral sends $f(-y)$ to $f(y)$. So it suffices to bound $f(y)$ for $y \geq 0$.

Let $y \geq 0$ and restrict the integral to $x \in [-1, 0]$, which is legitimate as a lower bound because the integrand is positive everywhere. On that interval $0 \leq x^2 \leq 1$, so $1 + x^2 \leq 2$; and $x + y$ lies between $y - 1$ and y, so $(x+y)^2 \leq (1+y)^2$ and $1 + (x+y)^2 \leq 1 + (1+y)^2 = y^2 + 2y + 2$. Using $2y \leq y^2 + 1$,

$$(1 + x^2)\big(1 + (x+y)^2\big) \leq 2\,(y^2 + 2y + 2) \leq 2\,(2y^2 + 3) \leq 6(1 + y^2).$$

The interval $[-1, 0]$ has length 1, so

$$f(y) \geq \frac{1}{6(1+y^2)} \qquad \text{for all } y \in \mathbb{R}. \tag{2}$$

Step 3: convert the hypothesis into an algebraic statement. Combining (1) and (2),

$$\sum_{i,j=1}^{n} \frac{1}{1 + (a_i - a_j)^2} \;\leq\; 6 \sum_{i,j=1}^{n} f(a_i - a_j) \;\leq\; 6An. \tag{3}$$

Step 4: Cauchy and Schwarz against the reciprocals. Put $u_{ij} = 1 + (a_i - a_j)^2 > 0$. Applying the Cauchy and Schwarz inequality to the n^2 pairs of numbers $\sqrt{u_{ij}}$ and $1/\sqrt{u_{ij}}$,

$$n^4 = \Big(\sum_{i,j=1}^{n} 1\Big)^2 = \Big(\sum_{i,j=1}^{n} \sqrt{u_{ij}} \cdot \frac{1}{\sqrt{u_{ij}}}\Big)^2 \leq \Big(\sum_{i,j=1}^{n} u_{ij}\Big)\Big(\sum_{i,j=1}^{n} \frac{1}{u_{ij}}\Big).$$

By (3) the second factor is at most $6An$, and it is strictly positive, so

$$\sum_{i,j=1}^{n} \big(1 + (a_i - a_j)^2\big) \;\geq\; \frac{n^4}{6An} \;=\; \frac{n^3}{6A}.$$

Therefore the conclusion holds with

$$\boxed{B = \frac{1}{6A}}$$

which is positive by Step 0, and the proof is complete.

Layer 1: WHAT

Block 1. Expand the squared sum inside the integral into n^2 cross terms, and check that each converges.

Block 2. Use translation invariance: the (i, j) cross term depends only on $a_i - a_j$. Name it $f(a_i - a_j)$; the hypothesis becomes a pairwise sum bounded by An.

Block 3. Compare the two sides. The conclusion sums $1 + (a_i - a_j)^2$; so what you want from f is a lower bound of size $1/(1 + y^2)$, and nothing more. Decide not to evaluate f.

Block 4. Get that lower bound by shrinking the domain of integration to a fixed interval of length one, and bounding the denominator crudely there.

Block 5. You now hold $\sum 1/u_{ij} \leq 6An$ and want $\sum u_{ij} \geq Bn^3$. Reciprocal families with n^2 members: apply the Cauchy and Schwarz inequality.

Block 6. Read off $B = 1/(6A)$, and check $A > 0$ from the case $n = 1$ so that $B > 0$.

Juxtaposition. The nominal route treats the hypothesis as an analysis problem to be solved exactly: compute f by partial fractions or residues, obtain $2\pi/(4 + y^2)$, then do Steps 3 to 5 anyway. It reaches the same place after an extra page. The genuinely zombie route never expands at all, and instead tries to bound Φ_n pointwise, or tries to argue directly that the points a_i must be spread out; both lose the pairwise information the conclusion is made of.

Bifurcation point. The first bifurcation is the very first line: expand the square, or estimate it. Expanding is correct because the conclusion is indexed by pairs. The second bifurcation is Step 3: compute f, or bound f. Bounding is correct because the problem said "there is a constant", and that phrase forbids you from caring about any fixed factor.

Layer 2: HOW

Block 1. **Cue.** A square of a finite sum sitting inside an integral, with the conclusion indexed by pairs (i, j). **Why natural.** $\left(\sum_i c_i\right)^2 = \sum_{i,j} c_i c_j$ is the only identity that converts a sum of n objects into data about n^2 pairs, and pairs are what the conclusion is about. The check that each cross term converges is one line, using $\phi \le 1$ and $\int \phi = \pi$, and it is what makes term by term integration legal rather than hopeful. **Link to the next block.** Once you have n^2 integrals, ask what each one depends on.

Block 2. **Cue.** The integrand $\phi(x - a_i)\phi(x - a_j)$ and an integral over the whole line. **Why natural.** Integration over $\mathbb{R}$ is translation invariant, so a shift $x = t + a_i$ costs nothing and reveals that the cross term is a function of $a_i - a_j$ alone. Structurally this says: the L^2 energy of a sum of translates of a fixed profile is a quadratic form in the pairwise displacements, with kernel equal to the autocorrelation of the profile. Recognise that and the analytic hypothesis has become an algebraic one. **Link to the next block.** Name the kernel f and put the hypothesis and the conclusion side by side.

Block 3. **Cue.** The conclusion's summand $1 + (a_i - a_j)^2$ is the reciprocal of $\phi(a_i - a_j)$, and the problem promises only the existence of a constant. **Why natural.** A Lorentzian has tails of order $1/y^2$, and the overlap of two Lorentzians separated by y is of order $1/y^2$ as well, so $f(y) \asymp 1/(1 + y^2)$ is what you expect before computing anything. Since only the order matters, the exact value of f is dead weight. Deciding what not to compute is the single most valuable habit in this problem. **Link to the next block.** You need the lower half of that estimate; get it in the cheapest way available.

Block 4. **Cue.** "Bound an integral of a positive function from below." **Why natural.** For a positive integrand there is exactly one cheap lower bound: throw away part of the domain and estimate the rest crudely. The choice of $[-1, 0]$ rather than $[0, 1]$ is not aesthetic. With $y \ge 0$ fixed, on $[-1, 0]$ the shifted argument $x + y$ stays in $[y - 1, y]$, so $|x + y| \le 1 + y$ with no case analysis, while both factors of the denominator remain controlled. A fixed length interval also guarantees the length factor is 1, so no length bookkeeping is needed. **Link to the next block.** The hypothesis is now an inequality between two explicit finite sums.

Block 5. **Cue.** Two sums over the same n^2 index pairs whose summands are reciprocals, one bounded above, the other to be bounded below. **Why natural.** The Cauchy and Schwarz inequality in the form $\left(\sum 1\right)^2 \le \left(\sum u\right)\left(\sum 1/u\right)$ is precisely the statement that a family and its reciprocals cannot both be small. Equivalently it is the arithmetic mean and harmonic mean inequality. Nothing else can turn an upper bound for $\sum 1/u$ into a lower bound for $\sum u$, because no pointwise inequality does so. **Link to the next block.** Substitute the numbers and read off the constant.

Block 6. **Cue.** The problem asks for the existence of $B > 0$, not for its value. **Why natural.** The arithmetic is $n^4/(6An) = n^3/(6A)$, so $B = 1/(6A)$. The only thing left to verify is that $A > 0$, which the case $n = 1$ gives for free since the left side of the hypothesis is then $\pi/2$. Checking that your constant is genuinely positive is the sort of edge case a grader looks for. **Link to the next block.** Nothing follows; the chain closes.

Layer 3: WHY

The truth beneath this problem is a single identity. For any $\phi \in L^2(\mathbb{R})$ and any reals $a_1, \dots, a_n$,

$$\int_{-\infty}^{\infty} \Big(\sum_{i=1}^{n} \phi(x - a_i) \Big)^2 dx \;=\; \sum_{i,j=1}^{n} K(a_i - a_j), \qquad K(y) = \int_{-\infty}^{\infty} \phi(t)\phi(t + y)\, dt.$$

The energy of a superposition equals a quadratic form in the pairwise displacements, and its kernel is the autocorrelation of the profile. That form is positive semidefinite for every choice of points, since the left side is an integral of a square; equivalently, K is a positive definite function in the sense of Bochner, and its Fourier transform $|\widehat{\phi}|^2$ is non-negative. So the hypothesis of this problem is not an arbitrary analytic condition: it says that a specific positive semidefinite quadratic form in the configuration $(a_1, \dots, a_n)$ grows at most linearly in n, which is a precise way of saying the points do not clump.

The second truth is the reciprocal duality. The hypothesis kernel behaves like $1/(1+y^2)$ and the conclusion kernel is exactly $1+y^2$. Reciprocal quantities are compared only by the Cauchy and Schwarz inequality (or its special case, the arithmetic mean and harmonic mean inequality), which is sharp when all u_{ij} are equal. That sharpness is worth internalising: the bound Bn^3 is what you get when all pairwise distances are comparable, that is, when the points sit on an arithmetic progression of spacing of order $\sqrt{n}$ scaled correctly, and in fact the truth is n^4 for the extremal configurations, which is why the remark that this can be improved is available. The problem deliberately asks for the weaker n^3 because n^3 is exactly what the one line Cauchy and Schwarz gives.

The discernment mechanism is exponent arithmetic, and you should perform it before writing any mathematics. There are n^2 pairs. The hypothesis is of size n. Cauchy and Schwarz over N terms with an upper bound M on the reciprocal sum gives N^2/M. Here $N^2/M = n^4/n = n^3$. If the problem had asked for n^4 you would know immediately that Cauchy and Schwarz alone cannot do it and that a counting argument over unit intervals is required. Reading the exponent tells you which tool is intended, every time.

To train this: for each analysis problem you meet, before starting, write down the number of terms, the size of the hypothesis, and the size of the conclusion, and check whether the conclusion is the Cauchy and Schwarz consequence of the hypothesis. If it is, the entire problem is the translation step from analysis into algebra, and all your effort should go there. Second, train the habit of asking "what am I allowed not to compute": every appearance of the phrase "there is a constant" deletes a page of work if you let it.

Revolving concepts

The Cauchy kernel and the Poisson kernel. $\phi(x) = 1/(1+x^2)$ has $\int \phi = \pi$ and $\int \phi^2 = \pi/2$. Normalised, $P_t(x) = \frac{t}{\pi(t^2+x^2)}$ is the Poisson kernel for the upper half plane; it is a probability density (the Cauchy distribution with scale t) and it forms a semigroup under convolution, $P_s * P_t = P_{s+t}$. Hence the exact value $f(y) = \pi^2(P_1 * P_1)(y) = \pi^2 P_2(y) = 2\pi/(4+y^2)$, obtained without a single residue. See Book 1, field: Calculus and integrals, topic: Poisson and Cauchy kernels.

Fourier transform and Plancherel. With $\widehat{\phi}(\xi) = \int \phi(x) e^{-i\xi x} dx$ one has $\widehat{\phi}(\xi) = \pi e^{-|\xi|}$ for the Cauchy kernel, and Plancherel gives $\int |\Phi_n|^2 = \frac{1}{2\pi} \int |\widehat{\phi}|^2 \left|\sum_j e^{-i\xi a_j}\right|^2 d\xi$. This is the frequency space reading of the same identity, and it is the reason the pairwise form is positive semidefinite. See Book 1, field: Real analysis, topic: Fourier transform, Plancherel, positive definite functions.

Autocorrelation. $K = \phi \star \phi$ with $K(y) = \int \phi(t)\phi(t+y)dt$; K is even when ϕ is even, $K(0) = \|\phi\|_2^2$ is the maximum, and $\widehat{K} = |\widehat{\phi}|^2 \geq 0$. Every problem about the energy of a superposition of translates reduces to K. See Book 1, field: Real analysis, topic: convolution and autocorrelation.

The Cauchy and Schwarz inequality in reciprocal form. For positive $u_1, \dots, u_N$, $\left(\sum u_k\right)\left(\sum 1/u_k\right) \geq N^2$, with equality if and only if all u_k are equal. Equivalent formulations: the arithmetic mean and harmonic mean inequality; the Engel form $\sum c_k^2/u_k \geq (\sum c_k)^2/\sum u_k$. Learn the Engel form as the default, since it specialises to both. See Book 1, field: Inequalities, topic: Cauchy and Schwarz, Engel form, arithmetic and harmonic means.

Bounding integrals from below. For a non-negative integrand, restricting to any subdomain gives a valid lower bound; choose a subdomain of fixed measure on which the integrand is uniformly comparable to its scale. This is the standard replacement for exact evaluation and it is nearly always

enough when the target is an order of magnitude. See Book 1, field: Real analysis, topic: crude estimates and domain restriction.

On recursion. There is no recursion in this problem, and you should not manufacture one. The structure that plays the role a recursion usually plays, namely the thing to generalise, is the pairwise quadratic form. Its generalisation is stated in Layer 3 for an arbitrary profile $\phi \in L^2$, and the same identity with m factors gives $\int \left(\sum_i \phi(x - a_i)\right)^m dx = \sum_{i_1,\ldots,i_m} K_m(a_{i_1}, \ldots, a_{i_m})$ with K_m an m point correlation depending only on differences. The counting refinement used for the n^4 improvement does contain an implicit recursion, namely the partition of $\mathbb{R}$ into unit intervals with occupancy numbers q_z satisfying $\sum_z q_z = n$, and there the standard tool is Jensen's inequality in the form $\sum_z q_z^2 \geq n^2/k$ when exactly k of the q_z are nonzero.

Circuit Bank

- **WHEN** a square of a sum inside an integral, with a conclusion indexed by pairs **DO** expand into cross terms before doing anything else, and verify each converges **BECAUSE** the cross terms are the only place the pairwise information lives, and any pointwise estimate on the sum destroys it
- **WHEN** an integral of a product of translates of one fixed profile **DO** substitute to make it depend only on the difference of the translations, and name that kernel **BECAUSE** integration over the line is translation invariant, so the energy of a superposition is always a quadratic form in the pairwise displacements
- **WHEN** the phrase "there is a constant $B > 0$" in the conclusion **DO** throw away every fixed numerical factor and refuse to evaluate any integral exactly **BECAUSE** the statement is about orders of magnitude only, so exact constants are ink spent on nothing
- **WHEN** an upper bound on $\sum 1/u_k$ and a required lower bound on $\sum u_k$ **DO** apply Cauchy and Schwarz in the form $(\sum u)(\sum 1/u) \geq N^2$ **BECAUSE** no pointwise inequality relates a positive number to its reciprocal in that direction, so only a global comparison can work
- **WHEN** a required lower bound for an integral of a positive function **DO** restrict to a fixed interval and bound the integrand crudely there **BECAUSE** positivity makes domain restriction free, and a fixed length interval removes all length bookkeeping

Backtesting

1. First useful cue: a squared sum inside an L^2 integral, with a conclusion summed over ordered pairs (i, j). That mismatch of index sets orders the expansion immediately.
2. Method the expression suggested: expand into pairwise overlap integrals, recognise them as a translation invariant kernel $f(a_i - a_j)$, bound f below by $1/(6(1 + y^2))$, then apply Cauchy and Schwarz against the reciprocal family.
3. Where you would hesitate, and the thought to reinforce: at the temptation to compute f exactly. Reinforce with the phrase "there is a constant $B > 0$", which is a licence to be crude, and with the observation that a Lorentzian overlap must be of order $1/y^2$ because Lorentzian tails are.
4. The trap the question sets: bounding the integrand before expanding, or applying Cauchy and Schwarz to the integral rather than to the pair sums. Both are legal and both lose the problem.
5. What makes you recognise it faster next time: count the index pairs, note the hypothesis has size n, and observe that $(n^2)^2/n = n^3$ is exactly the exponent demanded. When the exponent matches the Cauchy and Schwarz output, the intended solution is Cauchy and Schwarz and the rest of the problem is translation.

6. The reusable decision rule: convert every analytic hypothesis about a superposition into a pairwise algebraic sum first, then compare the two index sums and let the exponent choose the inequality.

Mirror drills

1. Evaluate $\int_{-\infty}^{\infty} \frac{dx}{(1+x^2)(1+(x+y)^2)}$ exactly. *Answer:* $2\pi/(4+y^2)$; write the integrand as $\pi^2 P_1(x)P_1(-(x+y))$ with P_t the Cauchy density of scale t and use $P_1 * P_1 = P_2$.
2. Let $u_1, \dots, u_N$ be positive with $\sum_k 1/u_k \le M$. Show $\sum_k u_k \ge N^2/M$. *Answer:* Cauchy and Schwarz on $\sqrt{u_k}$ against $1/\sqrt{u_k}$; equality when all u_k agree.
3. Suppose $\int_{-\infty}^{\infty} \left(\sum_{i=1}^{n} e^{-(x-a_i)^2}\right)^2 dx \le An$ for all n. Prove $\sum_{i,j=1}^{n} e^{(a_i-a_j)^2/2} \ge Bn^3$ for some $B > 0$. *Answer:* the cross term is $\sqrt{\pi/2}\, e^{-(a_i-a_j)^2/2}$ by completing the square, so the hypothesis bounds $\sum_{i,j} e^{-(a_i-a_j)^2/2}$ above by a multiple of n; now apply the same reciprocal Cauchy and Schwarz.

15.12 B6. Wilson reflection and double roots of a truncated exponential

Let p be an odd prime. Show that for at least $(p+1)/2$ values of n in $\{0, 1, 2, \ldots, p-1\}$,

$$\sum_{k=0}^{p-1} k!\, n^k \qquad \textit{is not divisible by } p.$$

Read it out loud

Strip the costume off. The expression $\sum_{k=0}^{p-1} k!\, n^k$ is the value at $x = n$ of the polynomial $f(x) = \sum_{k=0}^{p-1} k!\, x^k$ over the field $\mathbb{F}_p$, since every $k!$ with $k \leq p-1$ is a residue and every n is an element of $\mathbb{F}_p$. The statement "for at least $(p+1)/2$ values of n the value is nonzero" is the complement of "for at most $p - (p+1)/2 = (p-1)/2$ values of n the value is zero". So the entire problem is:

f has at most $(p-1)/2$ roots in $\mathbb{F}_p$.

Now audit the numbers. The polynomial f has degree $p-1$, so Lagrange's bound already gives at most $p-1$ roots. You are asked for at most $(p-1)/2$, which is exactly half of the trivial bound. That factor of two is the whole problem, and there is exactly one mechanism in the theory of polynomials over a field that halves a root count: multiplicity. If you can produce a polynomial in which every root of f appears at least twice, and whose degree is at most p, then $2N \leq p$ forces $N \leq (p-1)/2$ because p is odd. Everything after this observation is the search for that auxiliary polynomial.

Multiplicity means derivatives, so ask what happens when you differentiate. $f'(x) = \sum_{k \geq 1} k \cdot k!\, x^{k-1}$, whose coefficients are worse than the ones you started with. Differentiation likes factorials in the denominator, not the numerator. The only tool in existence that says anything about $k!$ modulo p is Wilson's theorem, and the only symmetry the coefficient list $0!, 1!, \ldots, (p-1)!$ possesses is the reflection $k \mapsto p-1-k$, under which Wilson's theorem gives $k!\,(p-1-k)! \equiv (-1)^{k+1}$. Applying that reflection replaces $k!$ by $1/k!$ and reverses the polynomial. The reversed object is $g(x) = \sum_{k=0}^{p-1} x^k/k!$, the truncated exponential, whose derivative is itself up to one term. That is the object you can differentiate. So the route is forced by the requirements of the endgame, not discovered by luck.

The fields wired in are number theory (Wilson, Fermat), polynomial algebra over a finite field (root counting with multiplicity, formal derivatives), and the structure of $\mathbb{F}_p[x]$ as a ring of functions. Within thirty seconds you should have on paper: "at most $(p-1)/2$ roots", "half of deg f, so double roots", and "differentiate: I need $1/k!$, so use Wilson to reverse".

One more reading, and it is the one that tells you the problem cannot be attacked head on. You are asked for at least $(p+1)/2$ good values of n but you are never asked to name one. There is a reason. Taking $n = 1$ gives $\sum_{k=0}^{p-1} k!$, and the assertion that this is never divisible by an odd prime p is Kurepa's conjecture, a long standing open problem. So no argument that identifies specific good residues can exist at the present state of knowledge. The problem must be, and is, a counting argument that never inspects a single value.

Approaches, ranked worst to best

1. **Tabulate and look for a pattern.** *Cue:* small modulus, finite check. *Circuit:* experimentation. Compute $f(n)$ for $p = 3, 5, 7, 11$ and hope the good n are describable. They are not: the set of roots of f has no known description, and as noted above even the single value $n = 1$ is an open problem. Experimentation here produces nothing but consumed time.

2. **Find a closed form for $\sum k! n^k$.** *Cue:* a sum with a clean summand. *Circuit:* generating functions. The generating function does satisfy a differential equation: over $\mathbb{F}_p$ one checks directly that

$x^2f'(x) + (x-1)f(x) + 1 = 0$, because $p! \equiv 0$ kills the truncation error. But this is a formal relation, not an evaluation, and sums of factorials have no closed form in any useful sense. Ranks second worst because it looks productive and is not, although the differential relation itself is worth keeping and reappears below.

3. **Force double roots directly on f using $x^p - x$.** *Cue:* multiplicity is needed, and $x^p - x$ vanishes on all of $\mathbb{F}_p$. *Circuit:* modify a polynomial's derivative without changing its values. This is the right idea applied to the wrong polynomial, and understanding why it fails is what makes the ideal route inevitable. From $x^2f' + (x-1)f + 1 = 0$, at any root z of f we get $z^2f'(z) = -1$, and $z \neq 0$ since $f(0) = 1$, so $f'(z) = -z^{-2} = -z^{p-3}$. To make the roots of f into double roots of $h = f + c(x)(x^p - x)$ we need $h'(z) = f'(z) - c(z) = 0$ at every root, so c must agree with $-x^{p-3}$ on the root set. Since the roots are unknown, you must take $c(x) = -x^{p-3}$, and then $\deg h = 2p - 3$, so the count $2N \leq 2p - 3$ yields only $N \leq p - 2$. The bound is useless. The obstruction is precise: f' takes a different value at each root, so no single cheap corrector exists.

4. **Padé approximation and a Hankel determinant of factorials.** *Cue:* "at most t roots", so assume more and derive a linear dependence. *Circuit:* a rational function with small numerator and denominator forces a Hankel matrix to be singular. This is Elkies' solution. Suppose f vanishes at all but $m < t = (p-1)/2$ nonzero residues, build $Q(x)$ of degree m vanishing at the surviving points, deduce that the coefficients of $x^t, \ldots, x^{2t-1}$ in $f(x)Q(x)$ vanish, and conclude that the Hankel matrix $A = ((i+j+1)!)_{i,j=0}^{t-1}$ kills a nonzero vector. Then invoke $\det\big((i+j+n)!\big)_{i,j=0}^{m} = \prod_{k=0}^{m} k!\,(k+n)!$, which is nonzero modulo p because every factorial involved is a factorial of an integer less than p. Complete, beautiful, and far more expensive: it needs Padé style reasoning and a binomial matrix row reduction to evaluate the Hankel determinant. Rank it second because it is a real solution, not a sketch, but it costs three times the ink.

5. **Wilson reflection to the truncated exponential, then double roots via $x^p - x$.** *Cue:* coefficients $k!$ plus a target of exactly half the degree. *Circuit:* Wilson's reflection reverses factorials into reciprocal factorials; the truncated exponential is its own derivative up to the top term; $x^p - x$ tunes derivatives for free. Four short steps, one standard theorem, no machinery beyond the formal derivative. This is the ideal route.

The ideal solution

Convention. Read the sum as the value at $x = n$ of a polynomial: the $k = 0$ term is the constant $0! \cdot x^0 = 1$ for every n, including $n = 0$. This is the ordinary convention for the constant term of a polynomial and it is the only reading under which the expression is a polynomial function of n. (If one instead insists on $0^0 = 0$, then $n = 0$ is deleted from the count and the argument below delivers $(p-1)/2$ good values rather than $(p+1)/2$.)

Work in $\mathbb{F}_p$ throughout. Define

$$f(x) = \sum_{k=0}^{p-1} k!\,x^k, \qquad g(x) = \sum_{k=0}^{p-1} \frac{x^k}{k!},$$

both polynomials over $\mathbb{F}_p$; every $k!$ with $0 \leq k \leq p-1$ is a product of nonzero residues and is therefore invertible modulo p, so g is well defined. The claim to be proved is that f has at most $(p-1)/2$ roots in $\mathbb{F}_p$.

Lemma 15.12.1 (Wilson reflection)**.** *For $0 \leq k \leq p-1$ we have $k!\,(p-1-k)! \equiv (-1)^{k+1} \pmod p$.*

Proof. Write $(p-1)! = k! \prod_{m=k+1}^{p-1} m$. As m runs over $k+1, \ldots, p-1$, the residue $p-m$ runs over $1, \ldots, p-1-k$, and $m \equiv -(p-m)$, so $\prod_{m=k+1}^{p-1} m \equiv (-1)^{p-1-k}(p-1-k)!$. Since p is odd,

$(-1)^{p-1-k} = (-1)^k$. By Wilson's theorem $(p-1)! \equiv -1$, so $-1 \equiv (-1)^k k!\,(p-1-k)!$, which rearranges to the stated congruence. □

Lemma 15.12.2 (reversal)**.** *For every nonzero $z \in \mathbb{F}_p$ we have $f(z) = -\,g(-1/z)$. Consequently f and g have exactly the same number of roots in $\mathbb{F}_p$.*

Proof. By the previous lemma, $k! = (-1)^{k+1}/(p-1-k)!$ in $\mathbb{F}_p$. Hence for any $z \neq 0$,

$$f(z) = \sum_{k=0}^{p-1} \frac{(-1)^{k+1} z^k}{(p-1-k)!} = \sum_{j=0}^{p-1} \frac{(-1)^{p-j}\, z^{\,p-1-j}}{j!},$$

where $j = p-1-k$ and $(-1)^{k+1} = (-1)^{p-j}$. Since p is odd, $(-1)^{p-j} = -(-1)^j$, so

$$f(z) = -\,z^{p-1} \sum_{j=0}^{p-1} \frac{(-1)^j z^{-j}}{j!} = -\,z^{p-1}\, g(-1/z) = -\,g(-1/z),$$

using $z^{p-1} = 1$ by Fermat's little theorem. For the count: $f(0) = 1 \neq 0$ and $g(0) = 1 \neq 0$, so all roots of both polynomials are nonzero, and $z \mapsto -1/z$ is a bijection of $\mathbb{F}_p^\times$ onto itself. Hence $f(z) = 0$ if and only if $g(-1/z) = 0$, and the two root sets are in bijection. □

Lemma 15.12.3 (the truncated exponential is almost its own derivative)**.** $g'(x) = g(x) + x^{p-1}$ *in* $\mathbb{F}_p[x]$. *In particular, if $g(z) = 0$ for some $z \in \mathbb{F}_p$ then $z \neq 0$ and $g'(z) = 1$.*

Proof. Differentiating termwise,

$$g'(x) = \sum_{k=1}^{p-1} \frac{k\, x^{k-1}}{k!} = \sum_{k=1}^{p-1} \frac{x^{k-1}}{(k-1)!} = \sum_{j=0}^{p-2} \frac{x^j}{j!} = g(x) - \frac{x^{p-1}}{(p-1)!}.$$

By Wilson's theorem $(p-1)! = -1$ in $\mathbb{F}_p$, so $-x^{p-1}/(p-1)! = x^{p-1}$, giving $g' = g + x^{p-1}$. If $g(z) = 0$ then $z \neq 0$ because $g(0) = 1$, so $z^{p-1} = 1$ and $g'(z) = 0 + 1 = 1$. □

Proof of the required bound. Set

$$h(x) = x^p - x + g(x) \in \mathbb{F}_p[x].$$

Since $\deg g = p-1$ (its leading coefficient is $1/(p-1)! = -1 \neq 0$), h has degree exactly p with leading coefficient 1; in particular h is not the zero polynomial. Its formal derivative is

$$h'(x) = p\, x^{p-1} - 1 + g'(x) = -1 + g(x) + x^{p-1},$$

because $p = 0$ in $\mathbb{F}_p$.

Let $z \in \mathbb{F}_p$ satisfy $g(z) = 0$. By Fermat's little theorem $z^p = z$ for every element of $\mathbb{F}_p$, so $h(z) = z - z + g(z) = 0$. By the previous lemma $z \neq 0$ and $z^{p-1} = 1$, so $h'(z) = -1 + 0 + 1 = 0$. A common root of h and h' is a root of h of multiplicity at least two: writing $h = (x-z)q$ gives $h' = q + (x-z)q'$, so $h'(z) = q(z) = 0$ and hence $(x-z)^2$ divides h.

Let N be the number of roots of g in $\mathbb{F}_p$. The roots are distinct elements, so $\prod_z (x-z)^2$ divides the nonzero polynomial h, whence $2N \leq \deg h = p$. As p is odd, $N \leq (p-1)/2$.

By the reversal lemma, f also has exactly $N \leq (p-1)/2$ roots in $\mathbb{F}_p$. Therefore the number of $n \in \{0, 1, \ldots, p-1\}$ with $f(n) \neq 0$ in $\mathbb{F}_p$ is at least

$$p - \frac{p-1}{2} = \boxed{\frac{p+1}{2}},$$

which is exactly the assertion of the problem. □

Layer 1: WHAT

Block 1. Restate: the sum is a polynomial f of degree $p-1$ over $\mathbb{F}_p$, and the claim is that f has at most $(p-1)/2$ roots.

Block 2. Audit: $(p-1)/2$ is half the trivial Lagrange bound, so the mechanism must be multiplicity. You need an auxiliary polynomial of degree at most p in which every root of f is at least double.

Block 3. Multiplicity needs derivatives; f' has coefficients $k \cdot k!$, which are worse. Reverse the coefficient list using Wilson's reflection, turning $k!$ into $1/k!$ and f into the truncated exponential g.

Block 4. Verify the reversal is a genuine bijection on root sets: $f(z) = -g(-1/z)$ for $z \neq 0$, and neither polynomial vanishes at 0.

Block 5. Differentiate g: $g' = g + x^{p-1}$, so at any root, $g' = 1$. The roots of g are simple, which is the opposite of what you want, but the derivative there is the same constant at every root, which is what makes the repair cheap.

Block 6. Repair with $x^p - x$: adding it changes no value on $\mathbb{F}_p$ and shifts the derivative by exactly -1, converting $g' = 1$ into $h' = 0$. The degree rises only to p. Count: $2N \leq p$, so $N \leq (p-1)/2$.

Juxtaposition. The nominal route attacks f where it stands. It either tries to evaluate the sum (impossible, and in the single case $n = 1$ an open problem) or tries the multiplicity argument on f itself. The latter is the instructive failure: because $f'(z) = -z^{p-3}$ varies over the root set, the corrector must be the polynomial $-x^{p-3}$ rather than a constant, the auxiliary polynomial has degree $2p-3$, and the count collapses to the vacuous $N \leq p-2$.

Bifurcation point. The split is the line "I need a polynomial whose derivative at every one of its roots is the same constant". Once that requirement is written down, the reflection to $1/k!$ is not an idea, it is the only way to satisfy the requirement, because the truncated exponential is the unique polynomial family that reproduces itself under differentiation. A candidate who never writes that requirement will apply the correct circuit to the wrong polynomial and lose the problem while feeling that the method worked.

Layer 2: HOW

Block 1. **Cue.** "for at least $(p+1)/2$ values of n in $\{0, 1, \ldots, p-1\}$" combined with a summand that is a polynomial in n with integer coefficients. **Why natural.** Counting residues at which an expression is nonzero is never done by inspecting residues; it is done by counting the roots of the corresponding polynomial over $\mathbb{F}_p$, because Lagrange's theorem is the only quantitative statement in the subject. Passing to the complement converts "at least $(p+1)/2$ nonzero" into "at most $(p-1)/2$ zero", which is a bound of the shape that root counting produces. **Link to the next block.** Compare the required bound with the degree.

Block 2. **Cue.** $\deg f = p-1$ and the required bound is $(p-1)/2$, exactly one half. **Why natural.** Over a field, the sum of the multiplicities of the roots of a nonzero polynomial is at most its degree. Therefore the only way to beat the naive bound by a factor of two is to make each root count twice. This is not one option among several; there is no other source of a factor of two in polynomial root counting. **Link to the next block.** Multiplicity is detected by the derivative, so ask whether you can differentiate the polynomial you have.

Block 3. **Cue.** The coefficients are $k!$, and differentiation multiplies the coefficient of x^k by k. **Why natural.** Factorials in the numerator get worse under differentiation and factorials in the denominator get better: $\frac{d}{dx}\frac{x^k}{k!} = \frac{x^{k-1}}{(k-1)!}$. So you want the reciprocal coefficients. The only theorem about factorials modulo p is Wilson's, and the only symmetry available on the index set $\{0, \ldots, p-1\}$ is $k \mapsto p-1-k$; putting the two together gives $k! = (-1)^{k+1}/(p-1-k)!$, which reverses the polynomial into the truncated exponential. Nothing here is invention: you specify what you need

and Wilson is the unique supplier. **Link to the next block.** Reversing a polynomial changes the variable, so you must check the change of variable respects the root count.

Block 4. **Cue.** The reversal produces z^{p-1} and z^{-j}, both meaningful only for $z \neq 0$. **Why natural.** $z^{p-1} = 1$ on $\mathbb{F}_p^\times$ by Fermat, which collapses the prefactor and leaves the identity $f(z) = -g(-1/z)$. The map $z \mapsto -1/z$ is a bijection of $\mathbb{F}_p^\times$, and both f and g have constant term 1, so neither vanishes at 0 and no root is created or destroyed. Checking that the substitution is a bijection on the set you are counting over is the step a careless write-up skips, and it is exactly where marks are lost. **Link to the next block.** You now hold a polynomial designed to be differentiated. Differentiate it.

Block 5. **Cue.** g is the truncated exponential, and $\exp' = \exp$. **Why natural.** Termwise differentiation gives $g' = g - x^{p-1}/(p-1)!$, and Wilson converts the correction into $+x^{p-1}$. On the root set of g this reads $g'(z) = z^{p-1} = 1$. Read that carefully: it says the roots of g are all simple, which appears to be a defeat, but it also says the derivative takes the *same* value at every root. That constancy is the asset. A correction that is constant costs no degree. **Link to the next block.** You need to subtract 1 from the derivative without disturbing any value on $\mathbb{F}_p$. Ask which polynomials are invisible as functions.

Block 6. **Cue.** $x^p - x$ vanishes at every element of $\mathbb{F}_p$ but is not the zero polynomial. **Why natural.** The evaluation map $\mathbb{F}_p[x] \to \{\text{functions } \mathbb{F}_p \to \mathbb{F}_p\}$ has kernel exactly the ideal generated by $x^p - x$. That ideal is your free parameter: adding a multiple of $x^p - x$ changes nothing about values and everything about derivatives, since $(x^p - x)' = -1$. Adding it once shifts every derivative by -1, which is precisely what turns $g'(z) = 1$ into $h'(z) = 0$, and it raises the degree only from $p-1$ to p. The degree budget is the reason this works and the reason the same manoeuvre on f does not: there the required corrector is $-x^{p-3}$, not a constant, and the degree balloons to $2p-3$. **Link to the next block.** With h of degree p and every root of g doubled, the count $2N \leq p$ closes the problem.

Layer 3: WHY

Three truths carry this problem and each one generalises.

The first is Lagrange's theorem with multiplicity: over any field, a nonzero polynomial of degree d has at most d roots counted with multiplicity. It is the only counting device available, so every problem of the form "the number of residues with property P is at most M" must be arranged into a polynomial whose degree is M, or into one of degree $2M$ with double roots, or $3M$ with triple roots. Reading the ratio between the degree you have and the bound you want tells you which multiplicity you are being asked to produce. Here the ratio is two, so you look for double roots, and the derivative is what detects them: z is a multiple root of h if and only if $h(z) = h'(z) = 0$, a fact that holds over any field, with the standard proof by writing $h = (x-z)q$.

The second is that over $\mathbb{F}_p$ polynomials and functions are different things, and the difference is exactly the ideal $(x^p - x)$. Every function $\mathbb{F}_p \to \mathbb{F}_p$ is a polynomial function, and two polynomials induce the same function if and only if they differ by a multiple of $x^p - x$. This means you have a free parameter of degree p available at every moment. You spend it to adjust derivative behaviour without touching values. The cost is degree, and since your entire argument is a degree budget, the question is always: how cheap is the corrector? The design decision in this problem is the choice of a polynomial whose derivative is a single constant on its own root set, because that makes the corrector a constant and the cost exactly p. The failed attempt on f is not a different idea, it is the same idea with a corrector of degree $p-3$, and it fails on budget alone. That comparison is the sharpest thing in the problem and you should be able to say it in one sentence in the arena.

The third is Wilson's theorem, which exists in your toolkit for exactly one purpose: it is the multiplicative structure of $\mathbb{F}_p^\times$ compressed into a statement about $(p-1)!$. Proof in one line:

pairing each element with its inverse in $\mathbb{F}_p^\times$ leaves only the self inverse elements ± 1 unpaired, so $(p-1)! \equiv -1$. Its most useful corollary is the reflection $k!\,(p-1-k)! \equiv (-1)^{k+1}$, which is the statement that the sequence of factorials is, up to signs, its own reversal in the reciprocal sense. Whenever a problem hands you factorials as coefficients modulo p, the reflection is the first thing you write, because there is nothing else to write.

The discernment mechanism, stated for reuse: compare the bound demanded with the degree you have; the ratio names the multiplicity you must manufacture; the manufacture requires a polynomial you can differentiate cleanly; if the one you are given cannot be differentiated cleanly, apply the unique available symmetry of its coefficients to convert it into one that can. To train it, take any problem asserting "at most half the residues satisfy ..." and force yourself to name, before doing anything, the auxiliary polynomial of degree at most p whose double roots you will count.

A final honest word on the shape of the result. The bound $(p-1)/2$ on the number of roots of the truncated exponential over $\mathbb{F}_p$ is far from what is believed to be true; heuristically the number of roots should be around 1 on average. The problem asks for a bound that a single degree count can prove, not for the truth. Recognising the gap between "provable in ten lines" and "true" is part of reading a Putnam B6 correctly: you are being asked for the cheapest rigorous statement, and you should stop the moment you have it.

Revolving concepts

Wilson's theorem. $(p-1)! \equiv -1 \pmod{p}$ for prime p, proved by pairing each element of $\mathbb{F}_p^\times$ with its multiplicative inverse and noting that only 1 and -1 are self inverse. The converse also holds: if $(n-1)! \equiv -1 \pmod{n}$ then n is prime. Corollary used here: $k!\,(p-1-k)! \equiv (-1)^{k+1}$, and in particular $\binom{p-1}{k} \equiv (-1)^k$. See Book 1, field: Number theory, topic: Wilson's theorem and factorial reflection.

Fermat's little theorem and the polynomial $x^p - x$. $z^p = z$ for all $z \in \mathbb{F}_p$, hence $x^p - x = \prod_{a \in \mathbb{F}_p} (x - a)$. The evaluation map from $\mathbb{F}_p[x]$ to the p^p functions $\mathbb{F}_p \to \mathbb{F}_p$ is surjective with kernel $(x^p - x)$, so $\mathbb{F}_p[x]/(x^p - x)$ is the ring of all functions. Its derivative is -1, which makes it the standard device for adjusting derivatives at prescribed points without changing values. See Book 1, field: Number theory, topic: polynomial functions over a finite field.

Formal derivative and multiplicity. Over any commutative ring the formal derivative is defined coefficientwise; over a field, z is a root of h of multiplicity at least two if and only if $h(z) = h'(z) = 0$, and more generally of multiplicity at least m if and only if the first $m-1$ Hasse derivatives vanish (over characteristic p you must use Hasse derivatives, not iterated ordinary derivatives, once m reaches p). The separability criterion $\gcd(h, h') = 1$ is the same fact in different clothes. See Book 1, field: Polynomials, topic: formal derivatives, multiplicity, separability.

Truncated exponential and the recursion to generalise. The coefficient sequence $c_k = 1/k!$ satisfies the first order recursion $c_k = c_{k-1}/k$, and the polynomial family $e_m(x) = \sum_{k=0}^{m} x^k/k!$ satisfies

$$e_m'(x) = e_m(x) - \frac{x^m}{m!},$$

the general statement of which the lemma above is the case $m = p-1$ over $\mathbb{F}_p$. This is the general principle: a first order recursion $c_k = c_{k-1}\lambda(k)$ on coefficients becomes a first order differential relation on the generating polynomial, with an error term supplied by the truncation. The mirror statement for the reversed sequence $c_k = k!$, satisfying $c_k = k\,c_{k-1}$, is $x^2 f'(x) + (x-1)f(x) + 1 = 0$, which over $\mathbb{F}_p$ holds exactly for the truncation because $p! \equiv 0$. Whenever coefficients obey a first order recursion, write down the corresponding differential relation before doing anything else: it is where all information about derivatives at roots comes from. See Book 1, field: Polynomials, topic: generating polynomials of first order recursions.

Reversal of a polynomial. For $\deg q = d$, the reversal is $x^d q(1/x)$, which reverses the coefficient list; combined with a scalar substitution $x \mapsto \lambda/x$ it is a bijection of $\mathbb{F}_p^\times$ onto itself and therefore

preserves root counts away from 0. Always confirm that the constant term and leading coefficient are nonzero before transporting a count across a reversal. See Book 1, field: Polynomials, topic: reciprocal and reversed polynomials.

Hankel determinants of factorials. $\det\big((i+j+n)!\big)_{i,j=0}^{m} = \prod_{k=0}^{m} k!\,(k+n)!$, proved by turning the matrix into $\big(\binom{i+j+n}{i}\big)$ and observing that decreasing n by one is a sequence of elementary row operations, so the determinant of the binomial version is independent of n and equals 1. Needed only for Elkies' alternative solution. See Book 1, field: Matrices and determinants, topic: Hankel and binomial determinants.

Knowledge gaps, named without disguise. To attempt this problem at all you must have installed: $\mathbb{F}_p$ as a field and Lagrange's root bound; the formal derivative and the multiple root criterion; Wilson's theorem with its reflection corollary; Fermat's little theorem in the polynomial form $x^p - x = \prod_a (x-a)$; and the fact that the kernel of the evaluation map is $(x^p - x)$. None of these appear in A-level or Further Mathematics. Two further items are optional context rather than requirements: Hasse derivatives, which are what you need if you ever want multiplicity three or higher in characteristic p; and the Artin and Hasse exponential together with the arithmetic of truncated exponentials modulo p, which is the natural home of the polynomial g and explains why its root structure is hard. Finally, be clear that Kurepa's conjecture, the assertion that $\sum_{k=0}^{p-1} k! \not\equiv 0 \pmod{p}$ for every odd prime, is open; the problem is soluble only because it never asks about a specific n.

Circuit Bank

- **WHEN** "for at least N of the residues modulo p, this expression is nonzero" **DO** pass to the complement and count the roots of the corresponding polynomial over $\mathbb{F}_p$ **BECAUSE** Lagrange's bound is the only quantitative tool for counting residues, and it counts roots, not non-roots
- **WHEN** a required root bound that is exactly half the degree **DO** manufacture an auxiliary polynomial of degree at most twice the bound in which every root is at least double **BECAUSE** multiplicity is the only mechanism in polynomial algebra that improves a root count by a fixed factor
- **WHEN** a need to change a polynomial's derivative at points of $\mathbb{F}_p$ without changing any of its values **DO** add a multiple of $x^p - x$, and count the degree that costs **BECAUSE** the kernel of the evaluation map is exactly $(x^p - x)$, whose derivative is -1, so it is a free parameter for derivative data and a priced one for degree
- **WHEN** factorials appearing as coefficients modulo p **DO** apply Wilson's reflection $k!(p-1-k)! = (-1)^{k+1}$ and consider the reversed polynomial **BECAUSE** Wilson is the only theorem about factorials modulo p, and $k \mapsto p-1-k$ is the only symmetry the coefficient list has
- **WHEN** a substitution used to transport a count from one polynomial to another **DO** prove it is a bijection of the set being counted, and check the excluded point separately **BECAUSE** a count transported across a non-bijection is not a count, and the point $x = 0$ is where reversals always break

Backtesting

1. First useful cue: the target $(p+1)/2$ out of p, which complements to "at most $(p-1)/2$ roots", read against $\deg f = p-1$. The ratio of exactly two is the cue.
2. Method the expression suggested: produce double roots. Since f cannot be differentiated cleanly, use Wilson's reflection to reverse f into the truncated exponential g, differentiate that, and repair the derivative with the free polynomial $x^p - x$.

3. Where you would hesitate, and the thought to reinforce: at the appearance of $x^p - x$, which looks unmotivated. Reinforce it as the kernel of the evaluation map: it is the unique way to change derivatives while fixing all values, and its own derivative is -1, which is exactly the amount you need to subtract.
4. The trap the question sets: applying the multiplicity circuit to f itself. It is the right circuit, the algebra is valid, and the degree budget silently destroys the bound because f' is not constant on its root set. The trap punishes candidates who fire circuits without auditing degrees.
5. What makes you recognise it faster next time: the requirement audit. Before choosing a polynomial, write the property you need from it, here "the derivative must take one and the same nonzero constant value at every root", and only then ask which available transformation supplies it.
6. The reusable decision rule: a root bound equal to half the degree means double roots; double roots mean a derivative you control; if the given polynomial's derivative is uncontrolled, reverse the coefficients using the only available theorem about them and try again.

Mirror drills

1. Show that $\sum_{k=0}^{p-1} k!\, x^k$ and $\sum_{k=0}^{p-1} x^k/k!$ have the same number of roots in $\mathbb{F}_p$. *Answer:* Wilson's reflection turns one coefficient list into the other, giving $f(z) = -g(-1/z)$ for $z \neq 0$; both have constant term 1, and $z \mapsto -1/z$ is a bijection of $\mathbb{F}_p^\times$.
2. Prove that every root of $g(x) = \sum_{k=0}^{p-1} x^k/k!$ in $\mathbb{F}_p$ is a simple root. *Answer:* termwise differentiation and Wilson give $g' = g + x^{p-1}$; a root z is nonzero since $g(0) = 1$, so $g'(z) = z^{p-1} = 1 \neq 0$.
3. Let $q \in \mathbb{F}_p[x]$ have degree less than p, and suppose there is a single nonzero constant c with $q'(z) = c$ for every root z of q in $\mathbb{F}_p$. Show that q has at most $p/2$ roots in $\mathbb{F}_p$. *Answer:* put $h = q + c\,(x^p - x)$, so h has degree p, every root of q satisfies $h(z) = q(z) = 0$ and $h'(z) = c - c = 0$, hence is a double root of h; now count multiplicities against $\deg h = p$.

Chapter 16

Putnam 2010

16.1 A1. splitting 1 through n into equal sums

Given a positive integer n, what is the largest k such that the numbers $1, 2, \ldots, n$ can be put into k boxes so that the sum of the numbers in each box is the same? (When $n = 8$, the example $\{1, 2, 3, 6\}, \{4, 8\}, \{5, 7\}$ shows that the largest k is at least 3.)

Read it out loud

This is a partition question wearing an optimisation costume. "Largest k" tells you the answer is one closed formula in n, not a case description, because a Putnam optimisation answer is always exact. Strip the costume off and the whole content is: distribute $1, \ldots, n$ into k nonempty disjoint groups sharing one invariant, the sum.

The field wired into this is elementary combinatorics, the arithmetic series and a pigeonhole style counting bound, nothing more exotic. What is really being tested is whether you reach, inside thirty seconds, for the two things every equal-invariant partition problem needs: an upper bound on k, and a construction proving that bound is tight. The bracketed $n = 8$ example is a gift and a trap at once. It proves $k \geq 3$ for one specific n, but it is not shaped like the general construction, so do not try to reverse engineer the formula from it.

The first useful cue is that n itself is one of the numbers being distributed, and it is the largest of them. Within thirty seconds you should have written: whichever box holds n already has sum at least n. That one sentence is the entire upper bound.

Approaches, ranked worst to best

1. **trial and error on small n.** *Cue:* the bracketed hint for $n = 8$. *Circuit:* pattern spotting. (you can find the pattern by hand for $n = 1, \ldots, 8$, but a guess is not a proof of either half, the bound or the construction, so this only orients you.)
2. **pigeonhole on box sizes.** *Cue:* k boxes, n numbers, no box empty. *Circuit:* counting argument. (let j be the number of singleton boxes; every other box holds at least 2 numbers, so $n \geq j + 2(k - j) = 2k - j$, giving $j \geq 2k - n$. Once $k > (n+1)/2$ this forces $j \geq 2$, and two singleton boxes hold two different numbers, hence two different sums, a contradiction. Correct, but the variable j and the case split cost more ink than the problem needs.)
3. **the box containing n.** *Cue:* n is the largest number being distributed. *Circuit:* extremal element bound. (the box holding n already has sum at least n with no further argument; combined with the fixed total this delivers the bound in two lines and points straight at the construction, pairing numbers to a constant sum.)

The ideal solution

Write $S = 1 + 2 + \cdots + n = \dfrac{n(n+1)}{2}$. Suppose $1, \ldots, n$ are placed in k boxes, each with common sum s. No box is empty, since an empty box has sum 0, which would force every box to sum to 0, impossible once $n \geq 1$. Summing the k equal box sums recovers the total: $ks = S$.

The number n lies in some box, and every number in a box is a positive integer, so that box's sum is at least n: $s \geq n$. Then

$$k = \frac{S}{s} \leq \frac{S}{n} = \frac{n+1}{2},$$

and since k is an integer, $k \leq \left\lfloor \dfrac{n+1}{2} \right\rfloor$.

This bound is attained. If $n = 2m$, form the m pairs $\{1, 2m\}, \{2, 2m-1\}, \ldots, \{m, m+1\}$, each summing to $2m + 1$. If $n = 2m + 1$, form the singleton $\{n\}$ together with the m pairs

$\{1, n-1\}, \{2, n-2\}, \ldots, \{m, m+1\}$, each summing to n. Both constructions use $\lfloor (n+1)/2 \rfloor$ boxes, matching the bound, so

$$\boxed{k_{\max} = \left\lfloor \frac{n+1}{2} \right\rfloor = \left\lceil \frac{n}{2} \right\rceil .}$$

Layer 1: WHAT

Block 1: the box containing n has sum at least n. Block 2: multiply the common sum s by k to recover the fixed total S, combine with block 1 to bound $k \leq (n+1)/2$. Block 3: construct a partition attaining the bound, split by the parity of n. Block 4: state and box the closed form.

Against the nominal route, guessing from the $n = 8$ example or running the heavier pigeonhole on singleton boxes, the ideal route reaches the bound in two lines because it uses the single most informative fact available: n is already sitting in some box. The bifurcation point is exactly there. The zombie response to "largest k" is to start counting boxes and sizes in the abstract; the trained response is to ask what the single largest object already forces.

Layer 2: HOW

Block 1. **Cue.** The numbers being distributed are $1, \ldots, n$, and n is manifestly the largest. **Why natural.** Every number sharing a box with n is a positive integer, so it can only add to the sum; a box containing n has sum at least n with no further work. **Link to the next block.** This lower bound on the common sum s is exactly the quantity that controls k, since k and s are reciprocally tied through the fixed total.

Block 2. **Cue.** The sum of the numbers in each box is the same. **Why natural.** A partition sends every number into exactly one box, so the k equal box sums must add to the grand total S; this is counting the same numbers twice, once by box and once as a whole. **Link to the next block.** Substituting $s \geq n$ into $ks = S$ turns a statement about sums into a statement about the count k, the object the question actually maximises.

Block 3. **Cue.** A bound with no matching example proves nothing about what is achievable. **Why natural.** Whenever $1, \ldots, n$ needs a constant sum split, the reflex is Gauss's pairing, i with $n+1-i$, since these pairs are engineered to sum to $n+1$ regardless of i; the odd case leaves the unpaired middle term n sitting alone. **Link to the next block.** This construction meets the block 2 bound with equality, so the two halves of the proof now agree.

Block 4. **Cue.** Bound and construction coincide at $\lfloor (n+1)/2 \rfloor$. **Why natural.** A Putnam optimisation answer is complete only once both halves match exactly. **Link to the next block.** State the closed form.

Layer 3: WHY

The mechanism underneath is a general principle about partitioning a totally ordered set under an equal invariant constraint: the extremal element is already forced into some part, and its own magnitude is a ready made bound on that part's invariant, with no combinatorial search required. This converts a question about the number of parts, which you would otherwise have to search for, into a question about the size of one specific part, readable directly off the ordering of the set.

Train the discernment by asking, first, whenever a problem partitions $1, \ldots, n$ into groups sharing a common invariant: which group must the extreme element occupy, and what does that force. Pigeonhole is not wrong, the second ranked approach proves the same bound, but it answers a more general question, how many singleton boxes must exist, than the one asked, how large

can the common sum be, and answering the more general question costs more ink for the same conclusion. The examiners hand you n itself as manifestly the largest object being partitioned; recognising that the set is already ordered, and the largest element is non negotiable, is the entire skill.

Revolving concepts

Arithmetic series. $1 + 2 + \cdots + n = n(n+1)/2$ is the fixed total every equal sum partition must divide exactly. See Book 1, field: Sequences and recurrences, topic: arithmetic series.

Gauss pairing. Pairing i with $n + 1 - i$ splits $1, \ldots, n$ into pairs of constant sum $n + 1$; it resurfaces whenever a problem wants a constant sum or constant product decomposition of consecutive integers. See Book 1, field: Combinatorics, topic: pairing arguments.

Pigeonhole principle. If n objects sit in k nonempty boxes and $n < 2k$, some structural constraint is forced, here at least two singleton boxes. Keep this as a second method precisely because it generalises to partition problems where no single element dominates the invariant. See Book 1, field: Combinatorics, topic: pigeonhole principle.

Circuit Bank

- **WHEN** $1, \ldots, n$ is partitioned into groups sharing a common additive invariant, and you are asked for the largest or smallest number of groups **DO** find which group the largest, or smallest, element must occupy, and read off the bound it forces before any general counting **BECAUSE** the extremal element is already committed to some group, and its magnitude is the cheapest possible constraint on that group's invariant
- **WHEN** an upper bound on a count k falls out of a fixed total against a per group minimum **DO** immediately look for the matching construction that attains equality, splitting by parity if needed **BECAUSE** a largest k such that answer is a closed form, and a bound alone never proves a formula exact
- **WHEN** you need $1, \ldots, n$ split into pairs, or a near pairing, of constant sum **DO** pair i with $n + 1 - i$, leaving n or nothing unpaired according to parity **BECAUSE** this pairing is engineered so every pair sums to exactly $n + 1$, independent of i

Backtesting

1. First useful cue: n, the largest number being partitioned, must sit in some box.
2. Method the expression suggested: bound the common sum below by n, then use the fixed total to bound k.
3. Where you would hesitate, and the thought to reinforce: whether boxes can be empty; reinforce that an empty box forces every box to sum to zero, impossible once $n \geq 1$.
4. The trap the question sets: the bracketed $n = 8$ example looks like the general construction; pattern matching off it directly, without deriving the parity split, produces an unproven guess.
5. What makes you recognise it faster next time: any phrase of the form "distributed among groups with a shared sum" is your cue to check what the largest single object already forces.
6. The reusable decision rule: extremal element bounds its group's invariant; combine with the fixed total for an upper bound on the group count; construct a matching partition, typically by Gauss pairing, to prove the bound tight.

Mirror drills

1. State the largest k for splitting $1, \ldots, 15$ into equal sum boxes. *Answer:* $n = 15$ odd, $m = 7$, so $k = 8$: the singleton $\{15\}$ plus the pairs $\{1, 14\}, \ldots, \{7, 8\}$, each summing to 15.
2. Split $1, \ldots, n$ into boxes of equal sum where every box must hold exactly two elements. Why does the extremal element argument become trivial here? *Answer:* $k = n/2$ automatically; the size constraint removes all freedom, so there is nothing left to optimise.
3. Ask the same question for equal *product* rather than equal sum. Does the extremal bound still work cleanly? *Answer:* the box containing n still has product at least n, but the total product $n!$ does not distribute as evenly across a fixed per box minimum as $n(n+1)/2$ does; this is genuinely harder, not a free restatement.

16.2 A2. differentiability forces f to be linear

Find all differentiable functions $f : \mathbb{R} \to \mathbb{R}$ such that

$$f'(x) = \frac{f(x+n) - f(x)}{n}$$

for all real numbers x and all positive integers n.

Read it out loud

Say the equation back to yourself before touching it: the instantaneous rate of change of f at x equals the average rate of change of f over the window $[x, x+n]$, for every integer window length n, and this holds identically, not at some unspecified intermediate point. That last clause is the whole problem. The Mean Value Theorem already guarantees SOME point in $(x, x+n)$ where the instantaneous rate matches the average rate; this equation demands it happen at the left endpoint itself, for every positive integer n at once. A statement that rigid about a merely differentiable function should smell like it forces a very narrow family of solutions.

The field wired in is real analysis, the derivative as a limit of difference quotients, and functional equations, an unknown function constrained by an identity rather than by an initial value problem. What is really being tested is whether you notice the equation is not one condition but an infinite family indexed by n, and whether you have the discipline to ask which members of that family actually carry information, rather than trying to use "for all n" all at once.

The first useful cue, inside thirty seconds you should have written the $n = 1$ instance,

$$f'(x) = f(x+1) - f(x),$$

because it is the simplest member of the family and it converts the right side into the most standard object in analysis, a forward difference. Everything else in the solution is what you do with this one line and its $n = 2$ neighbour.

Approaches, ranked worst to best

1. **Taylor expansion in n.** *Cue:* $f(x+n)$ looks like it wants a series expansion. *Circuit:* formal power series matching. (assuming f is smooth enough to expand $f(x+n) = \sum_k f^{(k)}(x) n^k / k!$ and matching coefficients is an unjustified leap: the problem only grants f differentiable once, and even granting the expansion, matching a formal series against an exact identity does not by itself prove convergence or validity. Costs ink to patch and is never fully sound on the hypotheses given.)

2. **guess $f = cx + d$ from $n = 1$ and verify.** *Cue:* the $n = 1$ relation looks linear-shaped. *Circuit:* pattern matching plus sufficiency check. (spotting and verifying the answer is necessary but not sufficient, the problem says find all, and worse, the $n = 1$ relation alone genuinely does not pin f down, see Layer 3, so a write up that only ever invokes $n = 1$ cannot be a complete proof no matter how it is phrased.)

3. **combine $n = 1$ and $n = 2$, telescope, differentiate.** *Cue:* two consecutive integer witnesses of the same family. *Circuit:* minimal witness extraction. (telescoping the $n = 2$ relation into two $n = 1$ gaps forces periodicity of f' in two lines, and a further differentiation of the $n = 1$ relation itself collapses that periodicity to a constant. Uses only that f is once differentiable, and only two of the infinitely many hypotheses on offer.)

The ideal solution

Sufficiency. If $f(x) = cx + d$ for constants $c, d \in \mathbb{R}$, then $f'(x) = c$, and $(f(x+n) - f(x))/n = (c(x+n) + d - cx - d)/n = c$ for every positive integer n and every real x. So every such f works.

Necessity. Suppose f has the stated property. Write the $n = 1$ instance, valid for every real t,

$$f'(t) = f(t+1) - f(t). \tag{1}$$

and the $n = 2$ instance,

$$2f'(t) = f(t+2) - f(t). \tag{2}$$

Telescope the right side of (2) through (1) applied at t and again at $t+1$:

$$f(t+2) - f(t) = \big(f(t+2) - f(t+1)\big) + \big(f(t+1) - f(t)\big) = f'(t+1) + f'(t).$$

Comparing with (2), $2f'(t) = f'(t+1) + f'(t)$, so

$$f'(t+1) = f'(t) \qquad \text{for every } t \in \mathbb{R}. \tag{3}$$

This is a value level identity, established by pure algebra on (1) and (2); it does not require differentiating f'.

Now define $g : \mathbb{R} \to \mathbb{R}$ by $g(x) = f(x+1) - f(x)$. Since f is differentiable everywhere, so are the two shifted functions $x \mapsto f(x+1)$ and $x \mapsto f(x)$, hence so is their difference g, and by the chain rule applied to f itself, not to f',

$$g'(x) = f'(x+1) - f'(x).$$

By (3) this is 0 for every x, so g is constant on $\mathbb{R}$; write $g(x) = c$ for all x, where $c = g(0)$. But by (1), $g(x) = f'(x)$ identically, so $f'(x) = c$ for every x. A function with derivative c everywhere on $\mathbb{R}$ differs from cx by a constant, so

$$\boxed{f(x) = cx + d \text{ for some real constants } c, d.}$$

Layer 1: WHAT

Block 1: verify sufficiency of the guessed form, to know exactly what shape the proof must arrive at. Block 2: write the $n = 1$ relation as an identity in t. Block 3: write the $n = 2$ relation and telescope it through two copies of the $n = 1$ relation, forcing $f'(t+1) = f'(t)$. Block 4: rename the $n = 1$ relation's right side as $g(x) = f(x+1) - f(x)$, differentiate g directly from its definition, using block 3 to conclude $g' \equiv 0$, hence f' constant. Block 5: integrate to the final closed form.

Against the nominal route, expanding formally in powers of n or stopping once a plausible answer is verified, the ideal route uses only the two smallest witnesses of the hypothesis and never assumes more regularity than the problem grants. The bifurcation point sits exactly at block 3 to block 4: the zombie instinct is to try to differentiate the periodicity relation $f'(t+1) = f'(t)$ as though it were a statement about a differentiable function f'; the trained instinct is to notice that relation is a fact about function *values*, and to instead differentiate a different, provably differentiable object, g, whose derivative happens to equal exactly that difference of values.

Layer 2: HOW

Block 1. **Cue.** The right side of the equation, $(f(x+n) - f(x))/n$, is a difference quotient over an integer step. **Why natural.** A difference quotient is one thing above all others: cheap to instantiate at the smallest available step. Setting $n = 1$ costs nothing and immediately turns the abstract family into the single cleanest member. **Link to the next block.** The $n = 1$ relation becomes the definition you will keep returning to throughout the proof, both as an equation in its own right and, later, as the definition of the auxiliary function g.

Block 2. **Cue.** Two members of the same family, $n = 1$ and $n = 2$, both hold simultaneously for the same f. **Why natural.** Whenever a family of relations shares a variable, the value of combining two members is that their difference, or in this case their telescoped relationship, can eliminate the unknown function entirely and leave a statement purely about f'. **Link to the next block.** Telescoping $f(t+2) - f(t)$ through the $n = 1$ relation at t and at $t+1$ expresses the $n = 2$ right side using nothing but f', so the $n = 2$ equation collapses to a bare relation between $f'(t)$ and $f'(t+1)$.

Block 3. **Cue.** $f'(t+1) = f'(t)$ looks like a statement crying out to be differentiated. **Why natural.** It cannot be differentiated directly, since f' is only known to exist, not to be itself differentiable; the honest move is to notice that $f'(x)$ already equals $f(x+1) - f(x)$ by the $n = 1$ relation, so the object worth differentiating is the right side, $g(x) = f(x+1) - f(x)$, which is manifestly differentiable because f is. **Link to the next block.** Differentiating g by the ordinary chain rule on f produces $g'(x) = f'(x+1) - f'(x)$, and block 2's periodicity says this is zero, so g, and therefore f', is constant.

Block 4. **Cue.** f' is now known to be a constant c. **Why natural.** A function with everywhere derivative c is cx plus a constant, this is the single variable integration you have used since school. **Link to the next block.** Box the final family and match it against the sufficiency check from the start, closing the loop.

Layer 3: WHY

The deep question this problem hides is why exactly two witnesses, $n = 1$ and $n = 2$, suffice, when the hypothesis is stated for every positive integer n. The honest answer, beyond A level, comes from treating the $n = 1$ relation on its own as a linear delay differential equation and asking for its exponential modes $f(x) = e^{\lambda x}$. Substituting gives $\lambda e^{\lambda x} = e^{\lambda x}(e^{\lambda} - 1)$, so $\lambda = e^{\lambda} - 1$. Among real λ this forces $\lambda = 0$ (the function $e^{\lambda} - 1 - \lambda$ has derivative $e^{\lambda} - 1$, negative for $\lambda < 0$ and positive for $\lambda > 0$, so it is strictly minimised, at value 0, exactly at $\lambda = 0$), but a transcendental equation of this shape is known to admit infinitely many further roots off the real axis, in the complex plane. Each such root would contribute a genuine extra solution to the $n = 1$ relation alone, so $n = 1$ by itself does not logically force linearity, only real exponential rigidity, which is why the second ranked approach above is marked incomplete rather than merely inelegant.

Now bring in $n = 2$: the same substitution gives $2\lambda = e^{2\lambda} - 1$. Since $e^{\lambda} = \lambda + 1$ from the first equation, squaring gives $e^{2\lambda} = (\lambda + 1)^2 = \lambda^2 + 2\lambda + 1$, so the second equation becomes $2\lambda = \lambda^2 + 2\lambda$, that is $\lambda^2 = 0$, so $\lambda = 0$ is now the *only* root common to both equations, real or complex. This is the reason two witnesses, not one, are load bearing: $n = 1$ alone leaves room for oscillatory or growing complex modes, and it is specifically the second relation, $n = 2$, that annihilates every one of them simultaneously. The discernment mechanism to train: when a functional equation is imposed "for all n" but your proof only ever seems to need finitely many instances, do not treat that as a coincidence of the write up, ask what the extra instances are ruling out, since the answer usually reveals the true rigidity of the equation and tells you, in advance, the minimum witness set to reach for.

Revolving concepts

Mean Value Theorem. For f differentiable on $[a, b]$, some $c \in (a, b)$ has $f'(c) = (f(b) - f(a))/(b - a)$. The given equation is a far stronger demand, equality at the left endpoint itself, for every integer window, which is why it pins f down completely rather than merely constraining it. See Book 1, field: Real analysis, topic: Mean Value Theorem.

Method of differences and telescoping. Any statement about $f(x+n) - f(x)$ can be rewritten as a telescoping sum of n consecutive forward differences, $f(x+n) - f(x) = \sum_{i=0}^{n-1} \big(f(x+i+1) -$

$f(x+i)\big)$. This is the general device whenever a step n relation and a step 1 relation on the same function must be reconciled, and it generalises beyond $n = 2$: for any n, the same telescoping combined with the $n = 1$ relation forces $\sum_{i=0}^{n-1} f'(x+i) = nf'(x)$, a genuinely different looking constraint for each n worth checking against the periodicity already established. See Book 1, field: Sequences and recurrences, topic: method of differences.

Zero derivative implies constant. A function differentiable on all of $\mathbb{R}$ with $g' \equiv 0$ is constant, itself a corollary of the Mean Value Theorem applied to any two points. It is the final joint that closes the argument, turning a value level periodicity into an outright constant. See Book 1, field: Real analysis, topic: consequences of the Mean Value Theorem.

Circuit Bank

- **WHEN** a functional equation is imposed for all positive integers n, but you suspect it is over determined **DO** test whether the smallest one or two witnesses, $n = 1$ and $n = 2$, already force the conclusion, before trying to use the full family at once **BECAUSE** over determined hypotheses usually collapse to a minimal witness set, and identifying it shortens the proof and exposes the equation's true rigidity
- **WHEN** you have a value level identity such as $f'(t+1) = f'(t)$ but do not know f' is itself differentiable **DO** rename the original defining relation's right side as a new function g built from shifts of f, and differentiate g instead, using only that f is differentiable **BECAUSE** shifting and subtracting a differentiable function never demands extra regularity, so this route never smuggles in an assumption the problem did not grant
- **WHEN** a step n relation and a step 1 relation hold for the same function **DO** telescope the step n gap into n consecutive step 1 gaps **BECAUSE** $f(x+n) - f(x) = \sum_{i=0}^{n-1}(f(x+i+1) - f(x+i))$ is always available and converts an unfamiliar gap into gaps you already have a formula for

Backtesting

1. First useful cue: the $n = 1$ instance turns the right side into the forward difference $f(x+1) - f(x)$.
2. Method the expression suggested: write the $n = 1$ and $n = 2$ relations as identities in a dummy variable, then combine them.
3. Where you would hesitate, and the thought to reinforce: wanting to differentiate $f'(t+1) = f'(t)$ directly; reinforce that only f, not f', is known to be differentiable, so differentiate the shifted difference function g instead, whose derivative equals exactly that value level statement.
4. The trap the question sets: assuming every one of the infinitely many hypotheses, one per n, must be invoked, when in fact exactly two suffice and the rest are automatically verified once the closed form is known.
5. What makes you recognise it faster next time: seeing "derivative equals an average rate of change" phrasing is your cue to isolate the smallest one or two witnesses rather than reasoning about the family in the abstract.
6. The reusable decision rule: verify sufficiency of the guessed closed form first; for necessity, telescope two adjacent witnesses to extract a periodicity of the derivative, then differentiate the original defining relation's right side, not the derivative itself, to collapse that periodicity to a constant.

Mirror drills

1. Check directly, without redoing the proof, that $f(x) = 7x - 3$ satisfies the equation for $n = 5$. *Answer:* $f'(x) = 7$; $(f(x+5) - f(x))/5 = (7(x+5) - 3 - (7x-3))/5 = 35/5 = 7$, matching.
2. Solve the discrete analogue: find all $f : \mathbb{Z} \to \mathbb{R}$ with $\Delta f(x) = f(x+n) - f(x))/n$ where $\Delta f(x) := f(x+1) - f(x)$, again for all positive integers n. *Answer:* $n = 1$ is a tautology, but $n = 2$ gives $2\Delta f(x) = f(x+2) - f(x) = \Delta f(x+1) + \Delta f(x)$, forcing $\Delta f(x+1) = \Delta f(x)$, a constant first difference, so f is an arithmetic progression, the discrete shadow of $f(x) = cx + d$.
3. Suppose the hypothesis held only for $n = 1$ and $n = 3$, skipping $n = 2$. Does the telescoping argument still force $f'(t+1) = f'(t)$ directly? *Hint:* telescoping $n = 3$ through three $n = 1$ gaps gives $3f'(t) = f'(t) + f'(t+1) + f'(t+2)$, that is $2f'(t) = f'(t+1) + f'(t+2)$, a genuinely different looking relation from the $n = 1, n = 2$ case; do not assume any two witnesses behave identically without checking.

16.3 A3. why a bounded directional eigenfunction must vanish

Suppose that the function $h : \mathbb{R}^2 \to \mathbb{R}$ has continuous partial derivatives and satisfies the equation

$$h(x,y) = a\frac{\partial h}{\partial x}(x,y) + b\frac{\partial h}{\partial y}(x,y)$$

for some constants a, b. Prove that if there is a constant M such that $|h(x,y)| \leq M$ for all $(x,y) \in \mathbb{R}^2$, then h is identically zero.

Read it out loud

This is a partial differential equation wearing an analysis problem's clothes. Read the right side out loud: $a\,\partial h/\partial x + b\,\partial h/\partial y$, the same two constants a, b multiplying h_x and h_y respectively. That exact combination has a name, the directional derivative of h in the direction (a, b), so the equation says nothing more exotic than "the directional derivative of h in a fixed direction equals h itself." Once you have said that sentence to yourself, the problem has already told you what to do: directional derivatives become ordinary one variable derivatives the instant you restrict to the line running in that exact direction.

The field wired in is multivariable calculus, partial derivatives and the chain rule along a parametrised path, collapsing into one variable ODE theory the moment you make that restriction. What is really being tested is restraint: whether you try to solve the PDE in general, a formula for h valid across both variables, when the problem only asks you to prove h is identically zero, a far narrower target that one well chosen line through each point can hit directly, without ever writing down a general solution formula.

The first useful cue, inside thirty seconds you should have written: $ah_x + bh_y$ is $\nabla h \cdot (a, b)$, the directional derivative in direction (a, b); and the second, immediately after: restrict to $(x_0 + at,\ y_0 + bt)$ and see what the chain rule hands you for free.

Approaches, ranked worst to best

1. **guess a separated form $h(x,y) = X(x)Y(y)$.** *Cue:* two variables, one equation. *Circuit:* separation of variables. (nothing in the hypotheses forces h to separate; even after finding the separable solutions of the PDE you would still need a separate argument that no other bounded solution exists, more work for a narrower target than the problem asks for.)

2. **solve the PDE in full generality by the method of characteristics.** *Cue:* a first order linear constant coefficient PDE. *Circuit:* general characteristic solution formula. (correct in principle: one can derive $h(x,y) = e^t\Phi(\text{coordinate orthogonal to } (a,b))$ for an arbitrary function Φ along characteristic coordinates, then impose the bound on that formula. This works, but it spends ink deriving the entire two parameter family of solutions when the problem only needs pointwise vanishing, and setting up characteristic coordinates correctly is exactly where an algebra slip costs you marks.)

3. **restrict to the single characteristic line through an arbitrary point.** *Cue:* the matching constants a, b on both partials. *Circuit:* directional derivative reduction. (fix one point, follow the one line through it in direction (a, b), reduce to the ODE $\phi' = \phi$, and let the global bound do the rest. Reaches the conclusion pointwise, with no machinery beyond the ordinary chain rule and a single elementary ODE.)

The ideal solution

If $a = b = 0$, the given equation reads $h(x,y) = 0$ for every (x,y) directly, and there is nothing further to prove. So assume from now on that $(a,b) \neq (0,0)$, a genuine direction in the plane.

Fix an arbitrary point $(x_0, y_0) \in \mathbb{R}^2$, and let $L(t) = (x_0 + at,\, y_0 + bt)$ for $t \in \mathbb{R}$ be the line through (x_0, y_0) in direction (a, b). Since h has continuous partial derivatives, h is differentiable at every point, so the multivariable chain rule applies to the composition $\phi(t) := h(L(t))$, giving

$$\phi'(t) = a\,\frac{\partial h}{\partial x}(L(t)) + b\,\frac{\partial h}{\partial y}(L(t)).$$

The given equation, evaluated at the point $L(t)$, says the right side equals $h(L(t)) = \phi(t)$. So

$$\phi'(t) = \phi(t) \qquad \text{for all } t \in \mathbb{R}.$$

This is the most elementary linear ODE there is, with general solution $\phi(t) = \phi(0)e^t$.

Since $|h(x, y)| \le M$ for every $(x, y) \in \mathbb{R}^2$, in particular $|\phi(t)| \le M$ for every $t \in \mathbb{R}$. But $|\phi(t)| = |\phi(0)|e^t \to \infty$ as $t \to \infty$ unless $\phi(0) = 0$. A bounded function cannot equal an unbounded one, so $\phi(0) = 0$, that is,

$$h(x_0, y_0) = \phi(0) = 0.$$

Since (x_0, y_0) was an arbitrary point of $\mathbb{R}^2$, this holds everywhere, so

$$\boxed{h(x, y) \equiv 0 \text{ on } \mathbb{R}^2.}$$

Layer 1: WHAT

Block 1: dispatch $a = b = 0$ in one line. Block 2: fix an arbitrary point and build the line through it in direction (a, b). Block 3: restrict h to that line and apply the chain rule, converting the PDE into $\phi' = \phi$. Block 4: solve the ODE and use the global bound to force $\phi(0) = 0$. Block 5: universally quantify over the arbitrary starting point to conclude $h \equiv 0$.

Against the nominal route, deriving the full two variable solution formula for h by characteristic coordinates before imposing the bound, the ideal route never writes down a general solution at all. The bifurcation point sits at the very first reading of the equation: the zombie response to "a PDE with continuous partial derivatives" is to reach for general PDE solving machinery; the trained response is to notice that $ah_x + bh_y$ is definitionally a directional derivative, and that a directional derivative restricted to its own direction is just an ordinary derivative in disguise.

Layer 2: HOW

Block 1. **Cue.** The equation is stated "for some constants a, b," which explicitly includes the possibility $a = b = 0$. **Why natural.** If $a = b = 0$, substituting into the right side gives 0 with no calculation at all, so the conclusion is free; always clear a degenerate parameter case first in a "for some constants" problem, so the main argument can assume the parameters are not simultaneously trivial. **Link to the next block.** This lets everything that follows treat (a, b) as a genuine nonzero direction in the plane, which is exactly what is needed to define a line worth restricting to.

Block 2. **Cue.** The right side is $ah_x + bh_y$, the same constants multiplying each partial, which is the definitional fingerprint of a directional derivative $\nabla h \cdot (a, b)$. **Why natural.** Recognising this dictionary entry is the whole solve: a directional derivative is, by definition, the ordinary derivative of h restricted to a line running in exactly that direction, so the equation is quietly telling you which line to draw before you have drawn anything. **Link to the next block.** Fixing an arbitrary point and following the line $(x_0 + at, y_0 + bt)$ through it converts a two variable hypothesis into a one variable object you can differentiate directly.

Block 3. **Cue.** h has continuous, not merely existing, partial derivatives. **Why natural.** Continuity of the partials is exactly the standard sufficient condition for h to be differentiable at each point, which is what licenses the multivariable chain rule along the path $L(t)$; partial derivatives can exist everywhere without the function being differentiable if they fail to be continuous, so this hypothesis is not decorative, it is load bearing. **Link to the next block.** Applying the chain rule to $\phi(t) = h(L(t))$ reproduces exactly $ah_x + bh_y$ evaluated along the line, which by the given PDE equals $h(L(t)) = \phi(t)$ itself, so the chain rule computation and the hypothesis collapse onto the same expression.

Block 4. **Cue.** $\phi' = \phi$ is the simplest first order linear ODE available. **Why natural.** Its solution $\phi(t) = \phi(0)e^t$ is immediate, and exponential growth is the one behaviour categorically incompatible with a uniform bound over an unbounded domain. **Link to the next block.** The global bound on h restricts to a bound on ϕ over all $t \in \mathbb{R}$, so $\phi(0)$, which is exactly h at the original point, must be 0.

Block 5. **Cue.** The starting point (x_0, y_0) was never given any special property. **Why natural.** A statement proved for an arbitrary element of a set is proved for the whole set; this is ordinary universal quantification, but it is worth stating explicitly rather than leaving it implied. **Link to the next block.** Conclude $h \equiv 0$ on all of $\mathbb{R}^2$.

Layer 3: WHY

The truth underneath this problem is an eigenfunction rigidity principle. Along every line in direction (a, b), h behaves as an eigenfunction of the directional derivative operator $D_{(a,b)}$ with eigenvalue 1, and an eigenvalue 1 solution of a linear autonomous first order ODE is, with no exceptions, a genuine exponential, Ce^t, unless $C = 0$. Exponential growth is the unique obstruction to boundedness on an infinite ray, so a global bound annihilates the solution completely, one characteristic line at a time, and since the characteristic lines through every point of $\mathbb{R}^2$ in direction (a, b) sweep out the entire plane, the pointwise conclusion becomes a global one for free.

This generalises past the specific eigenvalue 1 given here. If the equation instead read $c \cdot \nabla h = \lambda h$ for a constant $\lambda \neq 0$ and a fixed direction c, the identical restriction gives $\phi' = \lambda\phi$, so $\phi(t) = \phi(0)e^{\lambda t}$; for $\lambda > 0$ the same argument forces $\phi(0) = 0$ by letting $t \to +\infty$, while for $\lambda < 0$ the blow up happens as $t \to -\infty$ instead, and boundedness on the whole line still forces $\phi(0) = 0$ by that end. Only $\lambda = 0$ escapes this, since then ϕ is itself constant and boundedness imposes no constraint. The discernment mechanism to train: whenever a first order linear PDE has constant coefficients and expresses a linear combination of the partials as a constant multiple of the function itself, restrict to the characteristic line through an arbitrary point in the coefficient direction; this always converts the PDE into an exactly solvable one variable ODE, and any global bound hypothesis can then be tested against that ODE's solution directly, with no need to ever write down h as a function of both variables at once.

Revolving concepts

Directional derivative. $D_v h(x_0, y_0) = ah_x(x_0, y_0) + bh_y(x_0, y_0) = \nabla h(x_0, y_0) \cdot v$ for $v = (a, b)$, equal to $\frac{d}{dt}\big|_{t=0} h(x_0 + at, y_0 + bt)$. Recognising an expression as a directional derivative is the single most useful pattern match in any two variable calculus problem carrying constant coefficients on both partials. See Book 1, field: Calculus and integrals, topic: directional derivatives.

Differentiability from continuous partial derivatives. If the partial derivatives of h exist in a neighbourhood of a point and are continuous there, h is differentiable at that point, and the chain rule applies to h composed with any differentiable path through it. This is the standard lemma that licenses every chain rule computation in this solution; existence of partials alone does not suffice,

the classical counterexample being a function with well defined partials at the origin that is not even continuous there. See Book 1, field: Real analysis, topic: differentiability in several variables.

Method of characteristics. The line $L(t) = (x_0 + at, y_0 + bt)$ is the characteristic curve of the PDE $ah_x + bh_y = h$ through (x_0, y_0); restricting any first order linear constant coefficient PDE to its characteristics is the general technique of which this solution is the minimal, single line instance. See Book 1, field: Calculus and integrals, topic: first order linear PDEs.

Exponential incompatibility with global boundedness. $Ce^{\lambda t}$ for $\lambda \neq 0$ is unbounded on at least one end of $\mathbb{R}$, so any hypothesis forcing an exponential solution, combined with a two sided bound, forces the constant C to vanish. See Book 1, field: Real analysis, topic: growth rates and boundedness.

Circuit Bank

- **WHEN** an equation contains $af_x + bf_y$ with the same constants multiplying each partial **DO** recognise it as the directional derivative $D_{(a,b)}f$ and restrict to the line through an arbitrary point in that direction **BECAUSE** this is definitionally what the expression means, and restricting collapses the two variable statement into a one variable ODE via the ordinary chain rule
- **WHEN** a restricted single variable function satisfies $\phi' = \lambda\phi$ for a nonzero constant λ, and the original function is globally bounded **DO** compare the exponential solution $\phi(0)e^{\lambda t}$ against the bound as t runs toward whichever end makes $e^{\lambda t}$ blow up, and conclude $\phi(0) = 0$ **BECAUSE** $Ce^{\lambda t}$ is unbounded on at least one end of the real line unless $C = 0$, and boundedness of the original function restricts to boundedness of every line through it
- **WHEN** a problem grants "continuous partial derivatives" rather than merely "partial derivatives exist" **DO** treat that continuity as licensing the multivariable chain rule along any path, not as decorative phrasing **BECAUSE** partials can exist at a point without the function being differentiable there if they are not continuous, and the chain rule genuinely needs differentiability

Backtesting

1. First useful cue: matching constants a, b multiply h_x and h_y, the directional derivative fingerprint.
2. Method the expression suggested: restrict h to the line through an arbitrary point in direction (a, b), converting the PDE into an ODE via the chain rule.
3. Where you would hesitate, and the thought to reinforce: whether $a = b = 0$ needs separate handling; reinforce that without it, direction (a, b) is not even genuine, and the case costs nothing since the equation then gives $h \equiv 0$ directly.
4. The trap the question sets: reaching for the general two variable solution formula of the PDE when the problem only needs pointwise vanishing, spending ink on machinery the proof never uses.
5. What makes you recognise it faster next time: any first order PDE of the form (constant coefficient linear combination of partials) equals (a multiple of the function) is your cue to restrict to the line in the coefficient direction.
6. The reusable decision rule: dispatch the degenerate parameter case first; for the genuine case, restrict to the characteristic line through an arbitrary point, reduce to $\phi' = \lambda\phi$, and compare the exponential solution against the given global bound to force the value at that point to zero.

Mirror drills

1. Confirm the argument still applies when $b = 0$, say $h = 2h_x$. *Answer:* the line is horizontal, $L(t) = (x_0 + 2t, y_0)$, $\phi(t) = h(L(t))$ still satisfies $\phi' = \phi$ by the same chain rule computation, so $h(x_0, y_0) = 0$ for every point exactly as before; $b = 0$ needs no special handling once $(a, b) \neq (0, 0)$.
2. Does the sign of a matter? Take $h = -h_x$ (so $a = -1, b = 0$). *Answer:* no; the chain rule reproduces $ah_x + bh_y$ along the line regardless of the sign of a, so $\phi' = \phi$ still holds and the same exponential blow up argument forces $h \equiv 0$; only the literal direction the line runs changes, not the eigenvalue.
3. Suppose instead $-h = ah_x + bh_y$, eigenvalue -1 rather than 1, still with $|h| \leq M$. Does h still vanish identically? *Answer:* yes; now $\phi' = -\phi$, so $\phi(t) = \phi(0)e^{-t}$, which blows up as $t \to -\infty$ instead of $t \to +\infty$; boundedness over all real t still forces $\phi(0) = 0$, only the end of the line doing the work has switched.

16.4 A4. Towers of ten factored by base swapping

Prove that for each positive integer n, the number

$$10^{10^{10^n}} + 10^{10^n} + 10^n - 1$$

is not prime.

Read it out loud

This is a primality claim about a number written in base 10, assembled from a three storey tower of tens with a -1 bolted onto the end. Said out loud, the fields announce themselves: elementary number theory, and inside that, divisibility read off a base 10 representation. Nothing else is in the room. No counting, no analysis, no algebra past congruences. And there is a standing rule about the command term. When a Putnam problem says *prove that this is not prime* about an object you could never write down, it is not asking for a factorisation. It is asking you to exhibit one divisor and to prove that the divisor is proper. Your entire job is to find the modulus.

Put $S = 10^n$ and $T = 10^{10^n}$ and look at the digits. The -1 never reaches the two leading ones; it eats into 10^n alone, and $10^n - 1$ is a run of n nines. So

$$N = 1\underbrace{0\cdots0}_{T-S-1}1\underbrace{0\cdots0}_{S-n}\underbrace{9\cdots9}_{n},$$

a string of $T + 1$ digits. That picture belongs on your paper inside thirty seconds, because a digit string made of two isolated ones and a tail of nines fires the schoolbook divisibility circuits at once: digit sums for 3 and 9, alternating digit sums for 11.

The digit sum is $1 + 1 + 9n \equiv 2 \pmod 3$, so 3 and 9 are dead and you lose nothing by checking. Now the alternating sum. Both S and T are even, being powers of ten with positive exponent, so the two isolated ones sit in even positions and contribute $+1$ each, while the n nines contribute $9(1 - 1 + 1 - \cdots)$, which is 9 for odd n and 0 for even n. Hence $N \equiv 2 + 9 \equiv 0 \pmod{11}$ when n is odd, and $N \equiv 2 \pmod{11}$ when n is even. Half the problem has fallen to a test you learnt at twelve.

What is really being tested is the next sentence, and it is a sentence about failure. The test died for even n for exactly one reason: casting out elevens is the statement $10 \equiv -1 \pmod{11}$, which turns 10^n into $(-1)^n$, and that is $+1$ and useless when n is even. So you need a modulus in which the ground floor of the tower still lands on -1. That is the whole problem: strip the powers of two out of n by changing base.

Approaches, ranked worst to best

1. **Compute it and factor it.** *Cue:* "prove it is not prime" misread as "produce the factorisation". *Circuit:* no circuit at all, pure consumption. Even at $n = 1$ the number has $10^{10} + 1$ digits, so no machine will ever hold it. This is not a method, it is a wish, and it ranks last.

2. **Hunt a polynomial identity in $x = 10^n$.** *Cue:* a sum of powers of ten ending in -1 smells like a factorisable form. *Circuit:* Sophie Germain, difference of two squares, cyclotomic factorisation. Setting $x = 10^n$ turns the tail into $x - 1$ but turns the middle term into 10^x, which is not a polynomial in x. The three floors are never simultaneously polynomial in one variable, so the route stalls on line one. Worth ten seconds, because problems such as "$m^4 + 4$ is composite for $m > 1$" fall to precisely this circuit, but not one second longer.

3. **Blind search over small primes.** *Cue:* no structural idea, so test $7, 13, 37, 41, 101, \ldots$ in turn. *Circuit:* brute force modular arithmetic. It finds 11 for odd n and 101 for $n \equiv 2 \pmod 4$, then stalls at $n = 4$ unless you happen to try 10001. Right answers with no reason attached: no reusable circuit, and no way to know you have covered every n.

4. **Casting out elevens.** *Cue:* two ones in even digit positions and a tail of nines. *Circuit:* $10 \equiv -1 \pmod{11}$, so base ten collapses to a sign. Instant, complete for odd n, dead for even n. It ranks second because its failure is informative: it isolates the one feature of n that mattered, the parity of n itself.

5. **Run the same digit test in base 10^{2^m}.** *Cue:* parity of n was the only obstruction, and parity is the $m = 0$ case of the 2 adic valuation. *Circuit:* when a sign argument dies on an even exponent, change base until the exponent becomes odd. The ideal route. Four lines, every n at once, and it contains the previous route as its own special case $m = 0$.

The ideal solution

Lemma 16.4.1. *For every integer $m \geq 0$ we have $2^m \geq m + 1$.*

Proof. Induction on m. For $m = 0$ both sides equal 1. If $2^m \geq m + 1$, then $2^{m+1} = 2 \cdot 2^m \geq 2m + 2 \geq m + 2$, since $m \geq 0$. □

Fix a positive integer n and put

$$N = 10^{10^{10^n}} + 10^{10^n} + 10^n - 1.$$

Write $n = 2^m k$ with $m \geq 0$ an integer and k a positive odd integer, so that 2^m is the exact power of two dividing n. Put

$$d = 2^m, \qquad M = 10^d + 1.$$

Then $10^d \equiv -1 \pmod{M}$ by construction, so for every nonnegative integer j,

$$10^{dj} \equiv (-1)^j \pmod{M}. \tag{$*$}$$

Everything now reduces to writing each of the three exponents as d times an integer and reading off the parity of that integer.

Ground floor. $n = dk$ with k odd, so $(*)$ gives $10^n \equiv -1 \pmod{M}$.

First floor. By the Lemma, $n \geq 2^m \geq m + 1$, so 2^{m+1} divides 2^n, which divides 10^n. Hence $10^n = dj$ with $j = 10^n / 2^m$ an even integer, and $(*)$ gives $10^{10^n} \equiv 1 \pmod{M}$.

Second floor. Also $10^n \geq n \geq m + 1$, so 2^{m+1} divides 2^{10^n}, which divides 10^{10^n}. Hence $10^{10^n} = dj'$ with j' even, and $(*)$ gives $10^{10^{10^n}} \equiv 1 \pmod{M}$.

Adding the four terms,

$$N \equiv 1 + 1 + (-1) - 1 = 0 \pmod{M},$$

so M divides N. It remains to check that M is a proper divisor. On the one hand $M = 10^d + 1 \geq 11 > 1$. On the other hand $d = 2^m \leq n$, so $M \leq 10^n + 1$; and since $10^n \geq n + 1$ for every $n \geq 1$,

$$N > 10^{10^n} \geq 10^{n+1} = 10 \cdot 10^n > 10^n + 1 \geq M.$$

Thus $1 < M < N$ and $M \mid N$, so N is composite. ■

Remark 16.4.2. For odd n this is $M = 11$, the schoolbook test. For $n \equiv 2 \pmod{4}$ it is $M = 101$. For $n \equiv 4 \pmod{8}$ it is $M = 10001 = 73 \times 137$, which is not itself prime, and that costs nothing: a proper divisor is not required to be prime.

Layer 1: WHAT

Block 1. Read the command term. "Not prime" for an uncomputable number means: exhibit a divisor M and prove $1 < M < N$.

Block 2. Write N as a digit string. Two isolated ones at the even positions 10^n and 10^{10^n}, and n nines at the bottom.

Block 3. Fire the base 10 tests. Digit sum kills nothing; the alternating sum gives $11 \mid N$ exactly when n is odd.

Block 4. Diagnose the failure. The only thing that broke for even n is the parity of the exponent in $10^n \equiv (-1)^n$.

Block 5. Repair by base swapping. Replace base 10 by $B = 10^{2^m}$, where 2^m is the exact power of two dividing n. Then $n/2^m$ is odd by construction, and $B \equiv -1 \pmod{B+1}$.

Block 6. Check the upper floors. Their exponents are B times an even number, because 10^n and 10^{10^n} carry gigantic powers of two; the only arithmetic you need is $2^m \geq m+1$.

Block 7. Size check, then conclude.

The nominal route runs Block 1, Block 2, Block 3, and then does the zombie thing: "eleven failed, try thirteen, try thirty seven, try one hundred and one". You will burn twenty minutes and finish with a case analysis you cannot close, because for each new n you would need a new prime found by luck.

The bifurcation point is Block 4, and it is one sentence long. Zombie: "11 does not always work, so I need a different prime." Ideal: "11 failed only because n was even, so I will keep the same argument and change the base until the exponent is odd." Same observation, opposite move. One asks for a new object, the other asks what the old object actually depended on.

Layer 2: HOW

Block 1. **Cue.** The words "is not prime" attached to a number with more digits than the universe has particles. **Why natural.** Primality is an existence statement about factors, so its negation is an existence statement you can satisfy by exhibiting one witness. Nothing else is possible here, since evaluation is off the table. **Link to the next block.** Once you have committed to finding a modulus, the presentation of N becomes the only source of candidate moduli, so you look at how N is written.

Block 2. **Cue.** N is written in base 10, and the sole subtraction is -1 against the smallest term. **Why natural.** A -1 against 10^n is the definition of a run of nines. Converting a formula into its digit string is free, and digit strings are the objects on which the base 10 divisibility tests act. **Link to the next block.** A digit string of ones and nines is exactly the input format for casting out threes and elevens.

Block 3. **Cue.** Two isolated ones and a block of nines. **Why natural.** The divisibility tests are not folklore, they are the statements $10 \equiv 1 \pmod 9$ and $10 \equiv -1 \pmod{11}$ applied to $\sum d_i 10^i$. Applying them costs two lines and either wins outright or tells you which structural feature is in charge. **Link to the next block.** The eleven test wins for odd n and dies for even n, so the next question is forced: which step used the parity?

Block 4. **Cue.** A method that works for half the cases and fails for the other half by a single sign. **Why natural.** Never abandon a method that half works. Localise the failure instead. Here 10^{10^n} and $10^{10^{10^n}}$ both reduced to $+1$ regardless of n, because their exponents are even for trivial reasons; only $10^n \equiv (-1)^n$ consulted the parity of n. **Link to the next block.** Fixing the sign of one term means changing the modulus, and the modulus you have is $10^1 + 1$, so the family to vary is $10^d + 1$.

Block 5. Cue. The modulus $11 = 10^1 + 1$ has an exponent slot in it. **Why natural.** In the family $M = 10^d + 1$ you always have $10^d \equiv -1$, hence $10^{dj} \equiv (-1)^j$. The ground floor now contributes -1 precisely when n/d is an odd integer, so choose d to be the exact power of two dividing n: that is the unique way to make the quotient odd while keeping d a divisor of 10^n. **Link to the next block.** Having chosen $d = 2^m$, you must now verify that the two upper floors still contribute $+1$, that is, that their exponents are d times an even number.

Block 6. Cue. Exponents of the form 10^n and 10^{10^n}, and a required divisibility by 2^{m+1}. **Why natural.** Ten is even, so 10^t is divisible by 2^t. The demand $2^{m+1} \mid 10^n$ therefore reduces to $n \geq m+1$, and $n \geq 2^m \geq m+1$ holds because exponentials beat linear functions from the very first term. Every floor above the ground is drowning in factors of two; the ground floor is the only delicate one. **Link to the next block.** With all three congruences in hand, the sum is 0 modulo M, and only the size check remains.

Block 7. Cue. A divisibility statement with no size statement is worth zero marks. **Why natural.** $M \mid N$ makes N composite only if $1 < M < N$; a grader will look for that line and nothing else at the end. Both inequalities are one step: $M \geq 11$, and $M \leq 10^n + 1 < 10^{10^n} < N$. **Link to the next block.** Done.

Layer 3: WHY

The mathematical truth under this route is the behaviour of the multiplicative order. For $\gcd(a, M) = 1$, the value of $a^E \bmod M$ depends only on E modulo $\mathrm{ord}_M(a)$. Choosing $M = a^d + 1$ engineers $\mathrm{ord}_M(a) = 2d$ with $a^d \equiv -1$, which means the exponent map restricted to multiples of d becomes the sign character $j \mapsto (-1)^j$. Every congruence in the solution is that one character evaluated three times. The schoolbook eleven test is the case $a = 10$, $d = 1$; you have been using this theorem since primary school without naming it.

The discernment mechanism for "prove this is composite" problems is a scan of exactly three engines, in this order. First, an algebraic identity: difference of squares, Sophie Germain $x^4 + 4y^4 = (x^2 + 2y^2 - 2xy)(x^2 + 2y^2 + 2xy)$, or a cyclotomic factorisation of $a^N \pm 1$. Second, a congruence exhibiting a small factor. Third, a crude parity or size argument. The presentation decides which engine fires: a number given as a polynomial in a variable points at the first, a number given by its base 10 digits with a lonely -1 points at the second. Here the tower makes the first engine impossible, since the exponents are not polynomial in any common variable, so the second engine is not a guess, it is the only survivor.

Train the discernment like this: whenever a method covers some cases and fails on others, write one sentence naming the exact step that consumed the hypothesis you lack. Not "eleven does not work", but "eleven consumed the parity of n and nothing else". That sentence converts a dead end into a parametrised family, because the offending step always has a parameter hiding in it, and here the parameter was the 1 in $10^1 + 1$.

The examiners nudged in three ways, and you should be able to name all three. They wrote the number in base 10 rather than as an abstract $a^{a^{a^n}}$, which is an instruction to look at digits. They put a -1 at the end, which is the only way to make an alternating sum land on a multiple of eleven rather than on 2. And they left n completely free, which guarantees that the easy half of the cases falls out immediately and the other half forces the generalisation. A problem whose easy special case solves itself is a problem asking you to generalise the special case, not to abandon it.

Revolving concepts

Multiplicative order. For $\gcd(a, M) = 1$, $\mathrm{ord}_M(a)$ is the least $t \geq 1$ with $a^t \equiv 1$, and $a^E \equiv a^{E \bmod \mathrm{ord}_M(a)}$. The order divides $\varphi(M)$ and, more sharply, divides the Carmichael function $\lambda(M)$. For $M = a^d + 1$

the order is exactly $2d$. See Book 1, field: Number theory, topic: multiplicative order.

The factorisation $a^d+1 \mid a^n+1$ for n/d odd. Predicated on $a^n+1=(a^d)^k+1$ with k odd and the identity $x^k+1=(x+1)(x^{k-1}-x^{k-2}+\cdots+1)$. Its contrapositive is the Fermat number statement: a^n+1 can be prime only if n is a power of two. See Book 1, field: Number theory, topic: Fermat numbers.

The 2 adic valuation. $v_2(n)$ is the exponent of the exact power of two dividing n, and $n=2^{v_2(n)}k$ with k odd is the canonical splitting. Every argument that consumes the parity of an exponent is secretly a statement about v_2, and the way to generalise such an argument is always to replace "odd" by "odd after dividing out 2^{v_2}". See Book 1, field: Number theory, topic: p adic valuation.

Base B representation and its two tests. Writing $N=\sum d_i B^i$ gives $N \equiv \sum d_i \pmod{B-1}$ and $N \equiv \sum (-1)^i d_i \pmod{B+1}$. This problem is the second test performed in base $B=10^{2^m}$: in that base N has digits 1, 1 and a tail of k copies of $B-1$, and the alternating sum is $1+1+(B-1)=B+1$. See Book 1, field: Number theory, topic: base representation.

The tower recursion, generalised. Define $t_0=n$ and $t_{j+1}=b^{t_j}$. Then $v_2(t_{j+1})=t_j\,v_2(b)$, so for even b the valuations explode: $v_2(t_1)\geq n$ and $v_2(t_2)\geq b^n$. The general moral, and it applies to every tower problem you will ever meet: only the ground floor of a tower has small arithmetic, and every floor above it is divisible by anything you care to ask for. The companion computational recursion is the standard one for evaluating a tower modulo M: to find $b^t \bmod M$, find $t \bmod \lambda(M)$ and recurse down the tower, handling $\gcd(b,M)\neq 1$ by lifting the exponent above $\log_2 M$. See Book 1, field: Number theory, topic: towers and Euler's theorem.

Circuit Bank

- **WHEN** a problem says "prove N is not prime" for an N you cannot compute **DO** exhibit one explicit divisor M, then write the line $1<M<N$ **BECAUSE** primality is an existence claim about factors, and the size line is where marks are actually lost
- **WHEN** the number is presented in base ten and ends in -1 or a run of nines **DO** compute the digit sum modulo 3 and 9 and the alternating digit sum modulo 11 before anything else **BECAUSE** those tests are the congruences $10\equiv 1 \pmod 9$ and $10\equiv -1 \pmod{11}$, they cost two lines, and they are free information
- **WHEN** a sign argument dies because an exponent turned out even **DO** re-run the identical argument in base a^{2^m}, where 2^m is the exact power of two dividing that exponent **BECAUSE** a^d+1 divides $a^{dj}-(-1)^j$, so the sign survives exactly when the quotient exponent is odd
- **WHEN** you meet a tower $b^{b^{b^n}}$ with b even **DO** treat every floor above the ground as divisible by any power of two you need, and check only the ground floor **BECAUSE** $v_2(b^t)=t\,v_2(b)$ outgrows any linear demand, so the delicacy is always at the bottom
- **WHEN** a method solves half the cases and fails on the rest **DO** name in one sentence the single step that consumed the missing hypothesis **BECAUSE** that step always contains a parameter, and varying the parameter is the generalisation you need

Backtesting

1. First useful cue: the -1 sitting against 10^n in a base ten expression, which converts the tail into n nines and makes the whole object a digit string.
2. Method the expression suggested: alternating digit sums, that is, congruences modulo 10^d+1 where $10^d\equiv -1$.
3. Where you would hesitate, and the thought to reinforce: at even n, when eleven dies. Reinforce this: a half working method is a parametrised method with the parameter still set to its default value. Find the parameter, here the d in 10^d+1, before you look for a different method.

4. The trap the question sets: the tower invites you to treat the upper floors as the hard part. They are the trivial part, since their exponents are divisible by every power of two you could want. All the arithmetic content sits in the single term 10^n.

5. What makes you recognise it faster next time: the shape $a^{a^{a^n}} + a^{a^n} + a^n - 1$ has four terms whose signs must sum to zero modulo something, and $1 + 1 - 1 - 1 = 0$ is the only way four terms of size one can cancel. Seeing that target sum in advance tells you that you need two even exponents and one odd multiple of d.

6. The reusable decision rule: to kill a number of the form (tower) + (small stuff) -1, look for a modulus in which the base has order $2d$, and choose d so that the smallest exponent is an odd multiple of d.

Mirror drills

1. Prove that $10^{10^n} + 10^n + 1$ is composite for every positive integer n. *Answer:* its digit sum is 3, so 3 divides it, and it plainly exceeds 3.

2. Let $a \geq 2$ and let $n = 2^m k$ with k odd. Prove that $a^{2^m} + 1$ divides $a^n + 1$, and deduce that $a^n + 1$ can be prime only when n is a power of two. *Answer:* $a^{2^m} \equiv -1$ modulo $a^{2^m} + 1$ and k is odd, so $a^n = (a^{2^m})^k \equiv -1$; if $k > 1$ the divisor is proper.

3. For which integers $b \geq 2$ does the proof above show that $b^{b^{b^n}} + b^{b^n} + b^n - 1$ is composite for all $n \geq 1$? *Answer:* for every even b, since then $v_2(b^t) \geq t$ makes the upper floors even multiples of 2^m; for odd b the ground floor argument survives but the upper floors are odd, so the sign pattern collapses and the proof gives nothing.

16.5 A5. The cross product cannot host a group

Let G be a group, with operation $*$. Suppose that

(i) G is a subset of $\mathbb{R}^3$ (but $*$ need not be related to addition of vectors);

(ii) For each $\mathbf{a}, \mathbf{b} \in G$, either $\mathbf{a} \times \mathbf{b} = \mathbf{a} * \mathbf{b}$ or $\mathbf{a} \times \mathbf{b} = \mathbf{0}$ (or both), where $\times$ is the usual cross product in $\mathbb{R}^3$.

Prove that $\mathbf{a} \times \mathbf{b} = \mathbf{0}$ for all $\mathbf{a}, \mathbf{b} \in G$.

Read it out loud

An abstract group has been dropped onto a subset of $\mathbb{R}^3$, and the only bridge between the algebra and the geometry is a single dichotomy: for every pair, either the group product equals the cross product, or the cross product is zero. Said out loud: this is a group theory problem with a vector algebra weapon, and the fields are the group axioms (identity, inverses, associativity, nothing deeper), the algebra of the cross product in $\mathbb{R}^3$, and one line of linear algebra about spanning. There is no structure theory here, no classification, no finiteness. You are given four axioms and one dichotomy, and that is the entire toolkit.

The first move, inside thirty seconds, is to rewrite (ii) so that it is usable. As a disjunction it says almost nothing. As an implication it is precise and strong:

$$\mathbf{x} \times \mathbf{y} \neq \mathbf{0} \implies \mathbf{x} * \mathbf{y} = \mathbf{x} \times \mathbf{y}.$$

In other words, wherever the cross product is alive, the mysterious group operation is completely known, and wherever the cross product is dead, the group operation is completely unknown. All your information lives on the nonzero cross products, which is exactly the set of pairs the conclusion says cannot exist. So the shape of the proof is forced: assume some pair has $\mathbf{a} \times \mathbf{b} \neq \mathbf{0}$, mine that assumption for all the information it gives, and hit a contradiction.

The second thing to write down, before any calculation, is the list of ways in which $\times$ fails to be a group operation, because the contradiction has to come from one of them. It is anticommutative, so $\mathbf{y} \times \mathbf{x} = -(\mathbf{x} \times \mathbf{y})$. It kills every vector against itself, $\mathbf{x} \times \mathbf{x} = \mathbf{0}$. It has no identity element, since $\mathbf{e} \times \mathbf{x}$ is orthogonal to $\mathbf{x}$ and can never equal $\mathbf{x}$ unless it is zero. And it is not associative: $\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2 \mathbf{y}$ for $\mathbf{x} \perp \mathbf{y}$, while $(\mathbf{x} \times \mathbf{x}) \times \mathbf{y} = \mathbf{0}$. Each of those four failures becomes a lemma. That is the whole problem, and if you write that list first, the solution is bookkeeping.

Third, note the free resources. A group hands you an identity $\mathbf{e}$ that commutes with everything, and inverses. In this setting "commutes" has a geometric consequence, because anticommutativity means a commuting pair cannot have a live cross product. You have not touched the problem yet and you already own a lemma.

Approaches, ranked worst to best

1. **Classify the possible sets G.** *Cue:* the conclusion describes all admissible G, so try to find them all. *Circuit:* no circuit, wishful classification. You would be trying to describe every subset of $\mathbb{R}^3$ carrying some unspecified group law. There is no traction, since $*$ is not assumed continuous, bilinear, or related to the linear structure in any way at all. Last place.

2. **Coordinates.** *Cue:* $\mathbb{R}^3$ invites a basis and nine unknowns. *Circuit:* coordinate bash. Write $\mathbf{a} = (a_1, a_2, a_3)$ and expand. The cross product becomes explicit and the group product stays an unknown function of six variables, so you have added notation and no information. The problem gives no bilinearity to exploit, which is exactly what a coordinate bash needs.

3. **Lengths and angles.** *Cue:* $|\mathbf{x} \times \mathbf{y}| = |\mathbf{x}||\mathbf{y}| \sin\theta$ suggests a size argument. *Circuit:* metric bookkeeping, extremal choice. You can extract facts such as "orthogonal nonzero elements

have nonzero cross product", and you will need that. But no size argument can finish, because nothing forces $*$ to interact with lengths; the group could be anything on the pairs where $\times$ dies.

4. **Squeeze the identity element first.** *Cue:* every group has an $\mathbf{e}$, and $\mathbf{e}$ commutes with everything. *Circuit:* use free structure before assumed structure. Anticommutativity turns commuting into $\mathbf{e} \times \mathbf{x} = \mathbf{0}$ for all $\mathbf{x} \in G$, and if G contains three independent vectors this forces $\mathbf{e} = \mathbf{0}$. This is the correct opening and it is genuinely free, but on its own it does not close: knowing $\mathbf{e} = \mathbf{0}$ is useless until you have produced elements whose squares are zero.
5. **Collide associativity with the vector triple product.** *Cue:* $*$ is associative, $\times$ is not, and the dichotomy forces them to agree on live pairs. *Circuit:* find where the two operations must agree, then evaluate the same expression two ways. The ideal route. Build an orthogonal triple inside G, prove that squares vanish, note that $\mathbf{e} = \mathbf{0}$, and then two elements are forced to commute that anticommutativity forbids. Every step is three lines and every step is an axiom failure of $\times$ cashed in.

The ideal solution

Throughout, $\cdot$ is the usual dot product, and "$\mathbf{u}$ is a multiple of $\mathbf{v}$" allows the multiple $\mathbf{0}$. Recall two standard facts about $\mathbb{R}^3$: nonzero orthogonal vectors have nonzero cross product, and the vector triple product identity $\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) = (\mathbf{u} \cdot \mathbf{w})\mathbf{v} - (\mathbf{u} \cdot \mathbf{v})\mathbf{w}$, whose special case $\mathbf{x} \perp \mathbf{y}$ reads

$$\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2 \, \mathbf{y}. \tag{$*$}$$

Hypothesis (ii) will be used only in the implication form: if $\mathbf{x} \times \mathbf{y} \neq \mathbf{0}$ then $\mathbf{x} * \mathbf{y} = \mathbf{x} \times \mathbf{y}$; in particular $\mathbf{x} \times \mathbf{y}$ then lies in G, by closure.

Lemma 16.5.1. *If $\mathbf{x}, \mathbf{y} \in G$ satisfy $\mathbf{x} * \mathbf{y} = \mathbf{y} * \mathbf{x}$, then $\mathbf{x} \times \mathbf{y} = \mathbf{0}$.*

Proof. Suppose $\mathbf{x} \times \mathbf{y} \neq \mathbf{0}$. Then $\mathbf{y} \times \mathbf{x} = -(\mathbf{x} \times \mathbf{y}) \neq \mathbf{0}$ as well, so hypothesis (ii) applies to both orders and gives $\mathbf{x} * \mathbf{y} = \mathbf{x} \times \mathbf{y}$ and $\mathbf{y} * \mathbf{x} = \mathbf{y} \times \mathbf{x} = -(\mathbf{x} \times \mathbf{y})$. The assumed commutation forces $\mathbf{x} \times \mathbf{y} = -(\mathbf{x} \times \mathbf{y})$, hence $\mathbf{x} \times \mathbf{y} = \mathbf{0}$, a contradiction. □

Lemma 16.5.2. *If $\mathbf{x}, \mathbf{y} \in G$ are nonzero and orthogonal, then $\mathbf{x} * \mathbf{x}$ is a multiple of $\mathbf{y}$.*

Proof. Since $\mathbf{x}, \mathbf{y}$ are nonzero and orthogonal, $\mathbf{z} := \mathbf{x} \times \mathbf{y} \neq \mathbf{0}$, so $\mathbf{z} = \mathbf{x} * \mathbf{y} \in G$, and $\mathbf{x}, \mathbf{y}, \mathbf{z}$ are nonzero and mutually orthogonal. Then $\mathbf{x} \times \mathbf{z} \neq \mathbf{0}$ as well, so

$$\mathbf{w} := \mathbf{x} \times \mathbf{z} = \mathbf{x} * \mathbf{z} = \mathbf{x} * (\mathbf{x} * \mathbf{y}) = (\mathbf{x} * \mathbf{x}) * \mathbf{y},$$

the last equality being associativity of $*$. On the other hand $(*)$ gives $\mathbf{w} = \mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2 \mathbf{y}$, a nonzero multiple of $\mathbf{y}$.

Now suppose $(\mathbf{x} * \mathbf{x}) \times \mathbf{y} \neq \mathbf{0}$. Then hypothesis (ii) gives $(\mathbf{x} * \mathbf{x}) * \mathbf{y} = (\mathbf{x} * \mathbf{x}) \times \mathbf{y}$, so $\mathbf{w} = (\mathbf{x} * \mathbf{x}) \times \mathbf{y}$ is orthogonal to $\mathbf{y}$, whence $\mathbf{w} \cdot \mathbf{y} = 0$. But $\mathbf{w} = -|\mathbf{x}|^2 \mathbf{y}$ gives $\mathbf{w} \cdot \mathbf{y} = -|\mathbf{x}|^2 |\mathbf{y}|^2 \neq 0$. Contradiction. Hence $(\mathbf{x} * \mathbf{x}) \times \mathbf{y} = \mathbf{0}$, which for $\mathbf{y} \neq \mathbf{0}$ says exactly that $\mathbf{x} * \mathbf{x}$ is a multiple of $\mathbf{y}$. □

Lemma 16.5.3. *If $\mathbf{x} \in G$ is nonzero and some nonzero $\mathbf{y} \in G$ is orthogonal to $\mathbf{x}$, then $\mathbf{x} * \mathbf{x} = \mathbf{0}$.*

Proof. As above $\mathbf{x} \times \mathbf{y} \in G$ is nonzero and orthogonal to $\mathbf{x}$. Apply Lemma 2 twice, first to the pair $(\mathbf{x}, \mathbf{y})$ and then to the pair $(\mathbf{x}, \mathbf{x} \times \mathbf{y})$: the vector $\mathbf{x} * \mathbf{x}$ is a multiple of $\mathbf{y}$ and also a multiple of $\mathbf{x} \times \mathbf{y}$. Those two vectors are nonzero and orthogonal to each other, so the only common multiple is $\mathbf{0}$. □

Now the proof. Suppose, for contradiction, that there exist $\mathbf{a}, \mathbf{b} \in G$ with $\mathbf{a} \times \mathbf{b} \neq \mathbf{0}$. Put $\mathbf{c} = \mathbf{a} \times \mathbf{b} = \mathbf{a} * \mathbf{b} \in G$. Then $\mathbf{a}, \mathbf{b}$ are independent and $\mathbf{c}$ is nonzero and orthogonal to both, so $\mathbf{a}, \mathbf{b}, \mathbf{c}$ are linearly independent and span $\mathbb{R}^3$.

Let $\mathbf{e}$ be the identity of G. It commutes with $\mathbf{a}$, $\mathbf{b}$ and $\mathbf{c}$, so Lemma 1 gives $\mathbf{e} \times \mathbf{a} = \mathbf{e} \times \mathbf{b} = \mathbf{e} \times \mathbf{c} = \mathbf{0}$. The map $\mathbf{v} \mapsto \mathbf{e} \times \mathbf{v}$ is linear and vanishes on a spanning set, so it vanishes identically on $\mathbb{R}^3$. A vector whose cross product with every vector is zero is zero, so

$$\mathbf{e} = \mathbf{0}.$$

Next, $\mathbf{b}$ and $\mathbf{c}$ are nonzero and orthogonal, so $\mathbf{p} := \mathbf{b} \times \mathbf{c} \neq \mathbf{0}$ and $\mathbf{p} = \mathbf{b} * \mathbf{c} \in G$, and $\mathbf{b}, \mathbf{c}, \mathbf{p}$ are nonzero and mutually orthogonal. Lemma 3 applies three times: to $\mathbf{b}$ with witness $\mathbf{c}$, to $\mathbf{c}$ with witness $\mathbf{b}$, and to $\mathbf{p}$ with witness $\mathbf{b}$. Hence

$$\mathbf{b} * \mathbf{b} = \mathbf{c} * \mathbf{c} = \mathbf{p} * \mathbf{p} = \mathbf{0} = \mathbf{e}.$$

So in the group G, each of $\mathbf{b}$, $\mathbf{c}$ and $\mathbf{p} = \mathbf{b} * \mathbf{c}$ is its own inverse. Therefore

$$\mathbf{b} * \mathbf{c} = \mathbf{p} = \mathbf{p}^{-1} = (\mathbf{b} * \mathbf{c})^{-1} = \mathbf{c}^{-1} * \mathbf{b}^{-1} = \mathbf{c} * \mathbf{b},$$

so $\mathbf{b}$ and $\mathbf{c}$ commute. Lemma 1 then forces $\mathbf{b} \times \mathbf{c} = \mathbf{0}$, contradicting $\mathbf{p} \neq \mathbf{0}$.

Hence no such pair $\mathbf{a}, \mathbf{b}$ exists, that is, $\mathbf{a} \times \mathbf{b} = \mathbf{0}$ for all $\mathbf{a}, \mathbf{b} \in G$. ■

Remark 16.5.4. The conclusion says that all elements of G are pairwise parallel, so G lies on a single line through the origin. That is sharp: take G to be the x axis with $*$ equal to vector addition. Every cross product is zero, so (ii) holds vacuously in its second branch.

Layer 1: WHAT

Block 1. Rewrite (ii) as the implication "live cross product forces $* = \times$", and negate the goal: assume $\mathbf{a} \times \mathbf{b} \neq \mathbf{0}$.

Block 2. Anticommutativity: a commuting pair in G cannot have a live cross product (Lemma 1).

Block 3. The identity is free, and it commutes with everything, so $\mathbf{e}$ has zero cross product against the spanning set $\mathbf{a}, \mathbf{b}, \mathbf{c}$, giving $\mathbf{e} = \mathbf{0}$.

Block 4. Manufacture an orthogonal triple inside G: $\mathbf{b}$, $\mathbf{c} = \mathbf{a} \times \mathbf{b}$, and $\mathbf{p} = \mathbf{b} \times \mathbf{c}$, all in G by the implication.

Block 5. Associativity evaluated against the vector triple product: $\mathbf{x} * (\mathbf{x} * \mathbf{y}) = (\mathbf{x} * \mathbf{x}) * \mathbf{y}$ while $\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2 \mathbf{y}$, which pins $\mathbf{x} * \mathbf{x}$ to a line, twice, hence to zero (Lemmas 2 and 3).

Block 6. Three squares equal the identity, so $\mathbf{b} * \mathbf{c} = (\mathbf{b} * \mathbf{c})^{-1} = \mathbf{c} * \mathbf{b}$, and Lemma 1 kills it.

The nominal route is Block 1 followed by coordinates: name nine unknowns, expand cross products, and try to solve. You will fill a page and learn nothing, because $*$ has no coordinate description to solve for. A second nominal route is to hunt for a slick one line contradiction from $\mathbf{x} \times \mathbf{x} = \mathbf{0}$ alone; it does not exist, because $\mathbf{x} * \mathbf{x}$ is unconstrained by (ii) precisely when $\mathbf{x} \times \mathbf{x}$ is zero, which is always.

The bifurcation point is the sentence after Block 1. Zombie: "what does G look like?" Ideal: "which group axioms can I collide with the failures of $\times$?" The first question asks about an object you cannot see. The second asks about a list you can write down in ten seconds.

Layer 2: HOW

Block 1. Cue. Hypothesis (ii) is a disjunction, and the conclusion is that the second branch always holds. **Why natural.** A disjunction is used by negating the branch you do not want: $\mathbf{x} \times \mathbf{y} \neq \mathbf{0}$ forces $\mathbf{x} * \mathbf{y} = \mathbf{x} \times \mathbf{y}$. That implication is the only channel through which the geometry can speak to the group, so all information lives on live pairs. Since the goal denies that live pairs exist, contradiction is the only available frame. **Link to the next block.** Assuming a live pair immediately produces new elements of G, because $\mathbf{a} * \mathbf{b}$ lies in G by closure and equals $\mathbf{a} \times \mathbf{b}$.

Block 2. **Cue.** $\times$ is anticommutative and $*$ has no commutativity assumption. **Why natural.** If $\mathbf{x}$ and $\mathbf{y}$ commute under $*$ and their cross product were live, then $*$ would have to equal $\times$ in both orders, and $\times$ reverses sign. Two vectors equal to their own negatives are zero. This is the cheapest of the four axiom failures, so it is proved first and used everywhere. **Link to the next block.** Commuting pairs are handed to you by the group: the identity commutes with everything.

Block 3. **Cue.** Every group has an identity, and you have not used it. **Why natural.** Free structure before assumed structure. Lemma 1 converts "commutes with" into "has zero cross product with", so the identity is orthogonally deficient against every element of G. Once G contains a spanning triple, linearity of $\mathbf{v} \mapsto \mathbf{e} \times \mathbf{v}$ upgrades three vanishings to all of $\mathbb{R}^3$, and only the zero vector kills everything. **Link to the next block.** Knowing $\mathbf{e} = \mathbf{0}$ is only useful if you can produce elements whose $*$ square is the zero vector, so the next task is to compute squares.

Block 4. **Cue.** You need orthogonal pairs inside G, and the cross product manufactures orthogonality for free. **Why natural.** $\mathbf{a} \times \mathbf{b}$ is orthogonal to both factors by definition, and it is in G whenever it is nonzero. Iterating gives a mutually orthogonal triple $\mathbf{b}, \mathbf{c}, \mathbf{p}$ inside G, which is the richest configuration the hypothesis can produce. **Link to the next block.** An orthogonal pair is exactly the input Lemma 2 needs.

Block 5. **Cue.** $*$ is associative, and $\times$ is not; a repeated factor is where non associativity shows. **Why natural.** Evaluate $\mathbf{x} * (\mathbf{x} * \mathbf{y})$ two ways. Read as $\mathbf{x} * \mathbf{z}$ with $\mathbf{z} = \mathbf{x} \times \mathbf{y}$ live, it equals $\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2\mathbf{y}$, which points along $\mathbf{y}$. Read as $(\mathbf{x} * \mathbf{x}) * \mathbf{y}$, it would be orthogonal to $\mathbf{y}$ if $(\mathbf{x} * \mathbf{x}) \times \mathbf{y}$ were live. Along $\mathbf{y}$ and orthogonal to $\mathbf{y}$ at once is impossible for a nonzero vector, so $\mathbf{x} * \mathbf{x}$ is pinned to the line through $\mathbf{y}$. Running the same argument with the orthogonal witness $\mathbf{x} \times \mathbf{y}$ pins it to a second, perpendicular line, and two perpendicular lines meet only at the origin. **Link to the next block.** Squares are the zero vector, and the zero vector is the identity, so you now have three involutions.

Block 6. **Cue.** $\mathbf{b}^2 = \mathbf{c}^2 = (\mathbf{b} * \mathbf{c})^2 = \mathbf{e}$ is the standard configuration that forces commuting. **Why natural.** In any group, $(gh)^{-1} = h^{-1}g^{-1}$; if g, h, gh are all involutions then $gh = (gh)^{-1} = h^{-1}g^{-1} = hg$. Nothing geometric is used here, it is pure group theory, and it delivers a commuting pair whose cross product you already know is nonzero. **Link to the next block.** Lemma 1 closes the contradiction.

Layer 3: WHY

The mathematical truth underneath is that $(\mathbb{R}^3, \times)$ is a Lie algebra, not an associative algebra. Its correct axiom is the Jacobi identity, and its associator is genuinely nonzero: for orthogonal nonzero $\mathbf{x}, \mathbf{y}$ you have $\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -|\mathbf{x}|^2\mathbf{y}$ while $(\mathbf{x} \times \mathbf{x}) \times \mathbf{y} = \mathbf{0}$. A group law is associative by fiat. So the dichotomy (ii) is an attempt to let an associative operation impersonate a non associative one, and the impersonation can only survive on the pairs where $\times$ degenerates. The proof consists of trapping $*$ into a configuration where it must agree with $\times$ on enough products that associativity and anticommutativity can be pointed at each other.

The discernment mechanism generalises past this problem. When you are handed an exotic structure defined by "an operation with axioms A agrees, sometimes, with a concrete operation B", never study the structure. Study the axiom failures of B, list them, and then engineer configurations that force agreement on the products where the failure lives. Here the list has four entries and three of them become lemmas: anticommutativity gives Lemma 1, non associativity gives Lemma 2, and the absence of an identity for $\times$ is what makes $\mathbf{e} = \mathbf{0}$ so damaging. The fourth, $\mathbf{x} \times \mathbf{x} = \mathbf{0}$, is the reason (ii) tells you nothing about squares, which is why squares had to be computed the long way round through associativity.

To train it: whenever an operation appears in a problem, write its axiom table before you write anything else. Commutative or anticommutative. Associative or not, and if not, what the associator is. Identity, or no identity. Zero divisors, or none. For $\times$: anticommutative, non associative with $\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) = (\mathbf{u} \cdot \mathbf{w})\mathbf{v} - (\mathbf{u} \cdot \mathbf{v})\mathbf{w}$, no identity, and zero divisors everywhere since parallel vectors annihilate. That table is the solution in compressed form.

What the examiners nudged toward: the phrase "but $*$ need not be related to addition of vectors" is there to stop you from assuming linearity, which is the same as telling you that only the group axioms are available. The "or both" in (ii) is there to stop you from treating the dichotomy as exclusive, which would be a false extra hypothesis. And placing the problem in $\mathbb{R}^3$ rather than in an abstract Lie algebra guarantees the vector triple product is available in closed form.

Revolving concepts

Group axioms and involutions. An element with $g * g = e$ is an involution and equals its own inverse. The key local lemma: if $g^2 = h^2 = (gh)^2 = e$ then $gh = (gh)^{-1} = h^{-1}g^{-1} = hg$, so g and h commute. Globalised: a group in which every element is an involution is abelian, and is a vector space over the field of two elements, hence has order a power of two when finite. See Book 1, field: Abstract algebra, topic: involutions and elementary abelian groups.

Inverse of a product. $(g * h)^{-1} = h^{-1} * g^{-1}$, with the order reversed. The reversal is the entire content of Block 6: without it the identity $gh = hg$ does not appear. See Book 1, field: Abstract algebra, topic: group axioms.

Cross product. $\mathbf{x} \times \mathbf{y}$ is orthogonal to both factors, $|\mathbf{x} \times \mathbf{y}| = |\mathbf{x}||\mathbf{y}| \sin\theta$, and $\mathbf{x} \times \mathbf{y} = \mathbf{0}$ if and only if the two are parallel or one is zero. Anticommutative, bilinear, and satisfying the Jacobi identity $\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) + \mathbf{v} \times (\mathbf{w} \times \mathbf{u}) + \mathbf{w} \times (\mathbf{u} \times \mathbf{v}) = \mathbf{0}$. See Book 1, field: Geometry, topic: vector product in three dimensions.

Vector triple product. $\mathbf{u} \times (\mathbf{v} \times \mathbf{w}) = (\mathbf{u} \cdot \mathbf{w})\mathbf{v} - (\mathbf{u} \cdot \mathbf{v})\mathbf{w}$. If this is not installed, install it now: it is the single computational fact this problem needs, and it is the standard way to convert a double cross product into a linear combination of the original vectors. See Book 1, field: Geometry, topic: triple products.

The recursion hidden in $\mathbf{x}\times$. Fix $\mathbf{x} \neq \mathbf{0}$ and $\mathbf{y}$ orthogonal to it, and define $\mathbf{w}_0 = \mathbf{y}$, $\mathbf{w}_{k+1} = \mathbf{x} \times \mathbf{w}_k$. The triple product identity gives the linear recursion

$$\mathbf{w}_{k+2} = -|\mathbf{x}|^2 \, \mathbf{w}_k,$$

whose characteristic roots are $\pm i|\mathbf{x}|$: on the plane orthogonal to $\mathbf{x}$, the map $\mathbf{w} \mapsto \mathbf{x} \times \mathbf{w}$ is $|\mathbf{x}|$ times a quarter turn, so the orbit is a four cycle up to scaling. Generalised statement: any linear map A with $A^2 = -\lambda I$ on a subspace produces the recursion $\mathbf{w}_{k+2} = -\lambda \mathbf{w}_k$, solved by $\mathbf{w}_{2k} = (-\lambda)^k \mathbf{w}_0$ and $\mathbf{w}_{2k+1} = (-\lambda)^k \mathbf{w}_1$, and is a rotation by a quarter turn after rescaling when $\lambda > 0$. Recognising that $\mathbf{x}\times$ is a rotation in disguise is what makes Lemma 2 predictable rather than surprising. See Book 1, field: Linear algebra, topic: maps with $A^2 = -I$.

Spanning and the kernel of $\mathbf{v} \mapsto \mathbf{e} \times \mathbf{v}$. For $\mathbf{e} \neq \mathbf{0}$ this linear map has kernel exactly the line through $\mathbf{e}$, so it cannot vanish on a spanning set. See Book 1, field: Linear algebra, topic: rank and kernel.

Circuit Bank

- **WHEN** a hypothesis reads "either P or Q" **DO** immediately rewrite it as "not Q implies P" and note where the information actually lives **BECAUSE** a disjunction is inert, an implication is a tool, and the branch you are told to prove is the branch you must assume fails
- **WHEN** the conclusion says some quantity is always degenerate **DO** assume one nondegenerate instance and mine it for new elements and new configurations **BECAUSE** a live instance is the only source of information the hypothesis can convert, so it is both the enemy and the fuel

- **WHEN** a group law is imposed on a set that already carries another operation **DO** write the axiom table of the other operation, commutativity, associativity, identity, zero divisors, before touching the problem **BECAUSE** every contradiction available to you is one of those axiom failures, so the table is the solution list
- **WHEN** you have $g^2 = h^2 = (gh)^2 = e$ **DO** write $gh = (gh)^{-1} = h^{-1}g^{-1} = hg$ and conclude that g and h commute **BECAUSE** inverses reverse order, so involutions convert squaring information into commuting information
- **WHEN** you need a nonzero vector orthogonal to two given ones **DO** take their cross product, and note that membership in the ambient structure is often free **BECAUSE** orthogonality is manufactured, not found, and in this problem the manufactured vector is automatically in G

Backtesting

1. First useful cue: the "either or" in hypothesis (ii), rewritten as "$\mathbf{x} \times \mathbf{y} \neq \mathbf{0}$ implies $\mathbf{x} * \mathbf{y} = \mathbf{x} \times \mathbf{y}$".
2. Method the expression suggested: contradiction driven by the axiom failures of $\times$, one lemma per failure, with associativity of $*$ as the crowbar.
3. Where you would hesitate, and the thought to reinforce: at "how do I ever compute $\mathbf{x} * \mathbf{x}$, given that (ii) says nothing about it?" Reinforce this: when an operation is opaque on a product, compute that product inside a longer associative chain where the neighbouring products are transparent. Associativity is a transport mechanism for information.
4. The trap the question sets: believing you must understand G as a set. You never learn a single element of G; you only learn relations. The problem is testing whether you can work with an object you cannot see.
5. What makes you recognise it faster next time: the phrase "$*$ need not be related to addition" is an explicit instruction that only the four group axioms are in play, and the presence of a concrete non associative operation next to an associative one is the standing signal to evaluate one expression two ways.
6. The reusable decision rule: to prove that a concrete operation cannot support an abstract structure, force the two operations to agree on a configuration large enough to expose the axiom the concrete operation violates.

Mirror drills

1. Prove that a group in which every element satisfies $g * g = e$ is abelian. *Answer:* $g * h = (g * h)^{-1} = h^{-1} * g^{-1} = h * g$.
2. Suppose $\mathbf{e} \in \mathbb{R}^3$ satisfies $\mathbf{e} \times \mathbf{v} = \mathbf{0}$ for three linearly independent $\mathbf{v}$. Show $\mathbf{e} = \mathbf{0}$. *Answer:* the map $\mathbf{v} \mapsto \mathbf{e} \times \mathbf{v}$ is linear, so it vanishes on all of $\mathbb{R}^3$; if $\mathbf{e} \neq \mathbf{0}$, pick $\mathbf{v}$ not parallel to $\mathbf{e}$ to get a nonzero value.
3. For $\mathbf{x}$ a unit vector, show that $\mathbf{y} \mapsto \mathbf{x} \times \mathbf{y}$ restricted to the plane orthogonal to $\mathbf{x}$ squares to $-I$, and say what that means geometrically. *Answer:* the triple product identity gives $\mathbf{x} \times (\mathbf{x} \times \mathbf{y}) = -\mathbf{y}$ there, so the map is a quarter turn in that plane.

16.6 A6. A shrinking ratio contradicts a convergent tail

Let $f : [0, \infty) \to \mathbb{R}$ be a strictly decreasing continuous function such that $\lim_{x\to\infty} f(x) = 0$. Prove that

$$\int_0^\infty \frac{f(x) - f(x+1)}{f(x)}\, dx$$

diverges.

Read it out loud

An integral over an infinite ray of a quantity that measures, at each point, the relative drop of f over one unit step. Say the integrand out loud in its true form:

$$\frac{f(x) - f(x+1)}{f(x)} = 1 - \frac{f(x+1)}{f(x)}.$$

That rewriting is compulsory and it is free, and it tells you two things at once. First, the problem is scale invariant: replacing f by λf for $\lambda > 0$ changes nothing, so no absolute size of f can ever appear in the proof, only ratios. Second, the integrand lies strictly between 0 and 1, so the integral of a positive function is at stake and divergence can only mean divergence to $+\infty$.

The fields are real analysis and nothing else: improper Riemann integrals, monotone functions, and the theory of infinite products hiding one layer down. There is no formula for f and there never will be, so every step must be soft: comparison, telescoping, monotonicity, contradiction. If you catch yourself trying to compute, you have misread the problem.

Spend sixty seconds on reconnaissance, because it calibrates everything. For $f(x) = e^{-cx}$ the integrand is the constant $1 - e^{-c}$ and the integral diverges linearly. For $f(x) = 1/(1+x)$ the integrand is $1/(x+2)$ and it diverges logarithmically. For $f(x) = 1/\log x$ on $[3, \infty)$ the integrand behaves like $1/(x \log x)$ and it diverges like $\log\log x$. The pattern is the message: the slower f decays, the smaller the integrand, and yet the divergence survives every time, arriving more and more slowly. So the theorem is tight, no crude bound will prove it, and the mechanism must be something that converts "f reaches zero" into "the total relative drop is infinite" without ever looking at the rate.

That mechanism is visible if you sample at the integers. Put $t_n = 1 - f(n+1)/f(n) \in (0, 1)$. Then

$$\prod_{n=0}^{N-1} (1 - t_n) = \frac{f(N)}{f(0)} \longrightarrow 0,$$

and an infinite product of factors $1 - t_n$ tends to zero exactly when $\sum t_n$ diverges. So at the integers the statement is a restatement of the infinite product criterion, which is the honest reason the problem is true. The first useful cue, on paper within thirty seconds, is the rewritten integrand together with the sentence "this is an infinite product in disguise". The difficulty, and the reason this sits at A6, is that the integral is not the integer sample: at a non integer x the integrand compares $f(x+1)$ with $f(x)$, a sliding window, and sliding windows are not controlled by integer windows.

Approaches, ranked worst to best

1. **Test cases and extrapolate.** *Cue:* no formula for f, so try the standard decays. *Circuit:* pattern spotting. Mandatory as sixty seconds of reconnaissance, worthless as a proof: it establishes nothing about a general f, which may be flat for a million units and then plunge. Last place as a route, first place as a calibration.

2. **Bound the integrand pointwise from below.** *Cue:* a divergent integral usually wants a divergent minorant. *Circuit:* comparison test. Hopeless, and knowing why is worth more

than the attempt. Given any $\varepsilon > 0$ you can build a strictly decreasing f whose relative drop over a single unit step is below ε on an interval as long as you please. No fixed minorant exists, so comparison must be applied to a rearranged or aggregated quantity, never to the raw integrand.

3. **Take logarithms and telescope.** *Cue:* the product $\prod f(x+n+1)/f(x+n)$ telescopes to $f(x+N)/f(x)$, which tends to zero. *Circuit:* convert products into sums with log. Correct in spirit, and it is the official first route. Write $-\log(1-t) \le t + 2t^2 \le 2t$ for $t \in [0, 1/2]$, so the integrand dominates $\frac{1}{2}\log(f(x)/f(x+1))$, whose integral telescopes to $\int \log(f(x+N)/f(x+n))$ and blows up. The obstruction is uniformity: no constant C satisfies $-\log(1-t) \le Ct$ for all $t \in [0, 1)$, so the bound is available only where the one step drop is at most a half, and the two regimes interleave in x in an uncontrolled way. Repairing that costs Lebesgue measure and the monotone convergence theorem, or a careful split at half integers. Correct, but expensive in both ink and installed theory.

4. **Block telescoping with a frozen denominator.** *Cue:* the denominator $f(x)$ is the only obstacle to telescoping the numerator. *Circuit:* replace a varying denominator by its largest value on the block. For integers $a < b$, since $f(x+k) \le f(x+a)$ for $k \ge a$,

$$\int_a^b \frac{f(x) - f(x+1)}{f(x)} dx \ge \int_0^1 \frac{f(x+a) - f(x+b)}{f(x+a)} dx,$$

the numerator having telescoped. Then choose b so far out that $f(b) < f(a+1)/2$, which makes the last integrand at least $1/2$; iterate to get infinitely many disjoint blocks each contributing at least $1/2$. Fully elementary, about ten lines, and it is the route to write if you cannot see the next one. Very high ranking, and the safest under time pressure.

5. **Contradiction from the vanishing tail.** *Cue:* convergence of a positive integral is a statement about tails, and the integrand is built from a ratio. *Circuit:* assume convergence, and sandwich a ratio between two things the tail controls. The ideal. Six lines, no measure theory, no case analysis, no constructed sequence. The tail of a convergent integral bounds $f(a+1)/f(a)$ from above; a small ratio makes the integrand close to 1; an integrand close to 1 makes the integral diverge. The assumption of convergence eats itself.

The ideal solution

Step 0: the integral makes sense. First, $f(x) > 0$ for all $x \ge 0$. If $f(x_0) \le 0$ for some x_0, then for every $x > x_0$ strict decrease gives $f(x) < f(x_0) \le 0$, so $\lim_{x\to\infty} f(x) \le f(x_0+1) < 0$, contradicting the hypothesis that the limit is 0. Hence $f > 0$, the function

$$g(x) := \frac{f(x) - f(x+1)}{f(x)} = 1 - \frac{f(x+1)}{f(x)}$$

is continuous on $[0, \infty)$, and $0 < g(x) < 1$ because $0 < f(x+1) < f(x)$. Since $g > 0$, the function $T \mapsto \int_0^T g$ is strictly increasing, so the improper integral either converges to a finite limit or increases to $+\infty$; there is no oscillation to worry about.

Step 1: a shift identity. For real $0 \le a < T$, substituting $x \mapsto x+1$ in the second piece,

$$\int_a^T \big(f(x) - f(x+1)\big) dx = \int_a^T f - \int_{a+1}^{T+1} f = \int_a^{a+1} f - \int_T^{T+1} f.$$

Since f is decreasing and positive, $0 < \int_T^{T+1} f \le f(T) \to 0$ as $T \to \infty$. Therefore

$$\int_a^\infty \big(f(x) - f(x+1)\big) dx = \int_a^{a+1} f(x)\, dx < \infty. \tag{1}$$

This is telescoping wearing continuous clothes: the total of all one step drops from a onwards equals the mass of f over the single window $[a, a+1]$.

Step 2: assume convergence. Suppose, for contradiction, that $\int_0^\infty g$ converges, with value $L < \infty$. Write $T(a) = \int_a^\infty g(x)\,dx$ for the tail. Since $\int_0^a g \to L$ as $a \to \infty$, we get

$$T(a) \longrightarrow 0 \qquad (a \to \infty). \tag{2}$$

Step 3: the tail dominates the ratio. Fix any $a \geq 0$. Because f is strictly decreasing and continuous, $f(x) > f(a+1)$ for $x \in [a, a+1)$, so $\int_a^{a+1} f > f(a+1)$. Because $0 < f(x) \leq f(a)$ for $x \geq a$, and $f(x) - f(x+1) > 0$, we have the pointwise bound $\frac{f(x)-f(x+1)}{f(a)} \leq \frac{f(x)-f(x+1)}{f(x)}$ on $[a, \infty)$. Chaining these with (1),

$$0 < \frac{f(a+1)}{f(a)} < \frac{1}{f(a)} \int_a^{a+1} f(x)\,dx = \frac{1}{f(a)} \int_a^\infty \big(f(x) - f(x+1)\big)dx \leq \int_a^\infty \frac{f(x) - f(x+1)}{f(x)}\,dx = T(a).$$

Every integral here converges: the middle ones by (1), the last by the assumption in Step 2.

Step 4: the assumption eats itself. By (2) and the squeeze, $f(a+1)/f(a) \to 0$ as $a \to \infty$, hence

$$g(a) = 1 - \frac{f(a+1)}{f(a)} \longrightarrow 1.$$

So there is $A \geq 0$ with $g(x) \geq \frac{1}{2}$ for all $x \geq A$, and then for $T > A$,

$$\int_0^T g \geq \int_A^T g \geq \tfrac{1}{2}(T - A) \longrightarrow \infty,$$

contradicting convergence. Therefore the integral does not converge, and being an integral of a positive function, it diverges to $+\infty$. ■

Remark 16.6.1. Strict decrease is used only to guarantee $f > 0$ and to make one inequality strict. The same proof, with $<$ weakened to $\leq$ in Step 3, gives the result for any positive decreasing continuous f with limit 0. Continuity is used only so that the Riemann integrals exist.

Remark 16.6.2. The divergence can be arbitrarily slow. For $f(x) = 1/\log x$ on $[3, \infty)$ the integrand is $\big(\log(x+1) - \log x\big)/\log(x+1)$, of order $1/(x \log x)$, so the integral grows like $\log\log x$. No proof of this problem can produce a rate, which is exactly why every viable route is a contradiction or an aggregation argument.

Layer 1: WHAT

Block 1. Rewrite the integrand as $1 - f(x+1)/f(x)$; record scale invariance and $0 < g < 1$; prove $f > 0$ so that the integral is defined.

Block 2. Sample at the integers to see the mechanism: $\prod(1 - t_n) = f(N)/f(0) \to 0$, so this is the infinite product criterion in disguise. This block produces no proof, only the certainty that the theorem is about f reaching zero and nothing else.

Block 3. Assume convergence. Write down the one free consequence: the tail $T(a) \to 0$.

Block 4. Telescope the numerator over the whole ray: $\int_a^\infty (f(x) - f(x+1))dx = \int_a^{a+1} f$.

Block 5. Sandwich the ratio: $f(a+1)/f(a) < \frac{1}{f(a)} \int_a^{a+1} f \leq T(a)$, the first step by monotonicity of f on the window, the second by replacing the denominator $f(x)$ by the larger $f(a)$.

Block 6. Squeeze: the ratio tends to 0, so the integrand tends to 1, so the integral diverges. Contradiction.

The nominal route stops at Block 1 and attacks the integrand pointwise, looking for a minorant like c/x. It fails, and it fails silently, which is the worst kind of failure: you can spend twenty minutes producing bounds that are true and useless. The second nominal route goes through

logarithms, gets the telescoping right, and then hits the uniformity wall at $-\log(1-t) \le Ct$; that route is salvageable but only with measure theory.

The bifurcation point is Block 3, and it is a decision about what "converges" is for. Zombie: convergence is the thing to be refuted, so I must find a lower bound for the integral. Ideal: convergence is a hypothesis with consequences, and its consequence, a vanishing tail, is an upper bound for anything I can fit underneath the tail. Ask what fits underneath, and the answer is the ratio $f(a+1)/f(a)$, which is the very quantity the integrand is built from. That self reference is the whole solution.

Layer 2: HOW

Block 1. **Cue.** A quotient whose numerator is a difference of the same function at two points. **Why natural.** Dividing through gives $1 - f(x+1)/f(x)$, which exposes scale invariance and bounds the integrand in $(0,1)$. Scale invariance is not decoration: it tells you in advance that no argument using the size of f can work, which deletes half the possible attacks before you try them. **Link to the next block.** Once the integrand is a ratio, the natural object is a product, and products over a ray sample naturally at the integers.

Block 2. **Cue.** The step size is exactly 1, and $f(x+1)/f(x)$ telescopes when sampled at integers. **Why natural.** Multiplying the sampled ratios gives $f(N)/f(0)$, and the hypothesis $f \to 0$ says the product collapses to zero. An infinite product of factors in $(0,1)$ collapses to zero precisely when the complementary series diverges, so the integer version of the claim is a standard theorem. This is reconnaissance, not proof, but it fixes the mechanism in your head: the only hypothesis doing work is $f \to 0$. **Link to the next block.** The continuum version cannot be read off from the integer version, so you need an argument that never localises to a window. Contradiction is the frame that lets you use global information.

Block 3. **Cue.** "Prove the integral diverges", with a positive integrand. **Why natural.** For a positive integrand, convergence is equivalent to the tails tending to zero, and there is no other content in the word "converges". Assuming convergence therefore hands you a free family of small quantities indexed by a, and small quantities are useful only as upper bounds. So you should immediately go hunting for something to place underneath $T(a)$. **Link to the next block.** The integrand contains $f(x)$ in the denominator, so anything you place underneath $T(a)$ will involve ratios of values of f.

Block 4. **Cue.** A difference $f(x) - f(x+1)$ integrated over an unbounded interval. **Why natural.** Shifting the variable in the second half turns $\int_a^T f(x+1)dx$ into $\int_{a+1}^{T+1} f$, and the two integrals overlap on almost all of their range. What survives is the head $\int_a^{a+1} f$ minus the tail $\int_T^{T+1} f$, and the tail dies because $f \to 0$. This is the exact continuous analogue of $\sum_{k=a}^{b-1}(u_k - u_{k+1}) = u_a - u_b$, and you should reach for it the instant a one step difference appears under an integral sign. **Link to the next block.** You now have a finite, explicit value for the integral of the numerator alone, which is the only part of the integrand you can compute.

Block 5. **Cue.** The denominator $f(x)$ varies and is the only thing preventing the numerator's telescoping from being useful. **Why natural.** Monotonicity gives $f(x) \le f(a)$ on the tail, so replacing $f(x)$ by $f(a)$ weakens the integrand in the safe direction: $T(a) \ge \frac{1}{f(a)} \int_a^\infty (f(x) - f(x+1))dx = \frac{1}{f(a)} \int_a^{a+1} f$. And $\int_a^{a+1} f > f(a+1)$ because f exceeds its right endpoint value across the window. The chain closes on $f(a+1)/f(a)$. **Link to the next block.** The quantity now bounded by the tail is exactly the quantity appearing in the integrand.

Block 6. Cue. A positive quantity squeezed under something tending to zero, which also sits inside the integrand. **Why natural.** If $f(a+1)/f(a) \to 0$ then $g(a) \to 1$, so g is eventually at least $\frac{1}{2}$, and an integrand bounded below by a positive constant on an infinite ray has infinite integral. The contradiction is with the assumption that started the argument, so the assumption is false. **Link to the next block.** Nothing remains: positivity of g upgrades "does not converge" to "diverges to $+\infty$".

Layer 3: WHY

The mathematical truth beneath the problem is the infinite product criterion. For $t_n \in [0,1)$, the product $\prod(1-t_n)$ has a nonzero limit if and only if $\sum t_n$ converges; equivalently, the product collapses to zero exactly when the series diverges. The hypothesis $f \to 0$ says that the product of one step ratios collapses, so the total relative drop must be infinite. Everything else in this problem is the technical work of passing from the integers to the continuum.

That passage is the real difficulty, and it is worth naming precisely why it is hard. The integrand at x compares $f(x+1)$ to $f(x)$, a window that slides continuously. Knowing the drops across $[n, n+1]$ for every integer n tells you nothing about the drop across $[x, x+1]$ for irrational x, because f may do all of its descending in a tiny subinterval. The logarithm route feels natural and is natural, but it needs the inequality $-\log(1-t) \le 2t$, which holds only for $t \le \frac{1}{2}$, and the set of x where the one step drop exceeds a half can be a complicated set that interleaves with its complement. Kedlaya's official route pays for that with Lebesgue measure: it isolates a set of positive measure inside $[0,1]$ on which the bound is valid for all shifts, and then applies monotone convergence. Paul Allen's route refuses to localise at all, aggregating over blocks and freezing the denominator. Joshua Rosenberg's route, the one written above, refuses even to look at the integrand, and works only with global tails.

The discernment mechanism, and this is the transferable part: for "prove that this integral diverges" with an unspecified function, there are exactly three engines. First, exhibit a positive lower bound for the integrand on a set of infinite measure. Second, aggregate over blocks and show each block contributes at least a fixed amount. Third, assume convergence and mine it. Engine three is the cheapest whenever the integrand is built from a ratio, because convergence controls tails, tails control averages, and averages of a monotone function control ratios. Scan these three in order and commit within two minutes; the ranking above shows that engines two and three both close this problem while engine one cannot.

Train the reflex like this. Whenever you assume convergence of $\int_0^\infty g$ with $g \ge 0$, write two lines immediately, before thinking: the tail $\int_a^\infty g \to 0$, and g cannot be bounded below by a positive constant on any set of infinite measure. Those two lines are the entire content of the assumption. Then ask the only remaining question: which quantity built from the problem's data can I trap underneath the tail? Here the answer arrives by asking what the numerator integrates to on its own, which is what makes Step 1 the pivot of the proof.

The examiners nudged in several ways worth naming. The integrand is presented as $(f(x) - f(x+1))/f(x)$ rather than as $1 - f(x+1)/f(x)$, which hides the scale invariance and hides the product structure. The step size is 1 and the domain is $[0,\infty)$, so the integers are an available but treacherous scaffolding. "Strictly decreasing" is there to keep f positive and for no other reason. And $\lim f = 0$ is the only hypothesis with real content: drop it, replace it with $f \to c > 0$, and the integrand becomes integrable for the obvious examples, so any proof that does not use $f \to 0$ decisively is wrong.

Revolving concepts

Improper integrals of nonnegative functions. For $g \ge 0$ the map $T \mapsto \int_0^T g$ is increasing, so the improper integral converges if and only if these partial integrals are bounded, and convergence is equivalent

to $\int_a^\infty g \to 0$ as $a \to \infty$. There is no conditional convergence to worry about. See Book 1, field: Real analysis, topic: improper integrals and tails.

Telescoping, continuous form. The identity used in Step 1 generalises: if $f \geq 0$ is decreasing with $f \to 0$ then for any $h > 0$,

$$\int_a^\infty \big(f(x) - f(x+h)\big)dx = \int_a^{a+h} f(x)\,dx.$$

The discrete parent is $\sum_{k=a}^{b-1}(u_k - u_{k+1}) = u_a - u_b$ with $u_b \to 0$. The general recursion to install: whenever a problem contains the one step difference operator $(\Delta_h f)(x) = f(x) - f(x+h)$, its sum or integral over a ray telescopes to a head term, and the head term is a window average rather than a point value. See Book 1, field: Real analysis, topic: telescoping and difference operators.

Infinite products. For $t_n \in [0,1)$, $\prod(1-t_n)$ converges to a nonzero limit if and only if $\sum t_n < \infty$; if $\sum t_n = \infty$ the product tends to 0. The proof is log of the product, plus the two sided bound $t \leq -\log(1-t) \leq 2t$ valid on $[0, \frac{1}{2}]$, plus separate treatment of the case where t_n does not tend to 0. This is the theorem the problem is a disguise for. See Book 1, field: Real analysis, topic: infinite products.

Halving times. Given a positive quantity decaying to zero, define $a_0 = 0$ and let a_{n+1} be the first integer beyond a_n with $f(a_{n+1}) < f(a_n + 1)/2$. This recursion is well defined precisely because $f \to 0$, it produces infinitely many disjoint blocks, and any argument extracting a fixed amount of budget per block then diverges. Generalised: to prove a sum or integral is infinite when you know only that some quantity decays to zero, cut the domain at the halving times of that quantity and prove a uniform positive lower bound per piece. This is the engine of the block telescoping route. See Book 1, field: Real analysis, topic: stopping times and block decomposition.

Monotone function facts. A decreasing f with limit 0 is positive; on $[a, a+1]$ it satisfies $f(a) \geq f(x) \geq f(a+1)$, so $f(a+1) \leq \int_a^{a+1} f \leq f(a)$, and this two sided window bound is what converts integrals back into point values. See Book 1, field: Real analysis, topic: monotone functions and integral bounds.

Knowledge gaps, named without disguise. If any of the following is not installed, this problem is not approachable and no amount of cleverness will substitute. One: the tail characterisation of convergence for improper integrals of nonnegative functions, including the fact that a convergent integral forces the integrand below any positive constant on arbitrarily late sets. Two: infinite products and their equivalence with series, which is the actual theorem in play. Three: the inequality $-\log(1-t) \leq t + 2t^2$ on $[0, \frac{1}{2}]$ and, more importantly, the fact that no linear bound holds uniformly on $[0,1)$; that failure is the entire reason A6 is not A2. Four: for the official first route only, Lebesgue measure, sets of positive measure, and the monotone convergence theorem. You can score full marks without item four by taking the block route or the tail route, and you should know that in advance so that you do not stall waiting for measure theory you have not installed.

Circuit Bank

- **WHEN** an integrand is a quotient whose numerator is a difference of the same function **DO** divide through to expose $1 - f(x+h)/f(x)$, and say the words "infinite product" **BECAUSE** ratios multiply and telescope, so a sum or integral of $1 -$ ratio is a statement about a product collapsing
- **WHEN** you must prove an integral of a positive function diverges **DO** assume convergence and write the two consequences at once: the tail tends to zero, and the integrand cannot stay above a positive constant late in the ray **BECAUSE** those two lines are the complete content of convergence, and one of them will contradict something you can build
- **WHEN** a one step difference $f(x) - f(x+h)$ sits under an integral over a ray **DO** shift the second copy and telescope, obtaining the head window $\int_a^{a+h} f$ **BECAUSE** overlapping ranges

cancel, and the surviving boundary term dies exactly when $f \to 0$

- **WHEN** a varying denominator blocks a telescoping numerator **DO** freeze it at its largest value on the block using monotonicity, and check that the inequality moves in the direction you need **BECAUSE** a lower bound survives replacing the denominator by a bigger one, and telescoping is worth far more than the constant you lose
- **WHEN** the only hypothesis is that some quantity decays to zero **DO** cut the domain at halving times of that quantity and extract a fixed budget from each block **BECAUSE** infinitely many halvings exist precisely because the limit is zero, so any per block bound sums to infinity

Backtesting

1. First useful cue: the integrand rewritten as $1 - f(x+1)/f(x)$, which is scale invariant, lies in $(0,1)$, and telescopes multiplicatively when sampled at integers.
2. Method the expression suggested: assume convergence, use the vanishing tail as an upper bound, and trap the ratio $f(a+1)/f(a)$ underneath it using the telescoped numerator.
3. Where you would hesitate, and the thought to reinforce: after seeing that the integer sampled statement is the infinite product criterion, you will want to reduce the integral to that sample, and there is no such reduction. Reinforce this: when a discrete model of the problem is easy and the continuous version is not, do not force the reduction; find an argument that never localises. Global tails are the tool that never localises.
4. The trap the question sets: attempting a pointwise minorant for the integrand. No minorant exists, because f may be nearly flat over arbitrarily long stretches, so time spent here is unrecoverable.
5. What makes you recognise it faster next time: an integral whose integrand is 1 minus a ratio of the same function at shifted arguments is always a product in disguise; and the presence of $\lim f = 0$ as the only real hypothesis says that the proof must convert "reaches zero" into "infinitely many multiplicative halvings".
6. The reusable decision rule: to prove divergence of a positive integral built from ratios, assume convergence, telescope the numerator over the tail, freeze the denominator by monotonicity, and squeeze the ratio to zero; a ratio squeezed to zero drives the integrand to 1 and destroys the assumption.

Mirror drills

1. Let $t_n \in [0,1)$. Prove that $\prod_{n\geq 0}(1-t_n) = 0$ if and only if $\sum_n t_n = \infty$. *Answer:* if $t_n \not\to 0$ both statements hold trivially; otherwise take logarithms and use $t \leq -\log(1-t) \leq 2t$ for $t \in [0, \frac{1}{2}]$, valid for all large n.
2. Let f be positive, decreasing, with $f \to 0$. Prove $\int_a^\infty (f(x) - f(x+1))dx = \int_a^{a+1} f(x)dx$. *Answer:* shift the second copy to $\int_{a+1}^{T+1} f$, cancel the overlap, and note the boundary term $\int_T^{T+1} f \leq f(T) \to 0$.
3. Let $a_n > 0$ be decreasing with $a_n \to 0$. Prove $\sum_{n\geq 0}\left(1 - a_{n+1}/a_n\right)$ diverges. *Answer:* group indices into blocks on which a at least halves; within a block starting at N, $\sum(a_n - a_{n+1})/a_n \geq \sum(a_n - a_{n+1})/a_N \geq \frac{1}{2}$, and there are infinitely many blocks.

16.7 B1. power sums forced into a false inequality

Is there an infinite sequence of real numbers $a_1, a_2, a_3, \ldots$ such that

$$a_1^m + a_2^m + a_3^m + \cdots = m$$

for every positive integer m?

Read it out loud

Read this out loud and notice what kind of object you are being handed. It is not one equation in the unknowns $a_1, a_2, a_3, \ldots$, it is infinitely many equations sharing the same unknowns, one for every positive integer m, and every single one of them must hold at once. The fields wired into that sentence are real analysis, because you are silently being told each of these infinite sums converges, and inequalities, because the only honest way to relate several different power sums of one fixed sequence to each other is through an inequality that compares them.

Within thirty seconds you should have peeled off the first few instances explicitly, since a statement true for every positive integer m is never meant to be used in full generality on a first pass. Write

$$\sum a_i = 1, \qquad \sum a_i^2 = 2, \qquad \sum a_i^3 = 3, \qquad \sum a_i^4 = 4,$$

and look at what you now have: three consecutive exponents, 2, 3 and 4, each pinned to a fixed number, on the same underlying sequence. The moment three consecutive exponents like that appear together, the guiding light is Cauchy Schwarz, because a_i^3 splits exactly as the product $a_i \cdot a_i^2$ of an exponent 1 term and an exponent 2 term, and Cauchy Schwarz is precisely the inequality that bounds a sum of products by the two sums of squares of the factors, exponents 2 and 4, both of which you already have. Once you see that split, the rest of the problem is arithmetic, not insight.

Approaches, ranked worst to best

1. **explicit construction by trial.** *Cue:* the word sequence invites you to write one down. *Circuit:* guess and check on geometric or finitely supported sequences. (even correctly suspecting the answer is negative, trying candidate sequences can never prove nonexistence: ruling out infinitely many sequences by exhibiting failures in a handful of them is not a proof, and this route never terminates.)

2. **casework on $|a_i| \le 1$ against $|a_i| > 1$.** *Cue:* the exponents keep climbing while the targets m only grow linearly. *Circuit:* compare termwise growth rates directly. (if every $a_i^2 \le 1$ then $a_i^4 \le a_i^2$ termwise, giving $4 = \sum a_i^4 \le \sum a_i^2 = 2$, an immediate contradiction; otherwise some $a_k^2 > 1$, and you must further show a_k^{2m} eventually exceeds $2m$, forcing $\sum a_i^{2m} \neq 2m$ for m large. Correct, and close to the ideal route in spirit, but it costs two branches and an extra limiting argument where the winning route needs neither.)

3. **Cauchy Schwarz on the pair (a_i) and (a_i^2).** *Cue:* three consecutive integer exponents pinned to fixed values. *Circuit:* the general power sum bound $S_{p+q}^2 \le S_{2p} S_{2q}$. (one inequality, one numerical contradiction, no casework at all. This is the ideal route: the exponents in the hypothesis were already telling you which inequality to reach for.)

The ideal solution

No such sequence exists.

Suppose toward a contradiction that a sequence $a_1, a_2, a_3, \ldots$ satisfies $\sum_{i=1}^{\infty} a_i^m = m$ for every positive integer m. Taking $m = 2$ and $m = 4$ in particular,

$$\sum_{i=1}^{\infty} a_i^2 = 2, \qquad \sum_{i=1}^{\infty} a_i^4 = 4,$$

and both series have nonnegative terms, so both converge absolutely, with every partial sum bounded above by the full sum.

Fix a positive integer N. Applying the Cauchy Schwarz inequality to the finite real sequences $u_i = a_i$ and $v_i = a_i^2$ for $1 \le i \le N$ gives

$$\left(\sum_{i=1}^{N} a_i^3\right)^2 = \left(\sum_{i=1}^{N} u_i v_i\right)^2 \le \left(\sum_{i=1}^{N} u_i^2\right)\left(\sum_{i=1}^{N} v_i^2\right) = \left(\sum_{i=1}^{N} a_i^2\right)\left(\sum_{i=1}^{N} a_i^4\right).$$

Since both factors on the right are partial sums of nonnegative series bounded by their full sums,

$$\left(\sum_{i=1}^{N} a_i^3\right)^2 \le \left(\sum_{i=1}^{\infty} a_i^2\right)\left(\sum_{i=1}^{\infty} a_i^4\right) = 2 \times 4 = 8$$

for every N. By hypothesis, taking $m = 3$, the partial sums $\sum_{i=1}^{N} a_i^3$ converge to 3 as $N \to \infty$, so letting $N \to \infty$ in the bound above,

$$3^2 = 9 \le 8.$$

This is false, so no such sequence exists.

Layer 1: WHAT

Block 1: read off the specific instances of the hypothesis you actually need, at $m = 2, 3, 4$, rather than trying to use the hypothesis in full generality.

Block 2: notice the exponent identity $a_i^3 = a_i \cdot a_i^2$, which turns $\sum a_i^3$ into an inner product of the two sequences (a_i) and (a_i^2).

Block 3: apply Cauchy Schwarz to the Nth partial sums of those two sequences, bounding $\left(\sum_{i=1}^{N} a_i^3\right)^2$ by the product of the two known convergent totals, 2 and 4.

Block 4: pass to the limit $N \to \infty$, using the given convergence of $\sum a_i^3$ to 3, to obtain the false inequality $9 \le 8$ and close the proof.

The nominal route tries to construct or eliminate candidate sequences directly, or splits into cases on the size of $|a_i|$; the ideal route never touches an individual a_i at all, it only ever manipulates the three totals 2, 3, 4 through one inequality. The bifurcation point sits right after you write down $\sum a_i^2 = 2, \sum a_i^3 = 3, \sum a_i^4 = 4$: the question is whether you treat the a_i as unknowns to search over, or treat the three totals as data to feed into a known inequality relating power sums. Taking the second option is the entire difference in cost between the routes.

Layer 2: HOW

Block 1. **Cue.** The same sequence has its second, third and fourth power sums all pinned to fixed numbers. **Why natural.** Powers only relate to each other through multiplication, and $a_i^3 = a_i \cdot a_i^2$ is the one algebraic identity connecting the exponent 3 term to the exponent 1 and exponent 2 terms; writing it out is reading the exponents present, not insight. **Link to the next block.** Once a_i^3 is seen as the product $a_i \cdot a_i^2$, the two factors are exactly a pair of real sequences whose own sums of squares, $\sum a_i^2 = 2$ and $\sum (a_i^2)^2 = \sum a_i^4 = 4$, are the two other numbers the hypothesis already hands you.

Block 2. **Cue.** Two real sequences with known, finite sums of squares. **Why natural.** Cauchy Schwarz is the definitional fact that an inner product cannot exceed the product of the two norms, and $\sum u_i v_i$, $\sum u_i^2$, $\sum v_i^2$ are exactly an inner product and two squared norms once you name $u_i = a_i, v_i = a_i^2$. **Link to the next block.** Applying it produces an upper bound on $\left(\sum a_i^3\right)^2$ built purely from the two totals 2 and 4, entirely independent of what the hypothesis says $\sum a_i^3$ equals.

Block 3. **Cue.** The bound obtained, 8, is a fixed number that does not depend on the sequence at all. **Why natural.** The hypothesis independently supplies a second, competing value for the same quantity, 9; two honest derivations of one quantity that disagree is the cleanest possible contradiction, since no further casework is needed once both numbers are in hand. **Link to the next block.** This numerical clash is the whole proof; nothing about the individual terms a_i was ever pinned down or needed to be.

Block 4. **Cue.** The sum $\sum a_i^3$ that appears is an infinite sum, and the Cauchy Schwarz inequality as usually memorised is a finite statement. **Why natural.** A finite inequality does not automatically hold in a limit without justification, so you must apply it to partial sums, which are always finite, and only then let the index tend to infinity, using convergence you were already given. **Link to the next block.** This turns the numerical bound into a fully rigorous statement about the infinite sums, with no hand waving about an inequality obviously extending.

Layer 3: WHY

The underlying truth is that Cauchy Schwarz is a statement about the geometry of square summable sequences: any two sequences whose squares sum to something finite behave like vectors, and their inner product is capped by the product of their lengths. Every family of power sums of one fixed sequence is secretly a family of such inner products of the sequence with powers of itself, so different power sums of the same object are never independent data, they are linked by this one inequality, and more generally by Holder's inequality for exponent pairs beyond the square case.

The discernment mechanism to train is this: whenever a problem fixes several power sums of one sequence and asks about existence, before considering any construction, write the given totals as a short table indexed by exponent and scan for an exponent identity $p + q = r$ among the exponents present. If you find one, Cauchy Schwarz very likely closes the problem in one inequality, and the split needs the lower exponent in the pair to be an honest integer power of the sequence, here $p = 1, q = 2$, since a fractional power of a possibly negative real number is not defined; that restriction is why the trick uses exponents $2, 3, 4$ rather than, say, $1, 2, 3$, and it is worth noticing which exponent triples admit the split and which do not.

What the examiners were nudging toward is visible in the phrase "for every positive integer m": insisting the equality hold for the entire infinite family hands you far more than you need, since only three values close the argument. Whenever a problem supplies an infinite family of constraints and asks a yes or no existence question, check first whether a small finite subset already contradicts itself, because setters frequently oversupply constraints exactly when the intended answer is negative.

Revolving concepts

Cauchy Schwarz inequality. For real sequences $(u_i), (v_i)$ with $\sum u_i^2$ and $\sum v_i^2$ convergent, $\left(\sum u_i v_i\right)^2 \leq \left(\sum u_i^2\right)\left(\sum v_i^2\right)$, with equality exactly when (u_i) and (v_i) are proportional. This is the $p = q = 2$ instance of Holder's inequality $\sum |u_i v_i| \leq \left(\sum |u_i|^p\right)^{1/p} \left(\sum |v_i|^q\right)^{1/q}$ for conjugate exponents $1/p + 1/q = 1$. See Book 1, field: Inequalities, topic: Cauchy Schwarz and Holder.

Power sums of a sequence. Writing $S_m = \sum a_i^m$, the general lemma actually used here is $S_{p+q}^2 \leq S_{2p} S_{2q}$ for nonnegative integers p, q, obtained by applying Cauchy Schwarz to $u_i = a_i^p, v_i = a_i^q$. More

broadly, the moments of a fixed object are log convex in the exponent, a fact known in probability as Lyapunov's inequality; installing the general $S_{p+q}^2 \leq S_{2p} S_{2q}$ form directly is the efficient way to carry this problem's method into other years. See Book 1, field: Sequences and recurrences, topic: power sum and moment inequalities.

Series of nonnegative terms and the partial sum limit technique. A series of nonnegative reals converges exactly when its partial sums are bounded above, and any finite inequality that holds for every partial sum of two such series passes to the limit; this is the standard, necessary device for legally extending Cauchy Schwarz, the triangle inequality or AM GM from finite sums to infinite ones. See Book 1, field: Real analysis, topic: series convergence.

Newton's identities. For finitely many variables the power sums p_k and the elementary symmetric polynomials e_k satisfy the recursion $p_k - e_1 p_{k-1} + e_2 p_{k-2} - \cdots + (-1)^{k-1} e_{k-1} p_1 + (-1)^k k e_k = 0$, determining each p_k from the earlier power sums and the symmetric functions, or vice versa. Not load bearing here, but it is the standard orbiting machinery around power sums of a sequence and worth keeping distinct from the inequality route used in this solution. See Book 1, field: Polynomials, topic: Newton's identities.

Circuit Bank

- **WHEN** a problem fixes several power sums S_m of one real sequence and asks about existence or impossibility **DO** write the given values as a short table against their exponents and scan for a pair p, q with $2p, 2q$ and $p+q$ all present among the given exponents **BECAUSE** Cauchy Schwarz turns S_{2p} and S_{2q} into an upper bound on S_{p+q}^2, and this is almost always cheaper than direct construction or termwise growth comparisons
- **WHEN** you need to extend a finite algebraic inequality, Cauchy Schwarz, the triangle inequality or AM GM, to an infinite sum **DO** apply the finite inequality to the Nth partial sums, keep the bound valid for every N, then let $N \to \infty$ using whatever convergence the hypothesis already grants **BECAUSE** finite inequalities do not automatically survive the limit without this step, and skipping it costs rigour marks even when the final conclusion is correct
- **WHEN** an existence question hands you infinitely many simultaneous constraints, for every m, for every n, and similar phrasing **DO** check whether a small finite handful of the constraints already contradicts itself before attempting any construction **BECAUSE** setters who quantify over every instance almost always oversupply data when the intended answer is negative, and the redundant data is usually the fastest route to the contradiction

Backtesting

1. First useful cue: the exponents $2, 3, 4$ sitting among the for every m constraints, since $3 = 1 + 2$ while 2 and 4 are twice 1 and twice 2.
2. Method the expression suggested: Cauchy Schwarz applied to the pair of sequences (a_i) and (a_i^2), because a_i^3 is exactly their termwise product.
3. Where you would hesitate, and the thought to reinforce: doubting whether Cauchy Schwarz is valid for an infinite sum at all; reinforce by recalling the partial sum then limit fix rather than abandoning the approach.
4. The trap the question sets: quantifying over every positive integer m tempts you to think you must use all of that data, or construct an explicit sequence; you only ever need three specific values.
5. What makes you recognise it faster next time: seeing two or more power sums of one sequence fixed simultaneously should fire the Cauchy Schwarz reflex before construction is even considered.

6. The reusable decision rule: when a sequence's power sums are pinned at several exponents, hunt for an exponent identity among them and reach for Cauchy Schwarz or Holder before anything else.

Mirror drills

1. Is there a sequence of reals with $\sum a_i^2 = 3$, $\sum a_i^3 = 4$ and $\sum a_i^4 = 5$? *Answer:* no, by the identical argument, since Cauchy Schwarz would force $16 \leq 15$.
2. A sequence of nonnegative reals has $\sum a_i^2 = 5$ and $\sum a_i^4 = 5$. What is the largest possible value of $\sum a_i^3$, and what does achieving it tell you about the sequence? *Answer:* at most 5, by Cauchy Schwarz, with equality exactly when (a_i) and (a_i^2) are proportional, forcing every nonzero term to share one common value.
3. Is there a sequence of reals with $\sum a_i^4 = 2$, $\sum a_i^5 = 3$ and $\sum a_i^6 = 3$? *Answer:* no, the same split shifted to $u_i = a_i^2$, $v_i = a_i^3$ gives $9 \leq 6$, false.

16.8 B2. smallest integer side of a lattice triangle

Given that A, B, and C are noncollinear points in the plane with integer coordinates such that the distances AB, AC, and BC are integers, what is the smallest possible value of AB?

Read it out loud

Read this out loud: three lattice points, pairwise integer distances, not collinear, and you are asked for the smallest a single labelled side can be. The fields wired in are geometry, since you are working with lattice points and the triangle inequality, and number theory, since integer coordinates force $x^2 + y^2$ type expressions to be perfect squares wherever a distance is required to be an integer. This is a Heronian triangle question in disguise, a triangle with integer sides that also happens to sit on the integer lattice, and "smallest possible value of AB" is an extremal command term: you are not asked to classify every such triangle, only to pin down one number.

Within thirty seconds you should have done two things. First, hunt for a small example to know what you are aiming to prove is best possible: the right angled triangle with legs 3 and 4 and hypotenuse 5 sits on the lattice at $(0,0)$, $(3,0)$, $(0,4)$ with all three sides integers, giving a candidate answer of 3. Second, notice that AB being genuinely small, 1 or 2, is a very restrictive condition once you also demand A, B, C are not collinear, because the reverse triangle inequality $|AC - BC| < AB$ then leaves the other two sides almost no room to differ. That second observation, not the example, is the real key: it converts a geometry question into an integer rigidity question with almost no free choices left.

Approaches, ranked worst to best

1. **brute force coordinate search.** *Cue:* integer coordinates invites literally trying small lattice triangles. *Circuit:* systematic or hand enumeration of nearby lattice points. (even a search that correctly finds no working triangle with $AB \in \{1,2\}$ among the cases you tried never proves impossibility for the lattice points you did not check, so it cannot stand alone as a proof and costs unbounded ink.)

2. **Heron's formula together with the lattice area lemma.** *Cue:* integer sides make the triangle Heronian, and integer coordinates force twice the area to be an integer. *Circuit:* Heron's formula combined with the shoelace determinant fact. (for $AB = 1$ the two remaining sides are forced equal by the triangle inequality, and Heron's formula closes quickly via a difference of squares factorisation; but for $AB = 2$ the remaining sides can differ by 1, and Heron's formula then produces a quadratic Diophantine condition with no clean factorisation. This route terminates, unlike brute search, but one branch costs far more ink than the route below.)

3. **reverse triangle inequality rigidity plus the perpendicular bisector.** *Cue:* AB is forced tiny, leaving almost no room between $|AC - BC|$ and AB. *Circuit:* the equality case of the triangle inequality is exactly collinearity, which the hypothesis excludes. (place A and B using the fact that the only lattice points at distance 1 or 2 from a fixed lattice point are axis aligned; then $|AC - BC| < AB$ forces $AC = BC$ outright for $AB = 1$, and after one short parity check forces the same for $AB = 2$. In both cases C is pinned to the perpendicular bisector of AB, where integrality kills it in one further line. This is the ideal route: every step is forced, and the only casework is a single parity check.)

The ideal solution

The smallest possible value of AB is $\boxed{3}$, achieved by $A = (0,0)$, $B = (3,0)$, $C = (0,4)$: here $AB = 3$, $AC = 4$, $BC = 5$, all integers, and the three points are not collinear.

It remains to rule out $AB = 1$ and $AB = 2$.

Two facts do all the work. First, since A, B, C are not collinear, the reverse triangle inequality is strict:

$$|AC - BC| < AB,$$

because equality in $|AC - BC| \leq AB$ holds exactly when one of A, B, C lies on the segment joining the other two, which is a case of collinearity. Second, since A and B have integer coordinates and $AB \in \{1, 2\}$, the only lattice points at distance AB from A are obtained from A by moving along a single coordinate axis: the equation $x^2 + y^2 = 1$ forces $(x, y) \in \{(\pm 1, 0), (0, \pm 1)\}$ and $x^2 + y^2 = 4$ forces $(x, y) \in \{(\pm 2, 0), (0, \pm 2)\}$, since neither 1 nor 4 has any other representation as a sum of two integer squares. Translating A to the origin and applying a quarter turn about it, both of which preserve the lattice and every distance, you may therefore place $A = (0, 0)$ and $B = (AB, 0)$ without loss of generality.

Case $AB = 1$**.** Write $A = (0, 0)$, $B = (1, 0)$, $C = (x, y)$ with $x, y \in \mathbb{Z}$. Since $|AC - BC| < 1$ and AC, BC are integers, $AC = BC$. Then C lies on the perpendicular bisector of AB, the line $x = \frac{1}{2}$; but x is an integer, so $x \neq \frac{1}{2}$, a contradiction. So $AB = 1$ is impossible.

Case $AB = 2$**.** Write $A = (0, 0)$, $B = (2, 0)$, $C = (x, y)$ with $x, y \in \mathbb{Z}$. Then

$$BC^2 - AC^2 = \left((x - 2)^2 + y^2\right) - \left(x^2 + y^2\right) = 4 - 4x,$$

which is even, so $AC^2 \equiv BC^2 \pmod 2$. Since an integer and its square always share parity, $AC \equiv BC \pmod 2$, so $AC - BC$ is even. On the other hand $|AC - BC| < AB = 2$ forces $|AC - BC| \in \{0, 1\}$, and the only even value available is 0, so again $AC = BC$. Then C lies on the perpendicular bisector $x = 1$, so write $C = (1, y)$. Now $AC^2 = 1 + y^2$ must be a perfect square, say $AC = k$ for a positive integer k, so $k^2 - y^2 = 1$, that is $(k - y)(k + y) = 1$. Both factors are integers, so either $k - y = k + y = 1$, giving $y = 0, k = 1$, or $k - y = k + y = -1$, giving $k = -1$, impossible since $k > 0$. So $y = 0$ and $C = (1, 0)$, which lies on the line through A and B, contradicting noncollinearity. So $AB = 2$ is impossible.

Since $AB = 1$ and $AB = 2$ are both impossible and $AB = 3$ is achieved, the smallest possible value of AB is 3.

Layer 1: WHAT

Block 1: exhibit the $3, 4, 5$ triangle on the lattice to fix a target value and confirm 3 is achievable.

Block 2: extract the two governing facts, the strict reverse triangle inequality from noncollinearity, and the fact that lattice points at distance 1 or 2 from a fixed lattice point are forced onto the coordinate axes, which licenses a clean placement of A and B.

Block 3: run $AB = 1$: the strict inequality alone forces $AC = BC$, and the perpendicular bisector together with integer coordinates finishes it in one line.

Block 4: run $AB = 2$: the strict inequality only narrows $|AC - BC|$ to two candidate values, so a short parity argument on $AC^2 - BC^2$ is needed to knock out the value 1 before the same perpendicular bisector idea closes it, this time via a difference of squares factorisation rather than a bare fraction.

The nominal route reaches for Heron's formula and the lattice area identity, or for brute enumeration; the ideal route reaches for the reverse triangle inequality's equality case. The bifurcation point sits right after noticing AB is forced small: whether you reach for an area formula, which keeps both unknown sides present and produces a genuine two variable Diophantine condition in the $AB = 2$ branch, or for the reverse triangle inequality, which converts "the two other sides are close in value" directly into "they are equal," collapsing a whole degree of freedom in a single line.

Layer 2: HOW

Block 1. Cue. The problem asks for the smallest possible value of a specific side. **Why natural.** An extremal existence question is answered by pairing a construction, giving an upper bound on

how small the answer could be, with a proof that nothing smaller works; writing down a candidate construction first tells you exactly what you need to rule out. **Link to the next block.** The $3, 4, 5$ triangle fixes the target as 3 and tells you the remaining work is to eliminate $AB = 1$ and $AB = 2$ only.

Block 2. Cue. Noncollinear points with a very small integer distance between two of them. **Why natural.** Noncollinearity is precisely the strict case of the triangle inequality, so smallness of AB combined with noncollinearity is definitionally a statement that AC and BC are forced close together; separately, integer coordinates at a fixed small integer distance restrict to finitely many lattice directions, since $x^2 + y^2 = 1$ or 4 have only axis aligned solutions. **Link to the next block.** These two facts jointly let you place A, B concretely and immediately convert "AB small" into a rigid constraint on $AC - BC$.

Block 3. Cue. For $AB = 1$, the strict inequality gives $|AC - BC| < 1$ with both sides integers. **Why natural.** An integer strictly between -1 and 1 can only be 0, so $AC = BC$ is not a deduction to search for, it is forced immediately by integrality. **Link to the next block.** Equidistance from two fixed points is the definition of the perpendicular bisector, so C's x coordinate is pinned to a half integer, clashing with C being a lattice point.

Block 4. Cue. For $AB = 2$, the same inequality only narrows $|AC - BC|$ to $\{0, 1\}$, leaving one extra case to remove. **Why natural.** $AC - BC$ shares the parity of $AC^2 - BC^2$ exactly because an integer and its square always share parity, and $AC^2 - BC^2$ is directly computable from the coordinates as $4 - 4x$, visibly even; this rules out $|AC - BC| = 1$ without any further search. **Link to the next block.** With $AC = BC$ re established, the perpendicular bisector argument returns, now needing a factorisation $(k - y)(k + y) = 1$ rather than a bare non integer, because the bisector $x = 1$ is itself a lattice line this time.

Layer 3: WHY

The underlying truth is that noncollinearity is not a separate geometric condition to check at the end, it is the strict case of the triangle inequality itself, so any problem that pairs "noncollinear" with "a very small integer side" is handing you a rigidity statement for free: the other two sides are forced within one unit of each other before you do any work. The discernment mechanism to train is to test, immediately after reading a lattice distance problem, whether the given or requested side length is small enough that the triangle inequality alone, rather than an area formula, which treats all three sides symmetrically and loses this rigidity, pins down a relation between the other two sides. When it does, the triangle inequality route is essentially always cheaper, because it removes a degree of freedom outright rather than producing a further equation to solve.

The second truth worth training is about the lattice itself: a lattice point at a fixed small integer distance from another lattice point is not free to sit anywhere on the corresponding circle, it is constrained to the finitely many lattice points on that circle, and for the two smallest integer radii, 1 and 2, that finite set happens to be exactly the four axis aligned points, which is what licenses the WLOG placement rigorously rather than by appeal to intuition. This stops holding at distance 5, where $(3, 4)$ and its images also lie on the circle, so the placement trick is specific to small radii and you should check, not assume, that it applies before using it.

What the examiners were nudging toward is the phrase "noncollinear": it is doing far more work than excluding a degenerate picture, it is the entire mechanism that makes $AB = 1$ and $AB = 2$ impossible, and a candidate who treats it as a throwaway hypothesis rather than the engine of the proof will reach for Heron's formula and pay for it in the $AB = 2$ case.

Revolving concepts

Reverse triangle inequality and its equality case. For any three points, $|AC - BC| \leq AB$, with equality exactly when A, B, C are collinear with one of them between the other two; strict inequality is therefore equivalent to noncollinearity. See Book 1, field: Geometry, topic: triangle inequality.

Sum of two squares representations. An integer n can be written as $x^2 + y^2$ with x, y integers according to the prime factorisation of n in the Gaussian integers; the count and shape of such representations is why 1 and 4 admit only axis aligned lattice points at that distance from the origin, while 25 admits both axis aligned and genuinely diagonal ones. See Book 1, field: Number theory, topic: sums of two squares.

Heronian triangles and the lattice area lemma. A Heronian triangle has integer sides and rational area; a triangle with integer coordinate vertices automatically has an area equal to half an integer, by the shoelace determinant formula, so an integer sided lattice triangle is automatically Heronian. This is the mechanism behind the alternative, costlier approach, and it generalises the whole problem: no Heronian triangle at all, lattice or not, can have a side of length 1 or 2. See Book 1, field: Geometry, topic: Heronian triangles and Pick's theorem.

Lattice symmetries. Translations by integer vectors and rotations by a multiple of a right angle about a lattice point both preserve the set of lattice points and every pairwise distance; this is the precise justification for placing A at the origin and B on an axis, rather than an unexamined appeal to symmetry. See Book 1, field: Geometry, topic: lattice point symmetries.

Circuit Bank

- **WHEN** a problem pairs noncollinear points with a small integer side length **DO** apply the reverse triangle inequality immediately, before reaching for an area formula **BECAUSE** noncollinearity is exactly the strict case of the triangle inequality, so it directly restricts the other two sides once the given side is small, without introducing any new unknowns
- **WHEN** two lattice points are known to be equidistant from a third **DO** place the equidistant point on the perpendicular bisector and use integrality of coordinates or of the resulting distance to finish **BECAUSE** equidistance is the definition of the perpendicular bisector, and a lattice point restricted to a specific line is usually one short check away from a contradiction
- **WHEN** you need to compare the parity of two integer distances without computing either one **DO** compute the difference of their squares directly from coordinates and check its parity, since an integer and its square always share parity **BECAUSE** this avoids ever solving for the distances themselves, replacing a square root computation with a one line algebraic identity

Backtesting

1. First useful cue: the phrase noncollinear sitting next to a request for the smallest integer side.
2. Method the expression suggested: the reverse triangle inequality, since a small AB leaves almost no room between AC and BC.
3. Where you would hesitate, and the thought to reinforce: whether placing A and B on an axis is really justified for a general lattice point; reinforce by checking explicitly which lattice points sit at the given distance, rather than assuming rotational freedom.
4. The trap the question sets: reaching for Heron's formula on the strength of "integer sides" alone, which works but buries you in a quadratic Diophantine condition in the $AB = 2$ case that the triangle inequality route avoids entirely.

5. What makes you recognise it faster next time: any lattice distance problem where noncollinearity is stated explicitly should trigger a check of the reverse triangle inequality before any area formula.
6. The reusable decision rule: when a side length is small and the points are noncollinear, use the triangle inequality to force the other two sides equal, or nearly equal, before doing anything else.

Mirror drills

1. Show that no Heronian triangle, lattice or not, can have a side of length 1. *Answer:* the other two sides are forced equal by the triangle inequality, then Heron's formula reduces to a difference of squares equation with no integer solution.
2. Three lattice points have pairwise integer distances and are noncollinear; show $AB = 2$ still forces $AC = BC$ using only a coordinate parity argument, without first bounding $|AC - BC|$ by the triangle inequality. *Answer:* compute $AC^2 - BC^2$ from coordinates, observe it is always even, and conclude $AC \equiv BC \pmod 2$, so the odd difference is impossible regardless of the inequality bound.
3. Four lattice points, no three collinear, have all six pairwise distances integers; what is the smallest possible value of the shortest of the six distances? *Answer:* still at least 3, since each distance is a side of some noncollinear triangle among the four points and the identical argument applies; a matching example can be built by extending the $3, 4, 5$ triangle with a fourth suitably placed integer point.

16.9 B3. the threshold for redistributing balls to uniform

There are 2010 *boxes labeled* $B_1, B_2, \ldots, B_{2010}$*, and* 2010$n$ *balls have been distributed among them, for some positive integer* n*. You may redistribute the balls by a sequence of moves, each of which consists of choosing an* i *and moving* exactly i *balls from box* B_i *into any one other box. For which values of* n *is it possible to reach the distribution with exactly* n *balls in each box, regardless of the initial distribution of balls?*

Read it out loud

Read this out loud and notice the one detail that decides the whole problem: a move at index i does not let you choose how many balls to send, it forces you to send exactly i balls, and only out of box B_i itself. The fields wired in are combinatorics and counting, specifically the pigeonhole principle used not in its usual two pigeons one hole form but as a capacity bound: if every box is required to stay below its own label's threshold, there is a hard cap on the total number of balls that configuration can hold. The command term "for which values of n, regardless of the initial distribution" is asking for a threshold that works against an adversary who chooses the starting mess, so you need both a construction that succeeds above the threshold and a stuck configuration that fails below it.

Within thirty seconds you should have written down the quantity $\sum_{i=1}^{2010}(i-1)$, the total number of balls a distribution could hold while every single box B_i sits strictly below its own threshold i, since a box below threshold cannot originate any move at all and such a configuration is completely frozen. That sum is $\frac{2009\times 2010}{2} = 2010 \times 1004.5$, exactly half a box count short of a clean multiple of 2010, and this single number is doing all the work: below it a frozen starting configuration exists, and above it no configuration can freeze every box simultaneously, so some move is always available.

Approaches, ranked worst to best

1. **small case brute force.** *Cue:* 2010 boxes and an arbitrary n are unwieldy, so shrink both. *Circuit:* pattern hunting on a handful of boxes. (useful only for conjecturing the threshold; it proves nothing about 2010 boxes on its own and still needs a general argument bolted on afterwards.)

2. **hunting for a preserved residue or colouring invariant.** *Cue:* "for which n is it possible" reads like a classic invariant question. *Circuit:* parity or modular invariants under the move. (the move set here does not preserve any fixed modulus in a useful way, since a move at index i changes two box counts by exactly $\mp i$, and i itself varies from move to move; you can chase candidate invariants for a long time without finding one that pins down 1005 exactly, and time is lost before falling back to a direct counting argument.)

3. **capacity or pigeonhole obstruction for impossibility, paired with an explicit router box construction for sufficiency.** *Cue:* a move is only ever legal from a box holding at least its own label, and box B_1 is the unique box whose native move size is a single ball. *Circuit:* the frozen everywhere configuration bounds the threshold from below, and routing every ball through B_1 realises it from above. (the threshold falls out of one inequality, $\sum_{i=1}^{2010}(i-1)$ against 2010n, and the construction needs nothing more exotic than repeatedly emptying the largest eligible box through B_1. This is the ideal route.)

The ideal solution

The answer is $\boxed{n \geq 1005}$.

Why $n \le 1004$ fails. For $i = 1, \dots, 2010$, choose integers c_i with $0 \le c_i \le i-1$, so $c_1 = 0$, summing to exactly $2010n$; this is possible whenever $2010n \le \sum_{i=1}^{2010}(i-1) = \frac{2009\times 2010}{2} = 2010 \times 1004.5$, in particular whenever $n \le 1004$, since the allowances $[0, i-1]$ for $i = 1, \dots, 2010$ can realise any integer total from 0 up to that cap. Start with B_i holding c_i balls. Every box satisfies $B_i \le i-1 < i$, so no move is legal anywhere: a move at index i needs $B_i \ge i$. The distribution is frozen, and since $c_1 = 0 < n$, box B_1 can never reach its target of $n \ge 1$ balls. So for $n \le 1004$ the uniform distribution is not reachable from every start.

Why $n \ge 1005$ suffices. First, an invariant that holds at every distribution with total $2010n$, not only the initial one, since the total never changes under a move: some index i is always *eligible*, meaning $B_i \ge i$. Indeed if every $B_i \le i-1$ at once, the total would be at most $2010 \times 1004.5 < 2010 \times 1005 \le 2010n$, a contradiction.

Run the following three phases from an arbitrary starting distribution.

Phase 1, collect. While some index $i \ge 2$ is eligible, take the largest such i and move i balls from B_i into B_1. Each move strictly reduces the total held in boxes 2 through 2010 by at least 2, so after finitely many moves no index $i \ge 2$ remains eligible: every B_i for $i \ge 2$ now satisfies $B_i \le i-1$.

Phase 2, flush. Process the indices $j = 2010, 2009, \dots, 2$ in turn. Let v_j denote the value of B_j left by phase 1, unchanged while larger indices are processed since only B_1 and the current B_j are touched, so $0 \le v_j \le j-1$. Move 1 ball from B_1 to B_j, repeated $j - v_j$ times, bringing B_j up to exactly j; this makes B_j eligible, so move j balls from B_j back to B_1, leaving $B_j = 0$. Move on to the next smaller index.

This phase never asks B_1 for a ball it does not have, precisely because $n \ge 1005$. Right before processing index j, boxes $2010, \dots, j+1$ are already empty and boxes $2, \dots, j$ still hold their original phase 1 values, so

$$B_1 = 2010n - \sum_{k=2}^{j} v_k.$$

The largest single withdrawal you make before any repayment, while processing index j, is $j - v_j$ balls, so legality at this step needs

$$2010n - \sum_{k=2}^{j} v_k \;\ge\; j - v_j \quad\Longleftrightarrow\quad 2010n \;\ge\; j + \sum_{k=2}^{j-1} v_k.$$

Bounding every $v_k \le k-1$, the right side is at most $j + \frac{(j-2)(j-1)}{2}$, an expression that increases by exactly j each time j increases by 1, so it is largest over $j = 2, \dots, 2010$ at $j = 2010$, where it equals $2010 + \frac{2008\times 2009}{2} = 2010 + 2017036 = 2019046$. Since $2010n \ge 2010 \times 1005 = 2020050 > 2019046$, the required inequality holds at $j = 2010$, and therefore at every smaller j too. So phase 2 completes legally, leaving every box empty except $B_1 = 2010n$.

Phase 3, distribute. For $i = 2, \dots, 2010$ in turn, move 1 ball from B_1 to B_i, repeated n times for each i. Legal throughout, since B_1 decreases monotonically from $2010n$ down to $2010n - 2009n = n$ and never needs to supply more than it holds. At the end $B_i = n$ for every $i \ge 2$, and $B_1 = n$ too.

The target is reached from an arbitrary start, so $n \ge 1005$ suffices, and combined with the impossibility direction the answer is exactly $n \ge 1005$.

Layer 1: WHAT

Block 1: notice the move size at index i is pegged to the box's own label, so a box below its label is completely frozen, which immediately suggests a capacity bound: what is the largest total achievable while every box is simultaneously frozen.

Block 2: turn that into the exact threshold, $\sum_{i=1}^{2010}(i-1) = 2010 \times 1004.5$, splitting the integers n into $n \le 1004$, where a frozen starting configuration with this total exists, and $n \ge 1005$, where no configuration can freeze every box at once, so some move is always available.

Block 3: for sufficiency, notice box B_1 is the only box whose native move size is a single ball, so route everything through it: collect every eligible box into B_1 working from the largest index down, then re inject one ball at a time to flush whatever remains, then hand out exactly n balls to everyone.

Block 4: verify the flush phase never overdraws B_1 by checking the single worst case, the largest index processed first, against the very same threshold from block 2, closing the argument exactly at $n = 1005$ with no gap between the impossibility bound and the construction.

The nominal route hunts for a preserved residue or colouring invariant, or brute forces small cases; the ideal route never asks what is preserved, only what is achievable while every legality threshold simultaneously fails. The bifurcation point sits right after reading "moving exactly i balls": whether you treat this as a hint to search for an invariant, the usual reflex for "for which n" questions, or as a rigidity statement about which moves are even available, which is a plain counting question and the one that actually closes this problem.

Layer 2: HOW

Block 1. **Cue.** A move at index i moves exactly i balls, and is available only when $B_i \geq i$ already. **Why natural.** This is a threshold condition tied to each box's own label, and thresholds tied to labels are answered by asking what happens if every box sits just under its own threshold at once, which is a pure counting question, not a structural one. **Link to the next block.** That question is exactly $\sum_{i=1}^{2010}(i-1)$, the capacity of the frozen everywhere configuration.

Block 2. **Cue.** The capacity 2010×1004.5 sits exactly halfway between two multiples of 2010. **Why natural.** The total balls is always an integer multiple of 2010, namely $2010n$, so it either sits at or below 2010×1004, which realises the frozen cap or falls short of it, or it sits at 2010×1005 or above, which strictly exceeds the cap; there is no integer n landing on the cap itself, so the two cases meet with no ambiguous boundary to resolve. **Link to the next block.** This tells you $n = 1004$ is the last impossible value and $n = 1005$ is the first value where some move is always available, which is exactly the opening a construction needs.

Block 3. **Cue.** Among all 2010 boxes, only B_1's own move size is 1. **Why natural.** Any box with label $i \geq 2$ forces batches of i balls at a time, too coarse to land exactly on a target of n balls per box in general, so atomic control is only ever available at B_1. **Link to the next block.** This fixes the shape of the whole construction before a single move is written: collect everything reachable into B_1, then dispense it back out one ball at a time.

Block 4. **Cue.** The collect phase can strand up to $i-1$ balls in each box i that never became eligible on its own. **Why natural.** Those stranded balls still have to reach B_1 eventually, and the only way to move fewer than i balls out of box i is to first top it up to exactly i balls, borrowing the shortfall from B_1 before getting it back. **Link to the next block.** This makes the flush phase a sequence of loans against B_1, and the only remaining question is whether B_1 can always afford the loan, answered by the same capacity inequality from block 2, now checked at the single worst moment, the largest box, processed first, when the loan is biggest and nothing has yet been repaid.

Layer 3: WHY

The underlying truth is that this is a counting argument dressed as a token game: the impossibility direction and the sufficiency direction are both proved by reading the same inequality, $2010n$ against $\sum_{i=1}^{2010}(i-1)$, in opposite directions, which is why the threshold comes out as one clean integer rather than needing two unrelated techniques that happen to meet in the middle. The discernment mechanism to train is this: whenever a move in a combinatorial construction is legal only once a counter's value crosses a threshold tied to its own label, ask first what the maximum achievable total

is while every threshold fails at once, and compare that number to the actual total; the comparison usually is the answer, or is one integer step away from it.

The second truth worth training is about where fine control comes from: when the native move sizes in a construction are coarse and tied to position, the only place you ever get atomic, one unit control is wherever the move size happens to be 1, and any construction with an exact numerical target has no real alternative to routing through that one component. Recognising this before writing a single move is what turns the construction from a search into a fixed three step procedure.

What the examiners were nudging toward is the word "exactly" in "moving exactly i balls": it is the entire content of the problem, since a version of this question where you could move any number of balls up to i would collapse instantly to direct balancing with no threshold at all, and a candidate who treats that word as incidental phrasing rather than the whole obstruction will spend their time hunting for an invariant that does not exist. This problem is a small instance of chip firing, also called the abelian sandpile model when firing is forced once a threshold is crossed, a named area of algebraic combinatorics; recognising the family is what turns this from a from scratch construction into an application of a known shape.

Revolving concepts

Pigeonhole as a capacity bound. If nonnegative integers $x_1, \ldots, x_k$ satisfy $x_i \leq c_i$ for fixed caps c_i, then $\sum x_i \leq \sum c_i$; contrapositively, if the actual total of such a family exceeds $\sum c_i$, some x_i exceeds its own cap. This is the two directions of the whole proof, read forwards for the construction and backwards for the counterexample. See Book 1, field: Combinatorics, topic: pigeonhole as a counting bound.

Chip firing and the abelian sandpile model. A chip firing game places tokens on the vertices of a labelled structure and allows a vertex to fire, sending tokens elsewhere, once its token count reaches a threshold tied to that vertex; questions of termination and the shape of the eventual stable state form a developed area of algebraic combinatorics, of which this problem is a specific instance. See Book 1, field: Discrete structures and graphs, topic: chip firing games.

Triangular numbers and their parity relative to the count. $\sum_{i=1}^{k}(i-1) = \binom{k}{2} = k \times \frac{k-1}{2}$, an integer multiple of k exactly when k is odd, and a half integer multiple of k when k is even. This is why 2010, an even count, produces a threshold sitting exactly between two integers with no ambiguous boundary case to resolve, whereas an odd count lands the capacity exactly on a multiple of the count and needs a separate check at equality, as in mirror drill two. See Book 1, field: Combinatorics, topic: triangular numbers.

Greedy processing against a monotonic bound. When a construction must clear a sequence of requirements against a shared, partially replenished reserve, and the reserve needed for each requirement is monotonic across the processing order, it suffices to verify the single extremal requirement rather than checking every step; this recurs across scheduling and resource allocation constructions generally, not only here. See Book 1, field: Sequences and recurrences, topic: extremal case verification.

Circuit Bank

- **WHEN** a move or operation is legal only once a counter's current value crosses a threshold tied to its own label or index **DO** compute the total achievable while every threshold simultaneously fails, and compare it to the actual total **BECAUSE** if the actual total exceeds that capacity, some threshold must be exceeded somewhere at all times, which seeds both the impossibility bound below the capacity and the construction above it
- **WHEN** a construction needs to move an arbitrary, exact amount out of a structure whose native move sizes are coarse and tied to position **DO** find the one component whose native move size is the smallest possible, typically 1, and route everything through it **BECAUSE**

atomic control only ever exists at the finest native step size, so any construction with an exact numerical target has no alternative to funnelling through that component

- **WHEN** a multi phase construction repeatedly borrows from a shared reserve before repaying it **DO** check the reserve's sufficiency only at the single worst case, largest debt, earliest in the processing order, rather than summing every phase **BECAUSE** when the required bound is monotonic across the processing order, the extremal case implies every other case for free, and re deriving each phase separately wastes ink

Backtesting

1. First useful cue: the move size at index i is exactly i, not a free choice.
2. Method the expression suggested: a capacity or pigeonhole bound on the configuration where every box sits simultaneously below its own label.
3. Where you would hesitate, and the thought to reinforce: worrying whether B_1 can run out of balls midway through the flush phase; reinforce by checking only the single largest withdrawal against the same capacity inequality that gave the threshold.
4. The trap the question sets: 2010 boxes and an arbitrary n invite a hunt for a residue or colouring invariant that simply is not there, since the real obstruction is a plain counting bound.
5. What makes you recognise it faster next time: any competition problem where a move is described as "exactly" some label dependent amount should trigger the frozen configuration capacity check before anything else.
6. The reusable decision rule: when a move's legality depends on crossing a label tied threshold, compare the actual total to the capacity of the everyone frozen configuration; that comparison typically is the answer.

Mirror drills

1. With 6 boxes $B_1, \ldots, B_6$ and $6n$ balls under the identical move rule, for which n can you always reach n in every box? *Answer:* by the same bound, $\sum_{i=1}^{6}(i-1) = 15 = 6 \times 2.5$, so $n \geq 3$.
2. Redo the setup with 2011 boxes instead of 2010, so that $\sum_{i=1}^{2011}(i-1) = \frac{2010 \times 2011}{2} = 2011 \times 1005$ lands exactly on an integer multiple of the box count. *Answer:* at $n = 1005$ the frozen everywhere configuration is achievable with total exactly 2011×1005, so this boundary value is impossible too, and the threshold is $n \geq 1006$, one more than the naive halving suggests, because an odd box count removes the half integer gap that made 2010 unambiguous.
3. A row of N boxes holds $2N$ balls; the rule is "move exactly 2 balls from any box holding at least 2, into any other box," the same fixed amount regardless of the source box's label. For which starting distributions can you reach 2 balls in every box? *Answer:* exactly those where every box starts with an even number of balls: each move changes a box's count by ± 2, so every box's parity is invariant throughout, ruling out any start with an odd box; when every box starts even, repeatedly moving from a box above target to one at zero terminates at uniformity. This drill checks that removing the label dependence swaps the capacity obstruction for a parity obstruction, it does not remove every obstruction.

16.10 B4. Casoratian equal to one forces linear

Find all pairs of polynomials $p(x)$ and $q(x)$ with real coefficients for which

$$p(x)q(x+1) - p(x+1)q(x) = 1.$$

Read it out loud

Say the left side aloud before you touch a coefficient: p evaluated at x times q evaluated at $x+1$, minus the same product with the roles of p and q swapped. Two structural features are on the surface. The expression is bilinear, linear in p for fixed q and linear in q for fixed p. And it is antisymmetric: swap p and q and the whole thing changes sign. Bilinear and antisymmetric in two arguments is the definition of a 2×2 determinant, so inside thirty seconds you should have written

$$p(x)q(x+1) - p(x+1)q(x) = \det \begin{pmatrix} p(x) & p(x+1) \\ q(x) & q(x+1) \end{pmatrix} = 1.$$

That is not decoration. The instant the expression is a determinant you own a free move: adding a constant multiple of one row to the other does not change the value.

Read the columns next. They are the same two polynomials evaluated at x and at $x+1$, so the forward shift is wired into the question, not the derivative. This object has a name. It is the Casoratian of p and q, the discrete Wronskian, the exact analogue of $pq' - p'q$ with the derivative replaced by the shift. The fields embedded here are polynomial algebra, finite differences, and one line of 2×2 linear algebra. There is no analysis in this problem and you should refuse to import any.

What is really being tested is two reflexes. First, can you run a degree count on an expression whose top term cancels for free? The coefficient of x^{m+n} on the left is $p_m q_n - p_m q_n = 0$ automatically, so all of the information sits one degree down, at x^{m+n-1}, and you must compute that coefficient in general rather than in examples. Second, do you remember that the one configuration which kills a leading coefficient argument, namely $\deg p = \deg q$, is exactly the configuration a row operation dissolves? The problem is built out of that pairing.

One more free line before you start. The right side is a nonzero constant, so neither p nor q is the zero polynomial, and any common polynomial factor of p and q divides 1, hence $\gcd(p, q) = 1$. Record it. You will not need it on the winning route, but it is the fuel for the second best route and it costs nothing.

Approaches, ranked worst to best

1. **Guess the linear family and stop.** *Cue:* the equation looks as though only small solutions can survive. *Circuit:* pattern matching with no exclusion. Substituting $p = ax + b$, $q = cx + d$ takes one line and returns $bc - ad = 1$. That is the correct family and it is worth close to nothing, because the command term is *find all*. A family with no proof that nothing else works is a conjecture, and the Putnam pays for classification, not for examples.

2. **Full coefficient comparison.** *Cue:* a polynomial identity, so expand and equate coefficients. *Circuit:* binomial expansion of $p(x+1)$ and $q(x+1)$, then a descending induction on the minors $p_i q_j - p_j q_i$. This does close. The coefficient of x^k on the left is a double sum of binomial differences multiplied by the 2×2 minors $p_i q_j - p_j q_i$, and a descending induction shows every minor vanishes, which makes p and q proportional and the left side zero. Cost: a page of indices for a problem whose answer is two lines. It ranks low not because it is wrong but because ink is time and every index is an error site.

3. **Telescoping the ratio.** *Cue:* divide by $q(x)q(x+1)$ and a difference of a rational function appears. *Circuit:* convert to $\frac{p(x)}{q(x)} - \frac{p(x+1)}{q(x+1)} = \frac{1}{q(x)q(x+1)}$ and telescope over the integers. Informative but overweight. If $\deg q \geq 1$ the right side is summable, so $p(N)/q(N)$ converges as N runs through the integers, which forces $\deg p \leq \deg q$, and antisymmetry gives the reverse. You get $\deg p = \deg q$ cheaply, and then you still have to break the tie. You also imported convergence into a problem that is pure algebra, which is the wrong kind of borrowing.

4. **Shift and subtract.** *Cue:* the identity holds for every x, so it also holds at $x-1$. *Circuit:* eliminate the mixed terms by subtracting two instances, then run coprimality against a degree count. Subtracting the instance at $x-1$ from the instance at x gives $p(x)\big(q(x+1)+q(x-1)\big) = q(x)\big(p(x+1)+p(x-1)\big)$. Since $\gcd(p,q) = 1$, p divides $p(x+1)+p(x-1)$; that quotient has degree 0 and leading coefficient 2, so $p(x+1)+p(x-1) = 2p(x)$, the second difference of p vanishes, and p is linear. A first class route. It costs one extra idea (coprimality) and one extra object (the second difference) compared with the winner.

5. **Casoratian row reduction.** *Cue:* the left side is a determinant, and a determinant only ever needs a degree count plus a row operation. *Circuit:* compute the first surviving coefficient in general, then use invariance under a row operation to force the degrees apart. One lemma, one line of algebra, no induction, no divisibility, and it explains the answer rather than merely confirming it. This is the route to install.

The ideal solution

Write $D(p,q)(x) = p(x)q(x+1) - p(x+1)q(x)$ for the Casoratian of the pair.

Lemma 16.10.1. *Let p and q be nonzero polynomials of degrees m and n with leading coefficients p_m and q_n, and suppose $m+n \geq 1$. Then* $\deg D(p,q) \leq m+n-1$, *and the coefficient of x^{m+n-1} in $D(p,q)$ equals $(n-m)p_m q_n$.*

Proof. Write $p(x) = p_m x^m + p_{m-1}x^{m-1} + \cdots$ and $q(x) = q_n x^n + q_{n-1}x^{n-1} + \cdots$. Expanding the shift, $q(x+1) = q_n x^n + (nq_n + q_{n-1})x^{n-1} + \cdots$ and $p(x+1) = p_m x^m + (mp_m + p_{m-1})x^{m-1} + \cdots$, since only the top two terms of a binomial expansion can reach degree $n-1$ or $m-1$. Hence the coefficient of x^{m+n} in $D(p,q)$ is $p_m q_n - p_m q_n = 0$, and the coefficient of x^{m+n-1} is

$$\big(p_m(nq_n + q_{n-1}) + p_{m-1}q_n\big) - \big((mp_m + p_{m-1})q_n + p_m q_{n-1}\big) = (n-m)p_m q_n. \qquad \square$$

Lemma 16.10.2. *For every real constant t, $D(p,\, q - tp) = D(p,q)$.*

Proof. Expanding, $D(p, q-tp) = D(p,q) - t\big(p(x)p(x+1) - p(x+1)p(x)\big) = D(p,q)$. $\square$

Now suppose $D(p,q) = 1$ identically. Neither p nor q is the zero polynomial, since D would then be identically zero. Put $m = \deg p$ and $n = \deg q$.

Suppose first $m = n$. Set $t = q_n/p_m$ and $q_1 = q - tp$. By Lemma 2, $D(p,q_1) = 1$, so q_1 is not the zero polynomial, and by construction its leading term has been cancelled, so $n_1 := \deg q_1 < m$. Apply Lemma 1 to the pair (p,q_1), whose degrees are now different. Since $n_1 \neq m$ and $p_m \neq 0$ and the leading coefficient of q_1 is nonzero, the coefficient of x^{m+n_1-1} in $D(p,q_1)$ is nonzero, provided $m+n_1 \geq 1$. As $D(p,q_1) = 1$ has degree 0, this forces $m+n_1-1 = 0$, that is $m+n_1 = 1$. Note $m+n_1 = 0$ is impossible: it would mean $m = n_1 = 0$, and two nonzero constants have $D = 0$. From $m+n_1 = 1$ with $n_1 < m$ we get $m = 1$ and $n_1 = 0$, so $m = n = 1$.

Suppose instead $m \neq n$. Then $m+n \geq 1$, and Lemma 1 gives $\deg D(p,q) = m+n-1$ exactly, so $m+n = 1$ and $\{m,n\} = \{0,1\}$.

In both cases $\deg p \leq 1$ and $\deg q \leq 1$. Write $p(x) = ax+b$ and $q(x) = cx+d$ with a,b,c,d real, allowing a or c to vanish. Then

$$D(p,q) = (ax+b)\big(c(x+1)+d\big) - \big(a(x+1)+b\big)(cx+d)$$

$$= \big(acx^2 + (ac + ad + bc)x + bc + bd\big) - \big(acx^2 + (ac + ad + bc)x + ad + bd\big)$$
$$= bc - ad,$$

a constant. So the identity holds precisely when $bc - ad = 1$, and every such pair does satisfy it. The complete answer is

$$\boxed{p(x) = ax + b, \quad q(x) = cx + d, \quad bc - ad = 1}$$

with $a, b, c, d \in \mathbb{R}$. Equivalently, $\det \begin{pmatrix} p(0) & p(1) \\ q(0) & q(1) \end{pmatrix} = 1$: for polynomials of degree at most 1, the single instance of the identity at $x = 0$ already implies all the others.

Layer 1: WHAT

Block 1. Name the left side as the determinant of the 2×2 matrix with rows $\big(p(x), p(x+1)\big)$ and $\big(q(x), q(x+1)\big)$.

Block 2. Degree count in general: the x^{m+n} coefficient cancels for free, and the coefficient of x^{m+n-1} is $(n - m)p_m q_n$.

Block 3. Consequence: if $\deg p \neq \deg q$ then the Casoratian has degree exactly $m + n - 1$, so being equal to 1 forces $m + n = 1$.

Block 4. The blocked case $\deg p = \deg q$ is not a case, it is a degeneracy: a row operation $q \mapsto q - tp$ leaves the determinant untouched and strictly drops the degree, sending you back into Block 3.

Block 5. Both degrees are at most 1; substitute and read off $bc - ad = 1$; verify the family.

The nominal route splits from this at Block 1. The zombie move is to expand $p(x)q(x+1)$ and $p(x+1)q(x)$ in coordinates, collect the coefficient of x^k, and start comparing. That route is not wrong, it is expensive, and its expense is entirely caused by refusing to name what the expression is. The bifurcation point is the very first line you write: either you write out $\sum_i \sum_j$ with binomial coefficients, or you write det. Everything downstream is determined by that choice. Once the word determinant is on the page, you automatically ask the determinant question, which is not what are the coefficients but which operations preserve the value, and the answer to that question is the whole solution.

Layer 2: HOW

Block 1. Cue. Bilinear in p and q, and antisymmetric under swapping them. In A-level Further Maths you met exactly one bilinear antisymmetric gadget on pairs, the 2×2 determinant, and it is defined by those two properties together with its value on the identity. **Why natural.** Naming an expression by its structural type is not cosmetic. A determinant carries a group of symmetries with it (row operations, transposition, multilinearity), and symmetries are free information. An unnamed polynomial expression carries nothing. **Link to the next block.** Once it is a determinant you want to normalise the rows, but before normalising you must know what quantity the rows control. That quantity is degree, so you count degrees next.

Block 2. Cue. The identity is pinned to a constant, and constants are the degree 0 objects. Any argument that controls the degree of the left side finishes the problem, so compute the degree. **Why natural.** The natural first attempt, look at the top coefficient, dies immediately because the top coefficient is $p_m q_n - p_m q_n = 0$. That death is the signal, not a setback: whenever an antisymmetric combination annihilates its own leading term, the first surviving coefficient is where the content lives, so you compute it in full generality. This is a general habit: in any expression built to cancel, differentiate your attention one order down. **Link to the next block.** The formula $(n - m)p_m q_n$ has a visible zero set. It vanishes exactly when $m = n$. So the degree count settles everything except one case, and you now know precisely which case to attack.

Block 3. Cue. A nonzero coefficient at x^{m+n-1} pins the degree of the left side to $m+n-1$, and the right side has degree 0. **Why natural.** This is the cheapest kind of argument on the Putnam: an equation between polynomials is an equation between degrees first. You get $m+n=1$, so one of the two polynomials is a nonzero constant and the other is genuinely linear. **Link to the next block.** The only gap left is $m=n$. You do not attack it head on; you remember from Block 1 that you are holding a determinant.

Block 4. Cue. Equal leading degrees plus a determinant is the classic setup for row reduction. This is the first step of Gaussian elimination and the first step of the Euclidean algorithm, and both are triggered by the same cue: two objects with the same leading behaviour. **Why natural.** The row operation $q \mapsto q - tp$ with t constant leaves D unchanged, because the correction term is $t\,p(x)p(x+1)$ in both products and cancels. It is essential that t is constant: with a polynomial multiplier $r(x)$ the correction is $p(x)p(x+1)\big(r(x+1)-r(x)\big)$, which survives unless r is constant. That is why this is one row operation, not a full Euclidean descent. **Link to the next block.** Killing the leading term drops $\deg q$ strictly, which places you back in Block 3, and Block 3 then forces $\deg p = 1$.

Block 5. Cue. Both degrees are at most 1, so the whole problem is now four unknown real numbers and one identity. **Why natural.** At this point substitution is the correct move, not a fallback. Substituting into a bounded, tiny parametrisation is honest work; substituting into an unbounded family of coefficients was the zombie route. The verification also matters for marks: you must show the family really does satisfy the identity, since *find all* demands both inclusions. **Link to the next block.** Nothing follows; the classification is complete. Finish by noticing the answer is the determinant of the values at 0 and 1, which is the cleanest form to remember.

Layer 3: WHY

The mathematical truth underneath is that D is an alternating bilinear form on the vector space of polynomials, twisted by the shift. Alternating forms have a fixed behaviour you can rely on. They vanish on proportional pairs, they are unchanged when you add a multiple of one argument to the other, and they lose the top order term of any naive size estimate. All three of those facts were used, and none of them required knowing anything about p and q beyond their degrees.

The second truth is the degree formula $(n-m)p_m q_n$, and it deserves a moment. It says the Casoratian of two polynomials of the same degree is smaller than you would expect, and it is exactly this drop that permits nonzero constant Casoratians to exist at all. If the coefficient had been, say, $p_m q_n$ with no factor $(n-m)$, the answer to the problem would have been that no pairs exist. So the entire solution set of the problem is created by one degeneracy, and the row operation is the tool that measures how deep the degeneracy goes. The same phenomenon runs the continuous analogue: for the Wronskian, the coefficient of x^{m+n-1} in $pq' - p'q$ is again $(n-m)p_m q_n$, and the classification of $pq' - p'q = 1$ is again $p = ax+b$, $q = cx+d$ with $bc - ad = 1$. If you learn one of the two you have learned both.

The discernment mechanism, the thing that decides the route before you compute anything, is this: classify the expression by its symmetry type, then use the symmetry group as your move list. Antisymmetric in two objects means determinant means row operations. Symmetric in two objects means elementary symmetric polynomials means Newton or Vieta. Invariant under a shift means finite differences or periodicity. You are not searching for a clever idea; you are looking up which move list the structure hands you. To train it, take any competition expression and before solving it, write one line naming its symmetry type and one line naming the associated move list. That drill is fast, it is boring, and it is the difference between seeing the determinant in ten seconds and never seeing it.

The examiners nudged in three ways and you should be able to name all three. They wrote $q(x+1)$ rather than $q(y)$, which plants the shift. They put a nonzero constant on the right rather than 0, which plants coprimality and rules out the proportional pairs. And they asked for pairs of polynomials, which plants a two row object. A question asking for a single polynomial would not have had a determinant in it.

Revolving concepts

Casoratian (discrete Wronskian). For functions f, g define $C(f,g)(x) = f(x)g(x+1) - f(x+1)g(x)$. It is bilinear and alternating, it satisfies $C(f, g+tf) = C(f,g)$ for constant t, and for polynomials it vanishes identically if and only if p and q are proportional. That last claim is worth proving once: if $C(p,q) = 0$ then $p(x)q(0) - p(0)q(x)$ vanishes at every integer, since p/q is unchanged by the shift, so the polynomial $p(x)q(0) - p(0)q(x)$ has infinitely many roots and is identically zero. See Book 1, field: Polynomials, topic: Casoratian and Wronskian.

Second order linear recurrences and the Casoratian identity. If two sequences f_k and g_k both satisfy $u_{k+1} = \alpha_k u_k + \beta_k u_{k-1}$, then their Casoratian $C_k = f_k g_{k+1} - f_{k+1} g_k$ satisfies the first order recursion $C_{k+1} = -\beta_k C_k$, hence

$$C_k = C_0 \prod_{j=0}^{k-1} (-\beta_j).$$

This is the generalised recursion behind the problem: a constant Casoratian is exactly the case $\beta_j = -1$ for all j, and an alternating Casoratian is the case $\beta_j = 1$, which is the Fibonacci recursion and gives the Cassini identity $F_{k+1}F_{k-1} - F_k^2 = (-1)^k$. Whenever you meet a bilinear alternating combination of two solutions of the same linear recursion, compute how it transforms in one step; it always satisfies a first order recursion, which you can always solve as a product. See Book 1, field: Sequences and recurrences, topic: second order linear recurrences.

Finite differences. $\Delta f(x) = f(x+1) - f(x)$ lowers the degree of a polynomial by exactly one, so $\Delta^{k+1} p = 0$ if and only if $\deg p \le k$. The relation $p(x+1) + p(x-1) = 2p(x)$ that appears on the shift and subtract route is exactly $\Delta^2 p = 0$ evaluated one step back, hence p is linear. The general expansion is Newton's forward difference formula $p(x) = \sum_{k \ge 0} \Delta^k p(0) \binom{x}{k}$, the discrete Taylor series, which terminates for polynomials. See Book 1, field: Polynomials, topic: finite differences.

Coprimality and Bezout in $\mathbb{R}[x]$. $\mathbb{R}[x]$ is a Euclidean domain, so gcd exists and $\gcd(p,q) = 1$ if and only if $up + vq = 1$ for some polynomials u, v. Our identity is literally a Bezout relation with $u(x) = q(x+1)$ and $v(x) = -p(x+1)$, which is why the problem feels like a unimodularity condition. See Book 1, field: Polynomials, topic: Euclidean algorithm and Bezout.

Unimodular pairs. For integers, $bc - ad = 1$ is membership of $SL_2(\mathbb{Z})$ and is the engine behind continued fractions, Farey neighbours and the Stern and Brocot tree. The polynomial statement here is the same condition read one level up: the pair (p,q) is a unimodular pair for the shift. See Book 1, field: Matrices and determinants, topic: unimodular matrices.

Circuit Bank

- **WHEN** an expression is bilinear and antisymmetric in two objects **DO** write it as a 2×2 determinant and immediately list the operations that preserve it **BECAUSE** a determinant hands you a free normalisation, and normalising is always cheaper than computing
- **WHEN** the leading coefficient of an expression cancels identically **DO** compute the next coefficient in full generality, with letters, not examples **BECAUSE** a structure built to cancel stores all of its information in the first surviving coefficient
- **WHEN** a degree count is blocked only by the equal degree case **DO** use a row operation or a subtraction of leading terms to break the tie **BECAUSE** equal leading behaviour is a degeneracy to dissolve, not a case to fight

- **WHEN** a polynomial identity equals a nonzero constant **DO** write down at once that neither polynomial is zero and that they are coprime **BECAUSE** any common factor divides the constant, and this line is free
- **WHEN** a relation says a second difference or second derivative vanishes **DO** conclude the object is linear and stop **BECAUSE** $\Delta^2 p = 0$ and $p'' = 0$ each characterise degree at most one

Backtesting

1. First useful cue: the left side changes sign when p and q are swapped, and it is linear in each, so it is a 2×2 determinant with columns indexed by x and $x+1$.
2. Method the expression suggested: control the degree of the determinant, then remove the one obstruction to that control with a row operation.
3. Where you would hesitate, and the thought to reinforce: at the case $\deg p = \deg q$, where the leading coefficient argument dies. Reinforce this: the death of a leading coefficient argument at equal degrees is the standard prompt for elimination, in Gaussian elimination, in the Euclidean algorithm, and here.
4. The trap the question sets: *find all* invites you to produce the linear family from a one line substitution and feel finished. Producing the family is worth almost nothing; excluding everything else is the whole problem.
5. What makes you recognise it faster next time: the phrase Casoratian. Say it out loud when you see $f(x)g(x+1) - f(x+1)g(x)$, and the whole theory of alternating forms and second order recurrences arrives with it.
6. The reusable decision rule: classify by symmetry type first, then use the symmetry group as your move list, and only then compute.

Mirror drills

1. Find all pairs of real polynomials with $p(x)q'(x) - p'(x)q(x) = 1$. *Answer:* the same family, $p = ax + b$, $q = cx + d$ with $bc - ad = 1$. The degree lemma is identical: the coefficient of x^{m+n-1} in $pq' - p'q$ is $(n-m)p_m q_n$, and $q \mapsto q - tp$ again preserves the form.
2. Prove that $p(x)q(x+1) - p(x+1)q(x) = 0$ identically forces p and q to be proportional. *Answer:* the relation says p/q is invariant under the shift, so $p(x)q(0) - p(0)q(x)$ vanishes at every integer, hence identically.
3. Let f_k and g_k both satisfy $u_{k+1} = u_k + u_{k-1}$. Show $f_k g_{k+1} - f_{k+1} g_k$ alternates in sign with constant modulus, and deduce Cassini's identity $F_{k+1}F_{k-1} - F_k^2 = (-1)^k$. *Answer:* substituting the recursion gives $C_{k+1} = -C_k$; take $f_k = F_{k-1}$, $g_k = F_k$ and evaluate at $k = 1$.

16.11 B5. Iterative differential equation has no increasing solution

Is there a strictly increasing function $f : \mathbb{R} \to \mathbb{R}$ such that $f'(x) = f(f(x))$ for all x?

Read it out loud

This is a differential equation in which the right side is not a function of x and $f(x)$ in the usual sense. The unknown appears at two different arguments at once, at x on the left and at $f(x)$ on the right. That is an iterative differential equation, and the single sentence that describes it is: the growth rate of f at a point equals the value of f at the point f has already reached. It is self referential, and self reference is the whole content. Nothing in the standard toolkit of separable, linear, integrating factor or exact equations touches it, so do not reach for any of them.

The question is a yes or no existence question, so the first two minutes belong to scouting, not to proving. Push the three standard shapes through the relation. Linear: $f = ax + b$ gives $a = a^2x + ab + b$ for all x, so $a^2 = 0$, so $a = 0$ and then $b = 0$; the only linear solution is the zero function, which is not strictly increasing. Power law on a ray: $f = cx^\beta$ gives $c\beta x^{\beta-1} = c^{1+\beta}x^{\beta^2}$, which needs $\beta^2 = \beta - 1$; the discriminant is $1 - 4 = -3$, so there is no real exponent. Exponential: $f = ce^{kx}$ gives $cke^{kx} = ce^{kce^{kx}}$, a double exponential against a single one, impossible. Three natural shapes, three failures, and the middle failure is structural: the exponent equation $\beta^2 = \beta - 1$ is the golden ratio equation with the sign flipped and it has no real root. That is your evidence. Commit to proving the answer is no.

The fields embedded are monotonicity, differentiability and convexity on the real line, plus comparison for a first order differential inequality. What is really tested is bootstrapping: taking a relation you cannot solve and squeezing consequences out of it until one of them is impossible. Within thirty seconds you should have written three free facts. f strictly increasing gives $f' \geq 0$. The right side $f \circ f$ is a composition of increasing functions, hence increasing, so f' is increasing and f is convex. And f' increasing with $f' \geq 0$ gives $f'(x) > f'(x-1) \geq 0$, so $f' > 0$ everywhere. You are already holding a convex, strictly increasing function whose derivative is positive, and you have used nothing but the words in the question.

Approaches, ranked worst to best

1. **Keep hunting for a solution.** *Cue:* the question says is there, so try to build one. *Circuit:* ansatz search over the standard function library. Useful for sixty seconds as scouting and worthless after that. A failed search proves nothing, and the marks are all in the exclusion. If you are still substituting shapes at minute ten you have misread the command term.

2. **Monotonicity bookkeeping with no engine.** *Cue:* f is increasing, so collect sign facts. *Circuit:* convexity, fixed point chasing, tangent line bounds. You will correctly get $f' > 0$, f convex, f above its tangents. Then you stall, because none of those facts is quantitative enough to contradict anything: convex increasing functions exist in abundance. The lesson is that sign information alone never closes an existence problem; you need a rate.

3. **The bound $f(x) \leq x + 1$ by case analysis.** *Cue:* f convex means f lies above a tangent line, and $f' = f(f(x))$ ties that line back to f. *Circuit:* tangent line bound, then a two case argument on whether $f(x_0) < x_0$. This closes and it is the official first route: from $f(x_0) = b$ and $f'(x_0) = a > 0$ you get $f(x) \geq a(x - x_0) + b$, and feeding that back at the point b forces $f(x_0) \leq x_0 + 1$, hence $f' = f \circ f \leq 1$, which then contradicts the linear lower bound. Correct, and it costs a case split that is easy to garble in the arena.

4. **Invert the function.** *Cue:* f is strictly increasing, so it has an inverse. *Circuit:* differentiate $x = g(f(x))$ to get $g'(y) = 1/f(y)$, then compare integrals. Attractive, because $g' = 1/f$ converts the whole problem into the convergence of $\int^\infty dy/f(y)$: since g is unbounded above

on the range, that integral must diverge, so f cannot grow too fast. But to finish you still need a lower bound on the growth of f, which is the work of the winning route. It is the winner plus one extra layer.

5. **Differentiate, bootstrap the growth, then blow up.** *Cue:* the relation defines f', so it can be differentiated for free, and the chain rule turns the composition into a product of two positive increasing factors. *Circuit:* regularity bootstrap, then integrate a positive second derivative twice, then compare against $y' = \alpha y^2$. One direction, no case splits, and every step is a move you already own. This is the route to install.

The ideal solution

The answer is no: no such function exists. Suppose, for contradiction, that $f : \mathbb{R} \to \mathbb{R}$ is strictly increasing and satisfies $f'(x) = f(f(x))$ for all x.

Step 1: $f' > 0$ everywhere. Since f is strictly increasing and differentiable, $f' \geq 0$. Since f is strictly increasing, so is the composition $f \circ f$, and $f \circ f = f'$, so f' is strictly increasing. Hence for every x we have $f'(x) > f'(x-1) \geq 0$.

Step 2: f is twice differentiable, with f'' positive, continuous and increasing. The relation $f' = f \circ f$ exhibits f' as a composition of differentiable functions, so f' is differentiable and the chain rule gives

$$f''(x) = f'\big(f(x)\big)\, f'(x).$$

Both factors are positive by Step 1, so $f'' > 0$. Both factors are also increasing: f' is increasing by Step 1, and $f' \circ f$ is increasing as a composition of increasing functions. A product of two positive increasing functions is increasing, so f'' is increasing. It is continuous because $f' = f \circ f$ is continuous and f'' is a product of continuous functions.

Step 3: quadratic growth. Put $\alpha = f''(0)/4 > 0$, so $f''(x) \geq 4\alpha$ for all $x \geq 0$. Integrating the continuous function f'' from 0 to x gives, for $x \geq 0$,

$$f'(x) = f'(0) + \int_0^x f''(t)\, dt \geq f'(0) + 4\alpha x,$$

and integrating once more,

$$f(x) \geq f(0) + f'(0)x + 2\alpha x^2 \qquad (x \geq 0).$$

The difference $\big(f(0) + f'(0)x + 2\alpha x^2\big) - \alpha x^2 = \alpha x^2 + f'(0)x + f(0)$ is a quadratic with positive leading coefficient, so it is nonnegative for all sufficiently large x. Fix $M \geq 0$ with

$$f(x) \geq \alpha x^2 \qquad \text{for all } x \geq M.$$

Step 4: feed the growth back into the relation. Choose $T \geq M$ with $\alpha T^2 \geq M$; this is possible since $\alpha T^2 \to \infty$. For $x \geq T$ we have $f(x) \geq \alpha x^2 \geq \alpha T^2 \geq M$, so the bound of Step 3 may be applied at the point $f(x)$ itself:

$$f'(x) = f\big(f(x)\big) \geq \alpha f(x)^2 \qquad (x \geq T).$$

Also $f(x) \geq \alpha T^2 > 0$ for $x \geq T$, so $1/f$ is defined and differentiable there.

Step 5: the differential inequality has no global solution. For $x \geq T$,

$$\frac{d}{dx}\left(\frac{-1}{f(x)}\right) = \frac{f'(x)}{f(x)^2} \geq \alpha.$$

Integrating from T to x,

$$\frac{1}{f(T)} - \frac{1}{f(x)} \geq \alpha(x - T) \qquad (x \geq T).$$

Since $1/f(x) > 0$, the left side is strictly less than $1/f(T)$, so

$$\alpha(x - T) < \frac{1}{f(T)} \qquad \text{for every } x \geq T,$$

which is false as soon as $x > T + \dfrac{1}{\alpha f(T)}$. This contradiction shows no such f exists. The answer to the question is $\boxed{\text{no}}$.

Layer 1: WHAT

Block 1. Scout with the three standard shapes, notice the exponent condition $\beta^2 = \beta - 1$ has no real root, and commit to proving nonexistence.

Block 2. Harvest the free monotonicity: $f' \geq 0$, then $f' = f \circ f$ is increasing, then $f' > 0$.

Block 3. Differentiate the relation. The chain rule gives $f'' = f'(f)\, f' > 0$ and increasing, so the equation upgrades its own regularity and its own sign information.

Block 4. Integrate the bound $f'' \geq 4\alpha$ twice from a fixed base point to obtain $f(x) \geq \alpha x^2$ for large x.

Block 5. Substitute that growth back into the original relation at the point $f(x)$, obtaining $f' \geq \alpha f^2$ on a ray.

Block 6. Kill $f' \geq \alpha f^2$ by integrating $(-1/f)' \geq \alpha$: the left side stays negative while the right side grows without bound.

The nominal route is to treat this as a differential equation to be solved: separate variables (impossible, the right side is not a function of $f(x)$ alone in the required sense), try a power series about 0 (the coefficients of $f \circ f$ tangle instantly), or hunt for a closed form. The bifurcation point is the line where you stop asking what f is and start asking what f must satisfy. Concretely it is the decision to differentiate the relation in Block 3. Differentiating an equation you cannot solve feels wrong to a beginner, because differentiation usually makes things harder. Here it makes things easier, because the right side is a composition and the chain rule turns a composition into a product, and products of positive increasing things are positive and increasing. You are trading solvability for sign information, and sign information is what kills existence questions.

Layer 2: HOW

Block 1. **Cue.** The command term is is there, not find. That is a yes or no question, and on a yes or no question the first job is to decide which way to commit, cheaply. **Why natural.** Substituting the three standard shapes is the cheapest possible experiment, and their failure modes are informative rather than blank. The power law forces $\beta^2 = \beta - 1$, an honest quadratic with negative discriminant, so the failure is not an accident of the family; it is a real obstruction living in the exponents. **Link to the next block.** Having committed to no, you need a contradiction, and contradictions in analysis come from quantities that are forced to grow while being forced to stay bounded. So you start collecting growth information.

Block 2. **Cue.** Strictly increasing is a hypothesis, and every hypothesis is free information you must bank before working. **Why natural.** $f' \geq 0$ is immediate. The move worth naming is the next one: since f is increasing, $f \circ f$ is increasing, and the relation says $f \circ f$ *is* f'. So the equation itself upgrades monotonicity of f into monotonicity of f', which is convexity. Then $f'(x) > f'(x - 1) \geq 0$ upgrades a weak inequality into a strict one at no cost, and strictness matters because you will divide by f later. **Link to the next block.** Convexity is qualitative. To get a rate you need a second derivative, and the relation is begging to be differentiated because its right side is a composition of differentiable functions.

Block 3. **Cue.** The relation defines f' explicitly as $f \circ f$, and f is differentiable, so f' is differentiable. Any equation that expresses a derivative in terms of differentiable data can be differentiated again, and doing so is free. **Why natural.** The chain rule converts the composition into the product $f'(f(x))f'(x)$. Sign information multiplies well: positive times positive is positive, increasing times increasing is increasing when both are positive. So one differentiation turns a qualitative statement about f into a quantitative floor on f'', namely $f''(x) \geq f''(0) > 0$ for $x \geq 0$. This is the general phenomenon of regularity bootstrap: an equation of the form $f' = \Phi(f)$ with Φ smooth forces f to be smooth, one derivative at a time. **Link to the next block.** A positive floor on f'' is exactly what integrates into a growth rate, so integrate twice.

Block 4. **Cue.** A constant lower bound on a continuous second derivative on a ray. There is only one move: integrate, twice, from a fixed base point. **Why natural.** Integration converts pointwise sign information into global size information, which is the currency of contradiction arguments. Note the choice of base point 0: any fixed point works, and choosing one immediately avoids all the case analysis that the alternative routes carry. Note also that α is defined as $f''(0)/4$ and not as some abstract constant; you want the quadratic coefficient to be 2α so that after absorbing the linear and constant terms you are left with the clean statement $f(x) \geq \alpha x^2$. **Link to the next block.** You now have a growth rate for f. The relation $f' = f \circ f$ evaluates f at $f(x)$, and $f(x)$ is large when x is large, so the growth rate can be applied a second time, at the inner point.

Block 5. **Cue.** The right side of the relation is f evaluated at $f(x)$, and you have just proved a lower bound for f valid on a ray. Any bound valid on a ray can be applied at any point that lies in the ray, and $f(x)$ does lie there once x is large. **Why natural.** This is the self reference finally working for you rather than against you. The composition, which made the equation unsolvable, is exactly what lets you apply the growth bound twice and square it. The single technical care needed is to choose T so that $f(x) \geq M$, which is why T is chosen with $\alpha T^2 \geq M$, not merely $T \geq M$. **Link to the next block.** You now hold $f' \geq \alpha f^2$ on a ray with $f > 0$, which is a Riccati type differential inequality, and those are lethal.

Block 6. **Cue.** The inequality $y' \geq \alpha y^2$ with $y > 0$. This is the standard blow up template and it has a standard one line kill. **Why natural.** Divide by y^2 and recognise the left side as the derivative of $-1/y$. Integrating a lower bound on a derivative gives a lower bound on the increment, and the increment of $-1/y$ is bounded above by $1/y(T)$ because $-1/y$ is negative throughout. A bounded quantity forced to increase at a constant rate forever is the contradiction. Notice what is not needed: no ODE existence theory, no comparison theorem quoted from a book, no explicit solution. The substitution $y \mapsto -1/y$ is the entire content of the finite time blow up of $y' = \alpha y^2$. **Link to the next block.** Contradiction reached, so the supposed f does not exist and the answer is no.

Layer 3: WHY

The mathematical truth beneath this problem is that superlinear feedback cannot live on the whole line. Write the relation as: the rate of increase of f equals f evaluated further along its own graph. If f ever grows faster than linearly, then $f(x)$ is far to the right of x, so $f(f(x))$ is much larger than $f(x)$, so f' is much larger than f, so f grows even faster. The loop is unstable, and the precise form of the instability is $y' \geq \alpha y^2$, whose solutions satisfy $\frac{1}{y(T)} - \frac{1}{y(x)} \geq \alpha(x - T)$ and therefore cease to exist beyond $x = T + \frac{1}{\alpha y(T)}$. The general principle is Osgood's criterion: for $y' \geq g(y)$ with g positive, the solution blows up in finite time whenever $\int^{\infty} \frac{dy}{g(y)}$ converges. With $g(y) = \alpha y^2$ that integral converges, so the domain $\mathbb{R}$ demanded by the problem is impossible. Everything else in the solution exists only to prove that the feedback really is superlinear.

The second truth is the bootstrap. An equation of the form $f' = \Phi(f)$ where Φ is built from f itself is not an obstacle to regularity, it is a machine for producing it: f differentiable gives f' differentiable gives f twice differentiable, and by induction f is infinitely differentiable. Each differentiation also multiplies sign information, because the chain rule produces products, and products of positive increasing factors are positive and increasing. So the two things this equation is generous with are smoothness and signs, and the ideal solution spends both.

The discernment mechanism is a two question filter you should run on every existence question in analysis. Question one: can I solve it? If the unknown appears at more than one argument, or inside itself, the answer is essentially always no, and you must stop trying within a minute. Question two: what does the equation force in the limit? Existence on all of $\mathbb{R}$ is a strong hypothesis and it is usually the hypothesis that breaks, because equations that are locally harmless are globally lethal. To train the filter, take any functional or differential equation and immediately ask what it says as $x \to +\infty$ and as $x \to -\infty$, before doing anything else. Here the whole difficulty is at $+\infty$, and the diagnostic that told you so is that the right side is f evaluated at a point that runs off to the right.

The examiners nudged in three ways. They asked is there rather than find, which signals a decision problem. They wrote strictly increasing rather than monotone or continuous, which hands you exactly the sign facts needed to start the bootstrap; a weaker hypothesis would have made the problem harder, so the hypothesis is a gift and you must spend it. And they demanded the domain $\mathbb{R}$, all of it, which is the clue that the contradiction is global rather than local. Every one of those three words was placed to point at the route above.

Revolving concepts

Iterative differential equations. An equation in which the unknown is evaluated at an argument depending on the unknown, such as $f'(x) = f(f(x))$, is called an iterative differential equation. It is not an ordinary differential equation: the standard existence and uniqueness theorems do not apply, because the right side is not a function of $(x, f(x))$ in the required sense. The correct handling is the one used here: extract monotonicity, regularity and growth, and argue globally. See Book 1, field: Functional equations, topic: iterative differential equations.

Regularity bootstrap. If f is differentiable and $f' = f \circ f$, then by induction f is of class C^k for every k: assuming $f \in C^k$, the composition $f \circ f$ is C^k, so $f' \in C^k$, so $f \in C^{k+1}$. This is the general recursion in the problem, and the general form is: for $f' = \Phi(f)$ with Φ of class C^∞ on the relevant range, $f \in C^k \Rightarrow f \in C^{k+1}$, hence $f \in C^\infty$. The same induction is how one proves solutions of analytic differential equations are analytic. See Book 1, field: Real analysis, topic: bootstrapping regularity.

Monotonicity and convexity from composition. If u and v are increasing then $u \circ v$ is increasing; if u, v are positive and increasing then uv is increasing. A differentiable f with increasing f' is convex, and a convex function lies above each of its tangent lines, $f(x) \geq f(x_0) + f'(x_0)(x - x_0)$. This tangent line bound is the engine of the alternative route ranked third, and it is worth carrying because it converts a single derivative value into a global linear floor. See Book 1, field: Real analysis, topic: convexity.

Finite time blow up and Osgood's criterion. If y is positive and differentiable on $[T, \infty)$ with $y' \geq \alpha y^2$, then $(-1/y)' \geq \alpha$, so $\frac{1}{y(T)} - \frac{1}{y(x)} \geq \alpha(x - T)$ and hence $x - T < \frac{1}{\alpha y(T)}$, which bounds the domain. More generally, $y' \geq g(y)$ with $g > 0$ and $\int^\infty dy/g(y) < \infty$ forces blow up in finite time, by the same separation of variables. The recursion to remember is the family $y' = y^{1+\varepsilon}$: every $\varepsilon > 0$ blows up, and $\varepsilon = 0$ (linear growth, the exponential) does not. Superlinear is the threshold. See Book 1, field: Real analysis, topic: differential inequalities and blow up.

Inverse function differentiation. If f is strictly increasing and differentiable with $f' > 0$, its inverse g satisfies $g'(y) = 1/f'(g(y))$; combined with $f' = f \circ f$ this gives the clean relation $g'(y) = 1/f(y)$ on the range of f. Since g is unbounded above, $\int^\infty dy/f(y)$ must diverge, which is

a second and independent way to see that f cannot grow quadratically. See Book 1, field: Calculus and integrals, topic: inverse functions.

Circuit Bank

- **WHEN** a relation gives f' explicitly in terms of f composed with itself **DO** differentiate the relation immediately **BECAUSE** the chain rule turns the composition into a product, and products of positive increasing factors keep every sign you already own
- **WHEN** a question asks is there rather than find **DO** spend sixty seconds on linear, power and exponential ansatzes, then commit **BECAUSE** the failure mode of the power ansatz is usually an honest algebraic obstruction that tells you which way to argue
- **WHEN** you hold a constant positive lower bound on a second derivative **DO** integrate twice from a fixed base point **BECAUSE** that is the only mechanism that converts a pointwise sign into a global growth rate, and choosing a base point removes all case analysis
- **WHEN** you reach $y' \geq \alpha y^2$ with $y > 0$ **DO** write $(-1/y)' \geq \alpha$ and integrate **BECAUSE** the increment of $-1/y$ is bounded above by $1/y(T)$, so a constant rate of increase is immediately impossible, and this is the whole proof of finite time blow up
- **WHEN** an existence question specifies the domain $\mathbb{R}$ rather than an interval **DO** look for the contradiction at infinity, not locally **BECAUSE** equations that are harmless locally are frequently lethal globally, and the stated domain is the examiner naming where to look

Backtesting

1. First useful cue: the unknown appears at two arguments, x and $f(x)$, so this cannot be solved and must be attacked by extracting consequences.
2. Method the expression suggested: bank the monotonicity facts, differentiate the relation to get $f'' = f'(f)f' > 0$, integrate twice for quadratic growth, resubstitute to reach $f' \geq \alpha f^2$, and finish with blow up.
3. Where you would hesitate, and the thought to reinforce: at the decision to differentiate an equation you cannot solve. Reinforce this: differentiation is free whenever the right side is differentiable data, and here it converts a composition into a product, which is a strict gain in sign information.
4. The trap the question sets: spending the whole time trying to construct f, or proving only qualitative facts (convex, increasing, positive derivative) and expecting them to contradict something. They never will; you need a rate.
5. What makes you recognise it faster next time: the shape $f' = f \circ f$ is superlinear feedback, and superlinear feedback on all of $\mathbb{R}$ is impossible. Recognise the family $y' \geq \alpha y^2$ on sight.
6. The reusable decision rule: in an existence question about a function on all of $\mathbb{R}$, extract a growth rate first and look for the contradiction at infinity.

Mirror drills

1. Show there is no differentiable $f : \mathbb{R} \to \mathbb{R}$ with $f'(x) = 1 + f(x)^2$ for every real x. *Answer:* $\big(\arctan f(x)\big)' = f'/(1+f^2) = 1$, so $\arctan f(x) - \arctan f(0) = x$ for all x, while the left side is bounded by π.

2. Prove that any differentiable $f : \mathbb{R} \to \mathbb{R}$ with $f' = f \circ f$ is infinitely differentiable. *Answer:* induct: if $f \in C^k$ then $f \circ f \in C^k$, so $f' \in C^k$ and $f \in C^{k+1}$; the base case is $f \in C^1$ because $f' = f \circ f$ is continuous.
3. Let y be positive and differentiable on $[0, \infty)$ with $y' \geq y^{3/2}$. Show no such y exists. *Answer:* $\left(-2y^{-1/2}\right)' = y'y^{-3/2} \geq 1$, so $2y(0)^{-1/2} \geq x$ for every $x \geq 0$, which fails for large x.

16.12 B6. Entrywise powers matching matrix powers forever

Let A be an $n \times n$ matrix of real numbers for some $n \geq 1$. For each positive integer k, let $A^{[k]}$ be the matrix obtained by raising each entry to the kth power. Show that if $A^k = A^{[k]}$ for $k = 1, 2, \ldots, n+1$, then $A^k = A^{[k]}$ for all $k \geq 1$.

Read it out loud

The problem places two rival meanings of the word power side by side. The algebraic power A^k mixes everything: each entry of A^2 is a sum of n products of entries. The entrywise power $A^{[k]}$ mixes nothing at all: the (i, j) entry of $A^{[2]}$ depends only on the (i, j) entry of A and on no other entry. These two operations have essentially nothing in common, and the hypothesis says they agree for the first $n + 1$ exponents while the conclusion says they agree forever. So this is an induction whose base cases are handed to you on a plate, and the entire problem is the single question: what relation propagates them?

Now read the count. The matrix is $n \times n$ and the hypothesis runs to $n + 1$. The theorem that converts the size n into a linear relation among the powers of A is the Cayley and Hamilton theorem: A satisfies its own characteristic polynomial P, which is monic of degree n, so A^n is a fixed real linear combination of $I, A, \ldots, A^{n-1}$. That relation has an I in it, and an I is poison here, because the entrywise operation does not send the constant polynomial 1 to I; it sends it to the matrix of all ones. So you multiply the relation by A and obtain a relation among $A, A^2, \ldots, A^{n+1}$ with no constant term. There is your $n + 1$. The exponent range in the hypothesis is not decoration, it is the examiner writing down the degree of $xP(x)$ and daring you to notice.

The fields embedded are linear algebra (characteristic polynomial, Cayley and Hamilton), polynomial division in $\mathbb{R}[x]$, and linear recursions with constant coefficients. What is really tested is whether you can see both operations as evaluations of a polynomial and then use only the structure the two evaluations share. The matrix evaluation $g \mapsto g(A)$ is a ring homomorphism. The entrywise evaluation is a ring homomorphism too, but for the wrong multiplication (the entrywise product). The single property they share is linearity in g, so every line of the solution must be linear in g, and you are forbidden from multiplying the two sides of anything.

State the knowledge gaps honestly, because this problem is not approachable without them. You must have installed: the characteristic polynomial $P(x) = \det(xI - A)$ for general n and the Cayley and Hamilton theorem for general n (A-level Further Mathematics gives you at most the 2×2 and 3×3 cases, usually as a verification exercise, and that is not enough); division with remainder in $\mathbb{R}[x]$ and the notion of the ideal generated by a polynomial; and the idea that a matrix identity can be read entrywise as n^2 scalar identities. Without Cayley and Hamilton there is no finite linear relation among the powers, and with no finite relation there is no problem, only an infinite family of unrelated equations. Within thirty seconds you should have written $P(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$, then $A\,P(A) = 0$, then the observation that $xP(x)$ has degree $n + 1$ and zero constant term.

Approaches, ranked worst to best

1. **Direct entrywise induction.** *Cue:* it is an induction on k, so assume $A^m = A^{[m]}$ and prove the next case. *Circuit:* expand $(A^{m+1})_{ij} = \sum_\ell (A^m)_{i\ell} A_{\ell j}$ and hope. It dies on the first line. The hypothesis gives you $(A^m)_{i\ell} = A^m_{i\ell}$, so the sum becomes $\sum_\ell A^m_{i\ell} A_{\ell j}$, and there is no reason on earth for that to equal A^{m+1}_{ij}. It also uses only one of the $n + 1$ given cases, which is a guarantee that you have misread the question: an unused hypothesis is a misread hypothesis.

2. **Diagonalise and compare spectra.** *Cue:* powers of a matrix are controlled by eigenvalues. *Circuit:* $A = SDS^{-1}$, so $A^k = SD^kS^{-1}$. Two fatal defects. Nothing says A is diagonalisable, and, far worse, the entrywise power is not a spectral object at all: $A^{[k]}$ has no relationship to

the eigenvalues of A. You would be applying a tool to one side of an equation whose other side the tool cannot see.

3. **Classify the matrices satisfying the hypothesis.** *Cue:* understand the objects before proving anything about them. *Circuit:* small case exploration, 2×2 first. As a proof strategy this is a research programme, not a Putnam route, and you would spend the hour on examples. As scouting it is worth two minutes, because the example $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$ turns up quickly and it is exactly the object that shows the threshold $n+1$ is sharp. Scout, then leave.

4. **Cayley and Hamilton plus induction on the exponent.** *Cue:* the count $n+1$ names the degree of $xP(x)$. *Circuit:* show every entry is a root of $xP(x)$, then run the same order n recursion on both sides. This is a complete and short solution. Apply the hypothesis to the relation $A\,P(A) = 0$ to learn that every entry t of A satisfies $tP(t) = 0$; then for $m \geq n$ both A^{m+1} and $A^{[m+1]}$ are the same fixed linear combination of the preceding n terms of their own sequences, and induction closes it.

5. **Two evaluation maps agreeing modulo a common ideal.** *Cue:* both sides are the value of a polynomial, read in two different ways. *Circuit:* linearity in the polynomial, plus division with remainder. Identical content to the previous route, but stated so that the induction disappears into one line of polynomial division. It also answers the two questions the induction leaves unanswered: why you must multiply P by x, and why $n+1$ instances and not n. Install this one.

The ideal solution

For a polynomial $g(x) = \sum_{k \geq 0} c_k x^k$ with real coefficients, define two matrices:

$$E(g) = g(A) = \sum_{k \geq 0} c_k A^k, \qquad F(g) = \text{the matrix whose } (i,j) \text{ entry is } g(A_{ij}).$$

Both E and F are linear maps from $\mathbb{R}[x]$ to the space of $n \times n$ real matrices, since each is defined coefficient by coefficient. For $k \geq 1$ we have $E(x^k) = A^k$ and $F(x^k) = A^{[k]}$. They part company only at $k = 0$: $E(1) = I$ while $F(1) = J$, the matrix all of whose entries are 1. For that reason the whole argument is carried out inside

$$x\mathbb{R}[x] = \{\, g \in \mathbb{R}[x] : g(0) = 0 \,\},$$

the polynomials with zero constant term, which is an ideal of $\mathbb{R}[x]$.

Let $V = \operatorname{span}_{\mathbb{R}}\{x, x^2, \ldots, x^{n+1}\}$, the polynomials of degree at most $n+1$ with zero constant term. The hypothesis says $E(x^k) = F(x^k)$ for $k = 1, \ldots, n+1$, so by linearity

$$E(g) = F(g) \qquad \text{for every } g \in V. \tag{1}$$

The conclusion to be proved is that $E(g) = F(g)$ for every $g \in x\mathbb{R}[x]$, since that includes every monomial x^k with $k \geq 1$.

Step 1: every entry of A is a root of $xP(x)$. Let $P(x) = x^n + a_{n-1}x^{n-1} + \cdots + a_0$ be the characteristic polynomial of A. Then $xP(x)$ is a polynomial of degree $n+1$ with zero constant term, so $xP(x) \in V$. By the Cayley and Hamilton theorem $P(A) = 0$, hence

$$E\big(xP(x)\big) = A\,P(A) = 0.$$

By (1), $F\big(xP(x)\big) = 0$ as well. Reading that matrix identity entrywise, for every pair (i,j) and $t = A_{ij}$,

$$t\,P(t) = 0. \tag{2}$$

Step 2: both maps annihilate every multiple of $xP(x)$. Let $I_0 = xP(x)\,\mathbb{R}[x]$, the ideal of multiples of $xP(x)$. For $g = xP(x)h(x)$ we have

$$E(g) = A\,P(A)\,h(A) = 0,$$

because polynomials in A commute and $P(A) = 0$; and the (i,j) entry of $F(g)$ is $t\,P(t)\,h(t) = 0$ by (2). So E and F both vanish on I_0.

Step 3: V and I_0 together span $x\mathbb{R}[x]$. It suffices to show $x^m \in V + I_0$ for every $m \geq 1$. For $m \leq n+1$ this is immediate. For $m \geq n+2$, divide with remainder by the polynomial $xP(x)$, which is monic of degree $n+1$:

$$x^m = Q(x)\,xP(x) + R(x), \qquad \deg R \leq n.$$

The monomial x^m and the product $Q(x)xP(x)$ are both divisible by x, hence so is R, so R is a linear combination of $x, x^2, \ldots, x^n$ and lies in V. Therefore $x^m = Q(x)xP(x) + R(x) \in I_0 + V$.

Step 4: conclude. The map $E - F$ is linear on $\mathbb{R}[x]$, it vanishes on V by (1), and it vanishes on I_0 by Step 2. Hence it vanishes on $V + I_0$, which by Step 3 is all of $x\mathbb{R}[x]$. In particular, for every $k \geq 1$,

$$A^k - A^{[k]} = (E - F)(x^k) = 0,$$

that is, $A^k = A^{[k]}$ for all $k \geq 1$, as required. ■

Remark 16.12.1. The same content in the language of recursions makes the role of the exponent $n+1$ explicit. Cayley and Hamilton gives, for every $m \geq n-1$,

$$A^{m+1} = -a_{n-1}A^m - \cdots - a_0 A^{m+1-n},$$

so the sequence $(A^k)_{k\geq 1}$ satisfies a linear recursion of order n for free. The sequence $(A^{[k]})_{k\geq 1}$ satisfies the very same recursion for $m \geq n$ if and only if every entry t satisfies $t^{m+1-n}P(t) = 0$, and that holds as soon as $tP(t) = 0$. But the root condition (2) is not free: it is exactly what the last instance $k = n+1$ of the hypothesis pays for. That is why the threshold is $n+1$ and not n.

Remark 16.12.2. The threshold is sharp. Take $n = 2$ and $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$. Then

$$A^k = \begin{pmatrix} 1 & 2k \\ 0 & 1 \end{pmatrix}, \qquad A^{[k]} = \begin{pmatrix} 1 & 2^k \\ 0 & 1 \end{pmatrix},$$

and $2k = 2^k$ holds for $k = 1$ and $k = 2$ but fails for $k = 3$. So $A^k = A^{[k]}$ for $k = 1, \ldots, n$ while $A^{n+1} \neq A^{[n+1]}$, and the hypothesis cannot be weakened to $k \leq n$. Consistently with Step 1, here $P(x) = (x-1)^2$ and the entry 2 is not a root of $x(x-1)^2$.

Layer 1: WHAT

Block 1. Read the count: $n \times n$ with exponents up to $n+1$ means the characteristic polynomial multiplied by x.

Block 2. Name the two operations as evaluation maps E and F of the same polynomial, and record that the only property they share is linearity in that polynomial.

Block 3. Record the constant term obstruction, $E(1) = I$ but $F(1) = J$, and conclude that all work happens in the polynomials with zero constant term. This is the reason for multiplying P by x.

Block 4. Push $xP(x)$ through the hypothesis: $E(xP) = A\,P(A) = 0$, hence $F(xP) = 0$, hence every entry of A is a root of $xP(x)$.

Block 5. Both maps therefore kill every multiple of $xP(x)$.

Block 6. Division with remainder shows that the multiples of $xP(x)$ together with the span of $x, \ldots, x^{n+1}$ exhaust the polynomials with zero constant term, so the two maps agree everywhere on that ideal.

The nominal route is Block 1 replaced by multiplication: write out $(A^{m+1})_{ij}$ as a sum, substitute the inductive hypothesis, and stare. The bifurcation point is the very first line, and the question that decides it is not about matrices at all: it is whether you treat the number $n+1$ as an incidental range or as a name. On the Putnam a numerical parameter that matches the size of the object is always a name. Here it names $\deg\big(xP(x)\big)$, and once you have written $xP(x)$ the problem is over in four lines.

Layer 2: HOW

Block 1. **Cue.** The size is n, the exponent range is 1 to $n+1$, and the conclusion is about all exponents. A finite hypothesis producing an infinite conclusion always means a recursion, and a recursion among powers of a matrix of size n means the characteristic polynomial. **Why natural.** Cayley and Hamilton is the only standard supplier of a monic annihilating polynomial whose degree is known before you look at A. That is what makes it usable in a problem where A is arbitrary: you cannot inspect A, but you can name $\deg P = n$. **Link to the next block.** You have a relation among $I, A, \ldots, A^n$. To use it you must transport it across the entrywise operation, so you now examine what that operation does to a polynomial relation.

Block 2. **Cue.** Both A^k and $A^{[k]}$ are values at A of the same monomial x^k, read in two different ways. Whenever the same formal object has two interpretations, name both as maps and compare them, rather than comparing the outputs case by case. **Why natural.** Naming the maps forces you to ask what structure they have. E is a ring homomorphism for matrix multiplication; F is a ring homomorphism for the entrywise (Hadamard) product; they share linearity and nothing else. Knowing precisely what you are allowed to do is worth more than any computation, because it eliminates whole families of doomed attempts, including every attempt to multiply the hypothesis by itself. **Link to the next block.** Linearity means you should ask on which subspace the two maps agree, and that question immediately exposes the constant term.

Block 3. **Cue.** The hypothesis starts at $k = 1$, not $k = 0$. An exponent range that pointedly excludes 0 is a warning about constants. **Why natural.** $E(1) = I$ and $F(1) = J$ are different matrices whenever $n \geq 2$, so the two maps disagree on constants and no argument may involve one. Consequently the natural domain is $x\mathbb{R}[x]$, and any annihilating polynomial you intend to use must have zero constant term. The characteristic polynomial does not, so you multiply it by x; this costs one degree, which is precisely the $+1$ in the hypothesis. **Link to the next block.** You now know exactly which polynomial to feed to the hypothesis: $xP(x)$, of degree $n+1$, the largest degree the hypothesis covers.

Block 4. **Cue.** A polynomial lying in the span where the two maps are known to agree, and whose image under one of the maps you can compute. That is the definition of a usable test object. **Why natural.** $E(xP) = A\,P(A) = 0$ by Cayley and Hamilton, so the agreement forces $F(xP) = 0$, and F is entrywise, so a matrix identity becomes n^2 scalar identities $tP(t) = 0$. This is the one moment in the problem where information crosses from the matrix world to the scalar world, and it crosses because F is entrywise and for no other reason. **Link to the next block.** Knowing the entries are roots of $xP(x)$ makes F kill not just xP but every multiple of it, which is what turns a single instance into an infinite family.

Block 5. **Cue.** A polynomial that vanishes at every entry, and a matrix polynomial that annihilates A. Both statements upgrade to whole ideals for free, because both maps respect multiplication in the appropriate sense. **Why natural.** $E(xPh) = AP(A)h(A) = 0$ uses that polynomials in A commute, and $F(xPh)$ has entries $tP(t)h(t) = 0$ uses only (2). So the common kernel of E and F contains the

ideal generated by $xP(x)$. Whenever you have two linear maps that agree on a subspace and vanish on a common ideal, the natural next question is whether those two pieces span the domain. **Link to the next block.** The answer is supplied by division with remainder, which is the standard way to prove that a subspace plus an ideal is everything.

Block 6. Cue. Ideal generated by a monic polynomial, plus a claim that low degree spans the rest. That is the division algorithm, every time. **Why natural.** Dividing x^m by the monic polynomial $xP(x)$ leaves a remainder of degree at most n, and the remainder inherits divisibility by x from the two other terms, so it lands in V. The inheritance step is the only fiddly point and it is one line, but it must be written: without it the remainder might have a constant term and the argument would leak into the one place the two maps disagree. **Link to the next block.** The linear map $E - F$ vanishes on a spanning set of $x\mathbb{R}[x]$, so it vanishes identically there, and the conclusion is exactly the statement about x^k.

Layer 3: WHY

The truth beneath this problem is that raising to a power is evaluation. Both sides of the equation $A^k = A^{[k]}$ are values of the monomial x^k under a map from $\mathbb{R}[x]$ into the $n \times n$ matrices, and once you see the problem that way it stops being about matrices and becomes a statement about two linear maps agreeing on a quotient. Agreement of linear maps on a spanning set is agreement everywhere; that is the only theorem used at the end. All the specific content of the problem is spent identifying the right spanning set, and the right spanning set is the $n + 1$ monomials the hypothesis names together with an ideal that both maps annihilate.

Cayley and Hamilton earns its place for a single reason: it supplies a nonzero element of the kernel of E whose degree is known in advance. Any monic polynomial annihilating A would do; the minimal polynomial μ would even be sharper, and the argument run with μ shows that the exponents 1 to $\deg \mu + 1$ already suffice. But $\deg \mu$ depends on A and cannot be quoted in a statement about every A, whereas $\deg P = n$ can. So the choice of the characteristic polynomial is not a preference, it is forced by the fact that the problem quantifies over all A of size n.

The discernment mechanism is the one to carry away: when two operations on the same object must be compared, first identify the largest structure they share, then confine every line of the argument to that structure. Here the shared structure is linearity in the polynomial, and the moment you accept that, the doomed routes die on their own. Multiplying the hypothesis by itself is illegal because F is not multiplicative for matrix multiplication. Diagonalising is illegal because F does not see the spectrum. Taking traces or determinants is useless because neither is linear in the polynomial in a compatible way. The problem is easy once the move list is short, and the move list is short as soon as you name what the two maps are.

There is also a general principle worth extracting, because it recurs across the paper. When a problem says that a property holding for $k \leq N$ forces it for all k, look for a linear recursion of order roughly N that both sides satisfy, and then ask what the final instance is buying. Here the matrix side satisfies the recursion for free and the entrywise side does not; the last instance $k = n + 1$ is precisely the payment that puts the entrywise side under the same recursion. Being able to say what the last hypothesis buys is the difference between reproducing a solution and understanding it, and it is also how you would guess the sharp example without being told one exists.

The examiners nudged three times. They wrote $n \times n$ and then $n + 1$, which is the degree of $xP(x)$. They wrote for all $k \geq 1$ rather than $k \geq 0$, which quietly warns that $k = 0$ is different, and $k = 0$ is different exactly because $F(1) = J \neq I$. And they defined $A^{[k]}$ verbally, entry by entry, rather than by a symbol, which is the examiner insisting that you think of it as n^2 separate scalar operations, which is the reading that makes Step 1 possible.

Revolving concepts

Characteristic polynomial and the Cayley and Hamilton theorem. For an $n \times n$ matrix A over a commutative ring, $P(x) = \det(xI - A)$ is monic of degree n and $P(A) = 0$. The standard proof multiplies the adjugate identity $\operatorname{adj}(xI - A)\,(xI - A) = P(x)I$ out and compares coefficients of powers of x, the entries of the adjugate being polynomials in x of degree at most $n - 1$. Beyond A-level for general n. See Book 1, field: Matrices and determinants, topic: Cayley and Hamilton.

Minimal polynomial and the annihilator ideal. The set $\{g : g(A) = 0\}$ is an ideal of $\mathbb{R}[x]$, hence generated by a single monic polynomial μ, the minimal polynomial; μ divides P and $\deg \mu \le n$. Running the solution with μ in place of P shows that exponents 1 to $\deg \mu + 1$ suffice for a fixed A. See Book 1, field: Linear algebra, topic: minimal polynomial.

Division with remainder and ideals in $\mathbb{R}[x]$. For monic D of degree d and any g, there are unique Q, R with $g = QD + R$ and $\deg R < d$. Consequently $\mathbb{R}[x] = \operatorname{span}\{1, x, \ldots, x^{d-1}\} + D\,\mathbb{R}[x]$, and the same argument intersected with $x\mathbb{R}[x]$ gives the decomposition used in Step 3. See Book 1, field: Polynomials, topic: division algorithm and ideals.

Linear recursions in a vector space, and the window principle. If two sequences (X_k) and (Y_k) with values in the same vector space satisfy the same recursion $Z_{m+1} = c_1 Z_m + \cdots + c_n Z_{m+1-n}$ for all $m \ge K$, and $X_k = Y_k$ for the n consecutive indices $k = K + 1 - n, \ldots, K$, then $X_k = Y_k$ for all $k \ge K + 1 - n$. The proof is induction on k, one step at a time. Generalised: a solution of an order n linear recursion is determined by any window of n consecutive terms, so two solutions agreeing on one window agree forever; this is the discrete analogue of uniqueness for an order n initial value problem. See Book 1, field: Sequences and recurrences, topic: linear recursions and initial windows.

Hadamard product and entrywise functional calculus. The entrywise product $(A \circ B)_{ij} = A_{ij}B_{ij}$ makes the $n \times n$ matrices a commutative algebra whose identity is the all ones matrix J, and $A^{[k]}$ is the kth Hadamard power. The map $g \mapsto F(g)$ is the direct product of the n^2 scalar evaluation homomorphisms $g \mapsto g(A_{ij})$. Recognising F as n^2 evaluations at once is what turns the matrix identity $F(xP) = 0$ into the scalar statement that every entry is a root. See Book 1, field: Matrices and determinants, topic: Hadamard product.

Circuit Bank

- **WHEN** a hypothesis about an $n \times n$ matrix runs over exponents 1 to $n + 1$ **DO** write down $A\,P(A) = 0$ with P the characteristic polynomial **BECAUSE** $xP(x)$ is the monic polynomial of degree $n + 1$ with zero constant term, and the exponent range is the examiner naming it

- **WHEN** two different operations on the same object must be compared **DO** identify the largest structure both respect and confine every line to it **BECAUSE** the shared structure is the complete move list, and knowing it kills the doomed routes before you spend ink on them

- **WHEN** an entrywise operation appears anywhere **DO** check the constant polynomial separately **BECAUSE** entrywise evaluation sends 1 to the all ones matrix and not to the identity, so any relation you transport must have zero constant term

- **WHEN** two sequences satisfy the same linear recursion of order n and agree on n consecutive terms **DO** stop and conclude they agree forever **BECAUSE** a solution of an order n linear recursion is determined by any window of n consecutive values

- **WHEN** a finite hypothesis is claimed to force an infinite conclusion **DO** ask what the final instance of the hypothesis buys **BECAUSE** it is always the instance that installs the recursion on the side that does not get it for free, and naming it tells you where the sharp example lives

Backtesting

1. First useful cue: the exponent bound $n+1$ against a matrix of size n, which is $\deg\big(xP(x)\big)$ and nothing else.
2. Method the expression suggested: treat A^k and $A^{[k]}$ as two linear evaluations of x^k, feed $xP(x)$ through the hypothesis to learn that every entry is a root of $xP(x)$, and then note both evaluations kill the whole ideal generated by $xP(x)$.
3. Where you would hesitate, and the thought to reinforce: at the multiplication of P by x, which looks arbitrary. Reinforce this: the entrywise map sends 1 to the all ones matrix, so constants are the one place the two readings disagree, and multiplying by x is how you leave that place.
4. The trap the question sets: starting an induction on k using only the previous case, which uses one hypothesis instead of $n+1$ and stalls on the first line. An unused hypothesis is a misread hypothesis.
5. What makes you recognise it faster next time: the phrase two ways of reading the same polynomial. Every problem that compares an algebraic operation with an entrywise or coordinatewise one is asking you to find the ideal that both readings annihilate.
6. The reusable decision rule: when a numerical parameter in the hypothesis matches a structural invariant of the object, treat it as the name of a specific polynomial and write that polynomial down before doing anything else.

Mirror drills

1. For $A = \begin{pmatrix} 1 & 2 \\ 0 & 1 \end{pmatrix}$, decide for which k the identity $A^k = A^{[k]}$ holds, and explain why this does not contradict the problem. *Answer:* A^k has top right entry $2k$ and $A^{[k]}$ has 2^k, so it holds exactly for $k = 1, 2$; since $n = 2$ the hypothesis needs k up to 3, and it fails there.
2. Suppose A is $n \times n$ with $A^k = A^{[k]}$ for $k = 1, \ldots, n+1$. Prove that every entry of A is either 0 or a real eigenvalue of A. *Answer:* apply the entrywise reading to $xP(x)$ to get $tP(t) = 0$ for each entry t; a real root of P is an eigenvalue.
3. Let (X_k) and (Y_k) be sequences of matrices satisfying the same recursion $Z_{m+1} = c_1 Z_m + \cdots + c_n Z_{m+1-n}$ for all $m \geq K$. Show that agreement on the window $k = K+1-n, \ldots, K$ forces agreement for all $k \geq K+1-n$. *Answer:* induction on k: each new term is the same fixed combination of n terms already known to agree.

Chapter 17

Putnam 2009

17.1 A1. squares constrain f to vanish everywhere

Let f be a real-valued function on the plane such that for every square $ABCD$ in the plane, $f(A) + f(B) + f(C) + f(D) = 0$. Does it follow that $f(P) = 0$ for all points P in the plane?

Read it out loud

Say the hypothesis back to yourself: whatever f is, its values at the four corners of ANY square sum to zero, with no continuity, boundedness or measurability assumed. The question asks whether this single constraint, applied to every square at once, forces f to be the zero function.

The fields wired in are functional equations and finite linear algebra, dressed as plane geometry. No continuity is given, so any route through limits, integrals or "average over all squares through P" is dead before it starts; that absence is a cue that the intended route is a FINITE combination of finitely many instances of the hypothesis.

The first fact worth banking: every point P is the centre of infinitely many squares, of every size and orientation, so you get to CHOOSE the square through P. The second cue is the word "every": the constraint holds simultaneously for a square and for the smaller squares inside it, so expect to combine constraints at more than one scale.

Approaches, ranked worst to best

1. **guess $f \equiv 0$ and stop.** *Cue:* the hypothesis looks like it should force zero. *Circuit:* none, this is not a proof. Settles nothing: the question asks whether it FOLLOWS, and no argument has been given.
2. **average f over all squares centred at P.** *Cue:* "every square through P" suggests integrating over angle and radius. *Circuit:* continuity and symmetry. Illegal, since f is not assumed continuous or measurable, and it reroutes a purely algebraic problem through analysis for no gain.
3. **coordinate bash on a generic square through P.** *Cue:* P is some point, write a general square with P as a vertex and grind. *Circuit:* brute linear algebra over unfixed parameters. Workable but wasteful: a generic square carries two free parameters you end up specialising anyway, without any guidance on which specialisation to choose.
4. **one square centred at P, plus its midpoint subdivision.** *Cue:* every point is a centre, and every square contains smaller squares built from its own midpoints. *Circuit:* finite linear system from self-similar nested squares. The ideal route: the single most information-dense configuration, reused at three scales, gives a finite linear system with $f(P)$ isolated after one clean cancellation.

The ideal solution

The answer is yes, $f(P) = 0$ for every point P.

Fix an arbitrary point P. Choose any square $ABCD$, vertices in order, with centre P; such a square exists for every size and orientation. Let E, F, G, H be the midpoints of sides AB, BC, CD, DA.

Two facts about this configuration, each checked once by placing P at the origin and A, B, C, D at $(\pm r, \pm r)$ in a rotated frame: $EFGH$ is itself a square, centred at P; and each of $AEPH$, $BFPE$, $CGPF$, $DHPG$ is a square with P as a vertex.

Apply the hypothesis to all six squares:

$$0 = f(A) + f(B) + f(C) + f(D)$$
$$0 = f(E) + f(F) + f(G) + f(H)$$

$$0 = f(A) + f(E) + f(P) + f(H)$$
$$0 = f(B) + f(F) + f(P) + f(E)$$
$$0 = f(C) + f(G) + f(P) + f(F)$$
$$0 = f(D) + f(H) + f(P) + f(G).$$

Add the last four equations. Each of A, B, C, D appears once, each of E, F, G, H appears twice, P appears four times:

$$f(A) + f(B) + f(C) + f(D) + 2\big(f(E) + f(F) + f(G) + f(H)\big) + 4f(P) = 0.$$

The first bracket is 0 by the first equation, the second bracket is 0 by the second equation. What remains is $4f(P) = 0$, so $f(P) = 0$. Since P was arbitrary, $f(P) = 0$ for every point P: yes, it follows.

Layer 1: WHAT

Block 1. Fix an arbitrary P; recognise it as the centre of a square you get to choose. Block 2. Build one configuration reusing P at three scales: the big square $ABCD$, the mid square $EFGH$ sharing its centre, and four small squares meeting at P as a vertex. Block 3. Write the hypothesis for all six squares. Block 4. Sum the four small-square equations; the big and mid square sums appear inside, ready to cancel. Block 5. Cancel with equations 1 and 2, isolate $4f(P) = 0$.

The nominal route (coordinate bash) reaches the same answer by grinding a system with unfixed parameters before specialising. The bifurcation point is Block 2: the decision to introduce the midpoints at all, since nothing in the statement mentions them. They are earned because a second scale of square sharing centre P was needed, and side midpoints are the cheapest way to manufacture one.

Layer 2: HOW

Block 1. **Cue.** "For all points P" and "for every square" are both universally quantified over the same underlying objects. **Why natural.** $f(P) = 0$ is one scalar equation; producing it for free needs P to occur with nonzero net coefficient in some finite combination of the zero-equations, so P must literally be a vertex or centre of a square in your configuration. **Link to the next block.** Since P can be the centre of a square of any size, you build the configuration around P rather than searching for one.

Block 2. **Cue.** The one square you have, $ABCD$, has P as its centre, never as a vertex, so its equation never mentions $f(P)$. **Why natural.** A square's centre is never one of its own vertices, so you need squares with P as a VERTEX, built only from points already defined by $ABCD$. Bisecting each side and joining to P closes up into a smaller square, because a square is symmetric under a quarter turn about its centre, and that turn sends each vertex to midpoint to centre path to the next. **Link to the next block.** This is why E, F, G, H exist: forced by asking how to put P on a corner using only $ABCD$'s own points.

Block 3. **Cue.** $EFGH$ sits in the picture as a by-product of Block 2. **Why natural.** By the same quarter-turn symmetry, $EFGH$ is a square with centre P, so the hypothesis applies to it exactly as to $ABCD$. **Link to the next block.** This supplies the equation that cancels the "doubled" midpoint terms in the next block; without it those terms are stranded.

Block 4. **Cue.** Six zero-equations, nine unknowns; you want a combination leaving only $f(P)$. **Why natural.** Summing the four small-square equations accumulates P's coefficient to 4, since P appears in every one; every original vertex appears once and every midpoint appears in exactly

two of the four small squares, since the four tiles exactly cover $ABCD$. **Link to the next block.** That clean arithmetic is what equations 1 and 2 are built to cancel.

Layer 3: WHY

The deep fact is that a square's centre and a square's vertex can be the SAME point P across squares sharing vertices and edges, so one picture generates several instances of the hypothesis whose unknowns overlap almost completely. You are not solving for f as a function, only for the one combination of six equations, nine unknowns, that the self-similar structure of a square under midpoint subdivision hands you as $4f(P)$.

The discernment mechanism: when a functional equation is indexed by a family of figures, ask whether the family is closed under a subdivision that fixes a point of interest. Squares are closed under joining the midpoints, and the midpoint square shares the original centre, which tells you in advance how many independent equations you can generate and which point they pivot on. The examiners nudge away from analysis (no continuity, on purpose, kills any averaging route) toward finite linear algebra: the official remark that the same statement holds for any regular polygon confirms that nothing special about four sides was used beyond the midpoint figure being the same shape, centred at the same point, an invariance under rotation by the appropriate fraction of a turn.

Revolving concepts

Functional equations on a geometric index set: the unknown is a function constrained by linear relations indexed by figures. The method is always to manufacture a finite sub-family whose relations, summed with integer coefficients, isolate the wanted value. See Book 1, field: Functional equations, topic: linear constraints over geometric families.

Self-similarity of a regular polygon under midpoint subdivision: joining consecutive midpoints of a square gives a smaller square, rotated 45°, same centre; more generally the midpoint polygon of any regular n-gon is again a regular n-gon with the same centre. Recursion generalised: iterating, Q_{m+1} the midpoint polygon of Q_m, every Q_m shares the centre of Q_0, giving a nested family of hypothesis instances all pivoting on one centre. See Book 1, field: Geometry, topic: self-similar polygon subdivision.

Targeted elimination with more unknowns than needed: decide in advance which unknowns must die and which must survive, then choose the combination of equations accordingly, rather than summing and hoping. See Book 1, field: Linear algebra, topic: eliminating unwanted variables by design.

Circuit Bank

- **WHEN** a real-valued constraint is stated "for every [figure]" and you are asked about a single point **DO** ask whether that point can be simultaneously the centre of one instance of the figure and a vertex of a smaller instance built from it by internal subdivision **BECAUSE** overlapping the unknowns across two scales turns infinitely many instances of the hypothesis into one solvable finite system
- **WHEN** you have k zero-equations in more unknowns than you need **DO** write down which unknowns must stay alive and which must die before adding anything **BECAUSE** blind summation is the nominal route; targeted elimination is the ideal one
- **WHEN** a problem about a general point or function gives no continuity, boundedness or smoothness hypothesis **DO** rule out any route needing a limit, integral or average immediately **BECAUSE** the setter is signalling the intended route is purely algebraic

Backtesting

1. First useful cue: the hypothesis holds for every square, so you may pick the most convenient one.
2. Method the expression suggested: build a square centred at P, then look for smaller squares generated from it that put P on a vertex.
3. Where you would hesitate, and the thought to reinforce: introducing E, F, G, H, unmentioned in the statement; reinforce that new points are earned by asking what construction forces P onto a vertex using only existing data.
4. The trap the question sets: averaging or integrating over all squares through P, tempting given the word "every" but illegal with no continuity assumed.
5. What makes you recognise it faster next time: "for every [rigid figure]" applied to a point that is naturally both a centre and a vertex of that figure at different scales.
6. The reusable decision rule: when a point can play two structural roles in the same family of figures, build the configuration using both roles at once and sum with targeted cancellation.

Mirror drills

1. Suppose $f(A) + f(B) + f(C) = 0$ for every equilateral triangle ABC. Does $f \equiv 0$ follow? *Answer:* Yes; for any P, use an equilateral triangle centred at P together with the three smaller equilateral triangles meeting at P, and cancel exactly as above.
2. Suppose $f(A) + \cdots + f(F) = 0$ for every regular hexagon $ABCDEF$. Is $f \equiv 0$ forced? *Answer:* Yes; a hexagon centred at P and its midpoint hexagon, also regular with centre P, play the roles of $ABCD$ and $EFGH$ above.
3. If the hypothesis is only $f(A) + f(C) = f(B) + f(D)$ for every square (opposite corners equal, not full sum zero), does $f \equiv 0$ still follow? *Answer:* No; $f(x, y) = x + y$ satisfies it for every square, since a square's diagonals share a midpoint, but is not identically zero.

17.2 A2. the product fgh decouples the system

Functions f, g, h are differentiable on some open interval around 0 and satisfy the equations and initial conditions

$$f' = 2f^2gh + \frac{1}{gh}, \quad f(0) = 1,$$

$$g' = fg^2h + \frac{4}{fh}, \quad g(0) = 1,$$

$$h' = 3fgh^2 + \frac{1}{fg}, \quad h(0) = 1.$$

Find an explicit formula for $f(x)$, valid in some open interval around 0.

Read it out loud

Three coupled, nonlinear, first order differential equations in three unknowns, each mixing all three functions through a product term and a reciprocal-product term. Read out loud, this is not "solve a system of ODEs" in the generic sense, since that generic problem is not tractable by hand. The fields wired in are ordinary differential equations, specifically separable and linear first order ones, algebraic manipulation of the equations before any calculus is done, and the product rule read backwards.

The first cue to bank within seconds: every polynomial term on the right is the SAME triple product fgh times one of f, g, h ($2f^2gh = 2f \cdot (fgh)$, $fg^2h = g \cdot (fgh)$, $3fgh^2 = 3h \cdot (fgh)$), and every reciprocal term is a constant divided by the OTHER two functions, which is the same as the missing function divided by the full product fgh (for instance $1/(gh) = f/(fgh)$). Whenever three equations all revolve around the same triple product, that product, not the three functions separately, is the object to solve for first. The second cue: the question only asks for f, a strong hint that g and h never need to be isolated individually, only some combination of them does.

Approaches, ranked worst to best

1. **attack the system by generic ODE system theory.** *Cue:* "three coupled nonlinear ODEs" superficially looks like something to diagonalise. *Circuit:* eigenvalue methods for linear systems. Hopeless directly: the system is nonlinear, there is no matrix to diagonalise, and no exam expects a numerical solution.

2. **guess the symmetric ansatz $f = g = h$.** *Cue:* the three equations look superficially parallel in shape. *Circuit:* pattern matching on symmetry. Fails on inspection: the polynomial coefficients are $2, 1, 3$ and the reciprocal numerators are $1, 4, 1$, neither list constant across the equations, so setting $f = g = h = \varphi$ forces $\varphi' = 2\varphi^3 + 1/\varphi^2$ from the first equation and $\varphi' = \varphi^3 + 4/\varphi^2$ from the second simultaneously, an inconsistency. Quick to discard, but not a route to the answer.

3. **solve for f, g, h individually via separate substitutions.** *Cue:* each equation looks like it wants its own change of variable. *Circuit:* ad hoc substitution per equation. Never exploits the shared fgh structure, leaving three genuinely coupled nonlinear equations with no reduction in dimension; far more ink than necessary.

4. **multiply each equation by the complementary pair of the other two functions and sum.** *Cue:* the reciprocal terms are each (missing variable)/(full product), and three product-rule pieces summing to a single derivative is the signature of a hidden total derivative. *Circuit:* engineer $f'gh + fg'h + fgh' = (fgh)'$ on purpose. The ideal route: one substitution, $k = fgh$, collapses three coupled nonlinear equations into a single separable one, after which f is recovered from k by direct integration, with g, h never needing to be found.

The ideal solution

Multiply the first equation by gh, the second by fh, the third by fg:

$$f'gh = 2f^2g^2h^2 + 1,$$
$$g'fh = f^2g^2h^2 + 4,$$
$$h'fg = 3f^2g^2h^2 + 1.$$

Add the three lines. The left side is $f'gh + fg'h + fgh'$, exactly $(fgh)'$ by the product rule for three factors. The right side collects to $(2+1+3)f^2g^2h^2 + (1+4+1) = 6(fgh)^2 + 6$. Writing $k = fgh$,

$$k' = 6(k^2+1), \qquad k(0) = f(0)g(0)h(0) = 1.$$

This is separable and can be solved without any black-box existence-uniqueness theorem: since $k^2 + 1 > 0$ always, define $\varphi(x) = \arctan(k(x))$ where k is defined. Then

$$\varphi'(x) = \frac{k'(x)}{1+k(x)^2} = \frac{6(k(x)^2+1)}{1+k(x)^2} = 6$$

identically, so $\varphi(x) = 6x + \varphi(0) = 6x + \pi/4$ (since $\varphi(0) = \arctan 1 = \pi/4$), hence

$$k(x) = \tan\left(6x + \frac{\pi}{4}\right),$$

valid on the open interval around 0 where $6x + \pi/4 \in (-\pi/2, \pi/2)$, in particular on $(-\pi/24, \pi/24)$.

Return to the first original equation to recover f alone, using $gh = k/f$: since $2f^2gh = 2f \cdot (fgh) = 2fk$ and $1/(gh) = f/(fgh) = f/k$,

$$f' = 2fk + \frac{f}{k}, \qquad \text{i.e.} \qquad \frac{f'}{f} = 2k(x) + \frac{1}{k(x)},$$

valid near 0 since $k(0) = 1 \neq 0$. Writing $u = 6x + \pi/4$ (so $u' = 6$), this is $\frac{d}{dx}\ln f = 2\tan u + \cot u$. Two antiderivatives, checked by differentiating back: $\frac{d}{dx}\left[-\frac{1}{3}\ln\cos u\right] = \frac{1}{3}\tan u \cdot u' = 2\tan u$, and $\frac{d}{dx}\left[\frac{1}{6}\ln\sin u\right] = \frac{1}{6}\cot u \cdot u' = \cot u$ (both valid near $x = 0$, where u is near $\pi/4$, so $\sin u, \cos u > 0$ and no absolute values are needed). Adding,

$$\ln f(x) = -\frac{1}{3}\ln\cos u + \frac{1}{6}\ln\sin u + C = \frac{1}{6}\ln\left(\frac{\sin u}{\cos^2 u}\right) + C.$$

At $x = 0$, $u = \pi/4$, $\sin u = \cos u = \sqrt{2}/2$, so $\sin u/\cos^2 u = \sqrt{2}$. Since $f(0) = 1$, $0 = \frac{1}{6}\ln\sqrt{2} + C$, so $C = -\frac{1}{12}\ln 2$. Hence

$$\boxed{f(x) = 2^{-1/12}\left(\frac{\sin(6x+\pi/4)}{\cos^2(6x+\pi/4)}\right)^{1/6}}$$

on the open interval around 0 where this is defined and positive, for instance $(-\pi/24, \pi/24)$.

Layer 1: WHAT

Block 1. Notice all three right sides share the triple product fgh: polynomial terms are fgh times f, g, h; reciprocal terms are the missing variable divided by fgh. Block 2. Multiply each equation by exactly the other two functions, engineering the left sides into the three pieces of the product rule for $(fgh)'$. Block 3. Sum, collapse to a single separable ODE in $k = fgh$. Block 4. Solve for k directly via the arctan composition, no uniqueness theorem needed. Block 5. Substitute $gh = k/f$ back into the FIRST original equation only, reducing it to a relation for $\ln f$ in terms of the known $k(x)$. Block 6. Integrate, apply $f(0) = 1$, box the answer.

The nominal route, chasing f, g, h separately, never gets past a genuinely coupled nonlinear system. The bifurcation point is Block 2: multiplying by the complementary pair rather than trying to isolate one function's equation in isolation. Everything after that is forced computation.

Layer 2: HOW

Block 1. **Cue.** Three reciprocal terms, $1/(gh)$, $4/(fh)$, $1/(fg)$, each missing exactly one letter out of three. **Why natural.** A term missing exactly one of three letters is that missing letter divided by the product of all three, since multiplying and dividing by it changes nothing. Naming $k = fgh$ turns three different-looking fractions into f/k, $4g/k$, h/k, one unifying object. **Link to the next block.** Once k is visible in the denominators, check the polynomial terms too; it is visible there as well ($2f^2gh = 2fk$), which tells you k is the right substitution for the whole system, not just the fractions.

Block 2. **Cue.** You want $(fgh)' = f'gh + fg'h + fgh'$ on the left, three terms, one per equation. **Why natural.** Multiplying equation 1 by gh turns its f' into $f'gh$, the first piece of that expansion; the same move on equations 2 and 3 (multiply by fh, then fg) produces the other two pieces. This is reverse engineering: you know the sum you want on the left and multiply each equation by whatever makes its derivative term match one summand. **Link to the next block.** Once the left side is fixed as $(fgh)'$, simplify whatever the right side becomes in terms of $k = fgh$.

Block 3. **Cue.** Every polynomial term on the right becomes a multiple of k^2, and every reciprocal term becomes a bare constant, because the multipliers you chose are exactly the reciprocal terms' own denominators. **Why natural.** This is the payoff for choosing multipliers to match the reciprocal denominators, not just to build the product rule: the same multiplication that assembles $(fgh)'$ on the left also clears every fraction on the right. **Link to the next block.** Summing gives one ODE, $k' = 6k^2 + 6$, purely in k: three coupled unknowns collapsed to one.

Block 4. **Cue.** $k' = 6(k^2 + 1)$ is separable, and arctan is the antiderivative of $1/(1 + t^2)$. **Why natural.** Composing arctan with $k(x)$ and differentiating via the chain rule gives $k'(x)/(1 + k(x)^2)$, which the ODE says is identically 6. A function with constant derivative on an interval is linear on it, no separate existence-uniqueness theorem needed, since the derivative of $\arctan(k(x))$ was computed directly from the given equation, not guessed and verified. **Link to the next block.** This gives $k(x)$ completely as an explicit function of x.

Block 5. **Cue.** The question only wants f, whose equation contains g, h only through gh. **Why natural.** $gh = k/f$ directly from $k = fgh$, removing g and h from the first equation entirely without solving for them individually. **Link to the next block.** The result is f'/f expressed purely in terms of the already-known $k(x)$, a linear equation for $\ln f$.

Block 6. **Cue.** $f'/f = 2k + 1/k$ with k explicit. **Why natural.** $f'/f = (\ln f)'$, so integrating directly gives $\ln f$ up to a constant; the two pieces $2\tan u$ and $\cot u$ each match a standard logarithmic derivative once $k = \tan u$. **Link to the next block.** The constant of integration is fixed by $f(0) = 1$, the only place the specific numerical answer enters.

Layer 3: WHY

This system is tractable by hand because the setter chose the polynomial coefficients $(2, 1, 3)$ and reciprocal numerators $(1, 4, 1)$ so that multiplying by the complementary pair simultaneously builds a total derivative on the left and a function of a single new variable $k = fgh$ on the right. That is a structural coincidence to look for in any system of the shape (function)$'$ = (stuff)$\times$(product of all the functions) + (stuff)/(product of the other functions): the product rule for n factors always has exactly n terms, one per factor, and if the given system's n equations are shaped so that multiplying equation i by everything except factor i produces the ith term of that expansion, the whole system decouples through the combined product.

The discernment mechanism: whenever a system's nonlinear terms are all scalar multiples of the SAME product of all the unknowns, test whether multiplying by the complementary factors and summing produces a total derivative. The test costs only a few lines, while grinding the system unreduced can cost an unbounded amount.

A secondary truth worth stating cleanly, since the official solution leans on it without proof: a first order ODE $k' = F(k)$ with an explicit solution candidate through the given initial condition does not, by itself, guarantee that candidate is the only one; in general an existence-uniqueness theorem (Picard-Lindelof, using that F is locally Lipschitz) is what would rule out other branches through the same point. Here that heavier tool is unnecessary, because $\arctan(k(x))$ is differentiated directly FROM the given equation, which constructs the solution rather than merely verifying a guess.

Revolving concepts

Product rule for three or more factors: $(fgh)' = f'gh + fg'h + fgh'$, the direct generalisation of $(fg)' = f'g + fg'$. Recursion generalised: for n factors, $(f_1 f_2 \cdots f_n)'$ is the n-term sum $\sum_i f_1 \cdots f_i' \cdots f_n$; whenever a system of n coupled equations, each multiplied by the complementary product of the other $n-1$ functions, sums to this expansion, the full product decouples the system. See Book 1, field: Calculus and integrals, topic: product rule for n factors.

Separable ODEs of the shape $k' = a(k^2 + b^2)$: composing $\theta(x) = \arctan(k(x)/b)$ and differentiating gives $\theta' = ab$ constant, so $\theta(x) = abx + \theta(0)$, hence $k(x) = b\tan(abx + \arctan(k(0)/b))$ directly, with no uniqueness theorem required. Recognise this family on sight whenever a purely quadratic-in-the-unknown right side appears. See Book 1, field: Sequences and recurrences, topic: separable ODEs with rational-in-k^2 right sides.

ODE existence and uniqueness (Picard-Lindelof): an initial value problem $k' = F(x,k)$, $k(x_0) = k_0$, with F continuous and locally Lipschitz in k, has exactly one solution near x_0; this is the tool informally invoked whenever a guessed solution is checked against an initial condition. Beyond A-level: a real-analysis theorem, not part of A-level Further Mathematics, worth knowing exists even though the direct arctan-composition argument above avoids needing it. See Book 1, field: Real analysis, topic: existence and uniqueness for ODEs.

Linear first order ODEs by direct antidifferentiation: once f'/f is known explicitly as a function of x, $\ln f$ is recovered by direct antidifferentiation and f by exponentiating, the constant fixed by one initial condition; this is the integrating factor method in its simplest form, trivial because the equation is already separated. See Book 1, field: Calculus and integrals, topic: linear first order ODEs via integrating factor.

Circuit Bank

- **WHEN** a system of n equations for n functions has every nonlinear term as a multiple of the same product of all n functions, and every other term missing exactly one factor **DO** name the full product k, multiply equation i by everything except the ith function, and sum, expecting a total derivative k' on the left **BECAUSE** this is exactly the n-term expansion of the product rule, engineered to appear whenever the system has this shape
- **WHEN** after reducing a coupled system to one variable k, the question only asks for one of the original functions **DO** substitute the relation for k back into only the one original equation defining the wanted function, expressing the others as k divided by it **BECAUSE** solving for every original function separately is wasted work when only a combination of them was ever needed
- **WHEN** a separable ODE $k' = ak^2 + b$ appears with a, b the same sign **DO** compose $\arctan(k(x))$ (scaled if $b \neq 1$) and differentiate directly from the given equation rather than guessing and

invoking uniqueness **BECAUSE** direct construction is rigorous on its own and avoids relying on an unproven black-box theorem

Backtesting

1. First useful cue: every equation's polynomial term is the full product fgh times one function, every fractional term the missing function divided by fgh.
2. Method the expression suggested: multiply each equation by the complementary pair of functions to build $(fgh)'$ on the left and clear all fractions on the right.
3. Where you would hesitate, and the thought to reinforce: whether you may simply assert $k(x) = \tan(6x + \pi/4)$; reinforce that composing arctan with $k(x)$ and differentiating directly from the ODE constructs the solution, so no separate uniqueness theorem is required.
4. The trap the question sets: it looks like you need g and h explicitly to answer a question about f; you never do, since $gh = k/f$ is all the first equation needs.
5. What makes you recognise it faster next time: three or more coupled equations all sharing the same full product of unknowns, in both their polynomial and reciprocal terms.
6. The reusable decision rule: reduce coupled nonlinear systems built around a shared product by completing the product rule for that product, before attempting to isolate any single function.

Mirror drills

1. Two functions satisfy $p' = p^2q + 1/q$, $q' = pq^2 + 1/p$, $p(0) = q(0) = 1$. Find pq explicitly. *Answer:* multiply by q, p and add: $(pq)' = 2(pq)^2 + 2$, so $k = pq$ gives $k' = 2(k^2 + 1)$, $k(0) = 1$, hence $k(x) = \tan(2x + \pi/4)$ by the identical arctan-composition argument.
2. If $k' = 3(k^2 + 4)$ with $k(0) = 2$, find $k(x)$. *Answer:* write $k = 2\tan\theta$, so $k^2 + 4 = 4\sec^2\theta$ and $k' = 2\theta'\sec^2\theta$; the equation becomes $2\theta'\sec^2\theta = 12\sec^2\theta$, so $\theta' = 6$; at $x = 0$, $\tan\theta_0 = 1$ so $\theta_0 = \pi/4$; hence $k(x) = 2\tan(6x + \pi/4)$.
3. A function satisfies $y'/y = 2\tan(3x) + \cot(3x)$, $y(0) = 1$. Find $y(x)$ near 0. *Answer:* $\cot(3x) \to \infty$ as $x \to 0$, so y'/y is not even defined at $x = 0$; no such differentiable y with $y(0) = 1$ exists on an interval containing 0, showing the $\pi/4$ phase shift in the original problem, keeping $u = 6x + \pi/4$ away from 0 and $\pi/2$, was load-bearing, not decorative.

17.3 A3. cosine matrix collapses to rank two

Let d_n be the determinant of the $n \times n$ matrix whose entries, from left to right and then from top to bottom, are $\cos 1, \cos 2, \ldots, \cos n^2$ (the argument of cos always in radians). Evaluate $\lim_{n\to\infty} d_n$.

Read it out loud

Read the construction literally: row j, column k holds $\cos\big(n(j-1)+k\big)$, since $1, 2, \ldots, n^2$ are laid down row by row and row j starts right after row $j-1$ has used up $n(j-1)$ of them. Every entry is cos of an integer that splits as a row-dependent piece $n(j-1)$ plus a column-dependent piece k. That additive split, argument equals row thing plus column thing, fed through one fixed function, is the entire content of the problem; nothing about the specific numbers $1, \ldots, n^2$ matters beyond it.

The fields wired in are linear algebra, specifically determinants, rank, and the fact that a matrix with more rows than its rank is singular, and trigonometry, specifically angle addition and its cleanest packaging, the complex exponential. This is a determinant question that is not really about computing a determinant: an $n \times n$ determinant of transcendental entries cannot be computed directly, so the cue is to ask what structural property of the matrix forces it to vanish, rather than to compute it.

The first thing to write down within seconds: $\cos(a+b) = \cos a \cos b - \sin a \sin b$ turns a function of a SUM into a combination of PRODUCTS of a function of a alone and a function of b alone. A matrix whose (j,k) entry is a sum of products of row-only and column-only pieces is, by definition, a sum of outer products, hence of bounded rank regardless of n. Once n exceeds that bound, the determinant is not merely small, it is exactly zero.

Approaches, ranked worst to best

1. **cofactor expand for small n, look for a numeric pattern.** *Cue:* "evaluate a limit" suggests computing a few terms. *Circuit:* brute determinant expansion. Enormous ink even at $n = 4$ or 5, and never becomes a proof for all n; you would be guessing a pattern continues rather than showing why it must.

2. **add column 3 to column 1, use $\cos x + \cos y = 2\cos\frac{x+y}{2}\cos\frac{x-y}{2}$ to show the new column 1 is a multiple of column 2.** *Cue:* the entries are cos of a linear sequence, and sum to product identities are the standard tool for combining two cosines. *Circuit:* elementary column operations preserve determinant; proportional columns force determinant 0. Correct and reasonably short, but it requires spotting the specific pair of columns to combine, leaving the reader unsure why THAT pair works and treating a structural fact as an algebraic coincidence.

3. **write $\cos(a_j + b_k)$ as the real part of $e^{i(a_j+b_k)} = e^{ia_j}e^{ib_k}$, an outer product, hence rank 1 over $\mathbb{C}$ and rank at most 2 over $\mathbb{R}$ after taking the real part.** *Cue:* "argument is a sum of a row piece and a column piece" is the signature of an outer product once you pass to complex exponentials. *Circuit:* rank subadditivity under matrix addition; a matrix of rank less than its size is singular. The ideal route: it explains WHY the vanishing happens, a rank bound intrinsic to the additive-phase structure, nothing to do with the specific numbers, needs no case-specific column hunting, and applies uniformly to every $n \geq 3$ in one line.

The ideal solution

Index rows and columns by $j, k \in \{1, \ldots, n\}$. Since $\cos 1, \cos 2, \ldots, \cos n^2$ are written row by row, the entry in row j, column k is

$$M_{jk} = \cos\big(n(j-1)+k\big).$$

Set $a_j = n(j-1)$ and $b_k = k$, so $M_{jk} = \cos(a_j + b_k)$.

By Euler's formula, $M_{jk} = \mathrm{Re}\big(e^{i(a_j+b_k)}\big) = \mathrm{Re}\big(e^{ia_j}e^{ib_k}\big)$. Let $u \in \mathbb{C}^n$ have jth entry e^{ia_j} and $v \in \mathbb{C}^n$ have kth entry e^{ib_k}. The complex matrix uv^{T} has (j,k) entry $e^{ia_j}e^{ib_k}$ and rank exactly 1, since u, v are nonzero (every entry has modulus 1). Write $u = p + iq$, $v = r + is$ with $p, q, r, s \in \mathbb{R}^n$, so $p_j = \cos a_j$, $q_j = \sin a_j$, $r_k = \cos b_k$, $s_k = \sin b_k$. Then

$$M = \mathrm{Re}(uv^{\mathsf{T}}) = pr^{\mathsf{T}} - qs^{\mathsf{T}},$$

a difference of two real outer products, each of rank at most 1. Since rank is subadditive over sums, $\mathrm{rank}(M) \le \mathrm{rank}(pr^{\mathsf{T}}) + \mathrm{rank}(qs^{\mathsf{T}}) \le 1 + 1 = 2$.

So M has rank at most 2 for every n. For $n \ge 3$, an $n \times n$ matrix of rank at most $2 < n$ is singular, since its rows cannot all be independent once there are more of them than the rank, hence $\det M = d_n = 0$.

This holds for every $n \ge 3$, so the sequence $(d_n)_{n\ge 3}$ is identically zero; only $d_1 = \cos 1$ and $d_2 = \cos 1 \cos 4 - \cos 2 \cos 3$ can possibly be nonzero, and they play no role in a limit as $n \to \infty$. Therefore

$$\boxed{\lim_{n\to\infty} d_n = 0.}$$

Layer 1: WHAT

Block 1. Translate the row by row filling rule into a formula, $M_{jk} = \cos(a_j + b_k)$ with $a_j = n(j-1)$, $b_k = k$. Block 2. Recognise cos of a sum of a row piece and a column piece as the real part of a complex outer product, via Euler's formula. Block 3. Split into real and imaginary parts to bound the real rank by 2. Block 4. Invoke rank less than size implies singular, for all $n \ge 3$. Block 5. Note the limit is unaffected by the finitely many small n where the bound $2 < n$ fails.

The nominal route, the column-addition trick, reaches the same conclusion by constructing a specific dependency between two named columns using a sum to product identity, rather than exposing the rank bound that makes every suitably chosen triple of columns dependent once $n \ge 3$. The bifurcation point is Block 2: whether you see the matrix's defining feature as an algebraic coincidence between column 1 and column 3, or as an outer product in disguise, transparent as soon as you pass to $e^{i\theta}$.

Layer 2: HOW

Block 1. **Cue.** The problem states exactly how the matrix is filled: left to right, then top to bottom, with $1, 2, \ldots, n^2$. **Why natural.** Writing out row 1 $(\cos 1, \ldots, \cos n)$ and row 2 $(\cos(n+1), \ldots, \cos 2n)$ by hand makes the pattern $M_{jk} = \cos(n(j-1) + k)$ obvious: row j starts at $n(j-1) + 1$. **Link to the next block.** Once the entry is a formula in j and k, ask what kind of formula it is, in particular whether j and k interact additively or multiplicatively inside the cos.

Block 2. **Cue.** The argument $n(j-1) + k$ is additively separable: a function of j alone, plus a function of k alone. **Why natural.** $\cos(a + b)$ never factors as a product of a function of a and a function of b over the reals directly, but it does via $e^{i(a+b)} = e^{ia}e^{ib}$, which DOES factor exactly; since cos is the real part of $e^{i(\cdot)}$, complexifying is exactly what turns an additive split into a genuine product split. **Link to the next block.** A matrix whose complex version is a product of a column vector and a row vector is by definition a rank 1 outer product; this is the moment the determinant question quietly becomes a rank question.

Block 3. **Cue.** You need a statement about the real matrix M, but the clean rank 1 structure lives in the complex matrix uv^{T}. **Why natural.** Taking the real part of a sum is linear, so the real part of a rank 1 complex outer product is what you get by expanding $u = p + iq$, $v = r + is$ and collecting

real terms, producing exactly two real outer products added together. **Link to the next block.** Two rank 1 pieces added have rank at most 2, since rank is subadditive over sums, giving the bound actually needed.

Block 4. **Cue.** $n \geq 3$ but rank ≤ 2. **Why natural.** An $n \times n$ matrix is invertible exactly when its rank equals n; rank strictly less than n means the rows, or columns, are linearly dependent, precisely the condition for the determinant to vanish. This is definitional, not an extra theorem to prove. **Link to the next block.** The determinant is exactly 0, not merely small, for every single $n \geq 3$, a stronger and easier fact than any asymptotic estimate would have been.

Block 5. **Cue.** The problem asks for a limit as $n \to \infty$, but you have shown something about every $n \geq 3$. **Why natural.** A sequence that is exactly 0 from some point onward converges to 0 trivially; the finitely many terms d_1, d_2 before that point cannot affect a limit, since a limit depends only on the tail of a sequence. **Link to the next block.** This is why d_1 and d_2 never need to be computed at all to answer the question as posed; eventual vanishing is strictly stronger than any statement about the limit alone.

Layer 3: WHY

The deep fact is a general lemma worth extracting: if $M_{jk} = \cos(a_j + b_k)$ for ANY real sequences (a_j), (b_k), then M has real rank at most 2. Nothing about the specific values $a_j = n(j-1)$, $b_k = k$ was used anywhere; the problem's numbers $1, \ldots, n^2$ are decorative, and the only structural fact in play is the additive separability of the argument of cos. This is exactly why the column-manipulation solution works too: any two columns of M live in a 2-dimensional space spanned by $(\cos a_j)_j$ and $(\sin a_j)_j$, so any three columns, for $n \geq 3$, are automatically linearly dependent, and the official solution's choice of columns 1 and 3 is one convenient dependency among many available.

The discernment mechanism: whenever a matrix entry is a fixed function of a sum of a row-index quantity and a column-index quantity, and that function is cos, sin, $e^{(\cdot)}$, or built from these, suspect a low rank structure before attempting any determinant computation. The complex exponential is the natural lens because it is the one function where argument adds and value multiplies are the same fact, $e^{i(a+b)} = e^{ia}e^{ib}$; cos and sin inherit a rank bound of 2, not 1, because taking a real or imaginary part of a rank 1 complex object can double the real rank, never worse.

What the examiners were nudging toward: "from left to right and then from top to bottom" is doing real work, forcing the additive-in-(j,k) structure on the argument; a different filling order would break the rank bound argument directly, and the problem would likely be far harder or false. Recognising that the filling RULE, not the specific numbers, generates the structure is the single most transferable habit here.

Revolving concepts

Rank of a matrix and the outer product as the canonical rank 1 object: a matrix uv^{T} always has rank at most 1, and every rank 1 matrix has this form; rank is subadditive, $\operatorname{rank}(X+Y) \leq \operatorname{rank}(X) + \operatorname{rank}(Y)$, letting two rank 1 pieces combine to a rank ≤ 2 bound. Beyond A-level: rank as a formal invariant, and its subadditivity, is undergraduate linear algebra, not part of A-level Further Mathematics, where matrix work stops at determinants, inverses and simple eigenvalue problems. See Book 1, field: Linear algebra, topic: rank and outer products.

Euler's formula as a device for linearising angle addition: $e^{i\theta} = \cos\theta + i\sin\theta$, and $e^{i(a+b)} = e^{ia}e^{ib}$ converts the additive structure of angle sums into multiplicative structure, exactly what is needed to see a sum of angles as a product of two separate vectors' entries. This is A-level Further Mathematics content already; what is new is applying it to build a MATRIX rather than to simplify

a single trigonometric expression. See Book 1, field: Complex numbers, topic: Euler's formula as a rank-revealing substitution.

Rank deficiency and vanishing determinants, with the general growing-family template: an $n \times n$ matrix is singular exactly when its rank is less than n. Recursion generalised: whenever a family $M^{(n)}$, indexed by growing size n, is built from a fixed number r of underlying vector generators independent of n (here, two: $(\cos a_j)$ and $(\sin a_j)$), every $M^{(n)}$ with $n > r$ is automatically singular; this template reappears whenever a matrix's rows or columns are secretly built from a bounded-dimensional family of building blocks, for instance a matrix whose (j,k) entry is a fixed degree d polynomial in j and k, which has rank at most $d+1$ regardless of n. See Book 1, field: Matrices and determinants, topic: singularity from bounded generator count.

Sum to product identities: $\cos x + \cos y = 2\cos\frac{x+y}{2}\cos\frac{x-y}{2}$, the classical, pre-exponential tool the column-operation route uses, itself derivable from the same angle addition formula. See Book 1, field: Geometry, topic: sum to product trigonometric identities.

Circuit Bank

- **WHEN** a matrix entry is cos, sin, or $e^{i(\cdot)}$ of an expression splitting as a function of the row index plus a function of the column index **DO** pass to the complex exponential, write the entry as the real or imaginary part of $e^{ia_j}e^{ib_k}$, and read off a rank bound of 1 complex or 2 real **BECAUSE** additive structure inside a fixed exponential-type function is exactly multiplicative structure outside it, the definition of an outer product
- **WHEN** a determinant question grows with n and asks for a limit or eventual value, not a formula for every n **DO** check whether the matrix has bounded rank independent of n before attempting to compute the determinant directly **BECAUSE** once n exceeds that bound the determinant is not approximately zero, it is exactly zero, a stronger and cheaper fact to establish than any estimate
- **WHEN** a column-operation solution to a determinant problem works but the choice of columns looks arbitrary **DO** look for the underlying rank bound that makes many different column combinations work simultaneously **BECAUSE** seeing the general reason converts a lucky-looking identity into a transferable method

Backtesting

1. First useful cue: the matrix is filled row by row with consecutive integers, so the (j,k) entry's argument is a row piece plus a column piece.
2. Method the expression suggested: pass to $e^{i(\cdot)}$ to turn the additive split into a genuine outer product, then bound the real rank by 2.
3. Where you would hesitate, and the thought to reinforce: whether rank at most 2 really forces the determinant to be exactly 0 rather than just small; reinforce that singularity is a definitional consequence of rank less than size, with no approximation involved.
4. The trap the question sets: dressing up a rank argument as a determinant computation over specific, intimidating numbers 1 through n^2, tempting you to compute rather than look for structure.
5. What makes you recognise it faster next time: seeing "the determinant of a matrix whose entries are cos, or sin, or $e^{i(\cdot)}$, of a linear function of row and column index".
6. The reusable decision rule: before computing any growing-size determinant, check whether the entries are secretly built from a bounded number of vectors via an additive-to-multiplicative identity; if so, the determinant vanishes identically once n exceeds that bound.

Mirror drills

1. Let $N_{jk} = \sin\big((j-1) + 2k\big)$ for $j, k = 1, \ldots, 5$. What is $\det N$? *Answer:* 0; $a_j = (j-1)$, $b_k = 2k$ split additively, so N is the imaginary part of a rank 1 complex outer product, real rank at most 2, and $5 > 2$.
2. Let $M_{jk} = 2^{j+k}$ for $j, k = 1, \ldots, 4$. Without expanding, why is $\det M = 0$? *Answer:* $M_{jk} = 2^j \cdot 2^k$ is already a real outer product, rank exactly 1, and $4 > 1$.
3. Let $M_{jk} = \cos(jk)$ for $j, k = 1, \ldots, 4$ (a product, not a sum, inside the cosine). Does the rank 2 argument apply? *Answer:* No; jk does not split as a function of j plus a function of k, so e^{ijk} is not a product $e^{i\alpha_j} e^{i\beta_k}$ and the outer product structure is gone; this determinant is not forced to vanish by this method, showing the ORIGINAL problem's additive index structure was essential, not decorative.

17.4 A4. Rational closure blocked by a quadratic residue

Let S be a set of rational numbers such that

(a) $0 \in S$;

(b) if $x \in S$ *then* $x+1 \in S$ *and* $x-1 \in S$; *and*

(c) if $x \in S$ *and* $x \notin \{0,1\}$, *then* $\dfrac{1}{x(x-1)} \in S$.

Must S contain all rational numbers?

Read it out loud

This is a closure question. You are given a seed, 0, and two production rules, and you are asked whether the closure of the seed under those rules is all of $\mathbb{Q}$. Read it verbally: start at zero, you may slide by whole numbers, and you may apply the map $T(x) = 1/(x(x-1))$; does every rational get hit? The fields wired into the question are the arithmetic of reduced fractions, quadratic residues, and the general logic of least fixed points. Nothing else is present, which is itself information: there is no analysis here, no counting, no algebra beyond $\mathbb{Z}$.

The single most important structural observation comes before any computation. Rule (b) does not change the denominator of a fraction in lowest terms; it only shifts the numerator by the denominator. So rule (b) preserves both the denominator and the numerator modulo that denominator. Every change of denominator in this problem is done by rule (c), and by rule (c) alone. Thus there is exactly one object worth computing, and you should have it on paper within thirty seconds: write $x = p/q$ in lowest terms and simplify $T(p/q)$ into a single fraction. That is not a guess, it is the only rule capable of doing anything interesting, so you interrogate it.

There is also a meta cue in the phrasing. *Must S* contain all rational numbers? A *yes* costs you a constructive algorithm reaching an arbitrary p/q; a *no* costs you one explicit set and three bullet points of verification. Before committing, spend sixty seconds generating: 0 gives all integers, $x = 2$ gives $1/2$, $x = 3$ gives $1/6$, $x = 3/2$ gives $4/3$, $x = 1/6$ gives $-36/5$. Denominators $2, 3, 5, 6$ all appear, so the set is not obviously thin, but no scheme appears either. When a generation attempt produces plenty of examples but no descent, stop generating and start hunting for what the rules cannot do.

One more guiding light, and it is the strongest one in the problem. If rule (c) had been $x \mapsto 1/x$, the answer would be yes, and the proof would be the continued fraction algorithm: given p/q in lowest terms with $q \geq 2$, translate to r/q with $0 < r < q$, invert to q/r, and the denominator has strictly dropped, so induction on the denominator finishes. The setter wrote $1/(x(x-1))$ instead. That map has degree 2, not degree 1: it is not a Mobius transformation, so the transitive action of translations and inversion on $\mathbb{Q}$ is deliberately broken. When a setter breaks a structure you know to be transitive, they are telling you the answer is no.

Approaches, ranked worst to best

1. **Reverse Euclid.** *Cue:* two rules that look like "shift" and "invert", the exact generators of the continued fraction algorithm. *Circuit:* translations plus inversion act transitively on $\mathbb{Q}$, so run the Euclidean descent backwards. This is a false friend and it is expensive. The map is quadratic, so from $T(x) = y$ you cannot in general solve for a rational x: x satisfies $yx^2 - yx - 1 = 0$, whose discriminant $y^2 + 4y$ is rarely a square. The descent has no inverse step, and you can lose twenty minutes before noticing.

2. **Induction on the denominator.** *Cue:* a closure statement about $\mathbb{Q}$ with a natural complexity measure. *Circuit:* prove reachability by strong induction on the denominator of the target. It fails, but it fails informatively: try to reach p/q from something simpler and you find T sends

denominator q to numerator q^2, which is the wrong direction. Ranked above the first route only because the failure points straight at the invariant.

3. **Experimental hunt for a missing class.** *Cue:* a closure question that resists generation. *Circuit:* list a few hundred members of the orbit, sort by denominator, and look for a residue class that never appears. This does find the answer: with denominator 5 you only ever see numerators $\equiv 1$ or $4 \pmod 5$. It is legitimate and it is slow by hand, and it leaves you with a conjecture rather than a reason.

4. **Compute rule (c) on the general symbol.** *Cue:* rule (b) fixes the denominator, so only rule (c) can change it. *Circuit:* substitute $x = p/q$ in lowest terms, simplify to one fraction, and read what the output is forced to look like. The best route by a distance. One line of algebra gives $T(p/q) = q^2/(p(p-q))$ with numerator and denominator automatically coprime. The numerator of every output is a perfect square up to sign, and squares miss residue classes. From there the counterexample is forced rather than found.

The ideal solution

The answer is $\boxed{\text{No}}$.

Throughout, $T(x) = \dfrac{1}{x(x-1)}$, defined for $x \in \mathbb{Q} \setminus \{0, 1\}$.

Lemma 17.4.1. *A rational number has exactly two representations a/b with $\gcd(a,b) = 1$ and $b \neq 0$, namely a/b and $(-a)/(-b)$. In particular, if $a/b = c/d$ with $\gcd(a,b) = \gcd(c,d) = 1$ and $b, d \neq 0$, then $(c,d) = \pm(a,b)$.*

Proof. From $a/b = c/d$ we get $ad = bc$. Then $b \mid ad$, and $\gcd(a,b) = 1$, so $b \mid d$. Symmetrically $d \mid b$, so $d = \varepsilon b$ with $\varepsilon = \pm 1$. Cancelling b from $ad = bc$ gives $c = \varepsilon a$ with the same sign. □

Lemma 17.4.2. *Let $x \in \mathbb{Q} \setminus \{0, 1\}$ and write $x = p/q$ in lowest terms with $q \geq 1$. Put $D = p(p-q)$. Then $D \neq 0$, $\gcd(q^2, D) = 1$, and*

$$T(x) = \frac{q^2}{D}.$$

Proof. We have $x(x-1) = \dfrac{p}{q} \cdot \dfrac{p-q}{q} = \dfrac{p(p-q)}{q^2}$, so $T(x) = q^2/D$. Since $x \neq 0$ we have $p \neq 0$, and since $x \neq 1$ we have $p \neq q$, so $D \neq 0$. Finally $\gcd(p,q) = 1$ and $\gcd(p-q,q) = \gcd(p,q) = 1$, so $\gcd(D,q) = 1$ and therefore $\gcd(D, q^2) = 1$. □

Claim 17.4.3. *For every $x \in \mathbb{Q} \setminus \{0, 1\}$, the number $T(x)$ is not of the form $n + \frac{2}{5}$ with $n \in \mathbb{Z}$.*

Proof. Suppose $T(x) = n + \frac{2}{5} = \dfrac{5n+2}{5}$ for some integer n. With the notation of the previous lemma this says

$$\frac{q^2}{D} = \frac{5n+2}{5}.$$

The left side is in lowest terms, and so is the right side, because $\gcd(5n+2, 5) = \gcd(2, 5) = 1$. By the first lemma, $(5n+2, 5) = \pm(q^2, D)$, and in particular

$$5n + 2 = \pm q^2.$$

Reduce modulo 5: the left side is $\equiv 2$, so $q^2 \equiv \pm 2 \pmod 5$, that is, $q^2 \equiv 2$ or $q^2 \equiv 3 \pmod 5$. But every square is congruent to 0, 1 or 4 modulo 5, since $0^2, 1^2, 2^2, 3^2, 4^2 \equiv 0, 1, 4, 4, 1$. This is a contradiction, so no such x exists. □

Now put

$$B = \left\{ n + \tfrac{2}{5} : n \in \mathbb{Z} \right\}, \qquad S = \mathbb{Q} \setminus B.$$

Property (a): $0 \notin B$, since $0 - \frac{2}{5}$ is not an integer, so $0 \in S$. Property (b): B is by construction invariant under $x \mapsto x \pm 1$, hence so is its complement S. Property (c): let $x \in S$ with $x \notin \{0, 1\}$. By the Claim, $T(x) \notin B$, so $T(x) \in S$. (The Claim is stronger than required: no rational at all, whether or not it lies in S, is sent into B.)

So S satisfies (a), (b) and (c), and $\frac{2}{5} \notin S$. A set obeying the three conditions need not contain all rationals.

Remark 17.4.4. The same argument excludes $\frac{3}{5}$ as well, so $\mathbb{Q} \setminus \{n \pm \frac{2}{5}\}$ also works, and more generally: if $\gcd(m, n) = 1$ and neither m nor $-m$ is a quadratic residue modulo n, then $\mathbb{Q}$ minus the class of m/n satisfies (a), (b), (c). The proof is the same two lemmas, ending at $m + kn = \pm q^2$.

Layer 1: WHAT

Block 1. Read the two rules and notice that (b) preserves the denominator and the numerator modulo the denominator, so only (c) can move you between denominators.

Block 2. Spend sixty seconds generating; get $1/2$, $1/6$, $4/3$, $-36/5$; find no descent; switch from "reach everything" to "find what cannot be reached".

Block 3. Compute (c) on the general symbol: $T(p/q) = q^2/(p(p-q))$.

Block 4. Read the two gifts in that output: the fraction is already in lowest terms, and its numerator is a perfect square.

Block 5. Convert "numerator is a square" into an obstruction: fix a denominator n and ask which numerators are attainable modulo n; only $\pm$ squares are.

Block 6. Choose the smallest modulus for which some class is missed by both $+$ squares and $-$ squares. That is $n = 5$, class $m = 2$.

Block 7. Check the resulting bad set is translation invariant and misses 0, so (a) and (b) survive, and quote the Claim for (c).

The nominal route is Block 2 continued forever: generate, tabulate, hope for a pattern in the numbers. The ideal route is Blocks 3 to 7. The bifurcation point is exactly the line where you stop substituting numbers into T and substitute the *symbol* p/q instead. The zombie computes T on examples and reads decimals; you compute T once on the general reduced fraction and read the shape of the answer. Everything after that is forced.

Layer 2: HOW

Block 1. Cue. Two closure rules, one of which is a pure translation. **Why natural.** Translation by 1 sends p/q to $(p+q)/q$, which is still in lowest terms with the same denominator. So the pair (denominator, numerator modulo denominator) is an invariant of rule (b). A rule that preserves everything you can measure is not the rule to study. **Link to the next block.** If (b) cannot change the denominator, then the entire question is about the image of (c).

Block 2. Cue. The question says "must", which is an invitation to answer no. **Why natural.** Generating is cheap and settles the mood of the problem. You are not looking for the answer here, you are looking for evidence of a descent. Sixty seconds produced four denominators and no algorithm, which is the signature of a false statement. **Link to the next block.** With no descent in sight, the only remaining currency is an invariant, and an invariant must come from the general form of the rule.

Block 3. Cue. Rules stated on rationals, with the target also rational. **Why natural.** Every arithmetic property of a rational lives in its lowest terms representation, so you substitute $x = p/q$ with $\gcd(p, q) = 1$ and simplify to a single fraction. This is not exploratory; it is the definition-level move. The rule is a function of x, and you want to know what it does to the data (p, q) that carries all arithmetic meaning. **Link to the next block.** The algebra takes one line and produces $q^2/(p(p-q))$, which you then read rather than manipulate further.

Block 4. Cue. The output has a squared numerator and a product denominator. **Why natural.** Two facts fall out for free. First, $\gcd(p, q) = 1$ forces $\gcd(p - q, q) = 1$, so the output fraction is already reduced: no information is lost by reading its numerator. Second, that numerator is q^2. A perfect square is the most restrictive shape in elementary number theory, because squares occupy about half the classes modulo an odd prime. **Link to the next block.** The reduced form being unique means the numerator modulo the denominator is a well defined feature of the output, so "which classes can the output occupy" is a legitimate question with a finite answer.

Block 5. Cue. A quantity forced to be a square. **Why natural.** The circuit is automatic: squares miss residues, so ask which residues. Because the denominator D may be negative, moving to the standard positive denominator flips the sign of the numerator, so what is forced is that the numerator is $\pm$ a square modulo the denominator. Both signs must be excluded. **Link to the next block.** You now need a modulus n and a class m with neither m nor $-m$ a square modulo n.

Block 6. Cue. You need m and $-m$ to be non residues simultaneously. **Why natural.** That is possible only if -1 is itself a square modulo n, since otherwise multiplying by -1 swaps residues and non residues and one of $\pm m$ always lands in the residue class. For an odd prime p, -1 is a square exactly when $p \equiv 1 \pmod 4$. So the smallest candidate modulus is 5, where $-1 \equiv 4 = 2^2$ is a square and the non residues are 2 and 3. The smallest bad class is $2/5$. **Link to the next block.** The bad set is a class modulo 1, hence automatically translation invariant, which is precisely what rule (b) demands.

Block 7. Cue. The write up must verify three conditions, not tell a story. **Why natural.** A full marks answer to a "must" question is an object plus three checks. The only check with content is (c), and you prove the stronger statement that no rational maps into the bad class, which removes any need to track membership of S during the argument. **Link to the next block.** Nothing remains; state the conclusion.

Layer 3: WHY

The mathematical truth at the bottom is a fixed point statement. The set of rationals actually forced into S is the least fixed point of the operator $X \mapsto \{0\} \cup (X + 1) \cup (X - 1) \cup T(X \setminus \{0, 1\})$, obtained as the union of the iterates started at $\emptyset$. Least fixed points are hard to describe from below and easy to bound from above: any set that contains the seed and is closed under the rules contains the closure. So to answer no you never need to know what S must contain. You need one closed proper superset. That asymmetry is why the counterexample route is cheap and the "prove it is everything" route is expensive, and it is the reason to test for an invariant early.

The second truth is about which invariant to look for. Rule (b) is a translation, so any invariant must be a function of the rational modulo 1, that is, of the pair (reduced denominator, reduced numerator modulo denominator). That reduces the search space from all conceivable invariants to a finite question for each denominator. Then rule (c) supplies the answer: it manufactures numerators that are squares. Squares are the standard source of unattainable classes, and quadratic residues are the standard language for saying so.

The third truth is the sign issue, which is where a careless solution dies. $D = p(p-q)$ can be negative, so the reduced representation with a positive denominator carries numerator $-q^2$. That is exactly why you need -1 to be a square modulo n, and it is why $n = 3$ and $n = 4$ fail: modulo 3 the non residue is 2 and $-2 \equiv 1$ is a residue; modulo 4 the non residue among units is 3 and $-3 \equiv 1$ is a residue. Five is the first modulus that works, and the answer $2/5$ published by the competition is the smallest possible counterexample class. Knowing that means you would have found it in the arena without searching.

The discernment mechanism to train is a two question filter for closure problems. First: does any rule reduce a complexity measure, so that a descent proof is available? Here the candidate measure is the denominator, and T squares it into the numerator, so no. Second: does any rule force the output into a special shape? Here it forces a square, so yes. If the answer to the first is no and to the second is yes, the answer to the problem is no, and the shape tells you the invariant. The examiners nudged with the choice of $1/(x(x-1))$ over $1/x$: they took a transitive pair of generators and replaced one of them with a quadratic map, and the square in q^2 is the fingerprint of that replacement.

Revolving concepts

Lowest terms. A rational r has a representation p/q with $\gcd(p,q) = 1$ and $q \geq 1$, and it is unique. Uniqueness is the load bearing fact: it turns an equation between two fractions into an equation between two integer pairs, which is what converts the analytic looking statement $T(x) = n + 2/5$ into the arithmetic statement $5n + 2 = \pm q^2$. See Book 1, field: Number theory, topic: reduced fractions and Euclid's lemma.

Coprimality under shifting. $\gcd(p-q,q) = \gcd(p,q)$, because the ideal generated by $\{p,q\}$ equals the one generated by $\{p-q,q\}$. This is the same one line fact that makes the Euclidean algorithm work, and here it is what makes T output an already reduced fraction. See Book 1, field: Number theory, topic: gcd identities.

Quadratic residues. For an odd prime p, exactly $(p-1)/2$ of the nonzero classes modulo p are squares. Euler's criterion says $a^{(p-1)/2} \equiv \pm 1 \pmod{p}$ according as a is or is not a residue, and the Legendre symbol is multiplicative. The first supplement, $(-1/p) = 1$ exactly when $p \equiv 1 \pmod 4$, is the fact that selects 5 over 3 here. The second supplement, $(2/p) = 1$ exactly when $p \equiv \pm 1 \pmod 8$, is the natural next fact to install. See Book 1, field: Number theory, topic: quadratic residues.

Closure as recursion. The set generated by a seed under a family of maps is a recursion: $S_0 = \{0\}$ and $S_{k+1} = S_k \cup (S_k + 1) \cup (S_k - 1) \cup T(S_k \setminus \{0,1\})$, with $S = \bigcup_k S_k$. Generalised: for any seed A and any family of partial maps $\{f_i\}$ on a set X, the generated set is the least fixed point of $X \mapsto A \cup \bigcup_i f_i(X)$, computed as the union of the iterates. There are exactly two proof techniques for it. To prove the closure is everything, exhibit a complexity measure strictly decreased by some composite of the maps and induct. To prove it is not everything, exhibit one set containing A and closed under every f_i. Every problem of this family, including the reachability problems on lattices, on the integers under two affine maps, and on states of a token game, is attacked with this same pair. See Book 1, field: Logic and set theory, topic: least fixed points and invariants.

Transitivity of translation and inversion. The maps $x \mapsto x+1$ and $x \mapsto 1/x$ generate a group acting transitively on $\mathbb{Q} \cup \{\infty\}$, which is the group theoretic content of the continued fraction algorithm. Recognising when a problem is and is not that situation is worth more than any single computation here. See Book 1, field: Number theory, topic: continued fractions and $\mathrm{SL}_2(\mathbb{Z})$.

Circuit Bank

- **WHEN** a question asks whether a closure of a seed under given rules is everything **DO** spend sixty seconds hunting a descent, and if no rule reduces a complexity measure, switch to

constructing one closed proper superset **BECAUSE** the closure is a least fixed point, so a single closed superset settles the question, while a yes needs a full algorithm

- **WHEN** a rule acts on rationals **DO** substitute $x = p/q$ in lowest terms, simplify to a single fraction, and check whether the output is automatically reduced **BECAUSE** all arithmetic content of a rational lives in its reduced pair, and coprimality of the output is what makes its numerator modulo its denominator well defined
- **WHEN** a numerator or denominator is forced to be a perfect square **DO** ask which residue classes are then unattainable, that is, reach for quadratic residues **BECAUSE** squares occupy about half the classes modulo an odd prime, so they are the cheapest source of an impossibility
- **WHEN** you need a class m modulo n with both m and $-m$ non residues **DO** take $n \equiv 1 \pmod 4$, the smallest being $n = 5$ with $m = 2$ **BECAUSE** -1 is a square modulo n exactly then, so residue status is preserved by negation and one non residue does both jobs
- **WHEN** a problem uses a map that resembles a Mobius map but is not one **DO** check its degree before assuming transitivity **BECAUSE** translations together with $x \mapsto 1/x$ act transitively on $\mathbb{Q}$, so any departure from degree one is the setter switching that transitivity off

Backtesting

1. First useful cue: rule (b) is a translation, so it cannot change the denominator; therefore rule (c) carries the entire question.
2. Method the expression suggested: substitute $x = p/q$ in lowest terms into $1/(x(x-1))$ and read the output, which is $q^2/(p(p-q))$ with a squared numerator and automatically coprime parts.
3. Where you would hesitate, and the thought to reinforce: at the sign of D. Reinforce that a negative denominator flips the numerator's sign, hence both m and $-m$ must be non residues, hence the modulus must have -1 as a square, hence 5 and not 3.
4. The trap the question sets: the resemblance to the continued fraction generators, which pulls you into an hour of attempted descent. The map is quadratic; there is no inverse step.
5. What makes you recognise it faster next time: a "must it be everything" closure question in which one rule is a translation. That combination means the invariant is a class modulo 1, and the only work is to find which classes the other rule can output.
6. The reusable decision rule: to disprove a closure claim, do not track what is generated; compute the general form of the generating map and read off which shapes it can never produce.

Mirror drills

1. Let $S \subseteq \mathbb{Q}$ contain 0 and be closed under $x \mapsto x \pm 1$ and under $x \mapsto 1/x$ for $x \neq 0$. Show $S = \mathbb{Q}$. *Answer:* induct on the reduced denominator q. If $q \geq 2$, translate p/q to r/q with $0 < r < q$, then invert to q/r, whose denominator r is smaller; the base case $q = 1$ is the integers, all reachable from 0.
2. Show that no element of S in the problem above need have reduced form $p/5$ with $p \equiv \pm 2 \pmod 5$. *Answer:* the same claim, with $5n+2$ replaced by $5n \pm 2$; both force $q^2 \equiv \pm 2 \pmod 5$, and $2, 3$ are not squares modulo 5.
3. For which odd primes p is the congruence $x^2 \equiv 2 \pmod p$ solvable? *Answer:* exactly for $p \equiv \pm 1 \pmod 8$, by the second supplement to quadratic reciprocity.

17.5 A5. Product of all element orders in an abelian group

Is there a finite abelian group G such that the product of the orders of all its elements is 2^{2009}?

Read it out loud

A finite abelian group is handed to you with no other structure, and you are asked about a single global statistic: the product of $\mathrm{ord}(g)$ over every $g \in G$. Read the structure verbally: a multiplicative statistic of a finite abelian group, and a specific target value that is a power of 2. The fields wired in are the classification of finite abelian groups, elementary counting of elements by order, and the arithmetic of congruences. What is being tested is whether you convert a product into a sum, whether you know how to sum a function of $\mathrm{ord}(g)$ over a group, and whether you look for a congruence obstruction before attempting an exact solve.

The first two moves should be on paper within thirty seconds. First, take the 2-adic valuation: write $\prod_{g\in G} \mathrm{ord}(g) = 2^{k(G)}$ where $k(G) = \sum_{g\in G} v_2(\mathrm{ord}\, g)$, and the question becomes: is $k(G) = 2009$ possible? Products of orders are unmanageable; sums of valuations are additive over any partition of G, and partitions are the only tool you have. Second, notice that each $\mathrm{ord}(g)$ is one of the factors of the product 2^{2009}, so each $\mathrm{ord}(g)$ divides 2^{2009}. Every element has 2-power order for free, without invoking anything.

Then compute the statistic on the smallest examples, because a target of 2009 is far too large to reach by construction and the only realistic outcome is an obstruction. The trivial group gives $k = 0$; $\mathbb{Z}/2$ gives $k = 1$; $(\mathbb{Z}/2)^2$ gives $k = 3$; $\mathbb{Z}/4$ gives $k = 0 + 2 + 1 + 2 = 5$; $(\mathbb{Z}/2)^3$ gives $k = 7$; $\mathbb{Z}/2 \times \mathbb{Z}/4$ gives $k = 11$; $\mathbb{Z}/8$ gives $k = 17$. The list $0, 1, 3, 5, 7, 11, 17, \ldots$ has a visible hole: no value is congruent to 2 modulo 4. That is the shape of the answer. Note also that k is not additive across direct products, since $k(\mathbb{Z}/2 \times \mathbb{Z}/4) = 11 \neq 1 + 5$, which kills any attempt to build a group factor by factor.

Finally, read the year as data. $2009 = 7^2 \cdot 41$, and more to the point $2009 \equiv 1 \pmod 4$. The setter chose a target in the residue class that forces the tightest branch of the argument, and you should compute that residue before anything else.

Approaches, ranked worst to best

1. **Construct a group by assembling cyclic factors.** *Cue:* the question says "is there", which invites a construction. *Circuit:* choose exponents $e_1, e_2, \ldots$ and tune them until the statistic hits the target. The worst route, because the statistic is not additive over direct products: adding one $\mathbb{Z}/2$ factor doubles the number of elements and changes every order count. The search space has no local structure to exploit and you learn nothing from a failed attempt.

2. **Tabulate small cases and conjecture.** *Cue:* a statistic with an unfamiliar definition. *Circuit:* compute $k(G)$ for every abelian 2-group of small order and look at the list of attainable values. Genuinely useful and correctly ranked in the middle: it produces the observation that $k \equiv 2 \pmod 4$ never occurs and points you at the modulus 4. It is data, not proof, and by hand it costs several minutes.

3. **Structure theorem plus a filtration count.** *Cue:* "finite abelian" is a classification cue. *Circuit:* write G as a product of $(\mathbb{Z}/2^i)^{e_i}$, count elements of order at most 2^m as 2^{s_m} with $s_m = \sum_i \min(i, m)e_i$, and sum $k(G) = \sum_i i(2^{s_i} - 2^{s_{i-1}})$. Correct, and it is the published route. It costs a page of index juggling and delivers divisibility of $k(G) + 1$ by 2^{s_1}, which is more than the problem needs. Ranked below the best route purely on ink.

4. **Partition by generated cyclic subgroup, then read modulo 4.** *Cue:* you are summing a function of $\mathrm{ord}(g)$ over a group. *Circuit:* elements generating the same cyclic subgroup have the same order, so $\sum_g F(\mathrm{ord}\, g) = \sum_C \varphi(|C|)F(|C|)$ over cyclic subgroups C. The best route.

Each cyclic subgroup of order 2^h contributes $h\varphi(2^h) = h2^{h-1}$, which is divisible by 4 for every $h \geq 2$. So modulo 4 only the subgroups of order 2 survive, the count collapses to a single number, and that number is $2^e - 1$ because the elements of order dividing 2 form a vector space over $\mathbb{F}_2$. Four lines to the obstruction.

The ideal solution

The answer is $\boxed{\text{No}}$.

Suppose G is a finite abelian group with $\prod_{g \in G} \operatorname{ord}(g) = 2^{2009}$.

Step 1: every element has 2-power order. Each $\operatorname{ord}(g)$ is one of the factors in the product, and all factors are positive integers, so $\operatorname{ord}(g)$ divides 2^{2009} and is therefore a power of 2. Consequently $\prod_g \operatorname{ord}(g) = 2^{k(G)}$ with

$$k(G) = \sum_{g \in G} v_2(\operatorname{ord} g),$$

and G is a finite abelian 2-group. We must show $k(G) = 2009$ is impossible.

Step 2: sum over cyclic subgroups. Declare $g \sim h$ when $\langle g \rangle = \langle h \rangle$. This is an equivalence relation, and the class of g is exactly the set of generators of the cyclic group $\langle g \rangle$. A cyclic group of order 2^h has $\varphi(2^h)$ generators, each of order 2^h; for $h \geq 1$ that is 2^{h-1} generators, and for $h = 0$ the class is the identity, contributing 0. Writing c_h for the number of cyclic subgroups of G of order 2^h,

$$k(G) = \sum_{h \geq 1} h\, \varphi(2^h)\, c_h = \sum_{h \geq 1} h\, 2^{h-1} c_h.$$

Step 3: reduce modulo 4. For $h = 2$ the coefficient is $2 \cdot 2 = 4$, and for $h \geq 3$ it is $h2^{h-1}$ with 2^{h-1} divisible by 4. So every term with $h \geq 2$ vanishes modulo 4, and

$$k(G) \equiv c_1 \pmod 4.$$

Since a subgroup of order 2 has exactly one generator, c_1 is the number of elements of G of order exactly 2.

Step 4: count the elements of order 2. The set $G[2] = \{g \in G : g^2 = e\}$ is a subgroup because G is abelian, and every non identity element of it has order 2. Hence $G[2]$ is a vector space over $\mathbb{F}_2$, so $|G[2]| = 2^e$ for some integer $e \geq 0$, and

$$c_1 = 2^e - 1.$$

Step 5: the congruence forces $e = 1$. We need $k(G) = 2009 \equiv 1 \pmod 4$, so $2^e - 1 \equiv 1 \pmod 4$, that is $2^e \equiv 2 \pmod 4$. If $e = 0$ then $2^e - 1 = 0 \equiv 0$; if $e \geq 2$ then $2^e \equiv 0$ and $2^e - 1 \equiv 3$. So $e = 1$.

Step 6: $e = 1$ forces G cyclic. By the structure theorem for finite abelian groups, $G \cong \prod_{j=1}^{r} \mathbb{Z}/2^{a_j}\mathbb{Z}$ with each $a_j \geq 1$. Each factor contributes exactly one element of order 2 to $G[2]$ together with the identity, so $G[2] \cong (\mathbb{F}_2)^r$ and $e = r$. Hence $r = 1$ and $G \cong \mathbb{Z}/2^i\mathbb{Z}$ for some $i \geq 1$.

Step 7: evaluate k on a cyclic 2-group. A cyclic group of order 2^i has exactly one subgroup of each order 2^h with $0 \leq h \leq i$, so $c_h = 1$ for those h and $c_h = 0$ otherwise. Therefore

$$k(\mathbb{Z}/2^i\mathbb{Z}) = \sum_{h=1}^{i} h\, 2^{h-1} = (i-1)2^i + 1.$$

The closed form follows by induction: it holds for $i = 1$ since both sides equal 1, and

$$\big((i-1)2^i + 1\big) + (i+1)2^i = 2i \cdot 2^i + 1 = i \cdot 2^{i+1} + 1,$$

which is the formula for $i + 1$.

Step 8: solve. We need $(i-1)2^i + 1 = 2009$, that is

$$(i-1)2^i = 2008 = 2^3 \cdot 251,$$

with 251 prime. Taking 2-adic valuations, $v_2\big((i-1)2^i\big) = i + v_2(i-1) \geq i$, while $v_2(2008) = 3$; hence $i \leq 3$. But for $i = 1, 2, 3$ the left side is 0, 4, 16 respectively, none of which is 2008. Contradiction.

Therefore no finite abelian group has the product of the orders of its elements equal to 2^{2009}.

Remark 17.5.1. The same computation says more. Since $k(G) \equiv 2^e - 1 \pmod 4$, the value $k(G)$ is never congruent to 2 modulo 4, so 2^n with $n \equiv 2 \pmod 4$ is never such a product either. And when $n \equiv 1 \pmod 4$ the group is forced to be cyclic, so the only attainable such values are $n = (i-1)2^i + 1$, namely $1, 5, 17, 49, 129, 321, 769, 1793, 4097, \ldots$, and 2009 sits in the gap between 1793 and 4097.

Layer 1: WHAT

Block 1. Take v_2 of the target: the multiplicative question becomes $k(G) = \sum_g v_2(\operatorname{ord} g) = 2009$.

Block 2. Each order divides the product, so every element has 2-power order; no classification theorem is needed for this step.

Block 3. Compute k on the smallest groups and notice that no value is 2 mod 4; the obstruction will be a congruence modulo 4.

Block 4. Choose the summation technique: partition G by the cyclic subgroup each element generates, since ord is constant on those classes and the class sizes are φ values.

Block 5. Reduce modulo 4: every cyclic subgroup of order at least 4 contributes a multiple of 4, so only the elements of order 2 survive.

Block 6. Count those: $G[2]$ is an $\mathbb{F}_2$ vector space, so there are $2^e - 1$ of them.

Block 7. $2009 \equiv 1 \pmod 4$ forces $e = 1$, hence G cyclic by the structure theorem.

Block 8. Evaluate $k(\mathbb{Z}/2^i) = (i-1)2^i + 1$ and kill $(i-1)2^i = 2008$ with a valuation count.

The nominal route is Block 1, then straight to the structure theorem: write down the exponents e_i, derive the exact formula for $k(G)$ in terms of them, and try to solve $k = 2009$ over all admissible exponent vectors. That is a genuine solution and it is heavy. The bifurcation point is Block 3 into Block 5: the decision to compute the statistic modulo 4 rather than exactly. The zombie tries to solve an equation; you first ask which residues the left hand side can occupy, and the equation dissolves into a single very restricted case.

Layer 2: HOW

Block 1. Cue. A product over all elements of a group, equal to a prime power. **Why natural.** Products over large index sets carry no leverage, because nothing you know about groups is multiplicative in that way. Valuations turn products into sums, and sums split along any partition of the index set. Since the target is a power of 2, the valuation to take is v_2. **Link to the next block.** Once the object is a sum of $v_2(\operatorname{ord} g)$, the next question is automatically "how do I sum a function of the order over a group".

Block 2. Cue. The target 2^{2009} has only 2 as a prime factor. **Why natural.** Each $\operatorname{ord}(g)$ appears as a factor of the product, and a positive integer factor of 2^{2009} is a power of 2. This is one line and it removes all odd primes without Cauchy's theorem, without the structure theorem, and without the phrase "we may assume". **Link to the next block.** Now every order is 2^h, so the sum $k(G)$ is a sum over cyclic 2-subgroups waiting to be organised.

Block 3. Cue. A target of 2009, far beyond anything you can construct by hand. **Why natural.** When a target is large and the objects are rigid, the realistic answers are "no, by congruence" or "yes, by an explicit family". Small case data distinguishes the two in two minutes. The values $0, 1, 3, 5, 7, 11, 17$ have a gap at 2 mod 4, which names the modulus before you have proved anything.

Link to the next block. You now know the target modulus is 4, so you need a summation identity precise enough to be read modulo 4.

Block 4. Cue. A sum of a function of $\operatorname{ord}(g)$ over all g in a finite abelian group. **Why natural.** The order is constant on the set of generators of a fixed cyclic subgroup, and those sets partition G: every g generates exactly one cyclic subgroup. That gives $\sum_g F(\operatorname{ord} g) = \sum_C \varphi(|C|)F(|C|)$, the same partition that proves $\sum_{d|n} \varphi(d) = n$. This is the standard organisation of any order weighted sum, so it is a reflex, not an inspiration. **Link to the next block.** The coefficients $h\varphi(2^h) = h2^{h-1}$ are now explicit, so their residues modulo 4 can be read off.

Block 5. Cue. Coefficients $h2^{h-1}$ with a factor of 2^{h-1} growing. **Why natural.** Modulo 4, the terms $h \geq 3$ die because $4 \mid 2^{h-1}$, and $h = 2$ dies because $2 \cdot 2 = 4$. Only $h = 1$ survives, with coefficient 1. That is why 4 is the correct modulus and not 2: modulo 2 the same computation gives only "k is odd", which 2009 satisfies, so no obstruction. Modulo 8 the terms $h = 2, 3$ both survive and the identity becomes messier. Four is the unique modulus at which the sum collapses to one count. **Link to the next block.** The survivor is the number of elements of order 2, so that number must now be computed.

Block 6. Cue. The set of elements killed by 2, in an abelian group. **Why natural.** Because G is abelian, $\{g : g^2 = e\}$ is closed under the group operation, and every element of it satisfies $g^2 = e$, which is exactly the axiom that makes it a vector space over $\mathbb{F}_2$ with scalar multiplication by 0 and 1. Vector spaces over $\mathbb{F}_2$ have 2^e elements. No classification is needed here either. **Link to the next block.** The congruence now reads $k(G) \equiv 2^e - 1 \pmod 4$, and the target's residue class selects e.

Block 7. Cue. $2009 \equiv 1 \pmod 4$ and $2^e - 1 \in \{0, 1, 3\} \bmod 4$. **Why natural.** The three cases $e = 0$, $e = 1$, $e \geq 2$ give 0, 1, 3 respectively, so the target residue 1 pins $e = 1$ exactly. A group with a unique subgroup of order 2 has exactly one cyclic factor, which is where the structure theorem finally earns its place. **Link to the next block.** One family remains, and it is a one parameter family, so the problem is now an equation in a single integer.

Block 8. Cue. A one parameter equation $(i - 1)2^i = 2008$ with a large power of 2 on the left. **Why natural.** Do not test values of i; count valuations. The left side has $v_2 \geq i$ and the right side has $v_2 = 3$, so $i \leq 3$, and three arithmetic checks finish. Reading the equation 2-adically is faster than tabulating and it is immune to arithmetic slips. **Link to the next block.** Nothing remains; the contradiction is complete.

Layer 3: WHY

The first truth is that statistics summed over a group should be summed over the classes of an equivalence relation on which the summand is constant. Here ord is constant on the generator set of each cyclic subgroup, and those generator sets partition G exactly. That single identity,

$$\sum_{g \in G} F(\operatorname{ord} g) = \sum_d \varphi(d) F(d) c_d,$$

where c_d counts cyclic subgroups of order d, is the engine of the whole solution, and it is the same identity as $\sum_{d|n} \varphi(d) = n$ read backwards. Once you have it, the problem is no longer about groups; it is about a weighted count.

The second truth is that existence questions with an arithmetic target are killed by congruences, and the correct modulus is the one that annihilates every term but one. That is a search you can carry out consciously: the coefficient of the h-th family is $h2^{h-1}$, so ask for the largest modulus

dividing $h2^{h-1}$ for all $h \geq 2$ but not for $h = 1$. The answer is 4. Choosing a modulus by inspecting the coefficients is a repeatable procedure, not luck.

The third truth is that the residue class of the target does real work. Modulo 4 the attainable residues of k are 0, 1 and 3, corresponding to the trivial group, the groups with a unique involution, and everything else. A target congruent to 2 dies instantly. A target congruent to 3 leaves a huge family open and the problem would need a different idea. A target congruent to 1, which is what 2009 is, collapses the family to the cyclic groups, where an exact formula is available. The setter picked the year and then picked the residue class that makes the problem solvable in a page, and reading 2009 mod 4 in the first thirty seconds tells you which of the three worlds you are in.

The discernment to train is the pair of habits: convert products to sums with a valuation, and never solve an exact equation for a large target before checking small moduli. In this problem both habits are visible in the write up as a single line each, which is exactly why the solution looks short. What makes the route ideal is not that it avoids the structure theorem, since it uses it at Step 6; it is that it postpones the structure theorem until only one parameter is left, so the classification is applied to a question with a one line answer instead of to a search over exponent vectors.

Revolving concepts

Order of an element and p-adic valuation. $\operatorname{ord}(g)$ is the least $n \geq 1$ with $g^n = e$, and $g^m = e$ if and only if $\operatorname{ord}(g) \mid m$. The valuation $v_p(n)$ is the exponent of p in n; v_p turns products into sums, which is the only reason this problem is tractable. See Book 1, field: Abstract algebra, topic: element orders.

Generators of a cyclic group. A cyclic group of order n has $\varphi(n)$ generators, and exactly one subgroup of each order dividing n. Partitioning G by generated cyclic subgroup and counting gives $\sum_{d \mid n} \varphi(d) = n$ for cyclic G, and the weighted version used above in general. See Book 1, field: Abstract algebra, topic: cyclic groups and Euler's totient.

Elementary abelian groups. An abelian group in which every element satisfies $g^p = e$ is a vector space over $\mathbb{F}_p$, so it has p^e elements and $p^e - 1$ elements of order exactly p. This converts a counting question into a dimension. See Book 1, field: Abstract algebra, topic: elementary abelian p-groups.

Structure theorem for finite abelian groups. Every finite abelian group is a direct product of cyclic groups of prime power order, uniquely up to reordering. Consequence used here: the number of cyclic factors of a finite abelian 2-group equals $\dim_{\mathbb{F}_2} G[2]$, so a unique involution forces cyclicity. This is the one genuinely beyond A-level input, and it must be installed before any Putnam group theory problem. See Book 1, field: Abstract algebra, topic: classification of finite abelian groups.

The recursion behind $k(\mathbb{Z}/2^i)$, generalised. The counts satisfy $k_i = k_{i-1} + i2^{i-1}$ with $k_0 = 0$, a first order linear recursion with an inhomogeneous term of the form ix^{i-1}. The general solution method is to sum the inhomogeneity directly using

$$\sum_{h=1}^{i} hx^{h-1} = \frac{d}{dx}\left(\frac{x^{i+1}-1}{x-1}\right) = \frac{ix^{i+1} - (i+1)x^i + 1}{(x-1)^2},$$

which at $x = 2$ gives $(i-1)2^i + 1$. The same differentiate the geometric series move solves every $\sum hx^h$ and $\sum h^2x^h$ that appears in generating function work. For a general prime the same partition gives

$$k_p(\mathbb{Z}/p^i) = \sum_{h=1}^{i} h\left(p^h - p^{h-1}\right) = ip^i - \frac{p^i - 1}{p-1},$$

by Abel summation, and the case $p = 2$ is the formula used above. See Book 1, field: Sequences and recurrences, topic: summing $\sum hx^h$ by differentiating a geometric series.

Circuit Bank

- **WHEN** a problem gives a product over all elements of a group **DO** take the p-adic valuation to convert it into a sum, then sum over classes of an equivalence on which the summand is

constant **BECAUSE** products over large index sets have no structure, while sums split along any partition, and the natural partition here is by generated subgroup

- **WHEN** you must evaluate $\sum_{g \in G} F(\operatorname{ord} g)$ **DO** partition by generated cyclic subgroup to obtain $\sum_d \varphi(d) F(d) c_d$ with c_d the number of cyclic subgroups of order d **BECAUSE** the order is constant on generator sets, the generator sets partition the group, and a cyclic group of order d has $\varphi(d)$ generators
- **WHEN** you must show a specific large value of a statistic is unattainable **DO** compute the statistic modulo the smallest modulus that annihilates every family of terms but one **BECAUSE** existence questions with arithmetic targets die to congruences, and the right modulus is visible in the coefficients rather than guessed
- **WHEN** a finite abelian group has all elements killed by a prime p **DO** treat it as a vector space over $\mathbb{F}_p$ **BECAUSE** the group axioms plus $g^p = e$ are exactly the axioms of an $\mathbb{F}_p$ vector space, so the size is p^e and the count of elements of order p is $p^e - 1$
- **WHEN** the target number in a problem is the year **DO** compute its factorisation and its residues modulo 4 and 8 before starting **BECAUSE** the setter chose the year and then chose the statement that its residue class makes provable, so the residue tells you which branch of the argument is intended

Backtesting

1. First useful cue: the target is a prime power, so take v_2 and turn the product of orders into the sum $k(G) = \sum_g v_2(\operatorname{ord} g)$; and every order divides the product, so every element has 2-power order for free.
2. Method the expression suggested: sum over cyclic subgroups, since ord is constant on generator sets and each class has $\varphi(2^h) = 2^{h-1}$ elements.
3. Where you would hesitate, and the thought to reinforce: at the choice of modulus. Reinforce that modulo 2 the identity says only that k is odd, which 2009 is, so the modulus must be 4; the coefficients $h2^{h-1}$ tell you this directly.
4. The trap the question sets: the word "abelian" invites an immediate structure theorem expansion and a search over exponent vectors, which is a page of work before any obstruction appears.
5. What makes you recognise it faster next time: any statistic of the form "combine $\operatorname{ord}(g)$ over all g" should trigger the cyclic subgroup partition instantly, exactly as $\sum_{d \mid n} \varphi(d) = n$ does.
6. The reusable decision rule: before solving an exact equation for a large target, reduce the attainable values modulo a small number chosen to kill all but one family of terms.

Mirror drills

1. Show that no finite abelian group has the product of the orders of its elements equal to 2^n with $n \equiv 2 \pmod 4$. *Answer:* $k(G) \equiv 2^e - 1 \pmod 4$, which lies in $\{0, 1, 3\}$, never 2.
2. Compute the product of the orders of all elements of $\mathbb{Z}/2 \times \mathbb{Z}/4$. *Answer:* orders are $1, 2, 2, 2, 4, 4, 4, 4$, so the product is 2^{11}.
3. For an odd prime p, find the product of the orders of all elements of $\mathbb{Z}/p^i$. *Answer:* p^K with $K = \sum_{h=1}^{i} h(p^h - p^{h-1}) = ip^i - \frac{p^i - 1}{p - 1}$.

17.6 A6. Disproving a two variable mean value theorem

Let $f : [0,1]^2 \to \mathbb{R}$ be a continuous function on the closed unit square such that $\dfrac{\partial f}{\partial x}$ and $\dfrac{\partial f}{\partial y}$ exist and are continuous on the interior $(0,1)^2$. Let

$$a = \int_0^1 f(0,y)\,dy, \quad b = \int_0^1 f(1,y)\,dy, \quad c = \int_0^1 f(x,0)\,dx, \quad d = \int_0^1 f(x,1)\,dx.$$

Prove or disprove: there must be a point (x_0, y_0) in $(0,1)^2$ such that

$$\frac{\partial f}{\partial x}(x_0,y_0) = b - a \quad \textit{and} \quad \frac{\partial f}{\partial y}(x_0,y_0) = d - c.$$

Read it out loud

Read the hypotheses first, because they are a quotation. Continuous on a closed region, differentiable on the interior, conclusion asserting the existence of an interior point where a derivative equals a difference of boundary data: that is the mean value theorem, transposed to two variables. The fields wired in are multivariable calculus (partial derivatives, the fundamental theorem of calculus, Fubini), real analysis (the intermediate value theorem, connectedness, convexity), and counterexample construction. The command term "prove or disprove" means the setter is not telling you which, so your first job is to decide, and deciding is most of the problem.

Now read the four constants structurally rather than as four separate integrals. For f that is C^1 on a neighbourhood of the square,

$$b - a = \int_0^1 \big(f(1,y) - f(0,y)\big)dy = \int_0^1\int_0^1 \frac{\partial f}{\partial x}(x,y)\,dx\,dy,$$

by the fundamental theorem of calculus inside and Fubini outside, and symmetrically $d - c$ is the double integral of $\partial f/\partial y$. So $(b-a, d-c)$ is nothing other than the average value of the gradient ∇f over the square. The assertion says: the average of the gradient is attained as an actual value of the gradient at some interior point. Write that sentence down within thirty seconds; the entire problem is contained in it.

Once it is written, the decision is immediate. For a scalar continuous function on a connected set, the average is attained, because the image is an interval containing values on both sides of the average, and the intermediate value theorem finishes. For a vector valued function this argument dies, because the continuous image of a connected set in $\mathbb{R}^2$ is connected but need not be convex, while the average is only guaranteed to lie in the closed convex hull of the image. Connected sets in the plane can have holes, and an average can sit in a hole. The statement is therefore false, and the counterexample must be a function whose gradient image is a connected set with the average sitting in a hole.

There is a second confirmation available for free. The one variable version of this false statement is the mean value theorem for vector valued functions, and its standard failure is $\gamma(t) = (\cos t, \sin t)$ on $[0, 2\pi]$: the average velocity is zero, and γ' is never zero because cosine and sine have no common zero. That object is the template for the counterexample you are about to build.

Approaches, ranked worst to best

1. **Topological existence proof.** *Cue:* two equations in two unknowns with a demanded solution. *Circuit:* apply a degree, winding number or Poincare and Miranda type argument to the map $(x,y) \mapsto (f_x - (b-a), f_y - (d-c))$. The worst route, and dangerous because the cue is real. There is no control of the gradient on the boundary, so no hypothesis of the required type is available, and the statement is false in any case. This can absorb an hour and produce nothing.

2. **Iterated mean value theorem.** *Cue:* the hypotheses are literally the mean value theorem hypotheses. *Circuit:* apply the one variable mean value theorem in x to the y averaged function, then again in y. This does not prove the statement, but it is not wasted. It shows that each equation alone is always solvable: there is a point where $f_x = b - a$ and a point where $f_y = d - c$. The route fails precisely at synchronisation, which is exactly the information you need to know that simultaneity is the content and where to attack.
3. **Random search among small polynomials.** *Cue:* "prove or disprove" with a whiff of falsity. *Circuit:* write down a family such as $\alpha x^2 + \beta xy + \gamma y^2 + \dots$, compute the four constants, and hunt for a member with no solution. It works, and the published counterexample $3(1+y)(2x-1)^2 - y$ is of this shape. It costs four integrals per candidate and gives no reason, so you cannot reproduce it under pressure on a different problem.
4. **Normalise, then quote the standard pathology.** *Cue:* the conclusion is an equation $\nabla f = v$ for a fixed vector v, and the constants are the average of ∇f. *Circuit:* subtract the linear function $v \cdot (x, y)$ to reduce to $v = 0$, so that the claim becomes "f must have an interior critical point", then build a gradient that rotates without vanishing, copying $(\cos, \sin)$. The best route. The search space collapses to one design constraint, and the verification is four integrals that are zero by inspection.

The ideal solution

The assertion is $\boxed{\text{false}}$, and the following function disproves it.

Take

$$f(x, y) = y \sin(2\pi x).$$

It is a product of a polynomial in y and a smooth function of x, hence C^∞ on all of $\mathbb{R}^2$; in particular it is continuous on $[0,1]^2$ and its two partial derivatives exist and are continuous on $(0,1)^2$, so it satisfies every hypothesis.

The four constants. Since $\sin 0 = \sin 2\pi = 0$,

$$a = \int_0^1 y \sin 0 \, dy = 0, \qquad b = \int_0^1 y \sin 2\pi \, dy = 0,$$

and

$$c = \int_0^1 0 \cdot \sin(2\pi x) \, dx = 0, \qquad d = \int_0^1 \sin(2\pi x) \, dx = \left[\frac{-\cos 2\pi x}{2\pi} \right]_0^1 = 0.$$

Hence $b - a = 0$ and $d - c = 0$, so the assertion, applied to this f, claims the existence of $(x_0, y_0) \in (0,1)^2$ with

$$\frac{\partial f}{\partial x}(x_0, y_0) = 0 \quad \text{and} \quad \frac{\partial f}{\partial y}(x_0, y_0) = 0.$$

No such point exists. The partial derivatives are

$$\frac{\partial f}{\partial x}(x, y) = 2\pi y \cos(2\pi x), \qquad \frac{\partial f}{\partial y}(x, y) = \sin(2\pi x).$$

Suppose $(x_0, y_0) \in (0,1)^2$ satisfies both equations. From $\sin(2\pi x_0) = 0$ with $0 < x_0 < 1$ we get $2\pi x_0 = \pi$, so $x_0 = \frac{1}{2}$. Then $\cos(2\pi x_0) = \cos \pi = -1$, so the first equation reads $-2\pi y_0 = 0$, giving $y_0 = 0$. But $y_0 = 0$ is not in $(0,1)$. Contradiction.

Therefore there is no point of the open square at which both equations hold, and the assertion is disproved.

Remark 17.6.1. The gradient of this f never vanishes anywhere in the open square, because cos and sin have no common zero. For fixed $y \in (0,1)$, as x runs over $(0,1)$ the point $\nabla f(x, y) = (2\pi y \cos 2\pi x, \sin 2\pi x)$ traverses the ellipse with semi axes $2\pi y$ and 1 centred at the origin. So the

image of ∇f is a union of ellipses about the origin, a connected set with the origin missing from it: the average of the gradient sits in a hole of the gradient's image. That is the whole mechanism, and it is the two variable version of the circle $(\cos t, \sin t)$ that defeats the mean value theorem for vector valued functions.

Remark 17.6.2. Each equation separately is always satisfiable when f is C^1 on a neighbourhood of the square. Put $F(x) = \int_0^1 f(x,y)\,dy$; then F is differentiable with $F'(x) = \int_0^1 f_x(x,y)\,dy$, and the mean value theorem gives $\xi \in (0,1)$ with $\int_0^1 f_x(\xi,y)\,dy = F(1) - F(0) = b - a$. The continuous function $y \mapsto f_x(\xi,y)$ has average $b-a$, so by the intermediate value theorem it equals $b-a$ somewhere. The same works for the other partial. What fails is only that the two points need not coincide, and that is exactly what the counterexample exploits.

Layer 1: WHAT

Block 1. Recognise the hypotheses as the mean value theorem transposed to a square, and the conclusion as a statement about a common point.

Block 2. Convert the four constants: $b-a$ and $d-c$ are the averages of f_x and f_y over the square, so the claim is that the average of ∇f is a value of ∇f.

Block 3. Decide truth. The scalar version is true by the intermediate value theorem; the vector version fails because a connected planar image need not be convex, and the one variable failure $(\cos t, \sin t)$ is already on record. Commit to disproving.

Block 4. Normalise. Replacing f by $g = f - (b-a)x - (d-c)y$ shifts the gradient by a constant and makes both averages zero, so it suffices to find f with $a = b$, $c = d$ and no interior critical point.

Block 5. Design. Force the gradient to rotate without vanishing: use sin and cos in the same variable, so that the two partials cannot vanish together, and use a period over $[0,1]$ so that the boundary integrals vanish by periodicity.

Block 6. Land on $f(x,y) = y\sin(2\pi x)$; compute the four constants, all zero; solve the gradient system and find the only solution has $y_0 = 0$, outside the open square.

Block 7. Write up: hypotheses checked, constants computed, system eliminated in two lines.

The nominal route is Block 1 followed by an attempt to prove the statement, either by iterating the mean value theorem or by a fixed point argument, and it can consume the whole session because the hypotheses look so much like a theorem. The bifurcation point is the single line at Block 3: *is this even true?* Everything before it is shared; everything after it diverges completely. Note where the ideal route spends its effort: not on the proof, which is six lines, but on the two minutes of structural reading that convert the four constants into one sentence about averages.

Layer 2: HOW

Block 1. Cue. Continuous on a closed region, differentiable on the interior, conclusion asserting an interior point. **Why natural.** That triple is the signature of the mean value theorem, and a setter who writes hypotheses in a familiar shape is telling you which theorem the problem is about. Recognising the quotation does not tell you the statement is true; it tells you which true theorem to compare it against, and comparison is how you decide. **Link to the next block.** If this is a mean value theorem, then the four constants must be playing the role of $f(b) - f(a)$, so you should express them as something the derivative produces.

Block 2. Cue. Four boundary integrals, paired as differences $b-a$ and $d-c$. **Why natural.** The pairing is a difference of values of f on opposite edges, averaged over the other variable. A difference of values of f along a segment is an integral of the derivative, by the fundamental theorem of calculus, and averaging over the second variable is another integral, which Fubini lets you combine. The result is not a coincidence of notation: the problem is stated in terms of a, b, c, d

precisely to disguise the fact that $(b-a, d-c)$ is the average gradient. **Link to the next block.** A statement about an average being attained is a statement about the intermediate value theorem, so ask whether the intermediate value theorem is available.

Block 3. **Cue.** An average of a vector valued function claimed to be attained. **Why natural.** For scalar functions on a connected domain the average lies between the infimum and the supremum, and the image is an interval, so the average is attained. In $\mathbb{R}^2$ the only thing continuity gives is that the image is connected, and the average lies in the closed convex hull of the image, which is strictly larger than the image whenever the image is not convex. Since connected planar sets can enclose a hole, the claim is false in general and you should stop looking for a proof. **Link to the next block.** The failure mode is now explicit: build a gradient image that surrounds the target without containing it. The one variable model of that is the circle.

Block 4. **Cue.** The conclusion is $\nabla f = v$ for a fixed vector v determined by f. **Why natural.** Adding a linear function $\ell(x, y) = \alpha x + \beta y$ shifts the gradient by the constant (α, β) and shifts the four constants in a controlled way: with $\alpha = b - a$ and $\beta = d - c$, the new function $g = f - \ell$ has $b_g - a_g = 0$ and $d_g - c_g = 0$, and $\nabla f(P) = v$ if and only if $\nabla g(P) = 0$. So there is no loss in searching only among functions with balanced boundary data, and the target becomes the single point 0. **Link to the next block.** The design problem is now stated in its simplest form: a function on the square with balanced boundary averages and no interior critical point.

Block 5. **Cue.** Two conditions to arrange at once: boundary integrals vanish, gradient never vanishes. **Why natural.** Each has a standard supplier. Vanishing integrals come from a full period of a sine, since $\int_0^1 \sin 2\pi x \, dx = 0$ and $\sin 2\pi \cdot 0 = \sin 2\pi \cdot 1 = 0$, which kills three of the four constants outright. A nonvanishing gradient comes from making the two partials into cos and sin of the same argument, which have no common zero. Both suppliers are the same object, so use $\sin 2\pi x$ and let the second variable enter as the simplest possible factor. There is one more constraint worth naming: a continuous function on the compact square attains a maximum, and an interior maximum is a critical point, so every extremum of your counterexample must sit on the boundary. A factor of y pushes the extremes to the edges $y \in \{0, 1\}$, which is exactly what is needed. **Link to the next block.** The candidate $f(x, y) = y \sin(2\pi x)$ now has to be checked, not invented.

Block 6. **Cue.** A specific candidate. **Why natural.** Verification is mechanical and must be done in full: hypotheses, then the four constants, then the system. The elimination order matters for brevity: solve the partial that involves one variable only, namely $f_y = \sin 2\pi x = 0$, which pins $x_0 = 1/2$ immediately, and substitute into the other. Solving the messier equation first would cost more ink. **Link to the next block.** Nothing remains except to write it in the order a grader reads it.

Block 7. **Cue.** A "prove or disprove" answer. **Why natural.** Full marks needs three things and no more: the example, a confirmation that it satisfies the hypotheses (which candidates that are merely continuous can fail), and a complete demonstration that the conclusion fails for it. No motivation is required in the write up; the normalisation and the design story stay in your head. **Link to the next block.** Done.

Layer 3: WHY

The mathematical truth beneath this problem is the difference between connectedness and convexity. Continuity gives you exactly one thing about images: the continuous image of a connected set is connected. In $\mathbb{R}^1$, connected means interval, and an interval containing points on both sides of a number contains that number; that single sentence is the intermediate value theorem and it is

why every scalar mean value statement in sight is true. In $\mathbb{R}^2$, connected sets can be annuli, so the deduction "the average lies between the extremes, therefore it is attained" has no analogue. The average of a vector valued function lies in the closed convex hull of the image, and the convex hull can be strictly bigger. Every counterexample to a vector mean value statement is built by making the image wrap around its own average.

That explains why the one variable case is already broken. For $\gamma(t) = (\cos t, \sin t)$ on $[0, 2\pi]$ the average velocity is $(\gamma(2\pi) - \gamma(0))/2\pi = 0$, yet $|\gamma'(t)| = 1$ for every t. The two variable statement here is the same failure dressed in boundary integrals: the four constants are a disguise for the average gradient, and the counterexample makes the gradient run around a family of ellipses centred on that average.

The discernment mechanism to train is a three step filter for any "prove or disprove" whose hypotheses quote a known theorem. First, name the theorem being generalised. Second, ask what the generalisation upgrades: scalar to vector, one variable to several, pointwise to uniform, finite to infinite. Third, recall the standard pathology for exactly that upgrade. Scalar to vector upgrades of mean value statements are broken by rotation. One variable to several for the intermediate value theorem survives, because connectedness survives. Pointwise to uniform is broken by moving bumps. Finite to infinite is broken by the shift operator. Each upgrade has a stock counterexample, and knowing the stock list converts a research question into a lookup.

There is one more piece of reading worth doing, on the hypotheses themselves. The problem asks only for continuity on the closed square with continuous partials on the open square. Why weaken it there? Two reasons, both deliberate. It mirrors the mean value theorem's own hypotheses, strengthening the false suggestion that a mean value theorem is coming. And it legalises counterexamples whose partial derivatives blow up at the boundary, such as $x^{1/3}y^{2/3}$, widening the class of solutions the setter will accept. It also quietly blocks the laziest route to the identity $b - a = \iint f_x$: with partials only defined inside, the Fubini and fundamental theorem argument needs care about improper integrals. That is why the ideal solution never uses that identity in the proof and uses it only in the reading. Your written answer computes the four constants directly from a function that is smooth everywhere, so no interchange is ever justified in the arena.

Honest knowledge gaps for this problem, in the order they must be installed: partial derivatives and the gradient of a function of two variables; the double integral and Fubini's theorem for continuous functions on a rectangle; the fundamental theorem of calculus applied along a segment with the other variable fixed; differentiation under the integral sign; the mean value theorem for vector valued functions and why it fails; and the fact that the average of a continuous vector valued function lies in the closed convex hull of its image. None of this is in the A-level further mathematics core, and all of it is standard first year undergraduate analysis. Without partial derivatives and the connectedness versus convexity distinction, this problem is not approachable at all; with them, it is six lines.

Revolving concepts

Mean value theorem and its vector valued failure. For $f : [p, q] \to \mathbb{R}$ continuous on $[p, q]$ and differentiable on (p, q) there is ξ with $f'(\xi) = (f(q) - f(p))/(q - p)$. For $\gamma : [p, q] \to \mathbb{R}^n$ with $n \geq 2$ this fails, the standard witness being the circle. The correct replacements are the mean value inequality $|\gamma(q) - \gamma(p)| \leq (q - p) \sup |\gamma'|$ and the integral form $\gamma(q) - \gamma(p) = \int_p^q \gamma'(t)\, dt$. See Book 1, field: Real analysis, topic: mean value theorems.

The scalar multivariable mean value theorem, which does survive. For f scalar and differentiable on a convex open set, $f(P) - f(Q) = \nabla f(R) \cdot (P - Q)$ for some R on the segment PQ, because restricting to the segment reduces it to one variable. The contrast with the failed statement in this problem is the whole point: scalar output survives, vector output does not. See Book 1, field: Real analysis, topic: multivariable mean value theorem.

Averages, connectedness and convex hulls. If u is continuous and real valued on a connected

compact set K, its average over any probability measure on K is a value of u. If u maps into $\mathbb{R}^n$ with $n \geq 2$, the average lies only in the closed convex hull of $u(K)$. The gap between hull and image is where every counterexample of this type lives. See Book 1, field: Real analysis, topic: intermediate value theorem and convexity.

Boundary data as an integral of the gradient. For f that is C^1 on a neighbourhood of $[0,1]^2$,

$$b - a = \iint_{[0,1]^2} \frac{\partial f}{\partial x}\, dA, \qquad d - c = \iint_{[0,1]^2} \frac{\partial f}{\partial y}\, dA,$$

by the fundamental theorem of calculus in one variable and Fubini in the other. This is the translation that makes the problem readable, and it is a special case of the divergence theorem on a rectangle. See Book 1, field: Calculus and integrals, topic: Fubini and the fundamental theorem of calculus.

Interior extrema are critical points. A continuous function on a compact set attains its extremes; if an extreme is attained at an interior point where f is differentiable, the gradient vanishes there. Hence any counterexample must have all of its extreme values on the boundary of the square, which is a design constraint you can check by eye before computing anything. See Book 1, field: Calculus and integrals, topic: critical points and Fermat's condition.

Normalisation by a linear function. For fixed $v = (\alpha, \beta)$, the substitution $g = f - \alpha x - \beta y$ gives $\nabla g = \nabla f - v$ and changes the boundary constants by $b_g - a_g = (b - a) - \alpha$ and $d_g - c_g = (d - c) - \beta$. Choosing $v = (b - a, d - c)$ reduces any statement of the form "$\nabla f = v$ somewhere" to "$\nabla g = 0$ somewhere". The general principle is that affine substitutions are free and should be used to send the target of a problem to zero before you start constructing. See Book 1, field: Calculus and integrals, topic: normalisation and affine reduction.

Circuit Bank

- **WHEN** a proposed statement demands two conditions at a single common point **DO** check first whether each condition alone is achievable, and if so attack the simultaneity with a counterexample **BECAUSE** the coordinates of a vector valued mean value statement are each true separately, and only the shared point fails, so simultaneity is where the statement is breakable
- **WHEN** the hypotheses of a prove or disprove question quote a known theorem **DO** name the theorem, name what the statement upgrades, and recall the stock pathology for that upgrade **BECAUSE** scalar to vector is broken by rotation, pointwise to uniform by moving bumps, finite to infinite by the shift, and the stock list turns a decision into a lookup
- **WHEN** the conclusion is an equation $\nabla f = v$ **DO** subtract the linear function $v \cdot (x, y)$ to normalise v to zero **BECAUSE** gradients shift by constants under adding linear functions, so the search reduces to functions with balanced boundary data and no interior critical point
- **WHEN** you need a construction with prescribed vanishing integrals **DO** insert a full period of a sine in the relevant variable **BECAUSE** integrating a sine over a whole period gives zero and the endpoint values also vanish, which discharges several conditions with one factor
- **WHEN** you need a planar vector field that never vanishes **DO** make its two components cos and sin of the same argument **BECAUSE** cosine and sine have no common zero, and this is the same object that defeats the mean value theorem for vector valued functions
- **WHEN** you are hunting a function with no interior critical point on a compact region **DO** force every maximum and minimum onto the boundary, for instance by a monotone factor in one variable **BECAUSE** an interior extremum of a differentiable function is a critical point, so a counterexample with an interior extremum cannot exist

Backtesting

1. First useful cue: hypotheses continuous on the closed square and differentiable inside, with a conclusion at an interior point, which is the mean value theorem in disguise; the four constants are the average of the gradient.
2. Method the expression suggested: rewrite $b-a$ and $d-c$ as double integrals of the partials, so the claim reads "the average of ∇f is attained", and then test that claim against the intermediate value theorem.
3. Where you would hesitate, and the thought to reinforce: at the decision to stop proving. Reinforce that continuity gives connectedness of the image and nothing more, and that an average only lies in the convex hull, so a vector valued mean value claim is false until proved otherwise.
4. The trap the question sets: hypotheses that quote a real theorem so exactly that you spend the session trying to prove a false statement, aided by the fact that the scalar coordinates really are each achievable.
5. What makes you recognise it faster next time: any claim of the form "two partial derivative conditions hold at the same point" should trigger the circle $(\cos t, \sin t)$ within seconds, because that is the canonical obstruction to simultaneity.
6. The reusable decision rule: when a proposed theorem upgrades a scalar existence result to a vector valued one, look for the counterexample first and the proof second, and build the counterexample by normalising the target to zero and making the image wrap around it.

Mirror drills

1. Assume f is C^1 on a neighbourhood of $[0,1]^2$. Show there is a point of $(0,1)^2$ where $\partial f/\partial x = b-a$, with no claim about the other partial. *Answer:* put $F(x) = \int_0^1 f(x,y)dy$, apply the mean value theorem to F to get ξ with $\int_0^1 f_x(\xi,y)dy = b-a$, then the intermediate value theorem in y.
2. Produce a C^1 curve $\gamma : [0,1] \to \mathbb{R}^2$ with $\gamma(0) = \gamma(1)$ and $\gamma'(t) \neq 0$ for all t, and say what it disproves. *Answer:* $\gamma(t) = (\cos 2\pi t, \sin 2\pi t)$; it disproves the mean value theorem for vector valued functions, since the average velocity is zero.
3. Show the assertion of this problem is true whenever $f(x,y) = u(x) + v(y)$ with u, v continuous on $[0,1]$ and differentiable on $(0,1)$. *Answer:* $b-a = u(1)-u(0) = u'(\xi)$ and $d-c = v(1)-v(0) = v'(\eta)$ by the mean value theorem, and since $f_x = u'(x)$ and $f_y = v'(y)$ are decoupled, the point (ξ,η) works.

17.7 B1. rationals as quotients of prime factorials

Show that every positive rational number can be written as a quotient of products of factorials of (not necessarily distinct) primes. For example,

$$\frac{10}{9} = \frac{2! \cdot 5!}{3! \cdot 3! \cdot 3!}.$$

Read it out loud

Say the claim in your own words before you touch a pen: the set of positive rationals you can build by multiplying and dividing prime factorials is all of $\mathbb{Q}^+$. That is a closure statement, not a computation. The object on trial behaves as a sub-monoid of the multiplicative group $\mathbb{Q}^+$ that is closed under reciprocals, hence a subgroup, and you are asked to show it is the whole group. The moment you see "quotient of products," the algebra structure should light up: a nonempty subset of a group, closed under the group operation and under inverses, is automatically a subgroup, so proving it equals the whole group reduces to proving it contains a generating set.

The second cue is arithmetic, not algebraic: "factorials of primes" is a sparse generating set. $2!, 3!, 5!, 7!, \ldots$ are specific composite integers, not the primes themselves. The real question buried inside the statement is how you manufacture a bare prime p out of factorials of primes, when $p!$ is not p, it is p times a whole tail of smaller factors.

Within the first thirty seconds you should have written down the one identity that makes the whole problem move: $p! = p \cdot (p-1)!$, so $p = p!/(p-1)!$. That line is the seed. Everything else is bookkeeping: turning "p divides a/b to some power" into "multiply or divide by $p!/(p-1)!$ that many times," then noticing that $(p-1)!$ is a genuinely smaller object in the one sense that matters, its largest prime factor has dropped, even while its numerical size has exploded.

Approaches, ranked worst to best

1. **fiddle with small examples until a pattern appears.** *Cue:* the worked example $10/9$. *Circuit:* no reusable circuit, pattern search only. You can reproduce the given example by trial, but nothing you find for $10/9$ tells you what to do with, say, $17/13$. A construction that only works because the numbers were small is not a proof; every mark would be riding on a leap you cannot justify. Unlimited ink, no guarantee.

2. **fix a/b and try to reduce $\max(a,b)$ at each step.** *Cue:* the standard "induct on size" reflex. *Circuit:* induction on the wrong invariant. This looks honest, since induction is the right shape of proof for a closure statement. But the instant you multiply by $p!/(p-1)!$, $(p-1)!$ is enormous next to p, so the numerator or denominator can grow by orders of magnitude in a single step. An induction on $\max(a,b)$, or on $a+b$, is chasing a quantity that moves the wrong way. This is the exact bifurcation point of the problem: the zombie route keeps faith with "smaller numbers" as progress and stalls immediately.

3. **strong induction on the largest prime factor appearing in a or b.** *Cue:* the identity $p = p!/(p-1)!$ only ever introduces primes strictly below p. *Circuit:* extremal peeling on a correctly chosen invariant. Let the numbers grow as large as they like; track only the largest prime dividing a or b. Every application of $p!/(p-1)!$ strictly lowers that invariant, because $(p-1)!$ is built entirely from integers below p. The invariant is a nonnegative integer and it strictly decreases, so the induction terminates in finitely many steps no matter how large the actual numerator and denominator become. This is the ideal route: one identity, one strong induction, no case analysis beyond "does p divide the numerator or the denominator."

The ideal solution

Every positive rational number has a unique representation a/b with a, b positive integers and $\gcd(a, b) = 1$. Call the largest prime dividing a or b the *rank* of a/b, and give $1/1$ rank 1 (no prime present).

Prove the claim by strong induction on the rank.

Base case. Rank 1 forces $a = b = 1$, and $1 = 2!/2!$, a quotient of products of prime factorials.

Inductive step. Suppose a/b has rank p, a prime, and suppose every positive rational of rank strictly less than p is a quotient of products of prime factorials. Since $\gcd(a, b) = 1$, the prime p divides exactly one of a, b; say $p \mid a$ (if instead $p \mid b$, run the same argument on b/a and invert the resulting expression at the end, which is still a quotient of products of factorials). Write

$$a = p^k a', \qquad \gcd(a', p) = 1, \quad k \geq 1,$$

so that $a/b = p^k \cdot (a'/b)$. Because p was the largest prime dividing a or b, both a' and b have every prime factor strictly below p.

Now raise the seed identity $p = p!/(p-1)!$ to the power k:

$$p^k = \frac{(p!)^k}{\big((p-1)!\big)^k}.$$

Substituting,

$$\frac{a}{b} = (p!)^k \cdot \frac{a'}{\big((p-1)!\big)^k b}.$$

The factor $(p!)^k$ is already a product of prime factorials. The fraction $a'/\big((p-1)!\big)^k b$ is built entirely from a', b and $(p-1)!$, and every prime factor of each of these three integers is strictly below p: true of a' and b by construction, and true of $(p-1)!$ because it is a product of the integers 1 through $p-1$. So this fraction has rank strictly less than p.

By the induction hypothesis, $a'/\big((p-1)!\big)^k b$ equals some quotient of products of prime factorials, N/D. Then

$$\frac{a}{b} = (p!)^k \cdot \frac{N}{D},$$

which folds $(p!)^k$ into the numerator alongside N's own factorial factors, and is again a product of prime factorials over a product of prime factorials. This closes the induction: every positive rational number is a quotient of products of factorials of primes. ■

As a check that the machine actually runs, follow it on the example. For $10/9$: the rank is 5, since $5 \mid 10$, so $10/9 = 5 \cdot (2/9) = (5!/4!) \cdot (2/9)$, leaving a residual $2/(4! \cdot 9) = 1/108$ of rank 3, resolved by $3^{-3} = (2!)^3/(3!)^3$, leaving $1/2$ of rank 2, resolved directly by $2 = 2!$ since $1! = 1$. Multiplying the layers back together reproduces a valid quotient of prime factorials equal to $10/9$, confirming the construction is not merely an existence argument on paper.

Layer 1: WHAT

Block 1: recognise $\mathbb{Q}^+$ as a multiplicative group and the target set as automatically a subgroup, so "produces every rational" reduces to "reaches every prime." Block 2: write down $p = p!/(p-1)!$, the one non-obvious algebraic input. Block 3: read this identity as trading a bare prime p for a factorial of p, which is allowed, divided by a factorial of $p-1$, which is not itself a prime factorial, but whose prime factors are all smaller than p. Block 4: choose the induction invariant correctly: not the size of a, b, but the largest prime factor, and confirm it strictly decreases. Block 5: run the strong induction to completion, with $p = 2$, i.e. $1! = 1$, as the trivial floor.

The nominal route stalls at Block 3, because the untrained instinct is to induct on $\max(a, b)$, which the substitution makes worse, not better. The bifurcation point sits exactly there: the ideal

route abandons "smaller numbers" as the measure of progress and replaces it with "fewer distinct large primes," the one quantity the construction actually controls.

Layer 2: HOW

Block 1. **Cue.** The word "quotient" and "products" applied to a set built from prime factorials. **Why natural.** $\mathbb{Q}^+$ under multiplication is a group; a nonempty subset closed under products and reciprocals is a subgroup, and a subgroup containing a generating set of the whole group must be the whole group. Reframing "show every rational is representable" as "show this candidate subgroup is everything" turns an open-ended construction into a closure argument, far cheaper to police. **Link to the next block.** Since the primes generate $\mathbb{Q}^+$ multiplicatively, by unique factorisation, it suffices to reach every prime.

Block 2. **Cue.** You need a bare prime p out of an object called $p!$, which is p times a large tail. **Why natural.** Definitions are not negotiable: $p! = p \cdot (p-1) \cdots 1$, so dividing by $(p-1)!$ peels the tail off and leaves exactly p. There is no other identity connecting a prime to a factorial, so this is the only door available, and you take it. **Link to the next block.** The price of using this door is that $(p-1)!$ now sits in a denominator or numerator, and it is not itself a prime factorial, so you owe the reader a reason that is still acceptable.

Block 3. **Cue.** $(p-1)!$ is a specific, fully factored integer: $1 \cdot 2 \cdots (p-1)$. **Why natural.** Every prime factor of $(p-1)!$ is at most $p-1$, strictly below p. That is not a numerical accident, it is definitional: multiplying integers below p can never land on a prime factor of size p or larger. **Link to the next block.** This is exactly the property needed to set up a terminating induction: whatever measure counts "how many primes at or above p are still present" has just dropped by one such prime.

Block 4. **Cue.** The temptation to induct on the size of the numbers. **Why natural.** Reject it explicitly: $(p-1)!$ can dwarf a or b, so size increases. The only quantity guaranteed to decrease is the largest prime factor present, because the substitution was designed to kill exactly that prime and introduce nothing at or above it. **Link to the next block.** Strong induction on that invariant, bottoming out at rank 1, closes the proof.

Block 5. **Cue.** The base case must be checked honestly, not waved through. **Why natural.** $1 = 2!/2!$ is the smallest possible quotient of prime factorials, and it is the only rational of rank 1; the induction has nowhere lower to go, so this is a genuine floor. **Link to the next block.** None, the proof is complete.

Layer 3: WHY

The truth underneath this problem is that $\mathbb{Q}^+$, as a multiplicative group, is free abelian on the primes: every positive rational is a unique finite product $\prod p_i^{e_i}$ with $e_i \in \mathbb{Z}$. Proving a subset generates $\mathbb{Q}^+$ is therefore never about the raw arithmetic size of that subset, it is about whether the subset reaches every generator. The examiners chose factorials of primes as the generating gadget precisely because $p!/(p-1)! = p$ is a clean way to reach a generator, while guaranteeing that the natural first attempt, tracking the size of the numbers, fails, so that only a correctly chosen invariant, the largest prime factor, gets the induction to close.

The discernment mechanism is this: whenever an induction stalls because a substitution makes the objects bigger, ask what got structurally simpler even though the objects got numerically larger. Here, "structurally simpler" means "supported on strictly smaller primes," a well ordered quantity, so strong induction on it is valid and finite. Train this discernment by refusing to accept "the

numbers are growing, so induction on size cannot work" as a dead end; instead ask which other measurement of the same object the last step actually controlled.

Revolving concepts

- **Unique factorisation.** Every positive integer is a unique product of primes, so every positive rational has a unique representation as a finite product of prime powers with integer, possibly negative, exponents. This is why "reach every prime" is equivalent to "reach every positive rational." See Book 1, field: Number theory, topic: prime factorisation.
- $\mathbb{Q}^+$ **as a free abelian group.** Under multiplication, $\mathbb{Q}^+$ is isomorphic to a direct sum of one copy of $\mathbb{Z}$ per prime, via the exponent of each prime in a rational's factorisation. A nonempty subset closed under products and reciprocals is a subgroup; a subgroup containing a generating set is the whole group. See Book 1, field: Abstract algebra, topic: generators of an abelian group.
- **Strong induction on a non-standard, well founded invariant.** The recursion here is $p \mapsto (p-1)!$, and the generalised form of the technique is: whenever a substitution grows the objects but strictly reduces some other well ordered quantity, largest prime factor, number of distinct prime factors, degree of a polynomial, rank in a filtration, strong induction on that quantity is valid, regardless of what happens to size. The method is always the same: identify the quantity, show it is a nonnegative integer, show it strictly decreases, check the floor case by hand. This generalises to any "peel the extremal piece, recurse on the residue" argument, the same skeleton behind the Euclidean algorithm and behind proofs of unique factorisation itself.
- **Legendre's formula (background, not load bearing here).** The exact power of a prime q dividing $n!$ is the sum over $i \geq 1$ of $\lfloor n/q^i \rfloor$. It confirms, from a different angle, that $p!$ contains no prime factor above p and exactly one copy of p itself, which is why $p = p!/(p-1)!$ is exact with no residue to clean up. See Book 1, field: Number theory, topic: prime factorisation of factorials.

Circuit Bank

- **WHEN** a problem asks you to show a small generating set produces an entire multiplicative structure, all rationals, all integers under some operation **DO** reframe as a subgroup closure question: show the candidate set is closed under the group operations and contains a known generating set **BECAUSE** a nonempty subset of a group, closed under the operation and inverses, is automatically a subgroup, so "produces everything" reduces to "contains generators"
- **WHEN** an induction on the size of the numbers stalls because your substitution makes them larger **DO** look for a different, well ordered invariant that the same substitution strictly decreases, and reinduct on that **BECAUSE** induction only needs a well founded decreasing quantity, never the size of the objects specifically
- **WHEN** you need to produce a prime p from factorials **DO** write $p = p!/(p-1)!$ immediately **BECAUSE** it is the only algebraic bridge between a bare prime and a factorial, the first thing to try, not a discovery to search for

Backtesting

1. First useful cue: "quotient of products" applied to a sparse generating set, flagging a group generation argument before any arithmetic starts.

2. Method the expression suggested: peel the largest prime using $p = p!/(p-1)!$ and recurse.
3. Where you would hesitate, and the thought to reinforce: the moment $(p-1)!$ appears and looks like it made the problem worse; reinforce that you are tracking the largest prime factor, not size.
4. The trap the question sets: it tempts an induction on $a+b$ or $\max(a,b)$, the default competition move, which fails outright once you substitute.
5. What makes you recognise it faster next time: any phrase of the form "every element of a multiplicative structure can be built from a sparse generating set" should trigger an immediate check of what invariant the generating step actually controls, before you commit to an induction variable.
6. The reusable decision rule: when a substitution both grows the numbers and kills the largest bad feature, a prime, a degree, an exponent, induct on the bad feature, never on the size.

Mirror drills

1. Show that every positive rational number can be written as a quotient of products of primorials $p\# = 2 \cdot 3 \cdot 5 \cdots p$, the product of all primes up to p. *Answer:* peel the largest prime p using $p = p\#/q\#$, where q is the prime immediately below p, then induct on the largest prime factor exactly as above.
2. Take $a/b = 1/5$. Compute $\max(a,b)$ before and after applying the substitution for $p = 5$, and separately compute the rank, the largest prime factor present, before and after, to see directly why size is the wrong induction measure. *Answer:* before, max $= 5$, rank 5; the substitution gives $1/5 = 4!/5! = 24/120$, so max has jumped to 120, while the residual to induct on is $1/24$, of rank 3, strictly below 5. Size grew, rank fell; only rank is a valid measure here.
3. Show that every positive integer n can be written as a product of terms of the form $p!/(p-1)!$ for primes p, with repetition allowed. *Answer:* immediate from unique factorisation, since $n = \prod p_i^{e_i}$, and each p_i equals $p_i!/(p_i-1)!$.

17.8 B2. the range of achievable jump costs

A game involves jumping to the right on the real number line. If a and b are real numbers and $b > a$, the cost of jumping from a to b is $b^3 - ab^2$. For what real numbers c can one travel from 0 to 1 in a finite number of jumps with total cost exactly c?

Read it out loud

Strip the story away and look at the algebra sitting inside it: the cost of one jump from a to b is $b^3 - ab^2 = b^2(b-a)$. That is (value of a function at the right endpoint) times (width of the step). The instant you factor it that way, you are no longer looking at a game, you are looking at a term of a Riemann sum for $\int t^2\,dt$, sampled at the right end of each subinterval. A finite sequence of jumps $0 = x_0 < x_1 < \cdots < x_m = 1$ has total cost $\sum_{i=1}^{m} x_i^2(x_i - x_{i-1})$, which is exactly the right endpoint Riemann sum of $f(t) = t^2$ over the partition $\{x_i\}$ of $[0,1]$.

Once you see that, "for what c can you achieve total cost c" turns into "what is the set of values a right endpoint Riemann sum of t^2 can take, ranging over every finite partition of $[0,1]$." The fields wired into this problem are elementary real analysis, Riemann sums, monotonicity, the intermediate value theorem, dressed as a discrete combinatorial game; there is no calculus beyond recognising the shape.

The first thing worth writing down, inside thirty seconds, is the pair of bounds that monotonicity of t^2 on $[0,1]$ hands you for free: since t^2 is increasing, a right endpoint sum always overestimates the true integral $1/3$, and since $t^2 \le 1$ on $[0,1]$, the sum can never exceed 1. So the achievable set sits inside $(1/3, 1]$ before a single example is constructed. The rest of the problem is proving that inclusion is equality.

Approaches, ranked worst to best

1. **try two or three jumps and solve for the cost algebraically.** *Cue:* small, explicit numbers. *Circuit:* no reusable circuit, exploratory only. The two jump path $0 \to t \to 1$ has cost $1 + t^3 - t$, which ranges only over roughly $[0.615, 1]$, never reaching down near $1/3$ for any fixed small number of jumps, and gives no proof that every intermediate value is hit. A warm up, not a route to a proof.

2. **minimise the cost over all partitions of a fixed size m using calculus, then let $m \to \infty$.** *Cue:* "total cost" invites optimisation. *Circuit:* Lagrange multipliers on the simplex of partitions. Correct in spirit, wasteful in execution: an m variable optimisation, identification of the minimising, equally spaced partition, a limit as $m \to \infty$, and then a separate argument for every intermediate value of c.

3. **recognise the sum as a Riemann sum, extract both bounds from monotonicity, then build one continuous one parameter family of constructions and invoke the intermediate value theorem once.** *Cue:* $b^2(b-a)$ factoring as (function value)(width). *Circuit:* Riemann sum bound plus a single IVT sweep. The ideal route. Both bounds cost one inequality each. Achievability costs one explicit family and one appeal to continuity, no case analysis, no separate argument for each value of c.

The ideal solution

The answer is $c \in \left(\frac{1}{3}, 1\right]$.

Every achievable cost lies in this range. Fix a finite sequence $0 = x_0 < x_1 < \cdots < x_m = 1$ and write its total cost as

$$C = \sum_{i=1}^{m} x_i^2(x_i - x_{i-1}).$$

Upper bound. Since $0 < x_i \le 1$ for every i, $x_i^2 \le 1$, so

$$C \le \sum_{i=1}^{m} (x_i - x_{i-1}) = 1,$$

with equality only if $x_i^2 = 1$ on every step of positive width, forcing $x_i = 1$ for every $i \ge 1$, hence $m = 1$. So $C = 1$ exactly for the single jump $0 \to 1$, and $C < 1$ for every path with $m \ge 2$.

Lower bound. On (x_{i-1}, x_i), t^2 is strictly less than x_i^2, since t^2 is strictly increasing and $t < x_i$ there, so

$$\int_{x_{i-1}}^{x_i} t^2 \, dt < x_i^2 (x_i - x_{i-1}).$$

Summing over i gives $\int_0^1 t^2 \, dt < C$, that is $\frac{1}{3} < C$, for every finite partition, strictly, since every subinterval has positive width.

So every achievable cost lies in $\left(\frac{1}{3}, 1\right]$.

Every value in this range is achievable. The endpoint $c = 1$ is achieved by the single jump $0 \to 1$. Fix $c \in \left(\frac{1}{3}, 1\right)$.

For $a \in [0, 1]$ and a positive integer n, consider the path taking n equal jumps from 0 to a, through $a/n, 2a/n, \ldots, a$, followed by one final jump from a to 1. Its cost is

$$C(a, n) = a^3 \cdot \frac{(n+1)(2n+1)}{6n^2} + (1 - a),$$

using $\sum_{k=1}^{n} k^2 = n(n+1)(2n+1)/6$ for the first phase and $1^3 - a \cdot 1^2 = 1 - a$ for the last jump. For fixed n, $C(\cdot, n)$ is a polynomial in a, hence continuous on $[0, 1]$, with

$$C(0, n) = 1, \qquad C(1, n) = \frac{(n+1)(2n+1)}{6n^2}.$$

The right hand quantity tends to $\frac{1}{3}$ as $n \to \infty$, so choose n large enough that $C(1, n) < c$. Since $C(0, n) = 1 > c$ and $C(1, n) < c$, the intermediate value theorem applied to the continuous function $a \mapsto C(a, n)$ on $[0, 1]$ produces some $a \in (0, 1)$ with $C(a, n) = c$ exactly. The corresponding finite jump sequence, n equal steps from 0 to a followed by one step from a to 1, realises total cost exactly c.

Hence every $c \in \left(\frac{1}{3}, 1\right]$ is achievable, and combined with the bound above, the achievable set is precisely $\left(\frac{1}{3}, 1\right]$. ■

Layer 1: WHAT

Block 1: factor the cost of a single jump as $b^2(b - a)$ and recognise it as a Riemann sum term. Block 2: extract the upper bound $C \le 1$ from $x_i \le 1$, with an equality analysis pinning down exactly the one jump path. Block 3: extract the strict lower bound $C > 1/3$ from monotonicity of t^2 against the true integral. Block 4: build a single continuous, two parameter family of constructions, equal steps up to a, then one jump to 1, whose cost at $a = 1$ approaches $1/3$ and at $a = 0$ equals 1. Block 5: invoke the intermediate value theorem once, for a well chosen fixed n, to hit any target c exactly.

The nominal route grinds through small, fixed numbers of jumps and never escapes a bounded sub-interval of $(1/3, 1]$; it cannot see $1/3$, because no fixed finite m gets arbitrarily close to it. The bifurcation point is the decision to stop treating m as fixed and instead build a family indexed by a continuous parameter a, at fixed large n, so that continuity and the intermediate value theorem finish the job in one step instead of a case by case chase.

Layer 2: HOW

Block 1. **Cue.** The cost formula $b^3 - ab^2$ factors instantly as $b^2(b - a)$. **Why natural.** This is exactly the shape, sampled function value times subinterval width, that defines a Riemann sum; recognising it is definitional, not clever, once the factoring is written down. **Link to the next block.** A sum of jump costs over a partition of $[0, 1]$ is therefore literally a right endpoint Riemann sum of t^2.

Block 2. **Cue.** $t^2 \leq 1$ on $[0,1]$. **Why natural.** Bounding each sampled value by the global maximum of f on $[0,1]$ bounds the whole sum by $\sum(x_i - x_{i-1}) = 1$, a telescoping sum that collapses to the interval's length regardless of the partition. **Link to the next block.** Equality forces every contributing x_i to equal 1, possible only for the trivial one jump path, so the upper bound is attained exactly once, sharply.

Block 3. **Cue.** t^2 is strictly increasing on $[0,1]$. **Why natural.** An increasing function, sampled at the right endpoint of each subinterval, always overestimates its own integral there, strictly, whenever the subinterval has positive width; this is the definition of an upper Riemann sum for a monotonic function. **Link to the next block.** Since every one of the finitely many subintervals has positive width, the strict inequality survives summation, giving $C > 1/3$, with no partition ever reaching equality.

Block 4. **Cue.** The lower bound is a limit, $C \to 1/3$ as the partition refines, never an equality; every value strictly above $1/3$ is needed, not just values near it. **Why natural.** A single family with one free continuous parameter is cheaper than patching together infinitely many separate near misses: fix a as the point where "many small jumps" switches to "one final big jump," and let the many small jumps approach the integral $a^3/3$ while the final jump contributes the exact, non limiting quantity $1 - a$. **Link to the next block.** As a ranges over $[0,1]$ continuously, the limiting cost $a^3/3 + (1 - a)$ sweeps continuously from 1 down to $1/3$, so the target c sits somewhere on that sweep.

Block 5. **Cue.** You need an exact value c, not a limit. **Why natural.** For any fixed finite n, $C(a,n)$ is an honest polynomial, continuous in a, and its two endpoint values already straddle c once n is large enough; the intermediate value theorem then guarantees an exact a with $C(a,n) = c$, not merely an approximation. **Link to the next block.** No further step is needed; the explicit finite sequence attached to that a and n is the answer.

Layer 3: WHY

The mathematics underneath this problem is the standard fact that for a strictly increasing integrable function, the right endpoint Riemann sum is always a genuine overestimate of the integral, on every nondegenerate partition, and the amount of overestimate shrinks to zero as the mesh refines. That single fact supplies both halves of the answer: it forces every achievable cost above $1/3$, since the overestimate is always positive, and it is exactly what lets you build constructions whose cost approaches $1/3$ as closely as you like, since the overestimate shrinks to zero. The upper endpoint 1 is a separate, easier fact: $t^2 \leq 1$ on the domain, attained only at the right edge, so it bounds every sum and is matched only by refusing to subdivide at all.

The discernment mechanism for choosing the intermediate value theorem over a case by case covering argument is this: once a target set is known to be an interval bounded by a value you can hit exactly, $c = 1$, and a value you can only approach, $c \to 1/3$, look for one continuous family connecting those two behaviours, rather than constructing each c separately. Continuity of the cost function in the construction's free parameters is not an accident, it is guaranteed because the cost is a polynomial in the jump points; whenever a competition construction is built from polynomial, or otherwise continuous, data, the intermediate value theorem is available for free, and reaching for it before reaching for an explicit inverse formula is the disciplined habit worth training. The examiners nudged toward this by writing the cost as a cubic in a and b, structurally guaranteeing continuity of any construction built from it.

Revolving concepts

- **Riemann sums, upper and lower.** For a partition of $[a, b]$, the upper sum uses the supremum of f on each subinterval, the lower sum the infimum; both converge to $\int_a^b f$ as the mesh tends to zero, and for a monotonic function the right endpoint sum coincides with the upper sum, if increasing, or the lower sum, if decreasing. See Book 1, field: Real analysis, topic: Riemann integration.
- **Strict inequality between a monotone sum and its integral.** For f strictly increasing and continuous, every finite, nondegenerate right endpoint sum sits strictly on one side of $\int f$; equality is only a limiting statement, never attained by a finite partition. This is the source of the open endpoint $1/3$ in the answer.
- **The intermediate value theorem as a sweeping tool.** Whenever a construction depends continuously on a free real parameter, and its value at the two ends of the parameter's range straddles a target, the intermediate value theorem produces an exact hit. This generalises far beyond integrals: any competition existence claim of the form "for every c in an interval, some explicit finite object achieves c exactly" is a candidate for one continuous family plus one application of the theorem, rather than a separate construction per value. See Book 1, field: Real analysis, topic: continuity and the intermediate value theorem.
- **Closed forms for power sums.** $\sum_{k=1}^{n} k^2 = n(n+1)(2n+1)/6$ is the only piece of routine algebra in the whole solution, standard A-level Further Mathematics content; its only role here is to make $C(a, n)$ an explicit polynomial, so continuity and the endpoint values can be read off directly.

Circuit Bank

- **WHEN** a cost or sum is defined as a function of the right endpoint times the width of a step, over a partition of an interval **DO** recognise it immediately as a Riemann sum and extract bounds from monotonicity of the function before attempting any construction **BECAUSE** monotonicity gives strict, free inequalities in one line, and tells you in advance which endpoint of the achievable range is attained and which is only approached
- **WHEN** an achievable set is claimed to be an entire interval, with one endpoint attained exactly and the other only approached **DO** build a single family with one continuous real parameter connecting a construction that attains the good endpoint to one that approaches the other, then apply the intermediate value theorem once for a suitably chosen fixed discretisation **BECAUSE** continuity of the construction in its parameters is often free, since the data is polynomial, so one application of the theorem replaces an unbounded case analysis
- **WHEN** a bound of the form $C \leq 1$ or $C \geq$ some constant is proved via $f(t) \leq M$ or $f(t) \geq m$ on the whole domain **DO** immediately check when equality can hold, since competition answers usually want the exact attained endpoint identified, not only the bound **BECAUSE** a bound without its equality case is an incomplete answer to a "for what values" question, exactly what this problem asks

Backtesting

1. First useful cue: $b^3 - ab^2$ factors as $b^2(b - a)$, the signature shape of a Riemann sum term.
2. Method the expression suggested: bound the sum using monotonicity of t^2, then sweep the gap with a continuous family and one application of the intermediate value theorem.

3. Where you would hesitate, and the thought to reinforce: wanting a separate explicit construction for every target c; reinforce that one continuous family plus the theorem covers infinitely many targets at once.
4. The trap the question sets: small, fixed numbers of jumps look tractable and give a plausible looking sub-interval, but they never reach the true lower bound $1/3$, so stopping there gives a wrong, too narrow answer.
5. What makes you recognise it faster next time: any sum written as a value at the right endpoint times a step width, over a partition of a fixed interval, is a Riemann sum in disguise, regardless of the surface story, games, costs, journeys.
6. The reusable decision rule: bound first from monotonicity, then connect the two bounds with one continuous family and one sweep of the intermediate value theorem, rather than constructing each achievable value separately.

Mirror drills

1. For what c can a right endpoint Riemann sum of $f(t) = t$ over finite partitions of $[0, 1]$ equal c? *Answer:* $c \in \left(\frac{1}{2}, 1\right]$, by the identical bounding and sweeping argument, using $\int_0^1 t\,dt = \frac{1}{2}$.
2. For what c can a right endpoint Riemann sum of $f(t) = e^t$ over finite partitions of $[0, 1]$ equal c? *Answer:* $c \in (e - 1, e]$, since $\int_0^1 e^t\,dt = e - 1$ and the single jump value is $f(1) \cdot 1 = e$.
3. A rope of length 1 is cut once, at a point $x \in (0, 1)$, into two pieces of lengths x and $1 - x$. For which $c \in \left[0, \frac{1}{4}\right]$ does some cut give a product of piece lengths exactly c, and which single idea from this problem proves it? *Answer:* every $c \in \left[0, \frac{1}{4}\right]$, since $x(1 - x)$ is continuous on $[0, 1]$ with values 0 at $x = 0$ and $\frac{1}{4}$ at $x = \frac{1}{2}$, so the intermediate value theorem alone, with no Riemann sum needed, already finishes it.

17.9 B3. counting mediocre subsets via arithmetic progressions

Call a subset S of $\{1, 2, \ldots, n\}$ mediocre *if it has the following property: whenever a and b are elements of S whose average is an integer, that average is also an element of S. Let $A(n)$ be the number of mediocre subsets of $\{1, 2, \ldots, n\}$. (For instance, every subset of $\{1, 2, 3\}$ except $\{1, 3\}$ is mediocre, so $A(3) = 7$.) Find all positive integers n such that $A(n+2) - 2A(n+1) + A(n) = 1$.*

Read it out loud

Two separate structures are welded together here, and you should read them as two separate problems, to be multiplied against each other later, not as one tangled question. The first is a closure condition on subsets: whenever the integer average of two members exists, it is already a member. That is a midpoint condition, and the moment you see "closed under taking the midpoint, when it is an integer," the shape to think of is an arithmetic progression, because an arithmetic progression is exactly the family of sets where midpoints of members that ought to exist do exist and land back inside the set.

The second structure is arithmetical, not combinatorial: $A(n+2) - 2A(n+1) + A(n)$ is a second difference. Whenever a counting sequence is hit with a second difference, read it as "count something pinned by two new elements," because a first difference of a counting sequence very often counts the objects containing one distinguished new element, and a second difference then counts the objects containing two.

Within the first thirty seconds, write both of these down: a guess that mediocre sets are exactly the arithmetic progressions with an odd common difference, worth checking against the given example immediately, since $\{1, 3\}$ is precisely the one two term subset of $\{1, 2, 3\}$ with an even difference, and it is precisely the one excluded set; and the observation that $A(n+2) - 2A(n+1) + A(n)$ should turn, via a shift argument, into a count of mediocre subsets pinned at both ends of $\{1, \ldots, n+2\}$.

Approaches, ranked worst to best

1. **compute $A(n)$ for small n by hand and guess a recurrence.** *Cue:* the worked example $A(3) = 7$. *Circuit:* no reusable circuit, pattern spotting only. This confirms small cases but proves nothing, and the recurrence for $A(n)$ itself is not simple; you would be fishing for a closed form that never settles into a clean shape. Ink with no end in sight.

2. **attack $A(n+2) - 2A(n+1) + A(n)$ directly by comparing raw counts of mediocre subsets of the three ground sets.** *Cue:* "second difference" as a hint to compare set families. *Circuit:* brute comparison, no bijection. The right instinct, executed clumsily: comparing raw counts rather than building an explicit bijection wastes an enormous amount of casework and never isolates the clean combinatorial meaning of the second difference.

3. **prove the closure condition classifies mediocre sets exactly as odd difference arithmetic progressions, and separately reduce the second difference to a count of subsets pinned at both endpoints via a shift bijection, then multiply the two insights.** *Cue:* the midpoint closure condition together with the shape of a second difference. *Circuit:* structural classification plus telescoping bijection. The ideal route. Each half is a short, clean argument; combined, the problem collapses to an elementary divisor counting question with almost no casework.

The ideal solution

Step 1: classify the mediocre sets.

Claim 17.9.1. *A subset S of the integers with $|S| \geq 2$ is mediocre if and only if S is a finite arithmetic progression with odd common difference.*

Proof. If. Let $S = \{a, a+d, \ldots, a+md\}$ with d odd. Take two elements $a+id$, $a+jd$ with $i < j$; their average is $a + \frac{i+j}{2}d$. Since d is odd, this is an integer exactly when $i+j$ is even, and then, writing $\ell = (i+j)/2$, we have $i \leq \ell \leq j \leq m$, so $a + \ell d \in S$. So S is mediocre.

Only if. Let $S = \{s_1 < s_2 < \cdots < s_m\}$ be mediocre, $m \geq 2$. For $m = 2$: the only pair is (s_1, s_2), and its average, if an integer, lies strictly between s_1 and s_2, so it cannot lie in $S = \{s_1, s_2\}$; mediocrity forces $s_2 - s_1$ to be odd, which is exactly an arithmetic progression of length two with odd difference.

For $m \geq 3$, induct on m. Removing the top element gives $S' = \{s_1, \ldots, s_{m-1}\}$. Any two elements of S' have average at most $s_{m-1} < s_m$, so if their average is an integer it cannot equal s_m; since S is mediocre, that average lies in S, hence in $S \cap (-\infty, s_{m-1}] = S'$. So S' is mediocre, and by induction it is an arithmetic progression with odd common difference d'. Symmetrically, removing the bottom element gives a mediocre $S'' = \{s_2, \ldots, s_m\}$, an arithmetic progression with odd common difference d''.

The sets S' and S'' share the sub-list $\{s_2, \ldots, s_{m-1}\}$. When $m \geq 4$, this sub-list has at least two elements, and, being a run of consecutive terms inside both S' and S'', it is an arithmetic progression with difference d', from S', and also with difference d'', from S''. A set of at least two numbers determines its common difference uniquely, as (largest minus smallest) divided by (count minus one), so $d' = d'' =: d$. Hence $S = \{s_1, s_1 + d, \ldots, s_1 + (m-1)d\}$ is itself an arithmetic progression with odd common difference d. When $m = 3$, the same overlap argument, applied to $S' = \{s_1, s_2\}$ and $S'' = \{s_2, s_3\}$, directly forces $s_2 - s_1 = s_3 - s_2$, and that shared difference is odd by the $m = 2$ case, which anchors the induction at its base. □

Step 2: reduce the second difference to a count of doubly pinned mediocre sets.

For $n \geq 1$, the shift $T \mapsto T + 1$ is a bijection between the mediocre subsets of $\{1, \ldots, n\}$ and the mediocre subsets of $\{2, \ldots, n+1\}$: the average of $a+1$ and $b+1$ is an integer exactly when the average of a, b is, and it equals (the average of a, b) plus one, so it lies in $T + 1$ exactly when the original average lies in T. Hence $A(n)$ also counts mediocre subsets of $\{2, \ldots, n+1\}$, which are exactly the mediocre subsets of $\{1, \ldots, n+1\}$ that do not contain 1.

Write $f(n)$ for the number of mediocre subsets of $\{1, \ldots, n\}$ that do contain 1. The paragraph above gives

$$A(n+1) = A(n) + f(n+1).$$

Applying this once with n and once with $n+1$, then subtracting,

$$A(n+2) - 2A(n+1) + A(n) = f(n+2) - f(n+1).$$

Now split the mediocre subsets of $\{1, \ldots, n+2\}$ containing 1 according to whether they also contain $n+2$. Those that do not are exactly the mediocre subsets of $\{1, \ldots, n+1\}$ containing 1, counted by $f(n+1)$, so

$$f(n+2) = f(n+1) + g(n),$$

where $g(n)$ is the number of mediocre subsets of $\{1, \ldots, n+2\}$ containing both 1 and $n+2$. Combining the last two displays,

$$A(n+2) - 2A(n+1) + A(n) = g(n).$$

Step 3: count the doubly pinned sets. Any mediocre set counted by $g(n)$ contains both 1 and $n+2$, so it has at least two elements, and by the Claim it is an arithmetic progression with odd common difference d, first term 1, last term $n+2$. This forces $d \mid (n+2-1) = n+1$, and conversely every odd divisor d of $n+1$ gives exactly one such set, $\{1, 1+d, 1+2d, \ldots, n+2\}$. So $g(n)$ equals the number of odd divisors of $n+1$.

Step 4: solve. Write $n+1 = 2^a m$ with m odd. The odd divisors of $n+1$ are exactly the divisors of m, so $A(n+2) - 2A(n+1) + A(n) = 1$ if and only if m has exactly one divisor, if and only if $m = 1$, if and only if $n+1$ is a power of 2. Since n is a positive integer, $n+1 \geq 2$, so

$$n = 2^k - 1 \text{ for some integer } k \geq 1,$$

that is, $n = 1, 3, 7, 15, 31, \ldots$. ■

Layer 1: WHAT

Block 1: read the closure condition as a midpoint condition and guess the classification, arithmetic progressions with odd common difference, checking it instantly against the given example. Block 2: prove the classification by peeling the top and bottom element and reconciling the overlap. Block 3: read the second difference as a telescoping count, via two applications of a shift bijection, reducing it to counting mediocre sets containing both 1 and $n+2$. Block 4: apply the classification to that count directly: such sets biject with odd divisors of $n+1$. Block 5: translate "exactly one odd divisor" into "$n+1$ is a power of 2" and read off n.

The nominal route tries to find $A(n)$ itself in closed form and differentiate it twice, which is possible but heavier; it never needs the classification lemma at full strength and tends to bury the divisor count inside a messier sum. The bifurcation point is the decision to work with the second difference directly, as its own combinatorial object, rather than solving for $A(n)$ first and subtracting afterwards.

Layer 2: HOW

Block 1. Cue. "Whenever the average is an integer, it lies in the set." **Why natural.** This is definitional: the family of sets closed under taking midpoints, when the midpoint exists as an integer, is precisely the family with no gap that a midpoint could expose, and the only sets with that rigidity are arithmetic progressions, with the parity of the difference controlling whether midpoints exist at all. **Link to the next block.** Guessing the classification early lets you check it instantly against the one given data point, $\{1, 3\}$ excluded because its difference 2 is even, confirming the guess before a single line of proof is written.

Block 2. Cue. Proving a classification for a set of size m from a closure property. **Why natural.** Peel an extremal element, the top, or the bottom, and note that any average computed inside the remaining elements cannot equal the element removed, since it is out of range; this makes the smaller set inherit the closure property for free, so induction applies without extra hypotheses. **Link to the next block.** Removing the top, and separately the bottom, gives two overlapping smaller arithmetic progressions; forcing their shared difference to agree on the overlap is what promotes two progressions into one progression covering everything.

Block 3. Cue. $A(n+2) - 2A(n+1) + A(n)$, a second finite difference of a counting sequence. **Why natural.** A first difference of a counting sequence over a growing ground set very often counts exactly the objects using the newly available element, because every object either avoids it, and is counted by the smaller sequence, or uses it, and is exactly the difference; a second difference then counts objects using two newly available elements, by the same argument applied twice. **Link to the next block.** The two elements pinned by the second difference here are the two ends of the ground set, 1 and $n+2$, since each application of the shift bijection is anchored at a boundary.

Block 4. Cue. A mediocre set containing two specific pinned elements, 1 and $n+2$. **Why natural.** Once mediocre sets of size at least two are known to be exactly odd difference arithmetic

progressions, a set with two fixed endpoints has no freedom left: the common difference is forced by the two endpoints, and the only question is whether that forced difference is legal, an odd divisor of the gap. **Link to the next block.** Counting odd divisors of $n+1$ is now pure elementary number theory, not combinatorics.

Block 5. Cue. "Exactly one odd divisor." **Why natural.** The odd divisors of an integer are exactly the divisors of its largest odd factor, so having only one odd divisor, namely 1 itself, is the same statement as that odd factor being 1, that is, the whole integer is a power of 2. **Link to the next block.** This names every solution directly, $n = 2^k - 1$, with no further casework.

Layer 3: WHY

Underneath both halves of this problem is the same phenomenon: a condition phrased as closure under an operation applied to pairs pins down a classical, rigid family, and once that classification is in hand, a combinatorial counting question about the family collapses into ordinary number theory. Here the operation is "take the integer average," and the rigid family is "arithmetic progression with odd step"; the finite difference machinery is a separate, purely mechanical device for turning a numerical recurrence into a statement about one specific, small sub-family, the doubly pinned sets, rather than the whole family $A(n)$.

The discernment mechanism for the classification half is to notice that a closure condition defined pairwise, with no reference to size or position, almost always forces either total rigidity, a progression, a coset, a subgroup, or total freedom, any subset at all; the parity obstruction here, that an even difference two term set already breaks the condition, is the signal you are in the rigid case, not the free one. The discernment mechanism for the finite difference half is to recognise Δ^2 applied to a counting sequence as an instruction to find a bijection first and an algebraic identity second: a first difference of A isolates "contains the new top element," and running that isolation twice, once from each end of $\{1, \ldots, n+2\}$, is a completely generic technique that applies to any counting sequence defined on growing intervals, sets or lattice paths, whenever the underlying structure is translation invariant.

Both mechanisms trace back to the same extremal peeling idea used across this trio: strip an extreme element, recognise the property survives on what is left, and reconstruct the whole object from the pieces once each piece's structure is known.

Revolving concepts

- **Midpoint closure and arithmetic progressions.** A set closed under "take the average when it happens to be an integer" is, once it has at least two elements, forced to be an arithmetic progression with odd common difference; even difference fails because the midpoint of the extreme points then exists but sits strictly between them and outside a two term set. This is a discrete cousin of midpoint convexity from real analysis, where a function satisfying $f\left(\frac{x+y}{2}\right) \leq \frac{f(x)+f(y)}{2}$ for all real x, y, plus mild regularity, is forced to be convex; the same no gaps at the midpoint rigidity is at work. See Book 1, field: Discrete structures and graphs, topic: closure conditions and rigid families.
- **Finite differences of a counting sequence.** If $A(n)$ counts configurations on a ground set grown one element at a time, and the counted property is translation invariant, $\Delta A(n) = A(n+1) - A(n)$ counts configurations using the newly adjoined element, by splitting the larger family into "avoids it," matched bijectively to the smaller family, and "uses it," the difference. Iterating, $\Delta^k A(n)$ counts configurations pinned to use k specific new elements. Generalised recursion: for any counting sequence built this way, $\Delta^k A(n)$ always has this pinned at k points combinatorial meaning; the method is always the same shift bijection

applied k times, from whichever end, or ends, the ground set is being grown. See Book 1, field: Combinatorics, topic: finite differences and bijective counting.

- **Extremal peeling induction.** Strip the largest, or smallest, element of a finite structure, show the defining property survives on the remainder, apply the inductive hypothesis, then reconcile using an overlap between two different peelings, top and bottom here. This is the same skeleton as the strong induction used elsewhere in this trio, and it is the standard route for proving a local, pairwise closure condition forces a globally rigid structure. See Book 1, field: Discrete structures and graphs, topic: extremal and structural induction.
- **Counting divisors, odd divisors specifically.** Writing $n+1 = 2^a m$ with m odd, the divisors of $n+1$ split into those dividing m, the odd ones, and the rest; the odd divisors of $n+1$ are exactly the divisors of m, and m has exactly one divisor if and only if $m = 1$. This is the standard reduction whenever a problem asks specifically about odd divisors rather than all divisors. See Book 1, field: Number theory, topic: the divisor function.

Circuit Bank

- **WHEN** a subset is defined by a pairwise closure condition of the form: if two elements satisfy a relation, then a derived element is also a member **DO** guess immediately that the condition classifies a classical rigid family, a progression, a coset, an interval, and check the guess against any worked example given in the statement **BECAUSE** pairwise closure conditions with no reference to position or size almost always force total rigidity, and the given example is the cheapest possible confirmation
- **WHEN** a problem hands you a finite difference, first or second, of a counting sequence defined on a growing or shrinking ground set **DO** look for a bijection showing the first difference counts configurations that use one specific new element, then iterate it for higher differences **BECAUSE** finite differences of counting sequences are shift bijections in disguise, and unwinding them converts an opaque numerical recurrence into an explicit, smaller combinatorial family
- **WHEN** proving a classification lemma for structures of every size at least two **DO** induct by peeling the top and, separately, the bottom element, then force the two resulting sub-structures to agree on their overlap **BECAUSE** a single ended peel only proves half the rigidity; the overlap between two peels is what pins the whole object down uniquely
- **WHEN** a final answer reduces to a divisor count equal to exactly one **DO** write the relevant integer as a power of the excluded prime times the coprime remainder, since "exactly one divisor" is only ever true of the remainder equalling 1 **BECAUSE** this is the fastest, least error prone way to convert a divisor count condition into an explicit description of the integers satisfying it

Backtesting

1. First useful cue: the midpoint closure condition, immediately suggesting arithmetic progressions, confirmed against the one given exclusion, $\{1, 3\}$.
2. Method the expression suggested: treat $A(n+2) - 2A(n+1) + A(n)$ as a doubly pinned count via two shift bijections, not as a numerical sequence to solve for directly.
3. Where you would hesitate, and the thought to reinforce: the overlap argument in the classification proof, that two progressions sharing at least two points must share a common difference; reinforce by remembering a difference is determined by any two of its terms.

4. The trap the question sets: it invites you to find $A(n)$ in closed form first, a genuinely harder task than the problem needs, when the second difference alone is the clean object.
5. What makes you recognise it faster next time: any phrase of the shape "whenever a pairwise relation holds, a derived element is also in the set" is the fixed signature of a rigid family classification problem.
6. The reusable decision rule: classify the local closure condition first, reduce any finite difference of the counting sequence to a pinned sub-family second, then multiply the two results together.

Mirror drills

1. Prove that a subset S of $\mathbb{Z}$ with $|S| \geq 2$, closed under "if $a, b \in S$ and $a \equiv b \pmod 2$ then $(a+b)/2 \in S$," must be an arithmetic progression with odd common difference. *Answer:* identical proof to the Claim above, peeling the extreme elements of S and reconciling the overlap; the ground set being all of $\mathbb{Z}$ rather than $\{1, \ldots, n\}$ changes nothing.
2. Let $B(n)$ count the subsets of $\{1, \ldots, n\}$ containing no two consecutive integers. Find a bijective meaning for $B(n+1) - B(n)$, and use it to re-derive the recurrence for B. *Answer:* $B(n+1) - B(n)$ counts such subsets containing $n+1$; any such subset is $n+1$ together with a valid subset of $\{1, \ldots, n-1\}$ avoiding n, giving $B(n+1) - B(n) = B(n-1)$, that is, $B(n+1) = B(n) + B(n-1)$.
3. For which positive integers m does m have exactly three positive divisors, and why? *Answer:* exactly when $m = p^2$ for a prime p, since a divisor count of three forces $m = p^e$ for a single prime with $e + 1 = 3$.

17.10 B4. Circle averages and the dimension of balanced polynomials

Say that a polynomial with real coefficients in two variables, x, y, is balanced *if the average value of the polynomial on each circle centred at the origin is* 0*. The balanced polynomials of degree at most* 2009 *form a vector space V over $\mathbb{R}$. Find the dimension of V.*

Read it out loud

Read the command terms before you read the mathematics. *Vector space*. *Dimension*. *Find*. Those three words tell you that this is a linear algebra question wearing an analysis costume. The integral is not the object of study; the integral only exists to manufacture a linear functional, and your job is to count how many independent linear conditions that functional family imposes on a space whose dimension you already know. Say it out loud in that form: the ambient space is all real polynomials in x and y of degree at most 2009, which has dimension $\binom{2011}{2} = 2021055$, and the answer will be 2021055 minus the number of independent conditions. Once you have said that sentence, the problem has stopped being about circles.

Now read the condition. *The average on each circle centred at the origin is zero*. Two guiding lights are planted here. The first is *centred at the origin*: circles about the origin are the level sets of $x^2 + y^2$, so the natural coordinates are polar, and the natural grading on polynomials is by homogeneous degree, because a homogeneous polynomial of degree i scales by r^i when you scale the point by r. The second is *each*: the condition is indexed by a continuous parameter $r > 0$, so you appear to have uncountably many linear conditions on a finite dimensional space. Uncountably many conditions on a finite dimensional space must collapse, and the only question is how. That is the entire problem, and you should have it inside thirty seconds.

The fields wired into this question are: graded vector spaces and rank nullity from linear algebra; a single one variable integral over a period; parity, in the shape of the antipodal map $(x, y) \mapsto (-x, -y)$; and, behind the curtain, the mean value property of harmonic functions, which is the reason the answer is as clean as it is. What is really being tested is whether you will treat r as a nuisance variable to be integrated away, or as a formal variable whose powers separate the conditions for you.

Within thirty seconds you should have written

$$M_P(r) \;=\; \frac{1}{2\pi}\int_0^{2\pi} P(r\cos\theta, r\sin\theta)\, d\theta,$$

noted that arclength on the circle of radius r is $r\,d\theta$ and that the factor r cancels when you average, so this really is the average value, and then written $P = \sum_i P_i$ with P_i homogeneous. Everything after that is bookkeeping.

Approaches, ranked worst to best

1. **Coefficient bash.** *Cue:* the words "vector space" and "dimension" push you to set up a linear system. *Circuit:* write one linear equation per condition and row reduce. You expand P in the 2021055 monomials, compute $\int_0^{2\pi} \cos^a\theta \sin^b\theta\, d\theta$ for every $a + b \le 2009$, and stare at an infinite matrix indexed by r. Every individual step is legal and the ink is unbounded. Worse, without grading you have no handle on the rank at all, so the computation cannot end. This route fails not because it is wrong but because it never terminates.

2. **Fourier expansion on the circle.** *Cue:* an average over a circle is the zeroth Fourier coefficient. *Circuit:* convert $\cos^a\theta \sin^b\theta$ into a combination of $\cos k\theta$ and $\sin k\theta$ and read off the constant term. This is correct and it makes the parity fact obvious, since $\cos^a\theta \sin^b\theta$ with $a + b$ odd has only odd harmonics. It costs product to sum ink, and you still have to organise the count by

degree at the end, so you end up importing the grading anyway. Middle of the table: right idea, wrong order of operations.

3. **Harmonic decomposition.** *Cue:* "average over circles centred at the origin" is literally the mean value property. *Circuit:* write $P_i = \bigoplus_j (x^2+y^2)^j H_{i-2j}$ with H_k harmonic and use that harmonic polynomials of positive degree have zero circle average. This is the true reason the answer is what it is, it produces an explicit basis of V, and it generalises to spheres in $\mathbb{R}^m$ without a single new idea. It ranks below the next route only because in the arena it costs you the decomposition theorem, which you must either quote or prove, and you do not need it to get the number.

4. **Grade by degree, then one witness per even degree.** *Cue:* "polynomial", "each circle", "dimension". *Circuit:* homogeneous decomposition, a polynomial in r with infinitely many roots is zero, parity, rank nullity. Best. Three structural lines turn a continuum of conditions into 2010 conditions, parity kills 1005 of them for free, and each survivor is shown nontrivial by the single witness $(x^2+y^2)^{i/2}$. No integral is ever computed. This is the route that finishes in half a page.

The ideal solution

Let W be the space of real polynomials in x, y of degree at most 2009. Its monomials $x^a y^b$ with $a+b \le 2009$ number $\binom{2011}{2} = 2021055$, so $\dim W = 2021055$. For $i \ge 0$ let $W_i \subset W$ be the subspace of homogeneous polynomials of degree i, so $\dim W_i = i+1$ and $W = \bigoplus_{i=0}^{2009} W_i$; write $P = \sum_{i=0}^{2009} P_i$ for the corresponding unique decomposition.

The circle of radius $r > 0$ centred at the origin carries arclength $r\,d\theta$ and has total length $2\pi r$, so the average value of P on it is

$$M_P(r) = \frac{1}{2\pi} \int_0^{2\pi} P(r\cos\theta, r\sin\theta)\, d\theta.$$

Define the linear functional $\lambda_i : W_i \to \mathbb{R}$ by $\lambda_i(Q) = \frac{1}{2\pi}\int_0^{2\pi} Q(\cos\theta, \sin\theta)\, d\theta$. Homogeneity gives $Q(r\cos\theta, r\sin\theta) = r^i Q(\cos\theta, \sin\theta)$, so

$$M_P(r) = \sum_{i=0}^{2009} r^i\, \lambda_i(P_i) \qquad (r > 0).$$

The right hand side is a polynomial in r of degree at most 2009. A real polynomial vanishing at infinitely many points is the zero polynomial, so $M_P(r) = 0$ for all $r > 0$ if and only if $\lambda_i(P_i) = 0$ for every i. Hence

Claim 17.10.1. *$P \in V$ if and only if $\lambda_i(P_i) = 0$ for $i = 0, 1, \ldots, 2009$, and consequently $V = \bigoplus_{i=0}^{2009} \ker \lambda_i$ and $\dim V = \sum_{i=0}^{2009} \dim\ker \lambda_i$.*

The second sentence follows because the condition is imposed on each graded piece separately and the decomposition $W = \bigoplus W_i$ is direct.

Odd i. If i is odd and $Q \in W_i$ then $Q(-u,-v) = -Q(u,v)$. Substituting $\theta \mapsto \theta + \pi$ permutes the interval of integration modulo the period and sends $(\cos\theta, \sin\theta)$ to $(-\cos\theta, -\sin\theta)$, so $\lambda_i(Q) = -\lambda_i(Q)$, whence $\lambda_i(Q) = 0$. Thus $\lambda_i \equiv 0$ and $\dim\ker\lambda_i = i+1$.

Even i. Take $Q = (x^2+y^2)^{i/2} \in W_i$. On the unit circle $Q \equiv 1$, so $\lambda_i(Q) = 1 \neq 0$. Therefore λ_i is a nonzero linear functional on W_i, its image is all of $\mathbb{R}$, and rank nullity gives $\dim\ker\lambda_i = (i+1) - 1 = i$.

Summing, with the odd degrees $i = 1, 3, \ldots, 2009$ contributing $i+1$ each and the even degrees $i = 0, 2, \ldots, 2008$ contributing i each,

$$\dim V = \sum_{j=1}^{1005} 2j + \sum_{j=0}^{1004} 2j = 1005\cdot 1006 + 1004\cdot 1005 = 1011030 + 1009020.$$

Equivalently, λ_i is a nontrivial condition exactly for the 1005 even values $i = 0, 2, \ldots, 2008$, so $\dim V = \dim W - 1005 = 2021055 - 1005$. Either way,

$$\dim V = \boxed{2020050}.$$

A caveat worth clarifying: a circle has positive radius, so $r = 0$ was never in play, but even if you insisted on including the degenerate circle $\{(0,0)\}$ the condition $P(0,0) = 0$ is exactly $\lambda_0(P_0) = 0$, which is already on the list. The answer is unchanged.

Layer 1: WHAT

Block 1. Read "vector space" and "dimension" and reframe: V is the common kernel of a family of linear functionals, and the answer is $\dim W$ minus the number of independent conditions.

Block 2. Write the average as $M_P(r)$ and observe that the hypothesis is one linear condition for each $r > 0$, a continuum.

Block 3. Split P into homogeneous parts; homogeneity pulls the radius out of the integral as r^i.

Block 4. Recognise $M_P(r)$ as a polynomial in r; vanishing on $(0, \infty)$ forces every coefficient to vanish. The continuum has collapsed to 2010 conditions, one per degree.

Block 5. Parity: odd degree parts are killed by the antipodal map, so 1005 of those conditions are vacuous.

Block 6. Even degree: exhibit $(x^2 + y^2)^{i/2}$ as a witness that $\lambda_i \neq 0$, so each even condition cuts the dimension by exactly one.

Block 7. Count: $2021055 - 1005 = 2020050$.

The nominal route is Block 1 followed by an immediate assault on $\int_0^{2\pi} \cos^a\theta \sin^b\theta \, d\theta$ for all $a + b \leq 2009$, then an attempt to find the rank of the resulting system by hand. The bifurcation point is Block 3 into Block 4: the exact line where you decide whether r is a nuisance parameter to be integrated away, or a formal variable whose powers are the whole point. Treat r as formal and the problem is over in three lines. Treat r as a nuisance and you are still computing Wallis integrals at the end of the session.

Layer 2: HOW

Block 1. Cue. "The balanced polynomials of degree at most 2009 form a vector space V." The examiners told you it is a vector space. They did not have to. **Why natural.** A set defined by "some average vanishes" is a subspace precisely because averaging is linear; the sentence in the question is a hint that you should be hunting for the functional whose kernel this is, not for a description of the polynomials. **Link to the next block.** Once you are hunting for a functional, you must write the functional down, which forces Block 2.

Block 2. Cue. "on each circle centred at the origin". **Why natural.** Circles centred at the origin are the orbits of the rotation group, so the parametrisation $(r\cos\theta, r\sin\theta)$ is not a choice, it is the only description under which the family of circles becomes a single formula with a parameter. Averaging with respect to arclength is averaging in θ, because the arclength element $r\,d\theta$ has a constant density in θ. **Link to the next block.** You now have a function M_P of r, and a condition on all of $(0, \infty)$. Any time a condition is indexed by a continuous parameter you ask how the parameter enters, which is Block 3.

Block 3. Cue. P is a polynomial and the parameter enters as a scaling. **Why natural.** Scaling is exactly what homogeneous pieces are designed to absorb: $P_i(rv) = r^i P_i(v)$ is the definition. So decomposing by degree diagonalises the r dependence. Nothing else in your toolkit turns a scaling into a monomial. **Link to the next block.** Now $M_P(r) = \sum r^i \lambda_i(P_i)$, and you are staring at a polynomial identity in r.

Block 4. **Cue.** "$= 0$ for every r" with the left side a polynomial in r. **Why natural.** A nonzero real polynomial of degree d has at most d roots; infinitely many roots forces every coefficient to vanish. This is the standard device for converting a statement quantified over a continuum into finitely many statements, and it is the reason the problem is finite dimensional in the first place. **Link to the next block.** The conditions are now separated by degree, so you can treat each degree independently, and the cheapest thing to test on any linear condition is a symmetry.

Block 5. **Cue.** Odd degree, and an integral over a full period. **Why natural.** The circle is invariant under the antipodal map and a homogeneous polynomial of odd degree is anti invariant under it. An anti invariant function integrated over an invariant domain integrates to zero. You should test this reflex on every symmetric integral before computing anything. **Link to the next block.** Half the conditions are gone at zero cost. The remaining question is whether the even conditions are genuine, and to prove a linear functional is nonzero you produce one input it does not kill.

Block 6. **Cue.** Even degree i, and you want a polynomial whose circle average is visibly nonzero. **Why natural.** On the circle you want an integrand of constant sign, and the only homogeneous polynomial that is constant on the unit circle is a multiple of $(x^2 + y^2)^{i/2}$, which exists exactly when i is even. It is the object the problem is built around, since circles are its level sets. You do not need the value of the integral, only that it is not zero. **Link to the next block.** Rank nullity converts "nonzero functional" into "kernel of codimension one", and the count is immediate.

Block 7. **Cue.** A finite sum with a parity split. **Why natural.** There are 1005 even degrees in $\{0, 1, \ldots, 2009\}$ and 1005 odd ones. Count the conditions, not the kernels, and the arithmetic is one subtraction. **Link to the next block.** Done. The answer is 2020050.

Layer 3: WHY

The mathematical truth underneath is that averaging over the circle is the projection of a function onto its rotation invariant part. If you average a polynomial over the circle you are integrating it against the uniform measure on the orbit of the rotation group $SO(2)$, and this operator kills everything that is not invariant. So the number of conditions in degree i is exactly the dimension of the rotation invariant subspace of W_i, and that dimension is at most one.

Here is the proof, and it is short enough that you should keep it in the arena. Suppose $Q \in W_i$ is constant, say equal to c, on the unit circle. If i is odd then $Q(-v) = -Q(v)$ forces $c = -c$, so $c = 0$; then Q vanishes on the unit circle, and since every nonzero point of the plane is a positive multiple of a unit vector and $Q(tv) = t^i Q(v)$, Q vanishes identically. If i is even then $Q - c(x^2 + y^2)^{i/2}$ is homogeneous of degree i and vanishes on the unit circle, hence vanishes identically by the same scaling argument, so $Q = c(x^2 + y^2)^{i/2}$. In other words the invariants are generated by $x^2 + y^2$, the invariant subspace of W_i is one dimensional for even i and zero for odd i, and that single number is the entire content of the problem.

The discernment mechanism, the thing you actually train, is this. When a hypothesis is quantified over a continuous parameter, do not integrate the parameter away; isolate the dependence on it and ask what kind of function of the parameter you are looking at. If it is a polynomial, a power series, or an analytic function, the continuum of conditions collapses to the coefficients, and the collapse is usually the whole problem. You saw the same collapse in every generating function argument you have ever done: "equal for all x in an interval" means "equal coefficientwise".

The second trained instinct is cheaper still: to find the dimension of the kernel of a linear map you never need to describe the kernel. You need the rank, and the rank of a functional is one as soon as you can name a single input it does not kill. Candidates lose this problem by trying to write down a basis of V. You are not asked for a basis. Be lazy: name a witness, quote rank nullity, count.

What were the examiners nudging toward? Three things. They said "vector space" so you would look for linearity. They said "centred at the origin" so you would use polar coordinates and grading rather than a general position argument. They chose degree 2009, an odd number, so that the even and odd degrees split as 1005 and 1005 and the final arithmetic stays clean. The number 2009 is otherwise doing no work.

Revolving concepts

Average value on a curve. The average of F over a curve γ is $\frac{1}{\text{length}(\gamma)} \int_\gamma F\, ds$. For the circle of radius r about the origin, $ds = r\, d\theta$ and the length is $2\pi r$, so the average is $\frac{1}{2\pi} \int_0^{2\pi} F\, d\theta$ and the radius cancels. Note that the cancellation is special to circles about the origin; for a circle about another centre the density in θ is still constant but the polynomial no longer scales, which is why the problem specifies the centre. See Book 1, field: Calculus and integrals, topic: line integrals and averages.

Graded decomposition. $\mathbb{R}[x, y] = \bigoplus_{i \geq 0} W_i$ with $\dim W_i = i + 1$. In m variables $\dim W_i = \binom{i+m-1}{m-1}$ and the number of monomials of degree at most d is $N_m(d) = \binom{d+m}{m}$. That count satisfies the Pascal recursion $N_m(d) = N_m(d-1) + N_{m-1}(d)$, with $N_1(d) = d + 1$: split on whether the last variable appears. Every "how many monomials" count on the Putnam is this recursion or a stars and bars argument for the same thing. See Book 1, field: Combinatorics, topic: stars and bars.

Vanishing polynomials. A real polynomial with infinitely many roots is identically zero. The two dimensional version used above: a homogeneous polynomial vanishing on the unit circle vanishes identically, because $Q(tv) = t^i Q(v)$. See Book 1, field: Polynomials, topic: identity theorem.

Rank nullity and codimension. For linear $T : U \to \mathbb{R}$, either $T = 0$ and $\ker T = U$, or $T \neq 0$ and $\dim \ker T = \dim U - 1$. There is no third case, which is why one witness settles the matter. See Book 1, field: Linear algebra, topic: rank.

Harmonic polynomials and the mean value property. Let $\mathcal{H}_k$ be the harmonic homogeneous polynomials of degree k in two variables; $\mathcal{H}_0 = \langle 1 \rangle$ and $\dim \mathcal{H}_k = 2$ for $k \geq 1$, spanned by the real and imaginary parts of $(x + iy)^k$. Every homogeneous P_i decomposes as $P_i = \sum_j (x^2 + y^2)^j H_{i-2j}$ with $H_k \in \mathcal{H}_k$, and dimensions check: $2\lfloor i/2 \rfloor + [i \text{ even}] = i + 1$. Since a harmonic function has mean value equal to its value at the centre, and $H_k(0, 0) = 0$ for $k \geq 1$, only the j with $i - 2j = 0$ survives averaging. So P is balanced exactly when the coefficient of $(x^2 + y^2)^{i/2}$ in the decomposition of each even part P_i is zero, which exhibits an explicit basis of V and recovers the count. See Book 1, field: Real analysis, topic: harmonic functions.

Explicit circle moments. For nonnegative integers a, b,

$$\frac{1}{2\pi} \int_0^{2\pi} \cos^{2a}\theta \sin^{2b}\theta\, d\theta = \frac{\binom{2a}{a}\binom{2b}{b}}{4^{a+b}\binom{a+b}{a}},$$

and the average is 0 if either exponent is odd. You never need these values here, but they are the explicit form of λ_i, and they let you write a basis of $\ker \lambda_i$ by hand. See Book 1, field: Calculus and integrals, topic: Wallis integrals.

The generalisation you should carry away. In $\mathbb{R}^m$, with spheres about the origin and polynomials of degree at most d, the identical argument gives one condition per even degree at most d, so the balanced space has dimension $\binom{d+m}{m} - (\lfloor d/2 \rfloor + 1)$. For $m = 2$, $d = 2009$ that is $2021055 - 1005$. The invariants of the rotation group are generated by $x_1^2 + \cdots + x_m^2$ in every dimension, and that single fact is the whole engine.

Circuit Bank

- **WHEN** a hypothesis is quantified over all values of a continuous parameter **DO** isolate the parameter dependence and check whether it is polynomial or analytic; if so, set every

coefficient to zero **BECAUSE** a nonzero polynomial has finitely many roots, so a continuum of conditions collapses to finitely many coefficient conditions

- **WHEN** circles or spheres centred at the origin appear in a polynomial problem **DO** decompose into homogeneous parts **BECAUSE** homogeneity turns the radius into the scalar factor r^i, which separates the degrees and diagonalises the whole family of conditions
- **WHEN** you must show a linear condition is nontrivial **DO** name one explicit object the functional does not kill **BECAUSE** a nonzero functional has kernel of codimension exactly one, so a single witness converts the whole question into a subtraction
- **WHEN** you integrate over a domain invariant under a symmetry **DO** split the integrand into invariant and anti invariant parts and discard the anti invariant part **BECAUSE** anti invariant functions integrate to zero over invariant domains, which is usually half the work for one line of writing
- **WHEN** you are asked for the dimension of a space of objects, not for the objects **DO** compute the rank of the defining conditions and subtract **BECAUSE** describing the kernel costs a basis; counting the conditions costs a witness

Backtesting

1. First useful cue: "vector space" and "dimension", which reframe the problem as counting independent linear conditions, plus "each circle centred at the origin", which fixes polar coordinates and the grading by homogeneous degree.
2. Method the expression suggested: write the average as $M_P(r) = \sum_i r^i \lambda_i(P_i)$, treat r as a formal variable, and read off one condition per degree.
3. Where you would hesitate, and the thought to reinforce: the temptation is to start computing circle moments of monomials. Reinforce this instead: you never need the value of a linear functional, only whether it is zero, and the cheapest proof that it is not zero is one witness. Here the witness $(x^2 + y^2)^{i/2}$ is handed to you by the geometry.
4. The trap the question sets: the continuum of radii, which looks like infinitely many conditions and invites either despair or a numerical assault. The second trap is degree 2009, which invites you to think the size of the number matters; it does not, only its parity does.
5. What makes you recognise it faster next time: the pattern "linear condition indexed by a scaling parameter, applied to polynomials". Whenever scaling is the parameter, grade by degree first and read the conditions off the coefficients.
6. The reusable decision rule: a condition holding for all values of a continuous parameter is a coefficient identity in disguise; grade, separate, then count ranks with witnesses.

Mirror drills

1. Find the dimension of the space of real polynomials in x, y of degree at most d whose average on every circle centred at the origin is zero. *Answer:* $\binom{d+2}{2} - (\lfloor d/2 \rfloor + 1)$, by the identical grading argument; check it against $d = 2009$.
2. Show that the average of $x^a y^b$ over the unit circle is zero unless a and b are both even, and that it is strictly positive when both are even. *Answer:* apply $x \mapsto -x$ or $y \mapsto -y$, which preserves the circle and flips the sign when the corresponding exponent is odd; when both are even the integrand is nonnegative and not identically zero.

3. In $\mathbb{R}^3$, find the dimension of the space of polynomials in x, y, z of degree at most 4 whose average on every sphere centred at the origin is zero. *Answer:* $\binom{7}{3} - 3 = 35 - 3 = 32$, since the even degrees $0, 2, 4$ each contribute exactly one condition and the witness in degree i is $(x^2 + y^2 + z^2)^{i/2}$.

17.11 B5. A rational differential equation forced to diverge

Let $f : (1, \infty) \to \mathbb{R}$ be a differentiable function such that

$$f'(x) = \frac{x^2 - f(x)^2}{x^2\,(f(x)^2 + 1)} \qquad \text{for all } x > 1.$$

Prove that $\lim_{x\to\infty} f(x) = \infty$.

Read it out loud

A first order differential equation, rational in both variables, and no initial condition. The missing initial condition is the loudest cue in the question. You are not asked what f is, you are asked what f does, and the answer must be the same for every solution, so any route that tries to produce a formula is doomed before it starts. This is a qualitative question: the fields are real analysis (monotonicity from the sign of a derivative, the mean value theorem) and asymptotics, with a thin layer of algebra on top to find the right unknown.

Now read the right hand side out loud. "x^2 minus f squared, over x squared times f squared plus one." Two guiding lights. First, the numerator $x^2 - f^2$ vanishes exactly on the lines $f = \pm x$, so those two lines are the nullclines and they are what the whole picture will be organised around: above and below them the sign of f' flips. Second, and far more important, the denominator contains the factor $f^2 + 1$, and $f^2 + 1$ is exactly $\frac{d}{df}\left(f + \frac{1}{3}f^3\right)$. A denominator that is a polynomial in the unknown function is not an obstacle, it is an instruction: multiply by it.

So within thirty seconds you should have written two rearrangements, both trivial, both decisive:

$$f'(x)\left(1 + f(x)^2\right) = 1 - \frac{f(x)^2}{x^2}, \qquad f'(x) + \frac{1}{x^2} = \frac{x^2 + 1}{x^2\,(1 + f(x)^2)} > 0.$$

The first says that $g = f + \frac{1}{3}f^3$ has a derivative with no denominator at all and bounded above by 1. The second says that $f - \frac{1}{x}$ is strictly increasing, which is your lower bound on f, free of charge. Every remaining line of the solution is squeezed out of those two facts.

What is really being tested: whether you convert an equation you cannot solve into a statement about a quantity that is monotone. Monotone plus unbounded is the only cheap machinery for proving divergence, and finding the right monotone quantity is the entire skill.

Approaches, ranked worst to best

1. **Solve the equation.** *Cue:* "differentiable function such that $f' = \ldots$" reads like an invitation to integrate. *Circuit:* separate the variables or find an integrating factor. It is not separable: the term f^2/x^2 genuinely couples the variables, and the equation is not linear in f. The homogeneous substitution $f = xu$ produces $xu' + u = \frac{1-u^2}{1+x^2u^2}$, which is not autonomous and is worse than what you started with. Maximum ink, zero progress, and it also misses the point that no initial condition was given.

2. **Crude bounds and a monotone limit.** *Cue:* $|f'| \le 1$ style estimates. *Circuit:* trap f between barriers, deduce eventual monotonicity, then rule out a finite limit. This works and is the official first route: bound f below using $f' > -1/x^2$, prove $f(x) < x$ for large x by taking the least crossing point of f and x and contradicting it, conclude $f' > 0$ eventually, so a limit L exists, and if L were finite then $f' \to 1/(1 + L^2) > 0$ would force f to grow linearly. Correct, but it is three separate arguments plus a least element argument that is easy to write loosely and lose a point on.

3. **The nullcline pair $x + f$ and $x - f$.** *Cue:* the numerator vanishes on $f = \pm x$, so the natural coordinates are $x \pm f$. *Circuit:* show both combinations have a derivative of constant sign.

Compute: $(x+f)' = \frac{2x^2+f^2(x^2-1)}{x^2(1+f^2)} > 0$ and $(x-f)' = \frac{f^2(x^2+1)}{x^2(1+f^2)} \geq 0$. Both are one line and both are visibly signed, which gives $|f(x)| < x$ eventually with no least crossing point argument. Better than the previous route, but you still need a third argument to handle f itself, so it is not the shortest path.

4. **The antiderivative substitution $g = f + \frac{1}{3}f^3$.** *Cue:* the denominator $1+f^2$ is precisely $\frac{d}{df}\left(f+\frac{1}{3}f^3\right)$. *Circuit:* multiply through by the denominator so that the left side becomes an exact derivative. Best. One substitution removes every denominator involving f and produces $g' = 1 - f^2/x^2$, whose upper bound $g' \leq 1$ is immediate and whose limit is 1 as soon as you know f is sublinear. Upper bound, sublinearity, divergence of g and divergence of f all fall out in sequence, with no case analysis, no least element and no integration.

The ideal solution

Throughout, $x > 1$. Write f for $f(x)$ where no confusion arises. Two rearrangements of the given equation will be used repeatedly, and I refer to them as identity (i) and identity (ii):

$$\text{(i)} \qquad f'(x) + \frac{1}{x^2} = \frac{x^2 - f^2 + \left(1+f^2\right)}{x^2\left(1+f^2\right)} = \frac{x^2+1}{x^2\left(1+f^2\right)} > 0,$$

$$\text{(ii)} \qquad f'(x)\left(1+f(x)^2\right) = \frac{x^2 - f(x)^2}{x^2} = 1 - \frac{f(x)^2}{x^2}.$$

Step 1: f is bounded below. By identity (i), the function $x \mapsto f(x) - \frac{1}{x}$ has derivative $f'(x) + \frac{1}{x^2} > 0$ on $(1,\infty)$, hence is strictly increasing there. Therefore, for $x \geq 2$,

$$f(x) - \frac{1}{x} \;\geq\; f(2) - \frac{1}{2}, \qquad \text{so} \qquad f(x) \;>\; f(2) - \frac{1}{2} \;=:\; m.$$

Step 2: f grows at most like $x^{1/3}$. Put

$$g(x) = f(x) + \tfrac{1}{3}f(x)^3.$$

Then g is differentiable and, by the chain rule and identity (ii),

$$g'(x) = f'(x)\left(1+f(x)^2\right) = 1 - \frac{f(x)^2}{x^2} \;\leq\; 1.$$

Hence $x \mapsto g(x) - x$ has nonpositive derivative and so is non increasing on $(1,\infty)$; for $x \geq 2$ this gives $g(x) - x \leq g(2) - 2 =: K$. Combining with Step 1,

$$\tfrac{1}{3}f(x)^3 = g(x) - f(x) \;\leq\; (x+K) - m, \qquad \text{that is} \qquad f(x)^3 \;\leq\; 3x + D, \quad D := 3(K-m).$$

Set $X_0 = \max\{2, |D|\}$. For $x \geq X_0$ we have $3x + D \leq 4x$, so

$$f(x) \;\leq\; (4x)^{1/3} \qquad (x \geq X_0).$$

Step 3: $g'(x) \to 1$. For $x \geq X_0$ we have $m < f(x) \leq (4x)^{1/3}$, hence $f(x)^2 \leq \max\{m^2, (4x)^{2/3}\}$ and therefore

$$0 \;\leq\; \frac{f(x)^2}{x^2} \;\leq\; \frac{\max\{m^2, (4x)^{2/3}\}}{x^2} \;\longrightarrow\; 0 \qquad (x \to \infty).$$

By identity (ii), $g'(x) = 1 - f(x)^2/x^2 \to 1$.

Step 4: $g(x) \to +\infty$. Choose $X_1 \geq X_0$ with $g'(x) \geq \frac{1}{2}$ for all $x \geq X_1$. For $x > X_1$ the mean value theorem gives $\xi \in (X_1, x)$ with

$$g(x) - g(X_1) = g'(\xi)(x - X_1) \;\geq\; \tfrac{1}{2}(x - X_1),$$

so $g(x) \geq g(X_1) + \frac{1}{2}(x - X_1) \to +\infty$.

Step 5: $f(x) \to +\infty$. The map $\phi(t) = t + \frac{1}{3}t^3$ is strictly increasing on $\mathbb{R}$, since $\phi'(t) = 1 + t^2 > 0$, and $g = \phi \circ f$. Let $M > 0$. By Step 4 there is $X_2 \geq X_1$ with $g(x) > \phi(M)$ for all $x \geq X_2$. If some $x \geq X_2$ had $f(x) \leq M$, then $g(x) = \phi(f(x)) \leq \phi(M)$, a contradiction. Hence $f(x) > M$ for all $x \geq X_2$. As M was arbitrary,

$$\lim_{x \to \infty} f(x) = +\infty. \qquad \blacksquare$$

Remark 17.11.1. The same two rearrangements give the sharp asymptotic at no extra cost. From $g' \to 1$ and the mean value theorem applied to $g(x) - x$ one gets $g(x)/x \to 1$. Since $f \to \infty$, $g = \frac{1}{3}f^3\left(1 + 3f^{-2}\right) \sim \frac{1}{3}f^3$, so $\frac{1}{3}f(x)^3 \sim x$ and therefore

$$f(x) \sim (3x)^{1/3} \qquad (x \to \infty),$$

independently of which solution f you were handed. Numerically, solutions started anywhere between $f = -40$ and $f = 60$ near $x = 1$ all satisfy $f(x)/(3x)^{1/3} \to 1$; a solution that starts far below zero simply takes until $x \approx |f_0|^3/3$ to turn round, which is exactly what $g(x) - x$ non increasing predicts.

Layer 1: WHAT

Block 1. No initial condition and no closed form: the target is qualitative, so hunt for a quantity with a monotone or asymptotically constant derivative.

Block 2. The denominator $1 + f^2$ is the derivative of $f + \frac{1}{3}f^3$, so multiply through by it rather than dividing by it.

Block 3. That yields $g' = 1 - f^2/x^2$, an equation with no denominator, and immediately $g' \leq 1$.

Block 4. $g' \leq 1$ gives $g(x) \leq x + K$, which bounds f above by roughly $(3x)^{1/3}$, but only once f is known to be bounded below.

Block 5. The lower bound comes from the other rearrangement: $f' + \frac{1}{x^2} > 0$, so $f - \frac{1}{x}$ is increasing and f is bounded below by a constant.

Block 6. Feed the sublinear bound back into $g' = 1 - f^2/x^2$ to get $g' \to 1$, hence g eventually increases at rate at least $\frac{1}{2}$ and $g \to \infty$.

Block 7. $\phi(t) = t + \frac{1}{3}t^3$ is increasing, so $g \to \infty$ transfers to $f \to \infty$.

The nominal route is Block 1 followed by an attempt to separate variables, or by the estimate $|f'| \leq 1$ and a hope that boundedness of f' says something about f at infinity (it says nothing). The bifurcation point is inside Block 2: the line at which you decide whether to divide by $1 + f^2$ in order to isolate f', or to multiply by it in order to build an exact derivative. Dividing keeps the equation unsolvable. Multiplying converts it into a perturbation of $g' = 1$, which is the same as saying $f + \frac{1}{3}f^3 = x + c$, and that is the answer in disguise.

Layer 2: HOW

Block 1. **Cue.** A differential equation with no initial condition, and the demand is a limit. **Why natural.** If no initial condition is supplied then the conclusion must hold for every solution, so the conclusion cannot depend on any constant of integration. That rules out formula hunting on logical grounds alone, before you even test whether the equation is separable. **Link to the next block.** If you cannot produce f, you must produce a function of f whose behaviour you can control, so you go looking at the structure of the right hand side for a candidate.

Block 2. **Cue.** The denominator $x^2\left(f^2 + 1\right)$ splits into a factor depending only on x and a factor depending only on f. **Why natural.** When $f' = \frac{A(x)}{B(f)}$ the equation is separable and the correct unknown is $\int B$. Here $f' = \frac{1}{1+f^2}\left(1 - \frac{f^2}{x^2}\right)$ is that shape multiplied by a correction factor, so the

same substitution is still the right one; you simply have to carry the correction along. Since $\int(1+t^2)\,dt = t + \frac{1}{3}t^3$, the quantity to name is $g = f + \frac{1}{3}f^3$. **Link to the next block.** Multiplying by $1+f^2$ makes the left side $\frac{d}{dx}g$ and clears every f from the denominator, which is Block 3.

Block 3. Cue. $g' = 1 - f^2/x^2$, with no denominators anywhere. **Why natural.** The correction factor $1 - f^2/x^2$ is bounded above by 1 for free, since $f^2 \geq 0$. An inequality on a derivative is an inequality on the function via monotonicity, and you get it without integrating. **Link to the next block.** An upper bound on g is an upper bound on f, because $\phi(t) = t + \frac{1}{3}t^3$ is increasing, but converting it needs f bounded below, which is Block 5.

Block 4. Cue. $g(x) - x$ has nonpositive derivative. **Why natural.** This is the cleanest form of an integrated inequality: instead of writing $g(x) = g(2) + \int_2^x g'$ and estimating, you name the combination whose derivative has a fixed sign and quote monotonicity. Same content, less ink, and no integrability discussion. **Link to the next block.** You now have $\frac{1}{3}f^3 \leq x + K - f$, and to turn that into an upper bound on f you must control the $-f$ term, which is exactly a lower bound on f.

Block 5. Cue. You need a lower bound, and the equation gives $f' \geq -\frac{1}{x^2}$ with $\frac{1}{x^2}$ integrable at infinity. **Why natural.** The quantity $\frac{1}{x^2}$ is the derivative of $-\frac{1}{x}$, so $f - \frac{1}{x}$ has positive derivative. The identity $f' + \frac{1}{x^2} = \frac{x^2+1}{x^2(1+f^2)}$ is an algebraic identity, not an estimate, so there is nothing to check: put the two terms over a common denominator and the numerator is visibly positive. **Link to the next block.** With $f > m$ and $\frac{1}{3}f^3 \leq x + K - m$ you have $f = O(x^{1/3})$, and a sublinear bound is exactly what the correction term in g' needs.

Block 6. Cue. $g' = 1 - f^2/x^2$ and $f = O(x^{1/3})$, so the correction is $O(x^{-4/3})$. **Why natural.** This is bootstrapping: a crude bound on f, substituted back into the equation, sharpens the information about g', which in turn sharpens g and hence f. It is the standard loop in every asymptotic argument, and here one turn of the loop suffices. **Link to the next block.** $g' \to 1$ means g eventually increases at rate at least $\frac{1}{2}$, and the mean value theorem turns that into $g \to \infty$.

Block 7. Cue. You have $\phi(f) \to \infty$ with ϕ strictly increasing, and you want $f \to \infty$. **Why natural.** Monotone functions transfer limits to infinity in both directions: if f failed to exceed M at some large x, then $\phi(f)$ would fail to exceed $\phi(M)$ there. Write it as that one line contradiction rather than invoking an inverse function theorem. **Link to the next block.** Done. Note that you never needed f itself to be monotone, which is exactly what made this route shorter than the trapping routes.

Layer 3: WHY

The truth beneath the whole problem is that the equation is a perturbation of one you can solve exactly. Discard the term f^2/x^2 and you are left with $f' = \frac{1}{1+f^2}$, whose solutions satisfy $f + \frac{1}{3}f^3 = x + c$ exactly. That equation clearly forces $f \to \infty$ like $(3x)^{1/3}$. So the mathematical content of B5 is not "prove a limit", it is "prove the perturbation is harmless", and once you have said that sentence the route is forced: name the exactly solvable core, bound the perturbation, show it vanishes, transfer the conclusion.

The discernment mechanism is a two part reflex. Part one: when f' has a function of f in the denominator, the correct new unknown is the antiderivative of that function. Dividing to isolate f' is what a calculator does; multiplying to create an exact derivative is what a mathematician does, because $B(f)f' = \frac{d}{dx}\left[\int B\right](f)$ is the only way the chain rule can help you here. Part two: when the goal is a limit rather than a formula, you are always looking for a quantity with a signed derivative, because the mean value theorem is the only mechanism that converts local information (the sign

of f') into global information (behaviour at infinity). Everything else in analysis at this level is a repackaging of that.

How do you train it? Take every differential equation you meet and, before doing anything, write the right hand side in the form "a solvable core times a correction" and identify the correction's size. Then ask which combination of x and f has an unambiguously signed derivative. On this problem there are three good answers, $f - \frac{1}{x}$, $x - f$ and $g - x$, and the solution uses two of them. A route that uses none of them, such as trying to solve the equation, is a route with no mechanism for producing a conclusion at infinity.

One more structural point that is worth repeating, because it recurs across the whole Putnam. The problem says f is differentiable, not continuously differentiable. You do not need extra regularity here, since the argument is built on the mean value theorem alone. But note that the equation upgrades itself: f differentiable makes f continuous, which makes the right hand side continuous, which makes f' continuous, so f is automatically C^1; repeating the argument gives $f \in C^\infty$. That bootstrap is free and is a standard remark to have ready, though a full marks solution does not need it.

The examiners nudged in three ways. They put $f^2 + 1$ in the denominator, which is a pointer at $f + \frac{1}{3}f^3$ and is the only reason the cube appears in the asymptotic. They wrote the numerator as $x^2 - f^2$ rather than expanding, so the nullclines $f = \pm x$ are visible. And they set the domain to $(1, \infty)$ so that $1/x^2$ is small and integrable at infinity, which is what makes the lower bound in Step 1 a constant rather than a drifting function.

Revolving concepts

Monotonicity from the sign of the derivative. If F is differentiable on an interval and $F' \geq 0$ there, then F is non decreasing on that interval, and strict inequality gives strict monotonicity. The proof is the mean value theorem. This is the only tool used in the ideal solution, and it is the reason the solution avoids every integration. See Book 1, field: Real analysis, topic: mean value theorem.

The barrier trick. To convert a differential inequality $F'(x) \geq -h(x)$ into a bound, name H with $H' = h$ and observe that $F + H$ is non decreasing. Here $h(x) = 1/x^2$, $H(x) = -1/x$, and the conclusion is $f(x) \geq f(2) - \frac{1}{2} + \frac{1}{x}$. This replaces an integration of an inequality by a monotonicity statement, which is cheaper to justify. See Book 1, field: Real analysis, topic: differential inequalities.

The limit lemma. If F is differentiable on (a, ∞) and $F'(x) \to L$ then $F(x)/x \to L$; if in addition $L > 0$ then $F(x) \to +\infty$. Proof: apply the mean value theorem to $F(x) - Lx$ on $[X, x]$, where $|F' - L| < \varepsilon$ beyond X. This is the Cesaro and L'Hopital fact in the form you actually use. See Book 1, field: Real analysis, topic: asymptotics of derivatives.

Exact derivative substitutions. If $f' = \frac{A(x)}{B(f)}$, put $\Phi(t) = \int_0^t B$ and the equation becomes $\frac{d}{dx}\Phi(f) = A(x)$, hence $\Phi(f(x)) = \int A$. This is what separation of variables really is. When the equation is only approximately of that form, the substitution still helps: $\frac{d}{dx}\Phi(f) = A(x)\,(1 + \text{small})$. See Book 1, field: Calculus and integrals, topic: separable equations.

Bootstrapping, stated as a recursion. You have $y' = F(x, y)$ and a crude bound $|y| \leq C_0 x^{\alpha_0}$. Substituting it into the equation produces a bound on y', hence by integration a bound $|y| \leq C_1 x^{\alpha_1}$ with $\alpha_1 = T(\alpha_0)$ for some map T determined by F. Iterating gives $\alpha_{k+1} = T(\alpha_k)$, and the fixed point of T is the true growth exponent. Here $\alpha_0 = 1$ (from $|f| \leq x + C$), the map produces $\alpha_1 = 1/3$, and $1/3$ is already the fixed point, so one turn suffices; continuing the same loop on the constant rather than the exponent produces the full asymptotic expansion $f = (3x)^{1/3}\left(1 + c_1 x^{-2/3} + \cdots\right)$. The identical loop is how you extract asymptotics from every implicitly defined sequence or function. See Book 1, field: Real analysis, topic: bootstrapping and asymptotic expansions.

Nullclines and the phase plane. For $y' = F(x, y)$ the nullcline is the curve $F = 0$; the sign of F on either side of it dictates whether solutions rise or fall, and solutions can cross a nullcline only in the direction the sign change permits. Here the nullclines are $f = \pm x$, and the observation

that $(x-f)' \geq 0$ says that once $f < x$ the gap never closes. See Book 1, field: Real analysis, topic: qualitative theory of first order equations.

Existence and uniqueness. The right hand side $F(x,y) = \frac{x^2-y^2}{x^2(y^2+1)}$ is continuously differentiable in (x,y) for $x > 1$, so by the Picard and Lindelof theorem solutions exist locally and are unique for each initial value; the problem hands you one and asks for its behaviour. Regularity bootstraps: f differentiable implies f' continuous implies $f \in C^1$ implies $f \in C^2$, and so on, so $f \in C^\infty$. See Book 1, field: Real analysis, topic: existence and uniqueness for ODEs.

Circuit Bank

- **WHEN** f' has a polynomial in f in the denominator **DO** multiply through by that polynomial and recognise the left side as $\frac{d}{dx}\Phi(f)$ where Φ' is that polynomial **BECAUSE** $\Phi(f)$ is the natural unknown; the equation then has no denominator in f and its right side depends on x and on a bounded correction only
- **WHEN** a differential equation is given with no initial condition and the demand is a limit **DO** stop trying to solve it and hunt for a combination of x and f whose derivative has a fixed sign **BECAUSE** a conclusion that must hold for every solution cannot depend on a constant of integration, and monotone plus unbounded is the cheapest route to a limit
- **WHEN** you need a bound on f from a differential inequality $f' \geq -h(x)$ with h having an explicit antiderivative H **DO** state that $f + H$ is non decreasing **BECAUSE** monotonicity of an explicit combination replaces integrating an inequality and needs only the mean value theorem, so there is nothing for a grader to object to
- **WHEN** you have shown $F'(x) \to L > 0$ **DO** conclude $F \to +\infty$ and $F(x)/x \to L$ **BECAUSE** beyond some X you have $F' \geq L/2$, and the mean value theorem gives $F(x) \geq F(X) + \frac{L}{2}(x - X)$
- **WHEN** you have proved $\Phi(f) \to \infty$ for a strictly increasing Φ **DO** conclude $f \to \infty$ in one line **BECAUSE** if $f(x) \leq M$ at some large x then $\Phi(f(x)) \leq \Phi(M)$ there, contradicting the divergence of $\Phi(f)$

Backtesting

1. First useful cue: the factor $f^2 + 1$ in the denominator, which is the derivative of $f + \frac{1}{3}f^3$; the secondary cue is the absence of an initial condition, which tells you the answer must be solution independent.
2. Method the expression suggested: multiply by $1 + f^2$ to obtain the exact derivative $g' = 1 - f^2/x^2$, then run monotonicity arguments on $g - x$ and on $f - \frac{1}{x}$.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at the upper bound, where $\frac{1}{3}f^3 = g - f \leq x + K - f$ is useless until f is bounded below. Reinforce this: whenever a chain of bounds stalls, look for the missing bound in the other direction, and expect the equation to supply it through a different rearrangement of the same identity.
4. The trap the question sets: the equation looks separable at a glance, and the estimate $|f'| \leq 1$ looks useful. Neither is. A bounded derivative says nothing about the limit, and the coupling f^2/x^2 blocks separation. The second trap is assuming f is eventually monotone and using it without proof.
5. What makes you recognise it faster next time: the shape $\frac{\text{something in } x}{\text{polynomial in } f}$ on the right hand side. That shape means "integrate the denominator". Also the shape $x^2 - f^2$, which announces the nullclines $f = \pm x$.

6. The reusable decision rule: to prove divergence of an implicitly defined function, build a quantity whose derivative you can sign, prove that quantity diverges, and transfer back through a strictly monotone map.

Mirror drills

1. Let f satisfy $f'(x) = \frac{1}{1+f(x)^2}$ on $(0,\infty)$ with $f(0) = 0$. Find the asymptotic behaviour of f. *Answer:* multiply by $1 + f^2$ to get $\left(f + \frac{1}{3}f^3\right)' = 1$, so $f + \frac{1}{3}f^3 = x$ exactly and $f(x) \sim (3x)^{1/3}$.
2. Let F be differentiable on $(1,\infty)$ with $F'(x) \to L$. Prove $F(x)/x \to L$, and that $F \to +\infty$ when $L > 0$. *Answer:* given $\varepsilon > 0$ pick X with $|F' - L| < \varepsilon$ on $[X,\infty)$ and apply the mean value theorem to $F(x) - Lx$ on $[X, x]$; divide by x.
3. For the f of the problem, prove $f(x)/(3x)^{1/3} \to 1$. *Answer:* Steps 1 to 5 give $f \to \infty$ and $g' \to 1$; the previous drill gives $g(x)/x \to 1$, and $g \sim \frac{1}{3}f^3$ because $f \to \infty$, so $\frac{1}{3}f^3 \sim x$.

17.12 B6. Reaching every integer in a fixed number of steps

Prove that for every positive integer n, there is a sequence of integers $a_0, a_1, \ldots, a_{2009}$ with $a_0 = 0$ and $a_{2009} = n$ such that each term after a_0 is either an earlier term plus 2^k for some nonnegative integer k, or of the form $b \bmod c$ for some earlier positive terms b and c. [Here $b \bmod c$ denotes the remainder when b is divided by c, so $0 \leq (b \bmod c) < c$.]

Read it out loud

Read the quantifiers first, because they carry the whole difficulty. *For every positive integer n*, there is a sequence of *fixed length* 2010. So n is unbounded and the budget is not. That mismatch is the problem. Everything else, the two allowed moves and the bracketed definition of $b \bmod c$, is machinery.

Now read the machinery. You are building a straight line program: a list of values, each produced from earlier values by one of two instructions. Instruction one, add a power of two to an earlier term. Instruction two, reduce one earlier positive term modulo another. Ask the only question that matters: how much can one step change the description length of a number? Instruction one adds at most one binary digit. If you only ever add powers of two, then reaching n requires at least as many steps as there are ones in the binary expansion of n, and $n = 2^{3000} - 1$ has three thousand of them. So addition of powers of two is not the engine; it is the cheap way to build numbers with a small number of binary digits. The mod is the only instruction that can output something whose binary expansion is long and complicated when both inputs are simple. Therefore the mod must do essentially all of the work, and there will be at most a couple of them.

The fields wired in are elementary number theory (base representations, the uniqueness of the remainder, congruences) and polynomial algebra (the factor theorem and Taylor expansion in $\mathbb{Z}[x]$), with a light coating of computational complexity, since "sequence of fixed length from a fixed instruction set" is exactly the definition of a straight line program and the underlying model is that of addition chains.

Within thirty seconds you should have written the following. Which integers are free? Exactly those with few ones in binary. So fix $x = 2^h$ and think in base x: any integer whose base x digits are powers of two and few in number can be assembled by a handful of additions. Then ask the operative question: what pair of base x sparse integers has a remainder that reveals n?

Approaches, ranked worst to best

1. **Binary expansion.** *Cue:* "an earlier term plus 2^k" looks like an instruction to build n digit by digit. *Circuit:* greedily add the powers of two occurring in n. It produces n in exactly $s(n)$ steps where $s(n)$ is the binary digit sum, and $s(n)$ is unbounded. Dead for any n with more than 2009 ones. Its only value is diagnostic: it tells you precisely what the mod is for.

2. **Euclidean descent.** *Cue:* "$b \bmod c$" is the Euclidean algorithm step. *Circuit:* shrink large numbers towards n by repeated reduction. You can certainly steer numbers downwards, but the number of reductions needed grows with the numbers involved and there is no uniform bound. Same failure mode as the previous route, dressed differently, and it wastes more time because it feels productive.

3. **Primitive root construction.** *Cue:* 2 is a primitive root modulo 3^h, so every residue coprime to 3 is a power of two. *Circuit:* realise n as $2^c \bmod 3^{2h}$, and manufacture the modulus 3^{2h} as $2^{2hb} \bmod \left(3 + 2^b\right)$. This works and is genuinely short, but it imports the theorem that 2 is a primitive root modulo every power of 3, and it needs a separate patch when $3 \mid n$, since then n is not a power of two modulo 3^{2h}. Two knowledge costs and a case split.

4. **Euler inverse variant.** *Cue:* a remainder can be forced to be a quotient of sparse numbers if you can invert modulo the modulus. *Circuit:* use $x^{\varphi(m)-1} \equiv x^{-1} \pmod{m}$ to obtain $\frac{x^n+1}{x+1}$ as a single remainder, then reduce modulo $x+1$. Six steps rather than seven, and elegant, but it costs Euler's theorem and the observation that $\gcd(x, x^n + x + 1) = 1$. Worth knowing as a remark; not the route to write in the arena.

5. **Base x Taylor construction.** *Cue:* n must appear as a coefficient, and coefficients are exactly what a remainder modulo the square of a linear factor reveals. *Circuit:* $x^n + 1 \equiv n(x+1) \pmod{(x+1)^2}$ for odd n, then one further mod strips the base x digit. Best. Seven steps, entirely elementary, uniform in n, and the reduction to odd n is free because the instruction set is invariant under doubling. Nothing beyond the binomial theorem and the uniqueness of remainders is used.

The ideal solution

Call a finite sequence of integers $a_0, a_1, \ldots, a_N$ *legal* if $a_0 = 0$ and each later term is either $a_j + 2^k$ for some earlier index j and some integer $k \geq 0$, or $a_j \bmod a_l$ for earlier indices j, l with $a_j, a_l > 0$.

Lemma 17.12.1. *If $a_0, \ldots, a_N$ is legal then so is $2a_0, \ldots, 2a_N$.*

Proof. Certainly $2a_0 = 0$. If $a_i = a_j + 2^k$ then $2a_i = 2a_j + 2^{k+1}$, again an earlier term plus a power of two. If $a_i = a_j \bmod a_l$ with $a_j, a_l > 0$, write $a_j = qa_l + a_i$ with $0 \leq a_i < a_l$; then $2a_j = q(2a_l) + 2a_i$ with $0 \leq 2a_i < 2a_l$, so by uniqueness of the remainder $2a_i = (2a_j) \bmod (2a_l)$, and $2a_j, 2a_l$ are earlier positive terms. □

Lemma 17.12.2. *Let n be an odd positive integer and let x be an integer with $x > n$. Then*

$$(x^n + 1) \bmod (x+1)^2 = n(x+1).$$

Proof. Work in $\mathbb{Z}[x]$ and substitute $u = x + 1$, so $x = u - 1$. By the binomial theorem,

$$x^n + 1 = (u-1)^n + 1 = \sum_{k=0}^{n} \binom{n}{k} u^k (-1)^{n-k} + 1.$$

Because n is odd, the $k = 0$ term is $(-1)^n = -1$, which cancels the added 1; the $k = 1$ term is $nu(-1)^{n-1} = nu$; and every term with $k \geq 2$ is divisible by u^2. Hence there is an integer R with

$$x^n + 1 = n(x+1) + R(x+1)^2.$$

Since $1 \leq n < x + 1$ we have $0 \leq n(x+1) < (x+1)^2$, so by uniqueness of the remainder $n(x+1)$ is exactly the remainder of $x^n + 1$ on division by $(x+1)^2$. □

Lemma 17.12.3. *If $x > n \geq 1$ are integers then $n(x+1) \bmod x = n$.*

Proof. $n(x+1) = n \cdot x + n$ and $0 \leq n < x$. □

Proof of the statement. By Lemma 1 it suffices to construct a legal sequence of length 2010 ending in n for every *odd* n: for general n write $n = 2^s n_0$ with n_0 odd, take the sequence ending in n_0, and apply Lemma 1 exactly s times, which multiplies every term by 2^s and so ends in n while keeping $a_0 = 0$.

So let n be odd and choose $h \geq 1$ with $x := 2^h > n$. Define

$$\begin{aligned} a_0 &= 0, \\ a_1 &= 1 = a_0 + 2^0, \\ a_2 &= 2x + 1 = a_1 + 2^{h+1}, \end{aligned}$$

$$a_3 = (x+1)^2 = x^2 + 2x + 1 = a_2 + 2^{2h},$$
$$a_4 = x^n + 1 = a_1 + 2^{hn},$$
$$a_5 = n(x+1) = a_4 \bmod a_3,$$
$$a_6 = x = a_0 + 2^h,$$
$$a_7 = n = a_5 \bmod a_6,$$

and $a_i = a_{i-1} \bmod a_4$ for $8 \le i \le 2009$.

Every step is legal. The terms a_1, a_2, a_3, a_4, a_6 are each an earlier term plus a single power of two, as displayed. The term a_5 is Lemma 2 applied with this $x > n$, and its inputs $a_4 = x^n + 1$ and $a_3 = (x+1)^2$ are earlier and positive. The term a_7 is Lemma 3, with inputs $a_5 = n(x+1) > 0$ and $a_6 = x > 0$, both earlier. Finally each padding term is legal because $0 \le n < x^n + 1 = a_4$, so $n \bmod a_4 = n$, and the inputs $a_{i-1} = n > 0$ and $a_4 > 0$ are earlier and positive.

The sequence has $a_0 = 0$ and $a_{2009} = n$, as required. □

Remark 17.12.4. Seven steps suffice; the number 2009 in the problem is theatre, and the padding is the only place where it does any work at all. A variant of Lenstra, van Luijk and Della Torre reaches n in six: with n odd and $x = 2^h > n$, take $a_2 = x + 1$, $a_3 = x^n + x + 1$, $a_4 = x^{(n-1)(\varphi(a_3)-1)}$, $a_5 = a_4 \bmod a_3 = \frac{x^n+1}{x+1}$ and $a_6 = a_5 \bmod a_2 = n$. It is shorter because Euler's theorem lets one remainder perform a division, but it costs the theorem.

Layer 1: WHAT

Block 1. Compare budget with target: 2009 steps, n unbounded. Conclude that some instruction must be able to increase description length without bound, and identify it as the mod.

Block 2. Ask which numbers are free. Answer: those with few ones in binary. Fix $x = 2^h$ and work in base x, where a number is free exactly when it has few base x digits, each a power of two.

Block 3. Ask where an arbitrary n can appear as an output of a mod applied to sparse inputs. Since n enters algebra as an exponent, and remainders read off coefficients, the conversion of exponent into coefficient is differentiation: $\frac{d}{dx}x^n = nx^{n-1}$.

Block 4. Recall that the remainder of F modulo $(x-a)^2$ is $F(a) + F'(a)(x-a)$, the first order Taylor polynomial. So choose an F that is sparse, has $F(-1) = 0$ and $F'(-1) = \pm n$. That is $F(x) = x^n + 1$ with n odd, and the modulus is $(x+1)^2$, also sparse.

Block 5. Handle even n by the symmetry of the instruction set under doubling, which is exactly what removes the obstruction, since $x + 1$ divides $x^n + 1$ only for odd n.

Block 6. The remainder is $n(x+1) = nx + n$. Its base x digits are n and n, so one more reduction, modulo x, delivers n. This is where the condition $x > n$ earns its keep.

Block 7. Pad to index 2009 by repeatedly writing $n \bmod (x^n + 1)$.

The nominal route is Block 1 followed by writing n in binary and adding powers of two, or by an aimless Euclidean descent. The bifurcation point is inside Block 1: the moment you notice that the step count is independent of n and therefore refuse to spend even one step per binary digit. Once you have refused, the search space collapses to "which identity produces n from a mod of two sparse numbers", and there are only a few such identities in existence.

Layer 2: HOW

Block 1. Cue. "For every positive integer n" against "$a_0, \ldots, a_{2009}$". **Why natural.** A fixed length program with a fixed instruction set can only reach numbers whose complexity in that model is bounded. Addition of a power of two changes the binary description by one digit, so a program of pure additions has bounded digit sum. Since digit sums are unbounded, the mod is forced to be the load bearing instruction. Counting the resources before starting is the standard opening on any

construction problem with a budget. **Link to the next block.** If the mod carries the information, you must decide what its two inputs look like, which means deciding which numbers are cheap.

Block 2. **Cue.** "an earlier term plus 2^k". **Why natural.** The cheap numbers are exactly the sums of a few powers of two. Writing $x = 2^h$ turns "a few powers of two" into "a few base x digits", and a polynomial expression in x with small nonnegative coefficients, each a power of two, is cheap. This reframing is what lets you think about polynomials instead of integers, which is the only way the algebra becomes visible. **Link to the next block.** You are now looking for a polynomial identity: sparse F, sparse G, with F mod G revealing n.

Block 3. **Cue.** n appears in the problem only as a target value, but in any expression you can cheaply build it will appear as an exponent. **Why natural.** Exponents and coefficients are different slots. There are exactly two elementary devices that move a number from the exponent slot to the coefficient slot: differentiating x^n to get nx^{n-1}, and summing a geometric series with n terms so that the value at the special point is n. Both are the same fact seen from two sides, and both put n next to a linear factor. **Link to the next block.** Differentiation suggests a second order phenomenon at a root, and second order phenomena at a root are visible modulo the square of the corresponding linear factor.

Block 4. **Cue.** "$b \bmod c$" with c allowed to be any earlier term, so you may choose the modulus. **Why natural.** Division with remainder in $\mathbb{Z}[x]$ by $(x-a)^k$ returns the truncated Taylor expansion of the dividend at a: $F(x) \equiv F(a) + F'(a)(x-a) \pmod{(x-a)^2}$. Choosing $a = -1$ and $F(x) = x^n + 1$ makes $F(a) = 0$ for odd n and $F'(a) = n$, so the remainder is precisely $n(x+1)$. Both F and the modulus $(x+1)^2 = x^2 + 2x + 1$ are sparse, so both are cheap. Nothing here was guesswork: you asked for a sparse F with a double root structure at -1 and there is essentially one. **Link to the next block.** The polynomial identity has to be transported to integers, which requires the digits to be small enough not to interfere, and it also requires n odd.

Block 5. **Cue.** The identity fails for even n, because then $(-1)^n + 1 = 2 \neq 0$. **Why natural.** Before patching a case, check whether the problem has a symmetry that normalises it away. The instruction set here is homogeneous under doubling: adding 2^k becomes adding 2^{k+1}, and remainders scale, because $b = qc + r$ with $0 \leq r < c$ gives $2b = q(2c) + 2r$ with $0 \leq 2r < 2c$. So the set of reachable values is closed under multiplication by two, and it suffices to reach every odd n. **Link to the next block.** With n odd, run the identity, then convert the resulting integer $n(x+1)$ into n.

Block 6. **Cue.** The remainder $n(x+1) = nx + n$ written in base x. **Why natural.** Reducing modulo x deletes every digit except the last, and the last digit is n provided $n < x$. That proviso is exactly the condition under which the polynomial identity survives specialisation at $x = 2^h$: a polynomial identity with coefficients smaller than the base is a statement about digits, and digits do not carry. This is the same principle as Kronecker substitution, where polynomial multiplication is performed by multiplying large integers in a base big enough to stop carrying. **Link to the next block.** All that remains is bookkeeping: name the powers of two used, and pad.

Block 7. **Cue.** The index must reach exactly 2009. **Why natural.** Padding is not free, because there is no "do nothing" instruction. But $n \bmod c = n$ for any $c > n$, and you already own such a c, namely $x^n + 1$. The padding must be justified explicitly; a write up that says "pad the sequence" without naming the instruction is leaving a hole a grader is entitled to charge for. **Link to the next block.** The construction is complete and every step has been certified.

Layer 3: WHY

The underlying truth is that remainders compute Taylor coefficients, and that this survives specialisation to integers whenever the coefficients are smaller than the base. Both halves deserve to be stated plainly, because together they are the engine of a whole family of Putnam problems.

First half. In $\mathbb{Z}[x]$, dividing F by $(x-a)^k$ leaves the truncated Taylor polynomial $\sum_{j<k} \frac{F^{(j)}(a)}{j!}(x-a)^j$. For $k=1$ this is the factor theorem. For $k=2$ it is the statement that the remainder sees both the value and the derivative at a. So if you want a specific integer n to appear from an expression in which n is an exponent, reduce modulo the square of a linear factor at a root: the derivative brings the exponent down into a coefficient. That is why x^n+1 against $(x+1)^2$ is not a lucky pairing but the unique cheap pairing.

Second half. An integer written in base b is a polynomial evaluated at b, and the correspondence is faithful, meaning digits do not interact, exactly while all the coefficients lie in $[0,b)$. In the construction the relevant coefficients are 1, 2 and n, so the single condition $n < x$ makes every polynomial statement above a statement about the base x digits of an integer, and the integer remainder equals the specialised polynomial remainder. The same principle is behind Kronecker substitution, behind reading a generating function coefficient by evaluating at a large integer, and behind the standard proof that a polynomial with nonnegative integer coefficients can be reconstructed from two evaluations.

The discernment mechanism to train is a resource count followed by an identity search. Resource count: how much can each instruction change the complexity of the object, and how many instructions do you have? If the budget is finite and the target is unbounded, exactly one instruction is doing the work, and you should identify it in the first minute. Identity search: with the load bearing instruction identified, ask what identities produce your target as its output from cheap inputs, and go through your list. For remainders the list is short and you should have it memorised: the Taylor and factor theorem family ($x^n \pm 1$ against $(x\pm1)^2$), the geometric series family ($\frac{x^n-1}{x-1} \equiv n \pmod{x-1}$), Fermat and Euler, and orders and primitive roots. Every published solution to this problem uses one of those four.

Knowledge gaps, stated without cosmetics. If you do not already know that the remainder modulo $(x-a)^2$ is the first order Taylor polynomial, this problem is not approachable, because you will never generate the pairing x^n+1 with $(x+1)^2$ from scratch under time pressure. If you do not know Euler's theorem and the trick $x^{\varphi(m)-1} \equiv x^{-1} \pmod{m}$, the six step variant is invisible to you. If you do not know that 2 is a primitive root modulo every power of 3, and why lifting from modulus 9 suffices, the second published route is invisible to you. And if you have never met the notion of a straight line program or an addition chain, you will not perform the resource count in Block 1 with any confidence, and without that count you will spend the whole session writing n in binary. Install those four items before you claim this problem is solved.

Revolving concepts

Base b representation and digit extraction. Every nonnegative integer has a unique representation $N = \sum_j d_j b^j$ with $0 \le d_j < b$, and $N \bmod b^k$ keeps exactly the lowest k digits. An integer is cheap in this problem exactly when it has few nonzero binary digits, and choosing $x = 2^h$ makes "few base x digits, each a power of two" the operative notion of cheapness. See Book 1, field: Number theory, topic: base representations.

Division with remainder and its uniqueness. For integers b and $c > 0$ there are unique q and r with $b = qc + r$ and $0 \le r < c$. Uniqueness is what lets you certify a step: it is not enough that $x^n + 1 \equiv n(x+1)$ modulo $(x+1)^2$, you must also check $0 \le n(x+1) < (x+1)^2$, which is where $n < x$ is used. Half the available marks on a construction problem live in checks of this kind. See Book 1, field: Number theory, topic: the division algorithm.

Taylor expansion in $\mathbb{Z}[x]$ and Hensel lifting. For $F \in \mathbb{Z}[x]$ and $a \in \mathbb{Z}$, $F(x) = \sum_{j \ge 0} c_j (x-a)^j$

with $c_j = \frac{F^{(j)}(a)}{j!} \in \mathbb{Z}$, and $F \bmod (x-a)^k = \sum_{j<k} c_j (x-a)^j$. The recursion is $F_0 = F$, $c_j = F_j(a)$, $F_{j+1} = \frac{F_j - c_j}{x-a}$, which is synthetic division repeated. The same recursion over $\mathbb{Z}_p$ is Hensel lifting, where a simple root modulo p is refined to a root modulo p^k one digit at a time. Wherever you see "modulo a square", expect a derivative. See Book 1, field: Polynomials, topic: Taylor expansion and repeated division.

The two standard exponent to coefficient identities. For odd n, $x^n + 1 \equiv n(x+1) \pmod{(x+1)^2}$; for all n, $x^n - 1 \equiv n(x-1) \pmod{(x-1)^2}$. Both are the derivative statement in disguise, and both have geometric series proofs: $\frac{x^n-1}{x-1} = \sum_{j=0}^{n-1} x^j \equiv n \pmod{x-1}$ since each of the n terms is congruent to 1, and likewise $\frac{x^n+1}{x+1} = \sum_{j=0}^{n-1} (-1)^j x^{n-1-j} \equiv n \pmod{x+1}$ for odd n. Generalise the recursion: $\sum_{j=0}^{n-1} x^j$ satisfies $S_n = xS_{n-1} + 1$, and reducing that recursion modulo $x-1$ gives $S_n \equiv S_{n-1} + 1$, so $S_n \equiv n$, which is the cleanest proof and the one to carry. See Book 1, field: Number theory, topic: cyclotomic identities.

Euler's theorem and modular inversion by exponentiation. If $\gcd(x, m) = 1$ then $x^{\varphi(m)} \equiv 1 \pmod{m}$, so $x^{-1} \equiv x^{\varphi(m)-1}$. This is what allows a single remainder to perform a division and produces the six step variant. See Book 1, field: Number theory, topic: Euler's theorem.

Primitive roots modulo prime powers. If g is a primitive root modulo p^2 for odd prime p then it is a primitive root modulo p^h for every $h \geq 1$; 2 is a primitive root modulo 9, hence modulo every 3^h. Consequently every residue coprime to 3 is a power of 2 modulo 3^{2h}, which is the engine of the second published route. See Book 1, field: Number theory, topic: primitive roots.

Straight line programs and addition chains. An addition chain for n is a sequence starting at 1 in which every term is a sum of two earlier terms, ending at n; its length is between $\log_2 n$ and about $2\log_2 n$. The instruction "earlier term plus 2^k" is an addition chain with a restricted summand set, so its length is at least the binary digit sum. Adding the mod instruction changes the model completely, since a remainder can produce a number of high complexity from two of low complexity. The moral generalises: when a construction budget is fixed and the target is unbounded, find the instruction that is not complexity bounded. See Book 1, field: Discrete structures and graphs, topic: addition chains and complexity of representations.

Circuit Bank

- **WHEN** a construction must have length independent of an unbounded parameter **DO** identify which single instruction can increase the description length of the output without bound, and plan for the others to be mere preparation **BECAUSE** a bounded number of bounded increments is bounded, so every other instruction can only assemble simple objects
- **WHEN** you need an arbitrary integer n to appear as a value when n can only enter as an exponent **DO** look for n as a derivative, that is, reduce modulo the square of a linear factor, or as the number of terms of a geometric series **BECAUSE** differentiation and summation are the only cheap conversions of an exponent into a coefficient, and both place n beside a linear factor
- **WHEN** a problem asks for the Taylor data of a polynomial at a point **DO** divide by $(x-a)^k$ and read the remainder **BECAUSE** the remainder is exactly the degree $k-1$ Taylor polynomial at a, so remainders are a coefficient extraction device, not just a size reduction device
- **WHEN** you specialise a polynomial identity at an integer base b and want to read a coefficient with a mod **DO** check that every coefficient involved is smaller than b **BECAUSE** only then are the coefficients the base b digits with no carrying, so the integer remainder equals the specialised polynomial remainder
- **WHEN** an instruction set is invariant under a symmetry such as scaling **DO** use the symmetry to normalise the target, for instance to the odd case **BECAUSE** the reduction is free, and it

frequently removes precisely the obstruction, here that $x+1$ divides x^n+1 only when n is odd

- **WHEN** a construction problem specifies an exact length **DO** state explicitly which legal instruction performs the padding **BECAUSE** there is no do nothing instruction, so unjustified padding is a genuine hole, and here $n \bmod c = n$ for any earlier $c > n$ closes it

Backtesting

1. First useful cue: the fixed length 2010 against the unbounded n, which immediately demotes the addition instruction to preparation and promotes the mod to the engine.
2. Method the expression suggested: work in base $x = 2^h$ so that cheap integers are sparse polynomials in x, then use the remainder modulo $(x+1)^2$ to convert the exponent n into the coefficient n, and one further remainder modulo x to read the digit.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at even n, where the identity collapses. Reinforce this: before patching a case, test whether the problem's own instruction set has a symmetry that normalises it, and here doubling is that symmetry.
4. The trap the question sets: the number 2009 makes the budget look generous, so you start spending steps. It is not generous, it is irrelevant, and the true budget is a small constant. The second trap is the padding, which looks free and is not.
5. What makes you recognise it faster next time: the shape "fixed number of operations, unbounded target", which is a complexity question in disguise; and the pairing of a sparse dividend with the square of a linear divisor, which always means a derivative.
6. The reusable decision rule: count the resources first, identify the one instruction with unbounded output complexity, then search your list of identities for one that produces the target from cheap inputs under that instruction.

Mirror drills

1. Prove that $x^n + 1 \equiv n(x+1) \pmod{(x+1)^2}$ for odd n, and state the analogue for $x^n - 1$. *Answer:* substitute $u = x+1$ and expand $(u-1)^n + 1$ by the binomial theorem, or Taylor expand at $x = -1$; the analogue is $x^n - 1 \equiv n(x-1) \pmod{(x-1)^2}$ for all n.
2. Show that if $a_0, \ldots, a_N$ is a legal sequence then so is the sequence of doubled terms, and deduce that it suffices to reach every odd n. *Answer:* $a_j + 2^k$ doubles to $a_j + 2^{k+1}$, and $b = qc + r$ with $0 \le r < c$ doubles to $2b = q(2c) + 2r$ with $0 \le 2r < 2c$.
3. With $x = 2^h > n$, verify by hand for $n = 5$, $h = 3$ that $a_4 = x^5 + 1 = 32769$, $a_3 = (x+1)^2 = 81$, $a_5 = a_4 \bmod a_3 = 45$ and $a_7 = 45 \bmod 8 = 5$. *Answer:* $32769 = 404 \cdot 81 + 45$ and $45 = 5 \cdot 9 = n(x+1)$, then $45 = 5 \cdot 8 + 5$; note where $n < x$ is used in each of the two reductions.

Chapter 18

Putnam 2008

18.1 A1. a cyclic zero condition forces a potential

Let $f : \mathbb{R}^2 \to \mathbb{R}$ be a function such that $f(x, y) + f(y, z) + f(z, x) = 0$ for all real numbers x, y, and z. Prove that there exists a function $g : \mathbb{R} \to \mathbb{R}$ such that $f(x, y) = g(x) - g(y)$ for all real numbers x and y.

Read it out loud

You are given one identity in three free variables, and asked to produce a two variable object (indeed, ultimately a one variable function g) that satisfies it. Say that out loud and the shape of the problem is already visible: you have a spare degree of freedom sitting inside the hypothesis that the conclusion does not need. Whenever a hypothesis is quantified over more variables than the conclusion mentions, your first reflex should be to spend that extra variable by fixing it to a convenient constant, because a universally quantified identity remains true under any substitution you like.

There is no continuity, no differentiability, not even monotonicity attached to f. That silence is not an oversight, it is the examiners telling you plainly that this is an algebra problem, not an analysis problem. f is an arbitrary function on $\mathbb{R}^2$ subject to one constraint, and the constraint is an identity, so the entire solution has to come from substitution and nothing else. If your hand drifts toward a derivative or an integral here, that is the bifurcation point where you have left the ideal route.

Structurally, the identity $f(x, y) + f(y, z) + f(z, x) = 0$ is a statement that summing f around any directed triangle $x \to y \to z \to x$ gives zero. That is the signature of a quantity that is secretly a difference of a potential, the discrete cousin of a closed one form being exact. The first thing worth writing down inside thirty seconds is the most symmetric substitution available, $x = y = z$, because it costs nothing and immediately tells you how f behaves on the diagonal, before you commit to any particular strategy for building g.

Approaches, ranked worst to best

1. **impose regularity and integrate.** *Cue:* the shape $f(x, y) = g(x) - g(y)$ looks like it wants a fundamental theorem of calculus argument. *Circuit:* treat f as if it were smooth, differentiate the identity with respect to one variable, and try to integrate the result. (why it ranks here: the problem never grants continuity or differentiability, so any step that assumes it is invalid on the general function the problem actually hands you; you would be proving a strictly weaker statement while believing you had proved the real one, which on the Putnam scale is worth close to nothing).

2. **scattered numerical substitution.** *Cue:* a functional equation invites you to plug in numbers. *Circuit:* try several unrelated triples (x, y, z), hope a pattern for f emerges. (why it ranks here: it can work by luck, but it wastes ink chasing substitutions that do not obviously combine into a general argument, and in the arena you cannot afford to discover the right three substitutions by trial when a structural read of the identity hands them to you directly).

3. **structural substitution: diagonal, then a fixed basepoint.** *Cue:* the identity is symmetric under cycling x, y, z and holds for all reals, so it is safe to fix any of them. *Circuit:* collapse two variables to extract a symmetry property of f, then freeze one variable to a basepoint to expose a candidate for g. (this is the ideal route: every substitution is chosen because it isolates exactly one unknown, so the whole proof is three lines of algebra with nothing left to verify by guesswork).

The ideal solution

Substitute $x = y = z = t$ for an arbitrary real t. The identity gives $f(t,t) + f(t,t) + f(t,t) = 0$, so $3f(t,t) = 0$ and therefore

$$f(t,t) = 0 \quad \text{for every real } t.$$

Next substitute $z = x$ in the original identity, keeping x, y free. This gives

$$f(x,y) + f(y,x) + f(x,x) = 0.$$

Using $f(x,x) = 0$ from the previous step, this collapses to $f(x,y) + f(y,x) = 0$, so

$$f(y,x) = -f(x,y) \quad \text{for all real } x, y.$$

This is genuine full antisymmetry of f, not merely a coincidence at the point 0; it holds for every pair of reals.

Now fix the basepoint $z = 0$ and let x, y be free. The identity gives

$$f(x,y) + f(y,0) + f(0,x) = 0.$$

By antisymmetry, $f(0,x) = -f(x,0)$, so substituting,

$$f(x,y) + f(y,0) - f(x,0) = 0, \qquad \text{hence} \qquad f(x,y) = f(x,0) - f(y,0).$$

Define $g : \mathbb{R} \to \mathbb{R}$ by $g(x) := f(x,0)$. The last display reads exactly

$$f(x,y) = g(x) - g(y) \quad \text{for all real } x, y,$$

which is precisely the claim. This proves existence.

It is worth recording uniqueness as well, since it costs one more line and sharpens your understanding of the construction. Suppose g and h both satisfy $f(x,y) = g(x) - g(y) = h(x) - h(y)$ for all x, y. Fixing $y = 0$ gives $g(x) - g(0) = h(x) - h(0)$ for every x, so $g(x) - h(x) = g(0) - h(0)$ is the same constant for every x. Hence any two valid choices of g differ by a constant, and the full family of functions that work is exactly $f(x,0) + c$ for $c \in \mathbb{R}$, the discrete analogue of two antiderivatives of the same function differing by a constant.

Layer 1: WHAT

Block 1: substitute $x = y = z$ to learn $f(t,t) = 0$ for every t. Block 2: substitute $z = x$ to upgrade this diagonal fact into full antisymmetry, $f(y,x) = -f(x,y)$. Block 3: substitute $z = 0$ to write $f(x,y)$ purely in terms of $f(\cdot,0)$, using antisymmetry to remove the awkward term $f(0,x)$. Block 4: name $g(x) := f(x,0)$ and read off the required identity; optionally record that g is unique up to an additive constant.

The nominal route reaches for calculus, treating $f(x,y) = g(x) - g(y)$ as an invitation to differentiate and integrate, or else scatters numerical substitutions hoping to stumble onto the pattern that Blocks 1 to 3 produce directly. The bifurcation point sits at the very first line you write: whether you assume any regularity of f at all, or whether you commit immediately to spending the three free variables of the hypothesis on structurally chosen values (the diagonal, then a basepoint) rather than on arbitrary ones.

Layer 2: HOW

Block 1. Cue. The identity is universally quantified over three independent reals, and $x = y = z$ is the single substitution that uses none of that independence, so it is the cheapest possible test and should always be tried first on a multivariable identity. **Why natural.** Setting all three variables equal collapses three unknown values of f into three copies of the same value $f(t,t)$, turning the identity into a one term equation for free. **Link to the next block.** Having learned f vanishes on the diagonal, the natural next move is to see what else collapsing only two of the three variables reveals, since the diagonal fact must be a special case of something more general.

Block 2. **Cue.** The identity is unchanged if you set $z = x$ rather than $z = y$ or any other coincidence, and doing so leaves you with a live pair x, y, which is exactly the shape of the object you are trying to describe. **Why natural.** Setting $z = x$ produces $f(x, y) + f(y, x) + f(x, x) = 0$, and Block 1 has already told you $f(x, x) = 0$, so the third term vanishes and you are left with a direct relation between $f(x, y)$ and $f(y, x)$. **Link to the next block.** Antisymmetry is precisely the tool you need to simplify the awkward term $f(0, x)$ that will appear once you fix a basepoint, since it lets you rewrite any $f(0, \cdot)$ term as $-f(\cdot, 0)$.

Block 3. **Cue.** You want a formula for $f(x, y)$ using only a one variable object, so you should freeze the third slot z to a fixed value and see what falls out; $z = 0$ is the most convenient choice only because 0 is a familiar constant, any fixed real would do equally well. **Why natural.** Fixing $z = 0$ turns the three variable identity into a two variable one, $f(x, y) + f(y, 0) + f(0, x) = 0$, and antisymmetry from Block 2 converts the term $f(0, x)$ into $-f(x, 0)$, leaving an equation entirely in terms of $f(\cdot, 0)$ and $f(x, y)$. **Link to the next block.** Once every term is expressed through the single variable expression $f(\cdot, 0)$, that expression is visibly a candidate for g, and naming it is the only step left.

Block 4. **Cue.** The algebra in Block 3 already reads $f(x, y) = f(x, 0) - f(y, 0)$, so the function $g(x) := f(x, 0)$ is not invented, it is read directly off the page. **Why natural.** Defining g this way and observing the previous line is exactly $f(x, y) = g(x) - g(y)$ closes the existence claim with no further work; checking uniqueness up to a constant is one further substitution ($y = 0$) applied to two hypothetical representations. **Link to the next block.** None; this is the terminal block. Any generalisation (longer cycles, other domains) is deferred to the revolving concepts.

Layer 3: WHY

The identity $f(x, y) + f(y, z) + f(z, x) = 0$ says that summing f around any directed triangle returns to zero, whatever the three vertices. That is the discrete analogue of a closed differential form: on a connected space, a closed form is exact, meaning it can be written as the difference of values of a single potential function at the two endpoints. Here the connecting structure is trivial, every pair of reals is joined through the single basepoint 0, so the potential is recovered in one substitution rather than by any limiting or integrating process.

The discernment mechanism to train is this: whenever a functional identity is stated over more free variables than the conclusion needs, and the identity is symmetric under permuting or cycling those variables, treat the surplus variables as currency to be spent, not as further unknowns to solve for. The spending has an order: collapse variables against each other first (as in Blocks 1 and 2) to extract symmetry properties of the object, because those properties are usually needed to simplify what a basepoint substitution produces; only then fix a variable to a constant to expose the final formula. Doing the basepoint substitution first, before knowing the antisymmetry, would have left you staring at $f(0, x)$ with no way to relate it to $f(x, 0)$, forcing you to invent a second substitution mid-proof rather than having it ready.

The examiners nudge you toward this reading by giving you three completely free variables for a two variable conclusion and by giving you no regularity hypothesis whatsoever; a problem setter who wanted you to use calculus would have handed you continuity, since without it no calculus argument is even well posed. The total absence of hypotheses beyond the identity itself is therefore itself a cue, telling you the entire solution must be extractable by pure substitution.

Revolving concepts

Coboundaries and potential functions. A function f on pairs from a set X satisfying $f(x, y) + f(y, z) + f(z, x) = 0$ for all $x, y, z \in X$ is called a coboundary: there is a basepoint independent way

to write $f(x,y) = g(x) - g(y)$, obtained by fixing any basepoint $b \in X$ and setting $g(x) = f(x,b)$. This is the discrete, purely algebraic shadow of the statement that a closed one form on a simply connected domain is exact. See Book 1, field: Functional equations, topic: cyclic identities and potential constructions.

Antisymmetric two variable functions. A function with $f(y,x) = -f(x,y)$ for all x, y is called antisymmetric or skew symmetric; it automatically vanishes on the diagonal, $f(x,x) = 0$, since $f(x,x) = -f(x,x)$ forces this. Determinants of two vectors in the plane and the signed area of a triangle are the familiar antisymmetric objects from geometry; here antisymmetry is derived rather than assumed. See Book 1, field: Functional equations, topic: symmetry extraction by variable collapse.

Cauchy's functional equation, as a contrast. The equation $f(x+y) = f(x) + f(y)$ for $f : \mathbb{R} \to \mathbb{R}$ looks similarly innocent but is genuinely under determined by algebra alone: without continuity, monotonicity, or boundedness on some interval, it admits pathological solutions built from a Hamel basis of $\mathbb{R}$ over $\mathbb{Q}$, using the axiom of choice. The present problem needs no such hypothesis precisely because it hands you a full triangle of free variables to substitute into, giving pure algebra enough leverage to pin f down completely; recognising which functional equations have this extra leverage and which do not is the key discernment skill in this family. See Book 1, field: Functional equations, topic: when regularity hypotheses are load bearing.

Cyclic sums of any length. The three term identity generalises: for every $n \geq 2$ and all $x_1, \ldots, x_n \in \mathbb{R}$,

$$f(x_1,x_2) + f(x_2,x_3) + \cdots + f(x_{n-1},x_n) + f(x_n,x_1) = 0.$$

The general recursion and its solution: let $S(n)$ denote this n term cyclic sum. The base case $S(2) = f(x_1,x_2) + f(x_2,x_1) = 0$ holds by antisymmetry, and $S(3) = 0$ is the given identity. For $n \geq 4$, apply the three term identity to the triple (x_{n-1}, x_n, x_1) to get $f(x_{n-1},x_n) + f(x_n,x_1) = -f(x_1,x_{n-1}) = f(x_{n-1},x_1)$, which replaces the last two edges of the n cycle by a single shortcut edge, showing $S(n) = S(n-1)$. Induction downward to $S(2) = 0$ finishes the general claim. This shortcut chord reduction is the standard way to promote a local (three term) identity to a global (any length cycle) one, and it reappears whenever a relation is only given locally but you need it to hold along an arbitrarily long path. See Book 1, field: Functional equations, topic: telescoping identities and cycle reduction by induction.

Circuit Bank

- **WHEN** a functional identity is stated over more free variables than the conclusion mentions, with a cyclic or fully symmetric structure **DO** immediately test the most degenerate substitution, collapsing two or more of the variables to each other, before trying to solve for the object directly **BECAUSE** collapsing variables is the cheapest way to strip a cyclic identity down to a single variable fact, and it usually hands you a symmetry property for free that you will need later

- **WHEN** you need to build a one variable potential function g from a two variable difference type constraint $f(x,y)$ **DO** fix one variable to a convenient basepoint and define $g(x) := f(x, \text{basepoint})$ **BECAUSE** freezing a spare degree of freedom converts an identity that has more unknowns than equations into an explicit formula, and the resulting candidate only needs one line to verify

- **WHEN** a problem hands you a completely arbitrary function with no continuity, differentiability, or monotonicity stated **DO** treat every calculus tool as forbidden and solve purely by algebraic substitution **BECAUSE** a valid proof must work for every function satisfying the constraint as a set theoretic object, and any argument that silently assumes regularity is proving a different, weaker problem

- **WHEN** you have proved a short instance of a cyclic identity (length two or three) and suspect it extends to any cycle length **DO** state and prove the general n cycle claim by induction, reducing an n cycle to an $n-1$ cycle with one application of the base identity **BECAUSE** Putnam write ups reward the fully general closed form vanishing lemma, and the shortcut chord reduction is one line once you see it, so it is nearly free insurance against a grader asking for the general case

Backtesting

1. First useful cue: the identity holds for every real x, y, z, three independent degrees of freedom, while the conclusion only needs a one variable function, so before writing anything you already know you can afford to spend one whole variable freezing it to a constant.
2. Method the expression suggested: the cyclic symmetry under $x \to y \to z \to x$ in the constraint is the signature of a coboundary, and it suggests the potential function ansatz $g(x) = f(x, c)$ for a fixed c immediately, before any calculation.
3. Where you would hesitate, and the thought to reinforce: the natural hesitation is whether $f(0, x)$ and $f(x, 0)$ are related at all; reinforce the habit of collapsing the diagonal first, before touching a basepoint, since that always buys antisymmetry for free and removes the hesitation in advance.
4. The trap the question sets: nothing in the statement suggests f is continuous, and under time pressure the instinct is to reach for a derivative; the fix is recognising that a problem's silence about regularity is deliberate, not an omission you are meant to patch with an assumption.
5. What makes you recognise it faster next time: any functional equation whose left side is a sum of the same two variable function evaluated around a cycle of variables, equalling a constant, belongs to this family, and the diagonal collapse followed by a basepoint fix applies without exception.
6. The reusable decision rule: a cyclic sum of a two variable function that vanishes identically forces that function to be a coboundary; construct the potential by freezing one argument to a basepoint, after first extracting antisymmetry from a diagonal substitution.

Mirror drills

1. Let $h : \mathbb{R}^2 \to \mathbb{R}$ satisfy $h(x, y) + h(y, z) + h(z, x) = c$ for a fixed constant c and all real x, y, z. For which c does a function g with $h(x, y) = g(x) - g(y)$ exist, and what is g? *Answer:* setting $x = y = z = t$ forces $h(t, t) = c/3$ for every t, but any difference $g(t) - g(t)$ is always 0, so a representation of this shape forces $h(t, t) = 0$, which is only consistent with $c = 0$; when $c = 0$ the argument is identical to A1 and $g(x) = h(x, 0)$.
2. Suppose $f : \mathbb{R}^3 \to \mathbb{R}$ (now a genuinely three argument function) satisfies $f(x, y, z) + f(y, z, x) + f(z, x, y) = 0$ for all x, y, z. Must f decompose as a difference of two variable potentials the way A1's f does? *Answer:* no in general; A1's argument relied on f being built from pairs drawn out of a common domain, so that fixing one slot to a basepoint reduces the arity by exactly one and lands you back on a one variable object. Here cyclic invariance under permuting three genuinely independent arguments gives you an invariance property, not a telescoping decomposition, since there is no basepoint substitution that collapses three slots down to two while preserving the structure of the claim.
3. A function $k : \mathbb{Z}^2 \to \mathbb{Z}$ satisfies $k(m, n) + k(n, p) + k(p, m) = 0$ for all integers m, n, p. Show $k(m, n) = a(m) - a(n)$ for some $a : \mathbb{Z} \to \mathbb{Z}$, and that a is unique up to an additive constant. *Answer:* identical proof with basepoint 0, setting $a(m) := k(m, 0)$; uniqueness up to a constant

follows exactly as in the ideal solution's closing remark, by comparing two hypothetical representations at $n = 0$.

18.2 A2. a determinant game decided by pairing

Alan and Barbara play a game in which they take turns filling entries of an initially empty 2008×2008 *array. Alan plays first. At each turn, a player chooses a real number and places it in a vacant entry. The game ends when all the entries are filled. Alan wins if the determinant of the resulting matrix is nonzero; Barbara wins if it is zero. Which player has a winning strategy?*

Read it out loud

Strip away the game dressing and read what the winning condition actually is: Alan wants the final matrix to have full rank, Barbara wants it to have dependent rows. That single observation reclassifies this from a combinatorial game about 2008^2 individual real numbers into a linear algebra question about the row vectors of the matrix. You should have that translation written down within the first thirty seconds, because everything else follows from asking what structural property of the rows a player can force, not from thinking about entries one at a time.

Once you see it as a rank question, the next thing to notice is who is trying to force what. Alan cannot force independence directly, since Barbara gets to choose half of the entries and can always choose adversarially. Barbara, on the other hand, only needs to force some fixed linear relation among the rows, and a fixed relation is a much easier target than full rank is, because a relation just needs to hold once, robustly, no matter what real numbers Alan throws at the board.

The board size, 2008×2008, is chosen with a purpose: 2008 is even. Whenever a combinatorial game hands you an even sized board and a second moving player, the examiners are all but telling you to look for a pairing or mirroring strategy, since an even number of cells splits cleanly into pairs with nothing left over. The fields in play here are matrices and determinants for the target condition, and games and invariants for the mechanism that forces it; no case analysis on the actual sizes of Alan's numbers should appear anywhere in a correct solution, since the strategy has to work against every real number simultaneously.

Approaches, ranked worst to best

1. **heuristic guessing.** *Cue:* a game with a first mover suggests a first mover advantage. *Circuit:* guess an outcome from intuition about who moves first, without exhibiting a strategy. (why it ranks here: a competition proof of a game result requires an explicit strategy for the winning player together with a verification that it always succeeds; a correct guess with no construction earns essentially nothing, since the problem is asking you to prove a claim about all possible plays, not to state an opinion).

2. **small board exhaustion.** *Cue:* tiny cases such as 2×2 or 3×3 are checkable by hand. *Circuit:* play out or enumerate small boards, hoping the pattern generalises to 2008×2008. (why it ranks here: it costs enormous time relative to the insight gained, and worse, it risks a wrong generalisation, since a strategy that happens to work on a 2×2 board by coincidence need not reveal the mechanism that scales; you need the mechanism, not the outcome, and small cases do not hand you the mechanism for free).

3. **row copying pairing.** *Cue:* Barbara only needs two equal rows to force determinant zero. *Circuit:* pair each entry of row one with the entry directly below it in row two, and whenever Alan plays in either row, Barbara completes the pair with the identical value. (this is correct and legitimate, but it needs a second, separate rule to cover every cell outside rows one and two, since those cells must still be filled legally even though their values are irrelevant to the argument; the proof is correct but structurally two strategies stitched together, one load bearing and one not, and the underlying reason the determinant vanishes, namely two identical rows, is a comparatively narrow fact about the alternating property of the determinant rather than a fact about its rank in general).

4. **row negation pairing.** *Cue:* Barbara can force every row into a fixed hyperplane rather than merely duplicating one row. *Circuit:* pair the columns within every row into disjoint pairs, and whenever Alan writes a value x into a cell, Barbara writes $-x$ into its partner cell in the same row. (this is the ideal route: it is a single uniform rule applying to every move anywhere on the board with no special casing, it uses the evenness of 2008 exactly once and exactly where the problem places it, and the reason the determinant vanishes generalises immediately, more vectors than the dimension of the space confining them are automatically dependent).

The ideal solution

Claim. Barbara has a winning strategy.

Partition the 2008 columns into 1004 disjoint pairs, for instance column $2t-1$ with column $2t$ for $t = 1, \ldots, 1004$; any fixed pairing works equally well. This induces, inside every one of the 2008 rows, a partition of that row's 2008 cells into 1004 column pairs, and hence a partition of the entire grid of 2008^2 cells into $2008^2/2$ *cell pairs*, where a cell pair consists of two cells lying in the same row, in two columns that are partnered.

Barbara's strategy. Whenever Alan writes a real number x into a vacant cell, Barbara immediately writes $-x$ into the partner cell of that cell pair.

This is always legal. We show by induction on the number of completed rounds (one round is one move of Alan's followed by Barbara's response) that immediately after each of Barbara's moves, every cell pair is either both empty or both filled. This holds trivially before play begins, when every cell pair is empty. Suppose it holds after round k. On round $k+1$, Alan plays into some vacant cell; by the invariant, its partner cell was also vacant, so the partner cell remains a legal target for Barbara. Barbara writes $-x$ there, filling both cells of that pair and leaving every other pair exactly as it was. The invariant is restored, so it holds after round $k+1$. Since 2008^2 is even, the game consists of exactly $2008^2/2$ such rounds, and Barbara always has a legal reply until the board is completely filled.

The determinant vanishes. Fix a row i. Its 1004 cell pairs each contribute a value x from one cell and $-x$ from its partner, so their contribution to the row sum is $x + (-x) = 0$; summing over all 1004 pairs, the entries of row i sum to 0. Writing $r_i \in \mathbb{R}^{2008}$ for row i and $\mathbf{1} = (1, 1, \ldots, 1)$, this says $r_i \cdot \mathbf{1} = 0$ for every $i = 1, \ldots, 2008$. Hence every one of the 2008 row vectors lies in the hyperplane

$$H = \{v \in \mathbb{R}^{2008} : v \cdot \mathbf{1} = 0\},$$

which is the orthogonal complement of the line spanned by $\mathbf{1}$ and therefore has dimension 2007. A linear subspace of dimension 2007 cannot contain 2008 linearly independent vectors, since a subspace of dimension d contains no more than d linearly independent vectors. So the 2008 rows of the completed matrix are linearly dependent, which is exactly the condition

$$\det = 0.$$

Since Barbara's rule is fully explicit, is always legal by the invariant argument, and forces $\det = 0$ regardless of every real number Alan chooses at every stage, it is a genuine winning strategy for Barbara, independent of how Alan plays. Therefore

$$\boxed{\text{Barbara has a winning strategy.}}$$

Layer 1: WHAT

Block 1: translate the winning condition, det equals zero if and only if the rows are linearly dependent, so Barbara's task becomes forcing a dependence relation among the rows. Block 2: recognise that the cleanest way to force dependence against an adversary choosing arbitrary reals is to confine every row to one fixed hyperplane, since a subspace of dimension less than the number

of vectors automatically forces dependence, whatever those vectors turn out to be. Block 3: design a response rule that guarantees each row lands in the hyperplane of coordinates summing to zero, by pairing columns within a row and always answering a move with its negative in the paired cell. Block 4: verify the pairing rule is always legal, using an invariant maintained by induction on the number of completed rounds. Block 5: conclude det equals zero via the dimension bound, and note the argument holds for every possible sequence of Alan's moves, so it is a genuine winning strategy.

The nominal route reasons about individual entries, or guesses an outcome from who moves first, or exhausts tiny boards hoping a pattern emerges. The bifurcation point is the very first line: whether you translate the determinant condition into a linear algebra statement about the rows before searching for a strategy, or whether you start hunting for a strategy while still thinking entrywise. Everything downstream, the choice of hyperplane, the pairing rule, the legality proof, follows mechanically once that translation is made; none of it is available before it.

Layer 2: HOW

Block 1. **Cue.** The winning condition is stated purely in terms of the determinant of the finished matrix, a property of the rows as vectors, not a property of any individual entry. **Why natural.** Determinant zero is, by definition, equivalent to the rows being linearly dependent, so restating Barbara's goal in these terms costs nothing and immediately opens up the entire toolkit of linear algebra rather than leaving you stuck reasoning about 2008^2 separate real numbers. **Link to the next block.** Dependence is not something you can force by controlling individual numeric values against an adversary, but it is something you can force by confining every row to a common proper subspace, which suggests the next question, which subspace, and how do you confine a row to it turn by turn.

Block 2. **Cue.** You need a linear condition, satisfied by every row, that is forced regardless of the specific real numbers Alan chooses, which means the condition has to be additive or otherwise insensitive to the actual magnitudes involved. **Why natural.** A subspace of dimension d in $\mathbb{R}^{2008}$ can contain at most d independent vectors; if $d < 2008$ and all 2008 rows are confined to it, dependence is automatic with no further argument needed about the specific values, only about the confinement. **Link to the next block.** The simplest linear condition to force turn by turn on an adversarially filled row is that its coordinates sum to zero, since a sum is exactly the kind of quantity that a pairing and cancelling rule can guarantee.

Block 3. **Cue.** The board has 2008 columns, an even number, so every row can be split into whole pairs of columns with nothing left over; evenness is precisely what makes a pairing rule exact rather than approximate. **Why natural.** Pairing the columns of each row and always answering Alan's value x in one cell with $-x$ in its partner cell guarantees that every pair contributes exactly zero to the row sum, so the full row sum is a sum of zeros, forcing the row into the hyperplane orthogonal to $(1, \ldots, 1)$. **Link to the next block.** A rule stated this way is uniform, it says exactly the same thing regardless of which cell Alan chooses anywhere on the board, which is what makes its legality provable by a single invariant rather than by casework on where Alan plays.

Block 4. **Cue.** Before trusting that the rule always has somewhere legal to respond, you should ask explicitly what could go wrong, namely Barbara's target cell already being filled. **Why natural.** Proving the invariant, every cell pair is jointly empty or jointly filled immediately after each of Barbara's moves, by induction on the number of completed rounds, closes this gap completely: the invariant holding after round k guarantees Alan's move in round $k+1$ disturbs exactly one pair, which Barbara's response exactly repairs. **Link to the next block.** With legality secured for the entire length of the game, the only thing left is to read off the consequence for the determinant once the board is full.

Block 5. **Cue.** Every row now satisfies one fixed linear equation, membership in a 2007 dimensional hyperplane, and there are 2008 rows. **Why natural.** The dimension count, 2008 vectors inside a 2007 dimensional space, forces linear dependence by the basic fact that a subspace cannot contain more independent vectors than its dimension, which is exactly the condition for the determinant to vanish. **Link to the next block.** None; this is the terminal block. The strategy did not depend on any choice Alan made, so it succeeds against every possible sequence of his moves, which is what a winning strategy means.

Layer 3: WHY

The determinant is an alternating multilinear function of the rows, and it vanishes exactly when the rows fail to be independent. Against an adversary who is free to choose any real number at all, you cannot hope to control the actual value of the determinant, since Alan's freedom over magnitudes defeats any attempt to steer a numeric outcome directly. What you can control, turn by turn, is a linear relation, because linear relations are additive and are exactly the kind of property preserved by a pairing and cancelling response. This is the deep reason the problem reduces to a subspace confinement argument rather than a numeric one: rank and dependence are properties of the configuration of vectors, invariant under exactly the operations, negation and pairing, that a second player moving in response can enforce.

The discernment mechanism to train is recognising, as early as possible, when a combinatorial game's winning condition is secretly a linear algebraic one. The signal is a winning condition phrased about the whole configuration (rank, determinant, a system having a solution) rather than about any single entry; whenever you see that signal, translate immediately, then ask what fixed subspace or fixed relation the second (or sometimes first) player can force using a pairing argument, exploiting the board decomposing into an even number of symmetric slots. The examiners chose 2008, an even number, and phrased the win condition through the determinant rather than through any individual entry, precisely to make both of these signals available to you from the first reading.

It is also worth naming, as a matter of general strategy design, that a strategy proof for a game has two separate obligations that are easy to conflate: showing the response rule is always legal, and showing it always achieves the target condition. Confusing these, for instance checking legality only in a few example rounds and treating the target condition as obviously achieved, is the most common way a game solution loses marks; each obligation gets its own short, separately checkable argument here, legality by the round by round invariant, and the target condition by the direct row sum computation.

Revolving concepts

Determinant, rank, and dependence. For a square matrix, $\det = 0$ if and only if the rows, equivalently the columns, are linearly dependent; this converts any determinant vanishing question into a dependency question about vectors, and it is the single fact that opens up this entire problem. See Book 1, field: Matrices and determinants, topic: rank and dependency criteria for the determinant.

Dimension bounds on independent sets. In a vector space or subspace of dimension d, any set of more than d vectors is linearly dependent; equivalently a subspace of dimension d contains at most d linearly independent vectors. See Book 1, field: Linear algebra, topic: dimension and the exchange lemma.

Hyperplanes as orthogonal complements. For a fixed nonzero vector $u \in \mathbb{R}^n$, the set $\{v : v \cdot u = 0\}$ is a linear subspace of dimension $n-1$, the orthogonal complement of the line spanned by u. Here $u = (1, \ldots, 1)$ names the specific hyperplane that the pairing strategy forces every row into, turning an abstract dimension count into a concrete, checkable subspace. See Book 1, field: Linear algebra, topic: orthogonal complements and hyperplanes.

Pairing and mirroring strategies in combinatorial games. When a game's move set decomposes into an even number of symmetric slots that group into pairs, the second player can maintain the invariant that after every one of their turns each pair is jointly filled or jointly empty, by always responding inside the same pair as the opponent's most recent move. This lets the second player force any algebraic relation between paired slots, equal values, sums to zero, negatives, regardless of the opponent's actual choices. See Book 1, field: Games and invariants, topic: pairing and mirroring strategies on even sized boards.

Invariant maintenance as an induction on rounds. Proving a pairing strategy is well defined for the entire length of a game is, structurally, an induction on the number of completed rounds: state the invariant, verify it initially, then show that one round of play, the opponent's move together with the prescribed response, restores it. *General recursion and its solution:* for any pairing based strategy, let $I(k)$ be the invariant after round k. $I(0)$ holds trivially before any moves are made. If $I(k)$ holds, the opponent's move in round $k+1$ disturbs exactly one paired unit, and the prescribed response repairs it, so $I(k+1)$ holds. By induction, $I(k)$ holds for every k up to the last round, so the strategy is legal throughout and its target invariant, here the row sum condition, holds at the final round when the board is full. This template applies to any board filling or token based two player game where the total move count is even and the board decomposes into pairs. See Book 1, field: Games and invariants, topic: loop invariants for turn based strategies.

Circuit Bank

- **WHEN** a combinatorial game's winning condition is really a linear algebraic property of a matrix or a vector configuration, such as rank, determinant, or dependency **DO** translate the winning condition into the equivalent linear algebra statement before searching for a strategy **BECAUSE** the translated condition suggests concrete structural targets, equal rows, a shared hyperplane, that a pairing strategy can force, while the entrywise determinant formula suggests nothing constructive on its own
- **WHEN** the board or move set splits naturally into an even number of symmetric slots, such as pairable columns within each row, or pairable rows overall **DO** design a strategy where the second player always responds inside the same pair as the opponent's most recent move, and prove legality by an invariant maintained through the whole game rather than by checking a few sample rounds **BECAUSE** pairing keeps the invariant every pair is jointly empty or jointly filled trivially inductive, guaranteeing a legal response exists at every single turn until the board is full
- **WHEN** you want to force a specific linear relation, such as a sum to zero or two rows equal, onto data that an adversary is filling with arbitrary real numbers **DO** make the forcing rule purely algebraic and local, copying or negating a value into a fixed partner cell, rather than any rule that depends on the actual size of the numbers played **BECAUSE** a purely algebraic pairing rule works identically for every real number the adversary could ever choose, so no casework on the magnitude of Alan's numbers is ever needed
- **WHEN** you have more vectors than the dimension of the space you have confined them to **DO** state linear dependence as the immediate conclusion, with no further computation **BECAUSE** this is a one line consequence of the definition of dimension, and re deriving it from scratch inside a competition proof spends ink for no additional credit

Backtesting

1. First useful cue: the win condition is stated purely through the determinant of the completed matrix, never through any individual entry, signalling immediately that this is a linear algebra game rather than an arithmetic one.

2. Method the expression suggested: since determinant zero is equivalent to dependent rows, and Barbara moves second on an even sized board, a pairing or mirroring strategy is the natural mechanism for a second player to force a structural condition on the rows.
3. Where you would hesitate, and the thought to reinforce: the hesitation is whether Barbara always has a legal empty cell to respond into; reinforce proving this as an invariant maintained through the whole game by induction on the number of completed rounds, rather than assuming it is obvious.
4. The trap the question sets: it is tempting to assume the first player has the advantage by default, since Alan moves first and sets the tone; the trap is not questioning that assumption, when in a game about forcing a linear relation it is the second player who has the advantage, because the second player can always react to and cancel the most recent move.
5. What makes you recognise it faster next time: any $n \times n$ array filling game whose winning condition is determinant zero versus determinant nonzero, with n even, is solved by a pairing strategy for whichever player wants determinant zero, provided that player moves second.
6. The reusable decision rule: when the losing condition for a player is that the configuration lies in a proper subspace, and that player moves second on a board whose slots divide evenly, force the configuration into a fixed hyperplane using a pairing or negation response to every move of the opponent.

Mirror drills

1. Redo the game on a 5×5 board, so n is odd, under the same rules. Does the row negation pairing strategy still work directly? *Answer:* not directly, since a row of length 5 cannot be split into whole column pairs; but the row copying strategy, pairing row one with row two and mirroring exactly, still works whenever there are at least two rows to pair, giving two identical rows regardless of the parity of n, so Barbara wins for every $n \geq 2$ by that alternative strategy.
2. In the same 2008×2008 game, suppose instead Barbara wins whenever the determinant is a perfect square and Alan wins otherwise. Does the row negation pairing strategy still help Barbara? *Answer:* no; the pairing mechanism was tailored to force linear dependence, which pins the determinant at exactly zero, but it gives no control over the actual nonzero value of a determinant, so a restated target about the determinant's value, rather than its vanishing, needs an entirely different mechanism.
3. Two players alternately fill an 8×8 grid with values chosen from $\{1, -1\}$, Alan first, and Barbara wants the sum of all 64 entries to equal 0. Does Barbara have a winning strategy, and what is it? *Answer:* yes; since 64 is even, Barbara pairs the cells arbitrarily into 32 pairs, and whenever Alan writes a value v in one cell of a pair, Barbara writes $-v$, the opposite sign, in the other cell of that pair; every pair then contributes 0 to the total, forcing the grand total to be exactly 0, the same pairing and cancelling circuit as A2.

18.3 A3. gcd lcm swaps as sorting valuations by prime

Start with a finite sequence $a_1, a_2, \ldots, a_n$ of positive integers. If possible, choose two indices $j < k$ such that a_j does not divide a_k, and replace a_j and a_k by $\gcd(a_j, a_k)$ and $\operatorname{lcm}(a_j, a_k)$, respectively. Prove that if this process is repeated, it must eventually stop and the final sequence does not depend on the choices made.

Read it out loud

Read the problem as two separate claims stitched together by one shared operation. The first claim, that the process eventually stops, is a termination question: you need one quantity that only ever moves in one direction and cannot move forever. The second claim, that the final sequence does not depend on which legal pair you chose at each stage, is a confluence question: it says that however adversarially or however carefully you play this, you always land in the same place. These are logically independent statements, and you should resist any temptation to find a single argument that magically delivers both, because a proof that quietly uses one to justify the other is usually hiding a gap.

The operation itself is the real cue. Whenever gcd and lcm appear together acting on a pair, the identity $\gcd(a, b) \cdot \operatorname{lcm}(a, b) = ab$ should be the first thing you write down, since it is the reason the product of the whole sequence never changes, and it immediately supplies a uniform upper bound on every term forever after. But the deeper cue is that gcd and lcm are defined prime by prime, as the minimum and the maximum of exponents in a factorisation. That means the entire multiplicative process, once you fix a single prime p and look only at the exponent of p in each a_i, becomes a process on ordinary integers where a pair out of order gets replaced by its minimum and its maximum. That is exactly one comparator step of a sorting network, and once you see that, both halves of the problem separate cleanly: termination needs its own monovariant, because a sorting network with an adversarial firing order is not obviously finite in general, while confluence is close to automatic once you know that sorting by comparator exchanges always preserves which values are present, no matter the order the comparators fire in.

The field embedded here is number theory, through valuations and unique factorisation, married to a combinatorial idea, sorting networks, with the whole thing packaged as a termination and confluence question for a rewriting process. Within thirty seconds you should have written both the product identity and the phrase valuation, prime by prime, since those two observations are what the entire solution is built from.

Approaches, ranked worst to best

1. **empirical tracking.** *Cue:* small sequences with two or three terms are easy to compute by hand. *Circuit:* chase the actual integer values through several arbitrary sequences of moves, hoping to see why the endpoint is always the same. (why it ranks here: it conflates the multiplicative structure of several primes at once, obscures the actual reason uniqueness holds, and does not scale to a general proof no matter how many small examples you check).

2. **bad pair counting.** *Cue:* the stopping condition is that no out of order pair remains, so counting out of order pairs directly measures progress. *Circuit:* let a bad pair be a pair $j < k$ with a_j not dividing a_k, and show the count of bad pairs strictly decreases at every move by checking, case by case, that fixing one bad pair never creates a new one elsewhere. (this is a genuinely valid route to termination, and it gives an attractive explicit bound of at most $\binom{n}{2}$ moves, but the case analysis needed to show no new bad pair is ever created, splitting on whether a third index sits between, below, or above the two indices being fixed, costs more ink than a single scalar monovariant, for no gain on the harder uniqueness half of the problem, which still needs the prime by prime argument regardless).

3. **layered rightmost term stabilises first.** *Cue:* the last term of the sequence can only ever be the larger index in a legal move, so it can only grow. *Circuit:* show the last term is nondecreasing and bounded, hence eventually fixed, then repeat the argument one position to the left, and so on. (this is the standard route and its motivation is transparent, but write ups of it usually leave the induction from position n down to position 2 as an informal and so on, which a strict grader would want tightened into an explicit induction before granting full credit).
4. **single weighted product monovariant.** *Cue:* the layered idea says later positions matter more, and a strictly increasing sequence of weights is exactly how you compress a layered argument into one scalar. *Circuit:* define $Q = \prod_{i=1}^{n} a_i^i$ and show it strictly increases by an explicit factor at every move while staying bounded throughout. (this is the ideal route for termination: it captures the same rightward bias as the layered argument, but as one displayed inequality with an explicit bound, with no case split and no informal induction to tighten; uniqueness then has no competing family of approaches, the prime by prime valuation counting argument below is the natural and essentially forced route).

The ideal solution

Throughout, write $P = a_1 a_2 \cdots a_n$ for the product of the terms of the sequence at the start; this will turn out to be invariant.

A first invariant and a uniform bound. A single move replaces a_j, a_k ($j < k$) by $\gcd(a_j, a_k)$ and $\operatorname{lcm}(a_j, a_k)$. Since $\gcd(a_j, a_k) \cdot \operatorname{lcm}(a_j, a_k) = a_j a_k$, the product of all n terms is unchanged by every move, so it equals P at every stage of the process, forever. Moreover, at every stage every current term is at most P: if $x_1, \ldots, x_n$ are the current terms, with product P, then for any single index i, the remaining $n-1$ terms are positive integers, hence their product is at least 1, so $x_i \le P$. The same reasoning applied to any two current terms a_j, a_k gives $a_j a_k \le P$, and since $\operatorname{lcm}(a_j, a_k)$ divides $a_j a_k$ it satisfies $\operatorname{lcm}(a_j, a_k) \le a_j a_k \le P$ as well; combined with $\gcd(a_j, a_k) \le a_j \le P$, every value that can ever be written into the sequence, at any stage, is at most P.

Termination. Define

$$Q = \prod_{i=1}^{n} a_i^i,$$

using the current values of $a_1, \ldots, a_n$. Consider one move on positions $j < k$, and let $d = \gcd(a_j, a_k)$, $u = a_j/d$, $v = a_k/d$; then u, v are coprime positive integers and $\operatorname{lcm}(a_j, a_k) = duv$. The ratio of Q after the move to Q before the move is

$$\frac{d^j (duv)^k}{(du)^j (dv)^k} = \frac{d^{j+k} u^k v^k}{d^{j+k} u^j v^k} = u^{k-j}.$$

Since a_j does not divide a_k, we have $u = a_j/d \neq 1$, and u is a positive integer, so $u \ge 2$; since $k > j$, the exponent $k - j$ is a positive integer, so this ratio is at least $2^{k-j} \ge 2$. Hence Q is a positive integer that is at least doubled by every move. On the other hand, since every current term is at most P throughout the process, and each exponent i satisfies $i \le n$,

$$Q = \prod_{i=1}^{n} a_i^i \le \prod_{i=1}^{n} P^i = P^{n(n+1)/2}$$

at every stage. A positive integer that at least doubles at every step and stays below the fixed bound $P^{n(n+1)/2}$ can change only finitely many times, at most $\log_2\left(P^{n(n+1)/2}\right)$ times. So only finitely many moves can ever occur, and the process stops.

It remains to identify exactly when it stops: by definition the process stops precisely when no pair $j < k$ with $a_j \nmid a_k$ remains, that is, precisely when

$$a_j \mid a_k \quad \text{for every } j < k.$$

Uniqueness. Fix a prime p. For a positive integer x, write $v_p(x)$ for the exponent of p in the factorisation of x. Since gcd and lcm take the minimum and the maximum of exponents at every prime, a move on positions $j < k$ replaces the pair of valuations $\big(v_p(a_j), v_p(a_k)\big)$ by $\big(\min(v_p(a_j), v_p(a_k)), \max(v_p(a_j), v_p(a_k))\big)$, while every other position is untouched.

Claim. For every threshold $m \geq 0$, the count $c_m = \#\{\, i : v_p(a_i) \geq m \,\}$ is unchanged by every move.

Only positions j, k can change the count, so it suffices to check their combined contribution before and after. Write $x = v_p(a_j)$, $y = v_p(a_k)$ before the move. If $x, y \geq m$, both contribute before, and $\min(x, y) \geq m$ and $\max(x, y) \geq m$ both contribute after: 2 either way. If $x, y < m$, neither contributes before, and $\min(x, y) < m$ and $\max(x, y) < m$ (the larger of two values below m is still below m) neither contributes after: 0 either way. If exactly one of x, y is at least m, exactly one contributes before (1 total); after, $\min(x, y) < m$ (the smaller one, which is below m) does not contribute, while $\max(x, y) \geq m$ (the larger one) does: again 1 total. In every case the contribution is unchanged, proving the claim.

Now suppose the process terminates, in n terms $a'_1, \ldots, a'_n$. By the termination characterisation above, $a'_j \mid a'_k$ for every $j < k$, hence $v_p(a'_j) \leq v_p(a'_k)$ for every $j < k$ and every prime p; that is, for each fixed p the sequence $v_p(a'_1) \leq v_p(a'_2) \leq \cdots \leq v_p(a'_n)$ is nondecreasing. A nondecreasing sequence is recovered exactly from its threshold counts: because the sequence is sorted, the set of indices i with $v_p(a'_i) \geq m$, if nonempty, must be the top few indices $\{\, n - c_m + 1, \ldots, n \,\}$ (if some index qualifies, every later index qualifies too, since the sequence is nondecreasing), so

$$v_p(a'_i) = \max\{\, m : c_m \geq n - i + 1 \,\}$$

for each i, which depends only on the function $m \mapsto c_m$.

By the Claim, c_m is unchanged by every individual move, so by induction on the number of moves in any particular run, c_m at termination equals c_m for the original sequence, for every m, whichever legal moves were chosen along the way. Combined with the previous paragraph, this shows $v_p(a'_i)$ is the same value for every terminating run of the process, for every prime p and every index i. It suffices to check this for the finitely many primes dividing P, since a term's valuation at any other prime is identically 0 throughout. Since a positive integer is determined completely by its valuation at every prime, $a'_i = \prod_p p^{\,v_p(a'_i)}$ is therefore the same for every terminating run, for each i. This proves the final sequence does not depend on the choices made.

Layer 1: WHAT

Block 1: record the product identity $\gcd(a_j, a_k)\,\mathrm{lcm}(a_j, a_k) = a_j a_k$, giving an invariant product P and a uniform bound $x_i \leq P$ on every term forever. Block 2: for termination, define the weighted product $Q = \prod a_i^i$, show a single move multiplies it by an integer factor at least 2, and bound it above using Block 1, forcing finitely many moves. Block 3: characterise the stopping condition exactly as $a_j \mid a_k$ for all $j < k$. Block 4: for uniqueness, switch entirely to a fixed prime p and its valuations, and show a move replaces a pair of valuations by their minimum and maximum. Block 5: show the threshold count $c_m = \#\{i : v_p(a_i) \geq m\}$ is invariant under every move, for every m. Block 6: combine Block 3 (the final valuation sequence is sorted) with Block 5 (its threshold counts are choice independent) to recover the final valuation sequence uniquely for each prime, then recombine across primes by unique factorisation.

The nominal route tries to chase the raw integers a_i directly through several moves, which mixes the contributions of different primes together and hides why the endpoint is forced. The bifurcation point is the decision to decompose: termination is handled by compressing a layered, rightward biased improvement into a single weighted scalar, while uniqueness is handled by abandoning the integers entirely and working one prime at a time, because gcd and lcm are defined prime by prime and nowhere else do they have a simple description.

Layer 2: HOW

Block 1. **Cue.** gcd and lcm appear together acting on the same pair, and the identity relating their product to the original product of the pair is the single most standard fact attached to this combination. **Why natural.** Recording $\gcd(a_j, a_k)\,\mathrm{lcm}(a_j, a_k) = a_j a_k$ immediately gives an invariant, the total product P, and a moment's more work upgrades it to a uniform bound on every individual term at every stage, which is exactly the kind of control a termination argument needs. **Link to the next block.** A bounded uniform ceiling on every term is precisely what turns any candidate monovariant built from the a_i into a genuinely bounded quantity, so the next question is which monovariant to build.

Block 2. **Cue.** The move always replaces the term at the smaller index j by something no larger, and the term at the larger index k by something no smaller, so position matters, later positions are biased to grow. **Why natural.** Weighting each term by its position, $\prod a_i^i$, turns that positional bias into an exact algebraic statement: the ratio of Q after a move to Q before is u^{k-j} for an integer $u \geq 2$, hence at least 2, so Q strictly and substantially increases at every single move while Block 1's bound keeps it from increasing forever. **Link to the next block.** A bounded, strictly increasing positive integer can only change finitely many times, which is exactly the statement that the process stops; the only thing left is to say precisely what state it stops at.

Block 3. **Cue.** The process is defined to stop exactly when no legal move remains, and a legal move requires an out of order pair, so the stopping condition is the direct negation of that requirement. **Why natural.** No pair $j < k$ with $a_j \nmid a_k$ remaining means, by definition, $a_j \mid a_k$ for every $j < k$, and this is already the strongest possible divisibility statement about the final sequence, nothing further needs to be derived from it at this stage. **Link to the next block.** Divisibility between two positive integers is equivalent to an inequality between their valuations at every prime simultaneously, so this block is the bridge into the entirely separate, prime by prime argument that uniqueness needs.

Block 4. **Cue.** Uniqueness is a statement about the final integers themselves, but gcd and lcm have no simple additive or comparative description as integers; they do have an exact description as minimum and maximum once you pass to the exponent of a single fixed prime. **Why natural.** Fixing a prime p and tracking $v_p(a_j), v_p(a_k)$ converts the move into exactly one step of a sorting comparator, replace an out of order pair of numbers by their minimum, in position order, and by their maximum, which is a purely combinatorial operation with nothing multiplicative left in it. **Link to the next block.** Once the operation is recognised as a comparator step, the standard question for any comparator based process is what invariant survives every comparator firing, which is answered directly by counting how many values clear a threshold.

Block 5. **Cue.** A comparator step only touches two positions and replaces their values by the min and the max of the same two values, so any quantity that only cares how many values exceed a fixed threshold cannot be affected by the swap. **Why natural.** Checking the three cases, both above the threshold, both below, or exactly one above, shows the count of positions at or above any fixed threshold m is identical before and after a move, for every m and every prime p simultaneously, with no dependence on which of the two positions held the larger value. **Link to the next block.** An invariant count for every threshold is precisely the information needed to reconstruct a sorted sequence exactly, so the final step is purely about combining that reconstruction with the fact that the terminal valuation sequence is sorted.

Block 6. **Cue.** The terminal sequence is sorted in valuation, by Block 3, and its threshold counts are pinned down independently of any choices made, by Block 5. **Why natural.** A sorted sequence

is uniquely recoverable from its threshold counts, since the indices exceeding any threshold must be exactly the top few by sortedness, so the terminal valuation sequence at every prime is forced to be the same for every run; recombining across all primes dividing P via unique factorisation then pins down the terminal integers themselves. **Link to the next block.** None; this is the terminal block. Naming the wider frameworks this fits into, sorting networks and the structure theorem for finite abelian groups, is deferred to the revolving concepts.

Layer 3: WHY

The single structural truth underneath both halves of this problem is that gcd and lcm are, prime by prime, nothing but min and max of exponents. Every other fact used, the product identity, the bound on every term, the comparator description, the threshold counting, is a direct consequence of that one truth applied either multiplicatively across all primes at once, for the termination bound, or to a single prime in isolation, for uniqueness. Termination and uniqueness genuinely need different tools because they are different kinds of claims: termination is quantitative, it asks how long can this take, and needs a monovariant with an explicit bound; uniqueness is structural, it asks does the answer depend on the path, and needs an invariant that is exactly preserved, not merely bounded, at every step.

The discernment mechanism worth training is to recognise, the moment you see a process that repeatedly replaces a pair by its minimum and its maximum, that you are looking at one comparator of a sorting network, and that sorting networks have two separately provable properties: they terminate under any legal firing schedule if some potential function forces them to, and they are confluent, always landing on the same sorted output, because a threshold counting argument shows the underlying multiset of values is preserved by every single comparator firing regardless of order. Here that recognition has to be applied simultaneously at every prime, which is why decomposing into valuations first, before attempting either half of the proof, is the move that makes everything else routine rather than mysterious.

The examiners nudge you toward the prime decomposition by defining the operation entirely through gcd and lcm, objects with no natural description except through prime factorisations, and they nudge you toward a monovariant, rather than a direct combinatorial bound on the number of moves, by allowing the choice of which pair to fix at each stage to be completely adversarial, meaning no fixed schedule of moves can be assumed. A crude count of possible states is not available to you, so a genuine monovariant, one that improves under every conceivable choice, not just a convenient one, is the only tool suited to a process that a hostile or arbitrary chooser might run.

Revolving concepts

p-adic valuation. For a prime p and positive integer x, $v_p(x)$ is the exponent of p in the factorisation of x; $v_p(\gcd(a,b)) = \min(v_p(a), v_p(b))$ and $v_p(\operatorname{lcm}(a,b)) = \max(v_p(a), v_p(b))$ for every prime p, and since $\min(x,y) + \max(x,y) = x + y$, summing over all primes recovers $\gcd(a,b)\operatorname{lcm}(a,b) = ab$ without appealing to that identity directly. See Book 1, field: Number theory, topic: valuations and unique factorisation.

Monovariants and their construction by weighting position. A monovariant is a single quantity, attached to the current state of a process, that is bounded and moves in only one direction at every legal step, regardless of which legal step is chosen; boundedness together with strict monotonicity of an integer valued quantity forces termination. *General recursion and its solution:* when a process has a layered, lexicographic improvement, a later coordinate never gets worse and strictly improves whenever it changes, while earlier coordinates are secondary, choosing any strictly increasing weights $w_1 < w_2 < \cdots < w_n$ and combining via $\sum w_i x_i$ or $\prod x_i^{w_i}$ compresses that layered structure into one scalar monovariant, replacing an inductive argument carried out position by position with

a single displayed inequality; the exponent choice $w_i = i$ used here is the simplest such weighting, and the same compression applies to any finite process with a provable lexicographic improvement. See Book 1, field: Sequences and recurrences, topic: monovariants and weighted compression of a layered argument.

Sorting networks and confluence. A process that repeatedly finds any out of order pair, by position and value, and replaces it with its minimum and maximum in position order, applied in any sequence of choices, always terminates at the same fully sorted arrangement of the original multiset; this is the correctness property of comparator based sorting, and its standard proof is exactly the threshold counting argument used above, tracking how many elements clear each threshold. See Book 1, field: Discrete structures and graphs, topic: sorting networks and comparator based confluence.

Reconstructing a sorted sequence from its threshold counts. A nondecreasing sequence of n nonnegative integers is uniquely determined by the function $m \mapsto \#\{i : x_i \geq m\}$; this is the classical duality between a partition and its conjugate, the transpose of its Young diagram. See Book 1, field: Combinatorics, topic: partitions and conjugate partitions.

Invariant factors and the structure theorem for finite abelian groups. Every finite abelian group is isomorphic to a product $\mathbb{Z}/d_1\mathbb{Z} \times \cdots \times \mathbb{Z}/d_n\mathbb{Z}$ with $d_1 \mid d_2 \mid \cdots \mid d_n$, and this divisibility chain of invariant factors is uniquely determined by the group, independent of how it is presented. The final sequence of this problem is exactly the invariant factor decomposition of $\mathbb{Z}/a_1\mathbb{Z} \times \cdots \times \mathbb{Z}/a_n\mathbb{Z}$, computed by a purely combinatorial process instead of the row and column reduction usually used to compute a Smith normal form. This is genuinely beyond the installed toolkit and worth flagging rather than glossing: it is a full undergraduate algebra theorem, only its shadow is needed here. See Book 1, field: Abstract algebra, topic: Smith normal form and invariant factors.

Circuit Bank

- **WHEN** an operation replaces two numbers by their gcd and their lcm, or more generally by the componentwise minimum and maximum of some invariant attached to them **DO** decompose by prime, or by coordinate, and track the minimum and maximum of the exponents or coordinates separately for each prime **BECAUSE** gcd and lcm act independently and identically at every prime, so per prime bookkeeping turns a multiplicative question into a simple ordering question with nothing else mixed in
- **WHEN** you need termination of a process that can be run in any adversarial order rather than a fixed schedule **DO** hunt for a single scalar monovariant that is bounded and strictly improves at every legal move, preferring a weighted product or sum that captures a layered bias over an explicit layer by layer induction **BECAUSE** a single displayed inequality is faster to write and to check than an induction hidden behind an informal and so on, and it works under every possible ordering of moves by construction
- **WHEN** a claim has the form the final state does not depend on the choices made **DO** search for quantities attached to the current state, not to the sequence of moves, that are provably unchanged by every individual legal move, and check whether enough of them pin the state down completely **BECAUSE** if the invariants determine the state completely, as the threshold counts do here for every prime, then any two terminating runs share the same invariants and hence the same final state, with no need to compare the two runs move by move
- **WHEN** a sequence is repeatedly having an out of order pair, by position, replaced by its minimum and maximum in position order **DO** recognise it as a sorting network and recall that the final sorted order is already determined by the initial multiset, independent of which legal pair is fixed at each stage **BECAUSE** this is the defining correctness property of comparator based sorting, and citing it once justified saves re deriving confluence from scratch on every similar problem

- **WHEN** a problem asks you to prove two logically separate things, such as termination and then uniqueness of the outcome **DO** prove them in two clearly separated sections with two different invariants or techniques, never one argument claimed to deliver both at once **BECAUSE** conflating does it stop with does it always stop at the same place tends to produce a proof that silently assumes one claim while trying to establish the other

Backtesting

1. First useful cue: gcd and lcm are named together acting on the same pair, and the identity $\gcd(a,b)\,\mathrm{lcm}(a,b) = ab$ should be written down on sight, since it hands you both an overall invariant and a uniform bound on every term for the rest of the process.
2. Method the expression suggested: gcd and lcm are defined prime by prime as minimum and maximum of exponents, so the moment the problem asks whether the outcome is independent of choices, the suggested method is to track, separately for each prime, how the min and max operation moves valuations around, which is exactly sorting.
3. Where you would hesitate, and the thought to reinforce: the hesitation is usually how do you know this terminates when the moves can be chosen adversarially rather than in a fixed order; reinforce find one scalar quantity, strictly increasing and bounded, that does not care what order the legal moves are chosen in, and prefer a weighted product over a layered induction.
4. The trap the question sets: it is tempting to try to prove uniqueness by comparing two arbitrary sequences of choices move by move, an approach that looks like it needs to consider exponentially many cases; the trap is not noticing that several per prime, per threshold invariants pin down the entire final state at once, which sidesteps any need to compare move sequences directly.
5. What makes you recognise it faster next time: any process defined by replace a pair by the minimum and maximum of something, on a sequence, with a stopping condition phrased as fully sorted, is a comparator or sorting network problem, and both its termination and its confluence have standard, separately provable mechanisms, a monovariant, and threshold counting invariants respectively.
6. The reusable decision rule: a gcd or lcm replacement process reduces to sorting the valuation sequence at every prime simultaneously; prove termination with a strictly increasing bounded weighted product monovariant, and prove uniqueness with threshold count invariants, one prime at a time.

Mirror drills

1. A sequence of positive integers $b_1, \ldots, b_n$ repeatedly has two entries b_j, b_k, with $j < k$ and $b_j > b_k$, replaced by $\min(b_j, b_k)$ and $\max(b_j, b_k)$ whenever they are out of numeric order. Show this always terminates at the sorted sequence, independent of the order of swaps chosen. *Answer:* this is the identical circuit with the raw integers themselves playing the role of a single prime's valuations; termination follows from a position weighted sum or product monovariant, and uniqueness follows because every swap preserves how many entries meet or exceed any fixed threshold, forcing the final sorted arrangement of the original multiset to be the only possible outcome.
2. For positive integers a, b and a prime p, express $v_p(\gcd(a,b))$ and $v_p(\mathrm{lcm}(a,b))$ in terms of $v_p(a)$ and $v_p(b)$, and use this to reprove $\gcd(a,b)\,\mathrm{lcm}(a,b) = ab$ without quoting that identity directly. *Answer:* $v_p(\gcd(a,b)) = \min(v_p(a), v_p(b))$ and $v_p(\mathrm{lcm}(a,b)) = \max(v_p(a), v_p(b))$; since $\min(x,y) + \max(x,y) = x + y$ at every prime's exponent, the two sides of the claimed identity have the same valuation at every prime, hence are equal by unique factorisation.

3. A collection of chips is repeatedly merged by a rule that always preserves, for every threshold m, the number of chips with value at least m. If two different merge schedules both terminate at a totally sorted arrangement, must the two final arrangements be identical? *Answer:* yes, by exactly the uniqueness argument of A3: a sorted arrangement is uniquely reconstructed from its threshold counts, those counts are schedule independent invariants by hypothesis, so both schedules are forced to terminate at the same sorted arrangement.

18.4 A4. An iterated logarithm series that diverges

Define $f\colon \mathbb{R} \to \mathbb{R}$ *by*

$$f(x) = \begin{cases} x & \text{if } x \le e, \\ xf(\ln x) & \text{if } x > e. \end{cases}$$

Does $\displaystyle\sum_{n=1}^{\infty} \frac{1}{f(n)}$ *converge?*

Read it out loud

Say it to yourself: this is a series whose general term is not given by a formula at all, but by a recursion in which the argument is pushed down by a logarithm. The recursion is multiplicative: f at x is x times f at $\ln x$. The map $x \mapsto \ln x$ is a contraction that drags any starting point below e in finitely many steps, so the recursion is a finite product, and it terminates at a different depth depending on where x sits. That last clause is the whole problem. The domain of f is naturally cut into blocks by the tower $1, e, e^e, e^{e^e}, \ldots$, and on each block f is a different explicit product.

The fields wired in are real analysis (convergence of a positive series), calculus of iterated functions (the chain rule for $\ln \circ \ln \circ \cdots$), and the theory of the logarithmic scale of divergence, the Bertrand series. What is really being tested is whether you will unfold a recursion before you reach for a convergence test. Everybody knows the integral test. Almost nobody writes down the closed form on the second and third blocks before deciding what to do.

The first useful cue is the shape $f(x) = xf(\ln x)$ itself. Within thirty seconds you should have on your page:

$$f(x) = x \text{ on } (1, e], \qquad f(x) = x \ln x \text{ on } (e, e^e], \qquad f(x) = x \ln x \ln\ln x \text{ on } (e^e, e^{e^e}],$$

and the observation that the denominators are exactly the denominators of the classical Bertrand series $\sum \frac{1}{n \ln n \ln\ln n \cdots}$. Every one of those diverges. Your instinct at that point should already be that the answer is divergence, and the only remaining job is to find the cheapest honest proof.

The second cue is quieter and worth naming, because it is the one that makes the proof one line long. Take reciprocals in the recursion: $\frac{1}{f(x)} = \frac{1}{x} \cdot \frac{1}{f(\ln x)}$. Now read that as a statement about derivatives. If G is a function with $G' = 1/f$ on the previous block, then $\frac{1}{x} G'(\ln x)$ is the derivative of $G(\ln x)$. In other words the defining recursion of this problem *is* the chain rule for composing with ln. So $1/f$ is not merely comparable to the derivative of an iterated logarithm; it is exactly that derivative.

Approaches, ranked worst to best

1. **Ratio test or root test.** *Cue:* it is an infinite series, so run the standard tests. *Circuit:* default convergence testing. Both give exactly 1, because $f(n+1)/f(n) \to 1$ inside every block. You learn nothing and you have spent five minutes. Any test that only sees the term's exponential growth rate is blind to a series this close to the boundary.
2. **Comparison with $\sum \frac{1}{n \ln n}$.** *Cue:* the second block gives $f(x) = x \ln x$. *Circuit:* compare with a known Bertrand series. This ranks near the bottom for a reason worth burning in: the inequality points the wrong way. For $x > e^e$ we have $f(x) \ge x \ln x$, hence $\frac{1}{f(n)} \le \frac{1}{n \ln n}$, an upper bound by a divergent series. That carries zero information. To prove divergence you need a lower bound, so any comparison against a fixed Bertrand series must fail, because f eventually beats every one of them.
3. **Cauchy condensation.** *Cue:* the terms are positive and decreasing. *Circuit:* condense a monotone series. Correct, but you must condense once per logarithm and the bookkeeping multiplies. You are simulating the integral test by hand with worse constants.

4. **Block the sum along the tower and estimate each block by hand.** *Cue:* the closed form for f changes at $1, e, e^e, \ldots$. *Circuit:* partition the index set where the formula changes. This is the right skeleton. It ranks second only because to evaluate a block sum you will end up comparing with an integral anyway, so you may as well integrate from the start.
5. **Integral test, then substitute $y = \ln^k x$ on the kth block.** *Cue:* $1/f$ is the derivative of the kth iterated logarithm. *Circuit:* a multiplicative recursion under ln is the chain rule. Best. Each block contributes exactly 1, with no estimation, no epsilon, and no constants. An exact answer on every block is always cheaper than an inequality on every block.

The ideal solution

Write $\ln^1 x = \ln x$ and $\ln^k x = \ln(\ln^{k-1} x)$ for $k \geq 2$. Define the tower $t_0 = 1$ and $t_k = e^{t_{k-1}}$, so $t_1 = e$, $t_2 = e^e$, and so on.

Step 0: f is well defined on $[1, \infty)$. For $x > e$ the recursion refers to $f(\ln x)$ with $\ln x < x$. If some x had $\ln^k x > e$ for every k, the sequence $x_k = \ln^k x$ would be strictly decreasing and bounded below by e, hence convergent to some $L \geq e$ with $L = \ln L$, which is impossible since $\ln L < L$ for all $L > 0$. So every x falls below e after finitely many logarithms and the recursion closes.

Lemma 18.4.1. *For every $k \geq 1$ and every $x \in (t_{k-1}, t_k]$,*

$$f(x) = x \ln^1 x \, \ln^2 x \cdots \ln^{k-1} x$$

(the product being empty when $k = 1$).

Proof. Induction on k. For $k = 1$ the interval is $(1, e]$, where $x \leq e$ and $f(x) = x$ by definition. Let $k \geq 2$ and $x \in (t_{k-1}, t_k]$. Then $x > t_{k-1} \geq e$, so $f(x) = xf(\ln x)$, and $\ln x \in (t_{k-2}, t_{k-1}]$. By the induction hypothesis $f(\ln x) = \ln x \prod_{j=1}^{k-2} \ln^j(\ln x) = \prod_{j=1}^{k-1} \ln^j x$, since $\ln^j(\ln x) = \ln^{j+1} x$. Multiply by x. □

Lemma 18.4.2. *f is positive and increasing on $[1, \infty)$.*

Proof. Fix $k \geq 1$ and $x \in (t_{k-1}, t_k]$. For $1 \leq j \leq k-1$ we have $\ln^j x > \ln^j t_{k-1} = t_{k-1-j} \geq t_0 = 1$, so every factor in the Lemma is greater than 1 and is an increasing function of x on that interval. Hence f is positive and increasing on each block. The blocks fit together: at $x = t_{k-1}$ the two formulas differ only by the factor $\ln^{k-1} t_{k-1} = t_0 = 1$, so f is continuous at each tower point and therefore increasing on all of $[1, \infty)$. □

Lemma 18.4.3. *For every $k \geq 1$, $\displaystyle\int_{t_{k-1}}^{t_k} \frac{dx}{f(x)} = 1$.*

Proof. The function $L_k(x) = \ln^k x$ is differentiable wherever $\ln^{k-1} x > 0$, that is for $x > t_{k-2}$, and by the chain rule

$$L_k'(x) = \frac{1}{x \ln^1 x \cdots \ln^{k-1} x} = \frac{1}{f(x)} \qquad \text{on } (t_{k-1}, t_k]$$

by the first Lemma. Since $\ln^k t_{k-1} = \ln 1 = 0$ and $\ln^k t_k = 1$, the Fundamental Theorem of Calculus gives $\int_{t_{k-1}}^{t_k} \frac{dx}{f(x)} = L_k(t_k) - L_k(t_{k-1}) = 1$. □

Conclusion. By the second Lemma, $g(x) = 1/f(x)$ is positive and nonincreasing on $[1, \infty)$, so $g(n) \geq \int_n^{n+1} g(x)\, dx$ for every integer $n \geq 1$. Summing,

$$\sum_{n=1}^{N} \frac{1}{f(n)} \;\geq\; \int_1^{N+1} \frac{dx}{f(x)}.$$

By the third Lemma, $\int_1^{t_K} \frac{dx}{f(x)} = \sum_{k=1}^{K} 1 = K$, which is unbounded as $K \to \infty$. Hence the partial sums of the series are unbounded, and

$$\boxed{\sum_{n=1}^{\infty} \frac{1}{f(n)} \text{ diverges.}}$$

Layer 1: WHAT

Block 1. Unfold the recursion. Get the tower t_k and the closed form $f(x) = x \ln^1 x \cdots \ln^{k-1} x$ on the kth block, and check f is well defined at all.

Block 2. Take reciprocals and recognise $1/f$ as $\frac{d}{dx} \ln^k x$ on the kth block.

Block 3. Establish that f is positive and increasing across the block boundaries, so the integral test is available in the direction that proves divergence.

Block 4. Integrate over each block: the iterated logarithm carries the block onto $[0, 1]$, so every block contributes exactly 1. Infinitely many blocks, so the integral, and hence the sum, diverges.

Juxtaposition. The nominal route reads "series, positive terms" and runs down the checklist: ratio, root, comparison with $1/n$, comparison with $1/(n \ln n)$. It never unfolds the recursion, so it never sees that f outgrows every fixed Bertrand denominator, and it wastes its comparisons in the direction that cannot conclude. The ideal route treats f as an object to be computed before it is estimated.

The bifurcation point is the first line after reading the definition. Ideal: write f explicitly on $(e, e^e]$ and $(e^e, e^{e^e}]$. Zombie: write "let $a_n = 1/f(n)$" and start testing. Everything downstream is decided there.

Layer 2: HOW

Block 1. **Cue.** The recursion $f(x) = xf(\ln x)$ has an argument that strictly decreases and a base region $x \le e$. Any recursion with a contracting argument and a base case is a finite product, and its length depends on the starting point. **Why natural.** You cannot test a series whose term you cannot write. Unfolding costs two lines and converts an implicit definition into an explicit one. The tower $t_k = e^{t_{k-1}}$ is forced on you: it is precisely the set of points at which the number of logarithms needed changes, that is, the escape time of x under ln. **Link to the next block.** Once the closed form is on the page, the denominator $x \ln x \ln \ln x \cdots$ is staring at you, and that product is the most recognisable object in the theory of divergent series.

Block 2. **Cue.** A product of a variable and all of its iterated logarithms in a denominator. You have met that exact product as the derivative of $\ln \ln x$, of $\ln \ln \ln x$, and so on. **Why natural.** This is not a coincidence to be discovered; it is the recursion restated. Dividing $f(x) = xf(\ln x)$ through gives $\frac{1}{f(x)} = \frac{1}{x} \cdot \frac{1}{f(\ln x)}$, and $\frac{1}{x} G'(\ln x)$ is by the chain rule the derivative of $G \circ \ln$. So if $1/f$ is a derivative of something on one block, it is the derivative of that something composed with ln on the next. The recursion propagates antiderivatives. **Link to the next block.** Having an exact antiderivative means you will never need an inequality on a block, only an evaluation, so you now want an integral in your argument.

Block 3. **Cue.** You intend to use $\sum \ge \int$. That inequality has a hypothesis, and the hypothesis is monotonicity on the whole ray, not just inside blocks. **Why natural.** The two integral test inequalities are not symmetric: for nonincreasing positive g you get $\int_1^{N+1} g \le \sum_{n=1}^{N} g(n) \le g(1) + \int_1^N g$. Only the left one proves divergence of the sum from divergence of the integral. Choosing the wrong one is the single most common way to write a paragraph that proves nothing. Checking continuity at the tower points is not pedantry; a jump downward there would break the monotonicity you are

about to invoke. **Link to the next block.** With g nonincreasing and the block antiderivative known, the computation is now purely arithmetic.

Block 4. Cue. $\ln^k$ maps the kth block onto $[0,1]$, because t_{k-1} and t_k are by construction the points where the kth iterated logarithm equals 0 and 1. **Why natural.** The tower was defined to make this true; that is the payoff of choosing the blocks by escape time rather than by, say, powers of two. Each block therefore contributes exactly one unit of "logarithmic length", independent of k, and a sum of infinitely many ones diverges. Notice that the terms $1/f(n)$ are becoming absurdly small, of size roughly $1/t_k$ on the kth block; divergence is powered entirely by the fact that the blocks get long faster than the terms get small. **Link to the next block.** Nothing remains: unbounded partial sums is the definition of divergence for a positive series.

Layer 3: WHY

The truth beneath this problem is the structure of the logarithmic scale. For each fixed k the Bertrand series

$$\sum_n \frac{1}{n \ln^1 n \, \ln^2 n \cdots \ln^k n}$$

diverges, while inserting any exponent $1+\varepsilon$ on the last factor makes it converge. So the family of divergent series has no smallest member you can name at this level: for every convergent comparison you propose there is a divergent series just beneath it, and vice versa. This problem hands you the *diagonal* of that family. At each x it uses exactly as many logarithmic factors as are available before they drop below 1, and asks whether the diagonal falls on the divergent side. It does, and the exact block computation says why in the sharpest possible terms: each block contributes exactly 1, so the diagonal sits precisely at the divergence boundary, at the slowest possible rate, growing like $\ln^*$, the iterated logarithm of the upper limit.

The discernment mechanism is a general one, and it is worth stating as a policy rather than an observation. When a function is defined recursively by $F(x) = G(x)F(\sigma(x))$ with σ contracting into a base region, do not test, do not estimate, do not compare. Unfold. You get

$$F(x) = \Big(\prod_{j=0}^{k-1} G(\sigma^j(x))\Big) F_{\text{base}}(\sigma^k(x)), \qquad k = \text{escape time of } x,$$

and the natural partition of the domain is by escape time. Every question about F then becomes a question about that product on one block.

The second and sharper piece of discernment is reading a recursion as a differential statement. A multiplicative recursion with the factor $1/\sigma'(x)$ built into it, which is exactly what $f(x) = xf(\ln x)$ is since $\ln' x = 1/x$, is the chain rule written backwards. Train it like this: whenever you see $\frac{1}{h(x)} = \sigma'(x) \cdot \frac{1}{h(\sigma(x))}$, immediately write "so $1/h$ is the derivative of (previous antiderivative) $\circ\, \sigma$". The examiners nudged toward this by choosing ln and by putting the factor x in front, which is precisely $1/\ln'(x)$. Any other factor there and the problem would have a different answer; the factor x is the guiding light.

Revolving concepts

Iterated logarithm. $\ln^k x$ is defined for $x > t_{k-1}$, is increasing, and has derivative $\left(x \ln^1 x \cdots \ln^{k-1} x\right)^{-1}$. The tower $t_k = e^{t_{k-1}}$ is its "unit ladder": $\ln^k$ carries $[t_{k-1}, t_k]$ onto $[0,1]$. See Book 1, field: Real analysis, topic: iterated logarithms and growth scales.

Integral test. For g positive and nonincreasing on $[1,\infty)$: $\int_1^{N+1} g \le \sum_{n=1}^{N} g(n) \le g(1) + \int_1^N g$. Both halves come from $\int_n^{n+1} g \le g(n) \le \int_{n-1}^{n} g$. Know which half you need before you write it. See Book 1, field: Real analysis, topic: convergence tests.

Bertrand series and the Abel and Dini theorem. $\sum n^{-1}(\ln n)^{-\alpha}$ converges iff $\alpha > 1$; more generally $\sum \left(n \ln^1 n \cdots \ln^{k-1} n (\ln^k n)^{\alpha}\right)^{-1}$ converges iff $\alpha > 1$. Abel and Dini: for any divergent positive series $\sum a_n$ with partial sums s_n, the series $\sum a_n / s_n$ still diverges, so there is no slowest divergent series. See Book 1, field: Real analysis, topic: the logarithmic scale.

The generalised recursion. The recursion here is $F(x) = G(x)F(\sigma(x))$ with $\sigma(x) = \ln x$ and $G(x) = x$. In general, if σ is increasing with $\sigma(x) < x$ and the iterates enter a base region B after $k(x)$ steps, then $F(x) = \prod_{j=0}^{k-1} G(\sigma^j(x)) \cdot F_{\text{base}}(\sigma^k(x))$, and the domain splits by escape time. The special case $G = 1/\sigma'$ makes $1/F$ the derivative of the kth iterate of σ, since $\left(H \circ \sigma\right)'(x) = H'(\sigma(x))\sigma'(x)$. Discrete analogue: the recursion $u_x = x\, u_{\lfloor \ln x \rfloor}$ and the $\ln^*$ function of algorithm analysis, which counts exactly the escape time used above. See Book 1, field: Sequences and recurrences, topic: unfolding recursions with contracting arguments.

Circuit Bank

- **WHEN** a function is defined by a recursion whose argument strictly decreases into a base region **DO** unfold it into an explicit finite product and partition the domain by escape time **BECAUSE** an implicit definition cannot be estimated, and the block structure is invisible until the first two levels are written out
- **WHEN** a recursion has the form $\frac{1}{h(x)} = \sigma'(x) \cdot \frac{1}{h(\sigma(x))}$, for instance $h(x) = x\,h(\ln x)$ **DO** declare $1/h$ to be the derivative of an iterate of σ and integrate exactly **BECAUSE** that identity is the chain rule written backwards, so an exact antiderivative exists and no estimation is needed
- **WHEN** you are about to use the integral test **DO** first fix the direction you need, then verify monotonicity on the whole ray including the junction points **BECAUSE** the two integral test inequalities are not symmetric, and only $\sum \geq \int_1^{N+1}$ transfers divergence from the integral to the sum
- **WHEN** a comparison yields an upper bound by a divergent series or a lower bound by a convergent one **DO** discard it and reverse the search **BECAUSE** such a bound is logically empty, and recognising this in five seconds saves the ten minutes most candidates spend on it
- **WHEN** a denominator is a product of a variable and several of its iterated logarithms **DO** say Bertrand series out loud and expect the divergence boundary **BECAUSE** that family is the standard calibration scale for series near the boundary, and its behaviour is a fact you are expected to have installed

Backtesting

1. First useful cue: the recursion $f(x) = xf(\ln x)$ with a base case at e, which is an instruction to unfold, not to test.
2. Method the expression suggested: unfold into $x \ln x \ln\ln x \cdots$, recognise the Bertrand denominator, recognise $1/f$ as the derivative of an iterated logarithm, then integrate block by block.
3. Where you would hesitate, and the thought to reinforce: the moment you have the block formula and want an estimate. Do not estimate. Ask instead whether the block integrand has an exact antiderivative, because a recursion of this shape guarantees one.
4. The trap the question sets: the invitation to compare with $\sum \frac{1}{n \ln n}$. The comparison is true and useless, because $f(n) \geq n \ln n$ makes it an upper bound by a divergent series. The question is engineered so that every fixed comparison series is eventually beaten by f.

5. What makes you recognise it faster next time: the factor in front of the recursive call. It is x, and $x = 1/\ln'(x)$. A recursion whose multiplier is the reciprocal derivative of its argument map is a disguised chain rule, every time.
6. The reusable decision rule: unfold before you test; partition by escape time; and when the recursion is a chain rule, integrate exactly instead of estimating.

Mirror drills

1. Show that $\sum_{n\geq 16} \frac{1}{n \ln n \ln\ln n}$ diverges while $\sum_{n\geq 16} \frac{1}{n \ln n (\ln\ln n)^2}$ converges. *Answer:* the antiderivatives are $\ln^3 x$ and $-1/\ln\ln x$; the first is unbounded, the second is bounded.
2. Let $g(x) = x$ for $x \leq e$ and $g(x) = x(\ln x)g(\ln x)$ for $x > e$. Does $\sum 1/g(n)$ converge? *Answer:* it converges. Unfolding gives $g(x) = x\prod_{j=1}^{k-1}(\ln^j x)^2$ on the kth block, so the substitution $y = \ln^k x$ leaves $\int \frac{dy}{\ln^1 x \cdots \ln^{k-1} x} \leq 1/(t_1 t_2 \cdots t_{k-1})$ on that block, and those bounds are summable.
3. Prove the inequality used in the conclusion: for g positive and nonincreasing on $[1,\infty)$, $\int_1^{N+1} g(x)\,dx \leq \sum_{n=1}^{N} g(n)$. *Answer:* on $[n, n+1]$ we have $g(x) \leq g(n)$, so $\int_n^{n+1} g \leq g(n)$; sum over $n = 1$ to N.

18.5 A5. Regular polygon vertices force high polynomial degree

Let $n \geq 3$ be an integer. Let $f(x)$ and $g(x)$ be polynomials with real coefficients such that the points

$$(f(1), g(1)), (f(2), g(2)), \ldots, (f(n), g(n))$$

in $\mathbb{R}^2$ are the vertices of a regular n-gon in counterclockwise order. Prove that at least one of $f(x)$ and $g(x)$ has degree greater than or equal to $n-1$.

Read it out loud

Read the structure, not the words. You have one parameter, the integer k, running over 1 to n. You have two real polynomials evaluated at that parameter, and the resulting pair of numbers is being interpreted as a point of the plane. The condition on those n points is a rigid geometric one: they form a regular n-gon, and the labelling matches the counterclockwise order around it. The conclusion is about degrees.

Two real polynomials sampled at the same integers, with a planar condition on the pair, is the most reliable cue in the whole subject for one move: package them as a single complex polynomial $P(z) = f(z) + ig(z)$. The reason is not aesthetic. The condition "regular n-gon in counterclockwise order" has a one-line description in $\mathbb{C}$, namely $v_k = a\omega^k + b$ with $\omega = e^{2\pi i/n}$, and no clean description in real coordinates at all: in coordinates you are looking at $f(k) = \alpha + r\cos(\theta_0 + 2\pi k/n)$ and $g(k) = \beta + r\sin(\theta_0 + 2\pi k/n)$, two entangled trigonometric conditions instead of one algebraic one. The degree statement survives the packaging perfectly, because for real f_j, g_j the coefficient $f_j + ig_j$ vanishes if and only if both do. So $\deg P = \max(\deg f, \deg g)$, exactly, and the problem becomes: prove $\deg P \geq n-1$.

The fields embedded are polynomials (degree and root counting), complex numbers (roots of unity and plane similarities), and finite difference calculus (the shift operator acting on sequences). What is really being tested is whether you know the dictionary between polynomials and the difference operator: a polynomial sequence and a genuine geometric sequence are different species, and they can only agree on a bounded window of consecutive integers.

Within thirty seconds you should have written: $P = f + ig$, $\deg P = \max(\deg f, \deg g)$, and $P(k) = a\omega^k + b$ for $k = 1, \ldots, n$, where $\omega = e^{2\pi i/n}$ and $a \neq 0$. The problem is now "a polynomial agrees with a geometric sequence at n consecutive integers", and the number $n-1$ in the conclusion is the loud hint that you should be manufacturing $n-1$ roots of something.

Approaches, ranked worst to best

1. **Work in real coordinates with** cos **and** sin**.** *Cue:* the problem is stated with two real polynomials. *Circuit:* stay in the given coordinates. Last place. You end up needing Chebyshev-style identities, and every identity you derive is the real or imaginary part of the single complex identity you refused to write. Enormous ink, and the counterclockwise hypothesis becomes almost impossible to use cleanly.

2. **Suppose** $\deg \leq n-2$ **and interpolate.** *Cue:* a polynomial of degree at most $n-2$ is determined by $n-1$ of its values. *Circuit:* interpolation counting. The right instinct, since $n-1$ is one fewer than the number of data points, but on its own it is not a proof. You must find the specific linear functional that kills all polynomials of degree at most $n-2$ and does not kill the data. Without the complex packaging you will not see which functional to use.

3. **Induction on the number of sample points.** *Cue:* reduce the case of m points to $m-1$ points by differencing. *Circuit:* induction with the difference operator. Correct: if $Q(k) = t^k$ for $k = 1, \ldots, m$, then $R(z) = (t-1)^{-1}(Q(z+1) - Q(z))$ satisfies $R(k) = t^k$ for $k = 1, \ldots, m-1$ and has degree one less. Solid, but it costs a base case, an induction hypothesis and a check

that the degree really drops by exactly one, for a statement that a single operator identity settles outright.

4. **Finite differences in one shot.** *Cue:* degree at most $m-2$ is equivalent to the $(m-1)$st difference vanishing identically. *Circuit:* annihilate polynomials with Δ. Very good. The $(m-1)$st forward difference of Q at 1 equals $\sum_{i=0}^{m-1}(-1)^{m-1-i}\binom{m-1}{i}t^{1+i} = t(t-1)^{m-1} \neq 0$, which is impossible if $\deg Q \leq m-2$. Loses only to the next route because it needs a binomial evaluation.

5. **Kill the geometric part with the shift operator.** *Cue:* t^k is an eigenvector of the shift with eigenvalue t. *Circuit:* apply $E - tI$, then count roots. Best. Set $R(z) = Q(z+1) - tQ(z)$. Its leading coefficient is $(1-t)$ times that of Q, so it is nonzero and $\deg R = \deg Q$ exactly, while R vanishes at $1, 2, \ldots, m-1$ by construction. A nonzero polynomial with $m-1$ distinct roots has degree at least $m-1$. Three lines, no induction, no binomial sums, and it exposes the reason: $t \neq 1$ is what stops the operator from lowering the degree.

The ideal solution

Lemma 18.5.1. *Let $t \in \mathbb{C}$ with $t \neq 0$ and $t \neq 1$, let $m \geq 1$ be an integer, and let $Q \in \mathbb{C}[z]$ satisfy $Q(k) = t^k$ for $k = 1, 2, \ldots, m$. Then* $\deg Q \geq m-1$.

Proof. First, Q is not the zero polynomial, since $Q(1) = t \neq 0$. Let $d = \deg Q \geq 0$ and let $c \neq 0$ be its leading coefficient. Define

$$R(z) = Q(z+1) - t\,Q(z).$$

The polynomial $Q(z+1)$ has degree d and leading coefficient c, so R has degree at most d, and the coefficient of z^d in R is $c - tc = (1-t)c$, which is nonzero because $t \neq 1$ and $c \neq 0$. Hence R is a nonzero polynomial of degree exactly d.

For $k = 1, 2, \ldots, m-1$ both k and $k+1$ lie in $\{1, \ldots, m\}$, so

$$R(k) = Q(k+1) - t\,Q(k) = t^{k+1} - t \cdot t^k = 0.$$

Thus R has at least the $m-1$ distinct roots $1, 2, \ldots, m-1$. A nonzero polynomial has at most $\deg R$ distinct roots, so $d = \deg R \geq m-1$. □

Proof of the statement. Identify $\mathbb{R}^2$ with $\mathbb{C}$ by $(u, v) \mapsto u + iv$, and set

$$P(z) = f(z) + i\,g(z) \in \mathbb{C}[z].$$

Writing $f(z) = \sum_j f_j z^j$ and $g(z) = \sum_j g_j z^j$ with $f_j, g_j \in \mathbb{R}$, the coefficient of z^j in P is $f_j + ig_j$, which vanishes if and only if $f_j = g_j = 0$. Hence

$$\deg P = \max(\deg f, \deg g),$$

and it suffices to prove $\deg P \geq n-1$.

Let $v_k = P(k) = f(k) + ig(k)$ for $k = 1, \ldots, n$. By hypothesis $v_1, \ldots, v_n$ are the vertices of a regular n-gon listed in counterclockwise order. Let b be its centre and $r > 0$ its circumradius. The vertices of a regular n-gon lie on the circle of radius r about b at arguments forming an arithmetic progression of common difference $2\pi/n$, and listing them counterclockwise starting from v_1 means exactly that

$$v_k = b + re^{i(\theta_0 + 2\pi k/n)} = b + a\omega^k, \qquad a = re^{i\theta_0} \neq 0, \quad \omega = e^{2\pi i/n},$$

for $k = 1, \ldots, n$. Now set

$$Q(z) = \frac{P(z) - b}{a},$$

which is a polynomial with $\deg Q = \deg P$, since $a \neq 0$ and subtracting a constant cannot change the degree of a nonconstant polynomial and cannot lower a constant one below 0. By construction $Q(k) = \omega^k$ for $k = 1, \ldots, n$.

Since $n \geq 3$ we have $\omega \neq 1$, and $|\omega| = 1$ gives $\omega \neq 0$. Applying the Lemma with $t = \omega$ and $m = n$ yields $\deg Q \geq n - 1$. Therefore

$$\max(\deg f, \deg g) = \deg P = \deg Q \geq n - 1,$$

so at least one of f and g has degree at least $n - 1$. ■

Remark 18.5.2. The bound is sharp, and the same machinery says so. The unique polynomial Q of degree at most $n - 1$ with $Q(k) = \omega^k$ for $k = 1, \ldots, n$ has leading coefficient $\Delta^{n-1}Q(1)/(n-1)! = \omega(\omega - 1)^{n-1}/(n-1)!$, which is nonzero, so $\deg Q = n - 1$ exactly; its real and imaginary parts give f and g of degree at most $n - 1$ that do trace out a regular n-gon. So $n - 1$ is attained and cannot be improved.

Remark 18.5.3. Counterclockwise is used only to force the labels to run consecutively around the polygon, giving ω^k rather than $\omega^{\sigma(k)}$ for some permutation σ. If the order were clockwise, the same proof runs with $t = \omega^{-1}$, which is also neither 0 nor 1. If the labelling were an arbitrary permutation of the vertices the statement would be false in this form, which is precisely why that hypothesis is in the question.

Layer 1: WHAT

Block 1. Package the two real polynomials as one complex polynomial $P = f + ig$, and record that $\deg P = \max(\deg f, \deg g)$ exactly.

Block 2. Write the regular n-gon in counterclockwise order as $v_k = a\omega^k + b$, and normalise it away: $Q = (P - b)/a$ has the same degree and satisfies $Q(k) = \omega^k$.

Block 3. Apply the operator $Q \mapsto Q(z + 1) - \omega Q(z)$. Because $\omega \neq 1$ the leading coefficient survives, so the degree is unchanged.

Block 4. The new polynomial vanishes at $1, \ldots, n - 1$ by construction. Nonzero polynomial, $n - 1$ distinct roots, so degree at least $n - 1$. Push that back through Blocks 3, 2 and 1.

Juxtaposition. The nominal route keeps everything real: it writes $f(k) = \alpha + r\cos(\theta_0 + 2\pi k/n)$, tries to interpolate, and reaches for trigonometric identities to force a degree bound. It typically dies at the point where "regular n-gon" has to be turned into equations, because in coordinates that condition is a system, not a formula.

The bifurcation point is the sentence "the points are the vertices of a regular n-gon". Ideal: write $a\omega^k + b$. Zombie: write cos and sin. Everything after that is determined by which of those two lines you put on the page.

Layer 2: HOW

Block 1. Cue. Two real polynomials, sampled at the same integers, with a condition on the resulting planar points. **Why natural.** A pair of real numbers with geometric meaning is a complex number, and the plane's similarity structure is invisible in coordinates but algebraic in $\mathbb{C}$. The move costs nothing, because the degree bookkeeping is exact: $f_j + ig_j = 0$ with f_j, g_j real forces $f_j = g_j = 0$, so no cancellation between f and g can occur and $\deg P = \max(\deg f, \deg g)$. Note that the conclusion you want, "at least one of f, g has degree at least $n - 1$", is literally the statement $\max(\deg f, \deg g) \geq n - 1$, so the packaging is not just convenient, it is the exact shape of the target. **Link to the next block.** Now the hypothesis has to be written in $\mathbb{C}$, which is where the regular n-gon becomes a formula.

Block 2. **Cue.** "Regular n-gon", "counterclockwise order", and the index k running over consecutive integers. **Why natural.** A regular n-gon is by definition the orbit of one vertex under rotation by $2\pi/n$ about the centre. In $\mathbb{C}$, rotation by $2\pi/n$ about b is $z \mapsto b + \omega(z - b)$, so the orbit is $b + a\omega^k$. Counterclockwise order is exactly the statement that consecutive labels are consecutive under that rotation, so k appears in the exponent unpermuted. The normalisation $Q = (P - b)/a$ is then forced: the maps $z \mapsto az + b$ with $a \neq 0$ are precisely the orientation preserving similarities of the plane, they form a group, and none of them changes the degree of a polynomial. Normalising by a group that preserves the quantity you are measuring is always free. **Link to the next block.** The hypothesis is now the cleanest possible statement: a polynomial takes geometric values at consecutive integers.

Block 3. **Cue.** $Q(k) = \omega^k$ is a geometric sequence in k, and geometric sequences are eigenvectors of the shift. **Why natural.** On sequences, let E be the shift $(Eu)_k = u_{k+1}$. Then $Et^k = t \cdot t^k$, so $(E - tI)$ annihilates the geometric sequence exactly. On polynomials, E acts as $Q(z) \mapsto Q(z + 1)$, which preserves both degree and leading coefficient, so $E - tI$ multiplies the leading coefficient by $(1 - t)$. The single hypothesis $t \neq 1$ is what keeps that factor alive. This is the whole reason the argument works. Note also what the hypothesis $n \geq 3$ is and is not doing: $\omega \neq 1$ already holds for $n \geq 2$, so $n \geq 3$ is present only so that the word polygon means something. **Link to the next block.** You have a polynomial of the same degree as Q which vanishes wherever the geometric relation held.

Block 4. **Cue.** The target $n - 1$ is one less than the number of data points, and you now have a polynomial forced to vanish at $n - 1$ points. **Why natural.** The only elementary mechanism that bounds a degree from below is root counting: a nonzero polynomial of degree d has at most d distinct roots. So to prove "degree at least N" you manufacture a nonzero polynomial of that degree with N distinct roots. The differencing did the manufacturing: each of the $n - 1$ overlapping consecutive pairs $(k, k + 1)$ inside $\{1, \ldots, n\}$ produces one root. Counting the pairs is why the answer is $n - 1$ and not n. **Link to the next block.** Degrees were preserved at every stage, so the bound transfers straight back to $\max(\deg f, \deg g)$.

Layer 3: WHY

The mathematics underneath is the theory of linear recurrences, viewed as linear algebra on the space of sequences. Let E be the shift operator. Then:

1. The geometric sequence $u_k = t^k$ spans the eigenspace $\ker(E - tI)$.
2. The sequences $u_k = Q(k)$ with $\deg Q \leq d$ are exactly $\ker(E - I)^{d+1}$, that is, the generalised eigenspace at eigenvalue 1 of nilpotency degree $d + 1$. This is the operator form of the familiar fact that the $(d + 1)$st finite difference of a degree d polynomial vanishes.

So a polynomial sequence and a geometric sequence with ratio $t \neq 1$ live in generalised eigenspaces for different eigenvalues. They are transverse. They cannot be equal as sequences at all, and the quantitative form of that transversality is exactly this problem: they can agree on at most $d + 1$ consecutive terms, where d is the degree. Read the Lemma again with that in mind and it stops being an argument and becomes an accounting identity.

The discernment mechanism to install is this. When a problem gives you the values of a polynomial at consecutive integers, stop thinking of the polynomial as a function on $\mathbb{R}$ and start thinking of the sequence $u_k = P(k)$ together with the operators that act on it. Continuous tools (derivatives, calculus, curve sketching) are the wrong toolkit for data on a lattice; the shift and the difference are the right ones. The practical rule: whenever the data are of the form (polynomial) = (something with a multiplicative structure in k), hit both sides with $E - tI$ where t is the multiplier, and see what dies.

How to train it: memorise the dictionary in both directions. "Degree at most d" is "$\Delta^{d+1} \equiv 0$". "Geometric with ratio t" is "$(E - tI) \equiv 0$". "Polynomial of degree d times t^k" is "$(E - tI)^{d+1} \equiv 0$".

Then any hypothesis mixing polynomials and exponentials on the integers becomes a statement about which operator to apply first.

What the examiners nudged toward, and how: they chose a regular n-gon, which is the only common plane figure whose vertex set is literally a geometric sequence in $\mathbb{C}$. They said "counter-clockwise order" so that the exponent is k and not a permutation of k. They asked for $n-1$, one fewer than the number of points, which is the signature of a single differencing step, and they wrote the conclusion as "at least one of f and g", which is $\max(\deg f, \deg g)$, the exact invariant preserved by the complex packaging. Every hypothesis in the statement is a guiding light toward $P = f + ig$.

Revolving concepts

Complex packaging of a pair of real polynomials. If $f, g \in \mathbb{R}[z]$ then $P = f + ig \in \mathbb{C}[z]$ has $\deg P = \max(\deg f, \deg g)$, and $f = \frac{1}{2}(P + \bar{P})$, $g = \frac{1}{2i}(P - \bar{P})$ where $\bar{P}$ is the polynomial with conjugated coefficients. The packaging is a bijection between pairs of real polynomials and complex polynomials, and it preserves degree. See Book 1, field: Complex numbers, topic: complex form of planar conditions.

Roots of unity and regular polygons. $\omega = e^{2\pi i/n}$ satisfies $\omega^n = 1$ and $1 + \omega + \cdots + \omega^{n-1} = 0$ for $n \geq 2$; the vertices of the regular n-gon with centre b are $b + a\omega^k$, and the centroid of the vertices is b. The orientation preserving similarities of the plane are exactly $z \mapsto az + b$ with $a \neq 0$; the orientation reversing ones are $z \mapsto a\bar{z} + b$. See Book 1, field: Complex numbers, topic: roots of unity and plane transformations.

Finite difference calculus. $\Delta Q(z) = Q(z+1) - Q(z)$ lowers degree by exactly one and multiplies the leading coefficient by d; hence $\Delta^d Q$ is the constant $d!\,c$ for $\deg Q = d$ with leading coefficient c, and $\Delta^{d+1} Q \equiv 0$. The explicit expansion is $\Delta^p Q(x) = \sum_{i=0}^{p} (-1)^{p-i} \binom{p}{i} Q(x+i)$, so the leading coefficient of the interpolant through $p+1$ consecutive values is $\Delta^p Q(x_0)/p!$. See Book 1, field: Polynomials, topic: finite differences and Newton interpolation.

The generalised recursion. A sequence (u_k) satisfies a linear recurrence with constant coefficients and characteristic polynomial $\prod_i (x - \lambda_i)^{m_i}$ if and only if $u_k = \sum_i p_i(k) \lambda_i^k$ with $\deg p_i < m_i$. In operator form the solution space is $\ker \prod_i (E - \lambda_i I)^{m_i}$. Two consequences used here: polynomials of degree at most d are the case of the single root $\lambda = 1$ with multiplicity $d+1$, and geometric sequences are the case of a simple root $\lambda = t$. Since the generalised eigenspaces for distinct λ intersect trivially, a nonzero polynomial sequence is never geometric, and the finite version of that statement is the Lemma above. The same operator factorisation is the standard method for solving any linear recurrence, including the repeated-root case that A-level Further Maths handles only for order two. See Book 1, field: Sequences and recurrences, topic: linear recurrences and the shift operator.

Root counting. A nonzero polynomial over a field has at most deg distinct roots; equivalently, a polynomial of degree at most $N-1$ vanishing at N distinct points is identically zero. This is the only elementary source of lower bounds on degree. See Book 1, field: Polynomials, topic: factor theorem and root bounds.

Circuit Bank

- **WHEN** a pair of real polynomials is sampled at the same points and the pairs are read as points of the plane **DO** set $P = f + ig$ and record $\deg P = \max(\deg f, \deg g)$ **BECAUSE** planar similarity and rotation conditions are one-line algebra in $\mathbb{C}$ and a coupled system in coordinates, and the degree bookkeeping across the packaging is exact
- **WHEN** a regular n-gon appears with its vertices labelled in cyclic order **DO** write the vertices as $a\omega^k + b$ with $\omega = e^{2\pi i/n}$, then normalise to ω^k **BECAUSE** a regular polygon is by definition a rotation orbit, and $z \mapsto az + b$ with $a \neq 0$ is the full orientation preserving similarity group, which changes no degree

- **WHEN** a polynomial agrees with a geometric sequence t^k at consecutive integers **DO** apply $E - tI$, that is form $P(z+1) - tP(z)$ **BECAUSE** the operator annihilates the geometric part while preserving degree exactly when $t \neq 1$, so it converts data into roots for free
- **WHEN** you must bound a degree from below **DO** manufacture a nonzero polynomial of that same degree with many distinct roots **BECAUSE** root counting is the only elementary lower bound mechanism for degree, so every such proof is secretly a root count
- **WHEN** a hypothesis fixes an ordering or an orientation and you have not used it **DO** find the step where it removes a permutation or a sign ambiguity, and say so **BECAUSE** unused hypotheses on the Putnam are almost always load bearing, and here counterclockwise is what makes the exponent k rather than $\sigma(k)$

Backtesting

1. First useful cue: two real polynomials sampled at 1 to n with a planar condition on the pairs, which is the standing order to form $f + ig$.
2. Method the expression suggested: normalise the polygon to the roots of unity by an affine complex map, then annihilate the geometric sequence with the shift operator and count the roots that appear.
3. Where you would hesitate, and the thought to reinforce: at the claim $\deg P = \max(\deg f, \deg g)$. Reinforce it with the reason, that $f_j + ig_j = 0$ with both real forces both to vanish, so real and imaginary parts cannot cancel each other's leading terms. Without that line the packaging is not justified.
4. The trap the question sets: the invitation to expand in cos and sin, and the softer trap of forgetting to check $\deg R = \deg Q$ exactly. If you only claim $\deg R \leq \deg Q$, the root count gives nothing.
5. What makes you recognise it faster next time: the number $n - 1$ sitting one below the number of data points. That gap of exactly one is the signature of a single differencing step, just as a gap of two would signal a second difference.
6. The reusable decision rule: polynomial data on consecutive integers means work with the shift operator, not with calculus; and to bound a degree below, build roots.

Mirror drills

1. Show that no polynomial P of degree at most 8 satisfies $P(k) = 2^k$ for $k = 1, \ldots, 10$. *Answer:* $R(z) = P(z+1) - 2P(z)$ has the same degree as P, since the leading coefficient is multiplied by $1 - 2 = -1$, and vanishes at $1, \ldots, 9$, so $\deg P \geq 9$.
2. Compute the leading coefficient of the unique polynomial Q of degree at most $n - 1$ with $Q(k) = \omega^k$ for $k = 1, \ldots, n$, and deduce that the bound $n - 1$ in the problem is attained. *Answer:* $\Delta^{n-1}Q(1) = \sum_{i=0}^{n-1}(-1)^{n-1-i}\binom{n-1}{i}\omega^{1+i} = \omega(\omega-1)^{n-1} \neq 0$, so the leading coefficient is $\omega(\omega-1)^{n-1}/(n-1)!$ and the degree is exactly $n - 1$.
3. Suppose instead the points $(f(k), g(k))$ are the vertices of a regular n-gon listed in clockwise order. Does the conclusion change, and where exactly does the proof notice the difference? *Answer:* no change. You get $t = \omega^{-1}$, and the proof used only $t \neq 0$ and $t \neq 1$; the orientation enters solely in the identification of t.

18.6 A6. Producing every group element from a short sequence

Prove that there exists a constant $c > 0$ such that in every nontrivial finite group G there exists a sequence of length at most $c \log |G|$ with the property that each element of G equals the product of some subsequence. (The elements of G in the sequence are not required to be distinct. A subsequence *of a sequence is obtained by selecting some of the terms, not necessarily consecutive, without reordering them; for example, $4, 4, 2$ is a subsequence of $2, 4, 6, 4, 2$, but $2, 2, 4$ is not.)*

Read it out loud

Before any group theory, state the combinatorial shape. You are given a target set of size $n = |G|$. You must build an ordered list of length ℓ, and the objects you are allowed to hit the target with are the products of its subsequences, taken in the order they appear in the list. A list of length ℓ has exactly 2^ℓ subsequences, so it produces at most 2^ℓ elements, and therefore $\ell \geq \log_2 n$ is forced, always, for every group and every list. The problem asks for $\ell \leq c \log n$.

Stop and register what that means, because it is the single most important reading of this problem. You are being asked to come within a constant factor of an information theoretic bound that you can prove in one line. That is a very strong hint about the shape of the construction: if the counting bound is tight up to a constant, then a construction must gain a constant factor of the target per term. A constant factor per term means *doubling*. So before you know a single fact about G, you know the plan: find a way to append one element to a list and double the set of producible elements.

The fields wired in are finite group theory (cosets and the fact that right multiplication is a bijection), extremal and greedy combinatorics (grow an object until an obstruction appears), and a light touch of counting. The word "nontrivial" is not decoration: for the trivial group $\log |G| = 0$ and no positive length is allowed, so the hypothesis is there to make the statement true rather than to make the proof easier.

The first cue is the phrase "product of some subsequence". Interrogate the operation. If S is a list with producible set H, and you append one new element g at the *end*, then a subsequence of the new list either omits g, giving a product in H, or uses g, and then g is its last factor, so the product is hg for some $h \in H$. Hence the new producible set is exactly $H \cup Hg$. Within thirty seconds you should have on your page: $2^\ell \geq n$, the update rule $H \mapsto H \cup Hg$, and the observation that $|Hg| = |H|$ always, so doubling happens precisely when $H \cap Hg = \varnothing$.

Approaches, ranked worst to best

1. **Take a small generating set and multiply.** *Cue:* the elements of G are products of generators. *Circuit:* generation. Last. Every finite group has a generating set of size at most $\log_2 n$, so the size is right, but generation is the wrong notion here. A subsequence uses each term at most once and in the given order, so a list of k distinct generators produces at most 2^k elements, and there is no way to write a long word in the generators. To recover words you would have to repeat each generator many times, and the length blows up to something like n. Recognising within ten seconds that generation is not the operative notion is the whole value of this approach.

2. **Induct on a chain of subgroups with cyclic quotients.** *Cue:* finite groups are built from smaller ones. *Circuit:* induction on the index. If $H \trianglelefteq G$ with G/H cyclic of order p generated by the image of x, take a good list for H and append $x, x^2, x^4, \ldots, x^{2^s}$ with $2^{s+1} > p$. Subsequence products of the appended block give x^j for every $0 \leq j < p$, so the whole list produces Hx^j for all j, that is all of G, at a cost of $\lceil \log_2 p \rceil$ extra terms. Clean, and it even gives c close to 1. It fails in general: chief factors need not be cyclic, and a nonabelian simple group has no proper

normal subgroup at all. So this route settles the solvable case and stalls. Know it, do not bet the paper on it.

3. **Average over g and pigeonhole.** *Cue:* you want a good g, so compute the average of $|H \cup Hg|$. *Circuit:* double counting plus pigeonhole. Correct and quantitatively finer. Double counting gives $\sum_{g \in G} |H \cup Hg| = 2mn - m^2$ where $m = |H|$, so some g achieves $|H \cup Hg| \geq 2m - m^2/n$. In terms of the deficiency $s = 1 - m/n$ this is the recursion $s \mapsto s^2$, which reaches $s < 1/n$ after about $2 \log_2 n$ steps and gives length $\log_2 n + \log_2 \log n$, essentially optimal. It ranks below the next route only because it costs a double counting identity plus a page of inequality management, when exact doubling is available for free.

4. **Greedy exact doubling, then mirror the list.** *Cue:* doubling is exactly available while $H^{-1}H \neq G$. *Circuit:* grow until an obstruction appears, then exploit the obstruction. Best. Append any $g \notin H^{-1}H$ and the producible set doubles exactly. The greedy can only run $\log_2 n$ times, so it stops, and it stops precisely when $H^{-1}H = G$, which is exactly the hypothesis needed to finish. No inequalities, no averaging, no constants to chase.

The ideal solution

Throughout, G is a finite group with identity e and $n = |G| \geq 2$. For a finite sequence $S = (a_1, \ldots, a_k)$ of elements of G, say that S *produces* h if

$$h = a_{i_1} a_{i_2} \cdots a_{i_r} \quad \text{for some } 1 \leq i_1 < i_2 < \cdots < i_r \leq k,$$

and write $H(S)$ for the set of elements produced by S. We allow $r = 0$ inside the construction, with the empty product equal to e; the final sequence will produce every element by a nonempty subsequence, so this convention costs nothing and we check that explicitly at the end.

Lemma 18.6.1 (Update rule)**.** *If $S' = (a_1, \ldots, a_k, g)$ is S with g appended at the end, then $H(S') = H(S) \cup H(S)g$.*

Proof. A subsequence of S' either omits the last term, in which case it is a subsequence of S and its product lies in $H(S)$, or it includes the last term, in which case g is its final factor and its product is hg with $h \in H(S)$. Both types occur, so the produced set is exactly $H(S) \cup H(S)g$. □

Lemma 18.6.2 (Doubling criterion)**.** *Let $H \subseteq G$ be nonempty and $g \in G$. Then $H \cap Hg = \varnothing$ if and only if $g \notin H^{-1}H$, where $H^{-1}H = \{h_1^{-1}h_2 : h_1, h_2 \in H\}$. In that case $|H \cup Hg| = 2|H|$.*

Proof. $H \cap Hg \neq \varnothing$ means $h_1 g = h_2$ for some $h_1, h_2 \in H$, that is $g = h_1^{-1}h_2 \in H^{-1}H$; the two statements are contrapositives of each other. Right multiplication by g is a bijection of G, so $|Hg| = |H|$, and disjointness gives $|H \cup Hg| = 2|H|$. □

The greedy construction. Start with S_0 the empty sequence, so $H_0 = H(S_0) = \{e\}$ and $|H_0| = 1$. Given S_j with $H_j = H(S_j)$, if $H_j^{-1}H_j \neq G$ choose any $a_{j+1} \in G \setminus H_j^{-1}H_j$ and let $S_{j+1} = (a_1, \ldots, a_{j+1})$. By the two Lemmas, $H_{j+1} = H_j \cup H_j a_{j+1}$ and $|H_{j+1}| = 2|H_j|$.

By induction $|H_j| = 2^j$. Since $H_j \subseteq G$ we have $2^j \leq n$, so the process cannot run more than $\log_2 n$ times. Therefore it halts: there is an index $k \leq \log_2 n$ with

$$S = (a_1, \ldots, a_k), \qquad H = H(S), \qquad |H| = 2^k, \qquad H^{-1}H = G.$$

Note $k \geq 1$, because $H_0^{-1}H_0 = \{e\} \neq G$ as $n \geq 2$.

The mirrored sequence. Define the sequence of length $2k$

$$T = \left(a_k^{-1},\, a_{k-1}^{-1},\, \ldots,\, a_1^{-1},\, a_1,\, a_2,\, \ldots,\, a_k\right).$$

Claim 18.6.3. *T produces every element of G, and by a nonempty subsequence.*

Proof. Let $g \in G$. Since $H^{-1}H = G$ there are $h_1, h_2 \in H$ with $g = h_1^{-1}h_2$. Write

$$h_1 = a_{i_1}a_{i_2}\cdots a_{i_r}, \qquad i_1 < i_2 < \cdots < i_r,$$

which is possible since $h_1 \in H = H(S)$, and similarly $h_2 = a_{j_1}\cdots a_{j_s}$ with $j_1 < \cdots < j_s$. Inverting reverses the order of the factors:

$$h_1^{-1} = a_{i_r}^{-1}a_{i_{r-1}}^{-1}\cdots a_{i_1}^{-1}.$$

The first half of T lists $a_k^{-1}, a_{k-1}^{-1}, \ldots, a_1^{-1}$, that is, the inverses in *decreasing* index order. The factors of h_1^{-1} appear in exactly that order, since $i_r > i_{r-1} > \cdots > i_1$. So h_1^{-1} is the product of a subsequence of the first half of T, read left to right. Likewise h_2 is the product of a subsequence of the second half of T, read left to right, since $j_1 < \cdots < j_s$. Concatenating a subsequence of the first half with a subsequence of the second half is a subsequence of T, and its product is $h_1^{-1}h_2 = g$.

For the nonemptiness: if $g \neq e$ then $h_1 \neq h_2$, so at least one of h_1, h_2 differs from e and the corresponding selection is nonempty. If $g = e$, take instead $h_1 = h_2 = a_1$, which is legitimate since $a_1 \in H$ and $k \geq 1$; then $g = a_1^{-1}a_1$ is the product of the two term subsequence (a_1^{-1}, a_1) of T, the first term coming from the first half and the second from the second half. □

The length. T has length $2k \leq 2\log_2 n = \dfrac{2}{\ln 2}\ln n$. Hence

$$\boxed{c = \frac{2}{\ln 2} < 2.886 \text{ works for the natural logarithm, and } c = 2 \text{ for } \log_2.}$$

Since $n \geq 2$ we have $\log n > 0$, so the bound is a genuine positive length constraint, and the proof is complete. ■

Remark 18.6.4. The construction is within a factor 2 of optimal, and the counting bound of the opening paragraph proves it: a sequence of length ℓ has 2^ℓ subsequences, so producing all n elements forces $\ell \geq \log_2 n$. For $G = (\mathbb{Z}/2\mathbb{Z})^d$ the list of d standard basis vectors attains $\ell = \log_2 |G|$ exactly, so no $c < 1$ can work.

Layer 1: WHAT

Block 1. Count first: $2^\ell \geq n$, so $\ell \geq \log_2 n$ always. Target is $c \log n$, so the construction must gain a constant factor per term, that is, it must double.

Block 2. Find the update rule. Appending g at the end sends the producible set H to $H \cup Hg$, exactly.

Block 3. Find the exact condition for doubling: $|Hg| = |H|$ always, so doubling happens if and only if $H \cap Hg = \varnothing$, if and only if $g \notin H^{-1}H$.

Block 4. Run the greedy. It doubles $|H|$ each step and $|H| \leq n$, so it stops after at most $\log_2 n$ steps.

Block 5. Read the stopping condition, which is $H^{-1}H = G$, and recognise it as a gift rather than a failure: every g is now $h_1^{-1}h_2$.

Block 6. Realise $h_1^{-1}h_2$ as an ordered subsequence product by mirroring: reversed inverses, then the originals. Length $2k \leq 2\log_2 n$.

Juxtaposition. The nominal route asks "which elements should I put in the list?" and starts hunting for clever choices, structure theorems, or special cases (abelian first, then p-groups, then hope). It treats the problem as a search. The ideal route asks "what does appending one element do to the reachable set?" and treats the problem as an operator: it computes the update rule first and only then asks which g makes the update efficient. The choice of g then requires no cleverness at all, because any g outside $H^{-1}H$ works and the only question is whether such a g exists.

The bifurcation point is the very first move after reading the definition of subsequence product: compute $H(S$ with g appended$)$, or start choosing elements. Everything follows from that.

Layer 2: HOW

Block 1. **Cue.** A bound of the form $c \log |G|$, with c unspecified and log of unspecified base. Unspecified constant and base means the problem only cares about the order of magnitude. **Why natural.** Whenever the target is a logarithm of the size of the object, compute the trivial counting bound immediately, because a logarithm is exactly what a counting bound produces. Here a list of length ℓ has 2^ℓ subsequences, so $2^\ell \geq n$. The fact that the requested bound matches this lower bound up to a constant is the design of the problem speaking to you: it tells you no waste is permitted, so the construction must be exponentially efficient, so each term must multiply the reachable set by a constant factor. **Link to the next block.** Knowing you need doubling, the next question is forced: what exactly does one term buy?

Block 2. **Cue.** The definition of subsequence is spelled out at unusual length in the statement, with an explicit example distinguishing $4, 4, 2$ from $2, 2, 4$. When a problem defines its operation that carefully, the operation is the subject. **Why natural.** Order matters and each term is used at most once, so the only structural handle is the last term. Splitting the subsequences of S' by whether they use the last term is the unique decomposition that respects the order constraint, and it gives $H \cup Hg$ with no error terms. Getting an exact equality rather than an inclusion here is what makes the rest of the proof clean. **Link to the next block.** Now the question is purely about the size of $H \cup Hg$, which is a question about H and its translates.

Block 3. **Cue.** A union of a set and a translate, inside a group. **Why natural.** In a group, right multiplication by g is a bijection, so $|Hg| = |H|$ exactly; there is no loss. That single fact is why doubling is even conceivable, and it is the reason the problem is set in a group rather than in a monoid. Therefore $|H \cup Hg| = 2|H|$ if and only if H and Hg are disjoint, and unwinding the definition of disjointness gives the clean criterion $g \notin H^{-1}H$. Notice that you did not go looking for the set $H^{-1}H$; it was produced by asking when the update is efficient. This is the standard way the right object appears in a proof: not by inspiration, but by solving for the condition you need. **Link to the next block.** You have a criterion, so run it greedily until it fails.

Block 4. **Cue.** An operation that multiplies a bounded quantity by a fixed factor greater than 1. **Why natural.** Such a process must terminate, and the number of steps is the logarithm of the cap to the base of the factor. Here $|H_j| = 2^j \leq n$ gives $j \leq \log_2 n$ with no further work. Note the direction of the argument: you are not proving the greedy achieves something, you are proving it *cannot continue*, and it is the failure that you will harvest. **Link to the next block.** Termination means the criterion failed, so write down what failure says.

Block 5. **Cue.** The negation of "there exists $g \notin H^{-1}H$" is "$H^{-1}H = G$". **Why natural.** This is the quantum leap of the problem, and it is not a leap at all once you write both statements down side by side. The greedy was designed to fail, because its failure condition is a strong structural statement: every element of G factors as $h_1^{-1}h_2$ with both factors producible. In other words S does not produce all of G, but S together with the inverses of its products does. A greedy whose stopping condition is exactly the hypothesis you need is the standard architecture of extremal arguments; train yourself to ask, before running any greedy, what its failure would give you. **Link to the next block.** You need $h_1^{-1}h_2$ to be an ordered subsequence product, so build a list in which it is.

Block 6. **Cue.** Inversion reverses order: $(a_{i_1} \cdots a_{i_r})^{-1} = a_{i_r}^{-1} \cdots a_{i_1}^{-1}$. **Why natural.** The obstruction is the ordering rule, and the reversal identity is exactly the tool that matches it. If you lay the inverses down in decreasing index order, then any product h_1^{-1} reads off left to right in the order the reversal identity demands. Put the original list after it and h_2 reads off in its own natural

order. The concatenation of a subsequence of the first half with a subsequence of the second half is automatically a valid subsequence of the whole, since every position in the first half precedes every position in the second. The price is a factor of 2 in the length, which the problem's unspecified c has already told you is affordable. **Link to the next block.** Count: length $2k \leq 2\log_2 n$, and the constant is exhibited.

Layer 3: WHY

Three mathematical truths carry this proof, and each is worth extracting on its own.

The first is that in a group, translation is a bijection. $|Hg| = |H|$ for every subset H and every g: this is the reason a group is the right setting and the reason exact doubling is possible. In a monoid or a semigroup, Hg can collapse and the entire argument dies. Whenever you see a counting argument in a group, look for the place where translation invariance of cardinality is being used; it is almost always the load bearing step, and it is the same fact behind Lagrange's theorem, behind the orbit stabiliser theorem, and behind the pigeonhole step in the alternative averaging solution.

The second is the dichotomy attached to the difference set $H^{-1}H$. Either $H^{-1}H \neq G$, in which case H can be grown by a full factor of 2, or $H^{-1}H = G$, in which case H already covers G in the two sided sense. There is no middle case. This is the finite group analogue of the sumset dichotomy in additive combinatorics: a set either has small doubling, in which case it is structured, or it grows. Modern results on growth in groups, associated with Helfgott, Bourgain and Tao among others, are built on exactly this dichotomy in quantitative form, and the alternative averaging solution is a baby version of the same technology.

The third is the counting bound $2^\ell \geq n$. It is not a side remark; it is the design document of the problem. When a construction is requested at the scale of the counting bound, the construction is essentially forced to be a doubling process, because anything slower cannot reach n in $O(\log n)$ steps. Reading the target bound backwards to infer the mechanism is the single most transferable habit in this problem.

The discernment mechanism to train is the greedy design protocol. (1) Define the object you are growing and the quantity you are measuring. (2) Compute the exact update rule for one step, as an equality if possible. (3) Solve for the condition under which the update is maximally efficient. (4) Run the greedy while that condition holds, and bound the number of steps by the growth factor. (5) Write down what the failure of the condition says, and expect it to be the second half of the proof. Steps (3) and (5) are where the mathematics happens, and both are mechanical once you commit to writing exact conditions rather than searching for good choices.

What the examiners nudged toward: they defined subsequence at length, with an example, which tells you the ordering constraint is a real obstruction and must be addressed explicitly; that is the mirror. They left c and the base of the logarithm unspecified, which tells you constant factors are free; that is the licence to double the length. They said nontrivial, which is exactly the case $n = 1$ where $\log |G| = 0$ would make the statement false. Every clause is a guiding light.

Knowledge gaps to install before this problem is approachable

Be honest about this one, because it is an A6 and the entry cost is real. Nothing here is beyond A-level in difficulty, but the following are beyond A-level in content, and without them the problem is not merely hard, it is unreadable.

1. **Groups, subgroups, cosets.** A group is a set with an associative operation, an identity, and inverses. For a subset H and an element g, the right translate is $Hg = \{hg : h \in H\}$. Translation is a bijection of G, so $|Hg| = |H|$. If H is a subgroup, the translates are the right cosets, they partition G, and Lagrange's theorem follows: $|H|$ divides $|G|$. This problem uses the translate fact for arbitrary subsets, not just subgroups, which is a distinction worth noticing.

2. **Inversion reverses products.** $(xy)^{-1} = y^{-1}x^{-1}$, extended to $(x_1 \cdots x_r)^{-1} = x_r^{-1} \cdots x_1^{-1}$. This is the entire content of the mirror step.
3. **Product and difference sets.** For subsets $A, B \subseteq G$, $AB = \{ab\}$ and $A^{-1} = \{a^{-1}\}$. The set $H^{-1}H$ is the group analogue of a difference set. No structure theorem is needed, only the definitions.
4. **Normal subgroups, quotients and solvability**, if you want to understand why the induction route in the rankings fails in general. Not needed for the ideal solution.

Revolving concepts

Translation invariance and Lagrange. For any $H \subseteq G$ and $g \in G$, $|Hg| = |gH| = |H|$. For a subgroup H, cosets partition G and $|G| = [G : H]\,|H|$. See Book 1, field: Abstract algebra, topic: cosets and Lagrange's theorem.

The producible set and its update rule. For a sequence S over G, $H(S)$ is the set of ordered subsequence products. Appending g gives $H(S) \cup H(S)g$; prepending g gives $H(S) \cup gH(S)$. Both are exact equalities, and choosing which end to append to is a genuine design decision in problems of this type. See Book 1, field: Combinatorics, topic: subset sums and subsequence products.

Subset sums, the abelian shadow. In $(\mathbb{Z}/2\mathbb{Z})^d$ the same question is "find a short list of vectors whose subset sums cover the space", answered by any basis, with length exactly $\log_2 |G|$. In $\mathbb{Z}/n\mathbb{Z}$ the answer is the binary list $1, 2, 4, \ldots$, of length $\lceil \log_2 n \rceil$. Both are the abelian intuition that the general construction imitates. See Book 1, field: Combinatorics, topic: binary representation and covering by subset sums.

The generalised recursion. Two recursions appear, and both are worth stating in general form. The first is $m_{j+1} = 2m_j$ with cap n, giving $m_j = 2^j$ and stopping time $\log_2 n$; in general a greedy that multiplies a capped quantity by a fixed $\lambda > 1$ per step runs for at most $\log_\lambda(\text{cap}/m_0)$ steps, and this is the standard way to bound any doubling process. The second appears in the averaging solution: the deficiency $s_j = 1 - m_j/n$ obeys $s_{j+1} \leq s_j^2$, whose solution is $s_j \leq s_0^{2^j}$, doubly exponential decay, so $j \approx \log_2 \log_{1/s_0}(1/\varepsilon)$ steps suffice; this is the same recursion that governs Newton's method and repeated squaring, and recognising it saves you from trying to solve it by hand. See Book 1, field: Sequences and recurrences, topic: multiplicative and quadratic recursions.

Growth in groups. The dichotomy "either $H^{-1}H \neq G$ and H grows, or $H^{-1}H = G$" is the elementary ancestor of small doubling and small tripling results. See Book 1, field: Combinatorics, topic: growth and product sets in groups.

Circuit Bank

- **WHEN** a bound of the form $c \log |X|$ is requested with c and the base unspecified **DO** compute the counting lower bound first and compare **BECAUSE** if the target matches the counting bound up to a constant, the construction is forced to gain a constant factor per step, which means a doubling process and nothing slower
- **WHEN** you are building a set greedily **DO** write the exact update rule for one step before choosing anything, then solve for the condition under which the update is maximally efficient **BECAUSE** the right auxiliary object, here $H^{-1}H$, is produced by solving for the condition, not by inspiration, and an exact update rule costs nothing while an inclusion costs the whole proof
- **WHEN** a greedy process must terminate because a quantity is capped **DO** write down the negation of the continuation condition and expect it to be the hypothesis you need next **BECAUSE** extremal arguments are architected so that the failure of the growth step is itself a strong structural statement, so the stopping condition is the second half of the proof rather than an obstacle

- **WHEN** an ordering constraint stops you writing $x^{-1}y$ as an admissible product **DO** lay the inverses down in reverse order first and the originals afterwards **BECAUSE** inversion reverses products term by term, so a decreasing selection in the mirrored half is legal, and the cost is only a factor of two in length
- **WHEN** a counting argument takes place inside a group **DO** look for the step that uses $|Hg| = |H|$ and name it **BECAUSE** translation invariance of cardinality is almost always the load bearing fact, and the same argument in a monoid or semigroup would collapse

Backtesting

1. First useful cue: the requested bound is $c \log |G|$ and a length ℓ list has 2^ℓ subsequences, so the answer must be a doubling construction.
2. Method the expression suggested: compute the exact effect of appending one element, $H \mapsto H \cup Hg$; solve for exact doubling, $g \notin H^{-1}H$; greedy until it fails; harvest the failure $H^{-1}H = G$ with a mirrored list.
3. Where you would hesitate, and the thought to reinforce: at the moment the greedy stops, which feels like a dead end. Reinforce it with the protocol: a greedy is designed to fail, and its failure condition is the gift. Write the negation down before deciding you are stuck.
4. The trap the question sets: reading "product of some subsequence" as generation. Subsequences use each term at most once and in order, so a generating set is worthless by itself, and any approach based on words in generators will produce lengths of order $|G|$.
5. What makes you recognise it faster next time: the pairing of an information theoretic target with an ordered selection rule. That pairing has appeared before in covering codes, binary representations and subset sum constructions, and its answer is always doubling plus a device to handle the ordering.
6. The reusable decision rule: compute the update rule, solve for maximal efficiency, run the greedy, then use the stopping condition as a theorem.

Mirror drills

1. Prove that every finite group G has a generating set of size at most $\log_2 |G|$. *Answer:* greedily add any element outside the subgroup generated so far; the subgroup strictly grows, so by Lagrange its order at least doubles at each step, and the order is capped by $|G|$.
2. Show that the constant c cannot be taken smaller than 1 when log means $\log_2$. *Answer:* a list of length ℓ has 2^ℓ subsequences, so $2^\ell \geq |G|$; and in $(\mathbb{Z}/2\mathbb{Z})^d$ the d standard basis vectors achieve $\ell = d = \log_2 |G|$ exactly, so the bound is attained.
3. In $\mathbb{Z}/n\mathbb{Z}$ exhibit a list of length $\lceil \log_2 n \rceil$ whose subsequence sums are all residues, and say why the mirror step of the ideal solution is unnecessary here. *Answer:* take $1, 2, 4, \ldots, 2^{\lceil \log_2 n \rceil - 1}$; binary representation gives every value in $\{0, \ldots, n-1\}$. No mirror is needed because the group is abelian and the subset sums already form the full difference set, so the greedy never needs an inverse.

18.7 B1. rational chords force a rational centre

Let C be a circle in $\mathbb{R}^2$ whose centre is not a rational point. What is the maximum number of rational points that can lie on C? (A rational point is a point both of whose coordinates are rational numbers.)

Read it out loud

Say the sentence back to yourself in order: a circle, an irrational centre, then a count of rational points on the boundary. The command term is "maximum number", which always splits into two separate jobs: an upper bound, and a construction that meets it. The guiding light is the phrase "whose centre is not a rational point": the question is not really asking you to study circles, it is asking what an irrational centre can force about the points that sit on that circle. That phrasing is the cue to run the argument in reverse, asking not how many rational points fit, but what having several rational points would force onto the centre.

The field embedded here is elementary Euclidean geometry, but the mechanism doing the actual work is field arithmetic: the rational numbers are closed under addition, subtraction, multiplication, and division by anything nonzero. Any point built from rational data using only those four operations is again rational. A circle's defining property, equidistance from the centre, translates directly into perpendicular bisectors, and a perpendicular bisector of a rational segment is built from nothing but a midpoint (an average) and a rotated direction vector (a coordinate swap and a sign change). Both stay inside the rationals. Within thirty seconds you should have written: if three rational points lie on the circle, the centre lies on two perpendicular bisectors, both rational lines, so the centre is rational. That sentence, once written, is the whole problem; everything else is bookkeeping to make it rigorous.

Do not reach for the circle's algebraic equation $(x-a)^2+(y-b)^2=r^2$. It is correct but it drags in the radius, a genuinely new unknown you never need to touch. The moment a question ties a centre's rationality to points on the circle, the radius is irrelevant scaffolding: strip it away and work with distances pairwise, which is exactly what a perpendicular bisector already encodes.

Approaches, ranked worst to best

1. **numeric guessing.** *Cue:* trying explicit rational triples and checking whether a circle through them can have an irrational centre. *Circuit:* pattern spotting, not proof.: this only produces a conjectured answer of 2. It never closes the upper bound, since you cannot enumerate every rational triple, and it supplies no mechanism for the construction half.

2. **general circle equation.** *Cue:* writing $(x-a)^2+(y-b)^2=r^2$, or the equivalent 4×4 determinant, through three unknown points. *Circuit:* brute elimination of a, b, r.: fully rigorous, and in fact Kedlaya's own alternate remark, but it costs more ink: you carry the radius as a live unknown throughout, or expand a 4×4 determinant correctly under time pressure, either of which invites an arithmetic slip the leaner route never risks.

3. **perpendicular bisector routing.** *Cue:* "centre is not a rational point" tied to a count of points equidistant from it. *Circuit:* rational field closure under midpoint, rotation, and a nonsingular 2×2 linear solve.: the ideal route. It never introduces the radius, needs only two of the three hypothetical points at a time, and every step is a named field operation, so the rationality of the centre is visibly forced rather than argued about.

The ideal solution

The answer is $\boxed{2}$.

Upper bound: at most two. Suppose, for contradiction, that three rational points $P = (a, b)$, $Q = (c, d)$, $R = (e, f)$ all lie on a circle C with centre $O = (g, h)$. A line meets a circle in at most two points, so P, Q, R are pairwise distinct and noncollinear.

Because O is equidistant from P and Q, it lies on the perpendicular bisector of segment PQ. That bisector consists of every point (x, y) satisfying

$$(c - a)\left(x - \tfrac{a+c}{2}\right) + (d - b)\left(y - \tfrac{b+d}{2}\right) = 0,$$

a line $A_1x + B_1y = C_1$ with $A_1 = c - a$, $B_1 = d - b$, and C_1 built from rational sums and products of a, b, c, d; hence A_1, B_1, C_1 are all rational. The same construction on Q and R gives a second rational line $A_2x + B_2y = C_2$ through O.

Since P, Q, R are noncollinear, the vectors $(c - a, d - b)$ and $(e - c, f - d)$ are linearly independent. Rotating a vector by a quarter turn is an invertible linear map, so the corresponding normal vectors (A_1, B_1) and (A_2, B_2) are independent too, which means $A_1B_2 - A_2B_1 \neq 0$. The two rational lines therefore meet in exactly one point, and by Cramer's rule that point is

$$g = \frac{C_1B_2 - C_2B_1}{A_1B_2 - A_2B_1}, \qquad h = \frac{A_1C_2 - A_2C_1}{A_1B_2 - A_2B_1}.$$

Both coordinates are ratios of rational numbers with a nonzero rational denominator, so $O = (g, h)$ is rational. This contradicts the hypothesis that the centre of C is irrational. Hence no circle with an irrational centre carries three rational points.

Construction: two is achieved. Fix the two rational points $(0, 0)$ and $(1, 0)$. For any real t, the circle centred at $\left(\frac{1}{2}, t\right)$ with radius $\sqrt{\frac{1}{4} + t^2}$ passes through both, since each is at that exact distance from $\left(\frac{1}{2}, t\right)$. Choosing t irrational, say $t = 1$, gives a circle with centre $\left(\frac{1}{2}, 1\right)$, which is not a rational point, carrying two rational points on it.

Combining the bound and the construction, the maximum number of rational points on a circle whose centre is not rational is exactly 2.

Layer 1: WHAT

Block 1: read "equidistant from an irrational centre" as the true content of the problem, and translate it into perpendicular bisectors rather than the circle's algebraic equation. Block 2: show a perpendicular bisector of a rational segment is a rational line, since its coefficients come from a midpoint and a rotated direction vector, both field operations that keep rational inputs rational. Block 3: intersect two such rational lines, using noncollinearity of the three hypothetical points to guarantee the lines are not parallel, and read off a rational intersection point by Cramer's rule. Block 4: the contradiction closes the upper bound; supply the two point construction with an irrational centre to show the bound is attained.

The nominal route bifurcates from the ideal one immediately after "suppose three rational points lie on the circle": the zombie writer reaches for $(x - a)^2 + (y - b)^2 = r^2$, substitutes three points, and grinds through three simultaneous quadratics, or a 4×4 determinant, to isolate a and b, carrying the unnecessary unknown r the entire way. The ideal writer instead asks what a circle is, definitionally, and reaches for the perpendicular bisector, which never mentions r at all. That single choice, at the very first line of working, is the bifurcation point.

Layer 2: HOW

Block 1. Cue. "centre is not a rational point" set against "how many rational points lie on the circle". **Why natural.** A circle is defined as the locus equidistant from its centre; equidistance between two named points is precisely the perpendicular bisector construction, so any argument connecting the centre's rationality to points on the circle should be built from bisectors of chords, which need only

two points at a time, rather than the full circle equation, which needs all three points and the radius simultaneously. **Link to the next block.** Committing to bisectors means you now need the purely algebraic fact that a bisector of a rational segment has rational coefficients.

Block 2. Cue. midpoint and direction vector of a segment between two rational points. **Why natural.** the rational numbers are closed under addition, subtraction, multiplication, and division by a nonzero element; a midpoint is an average of rational coordinates, hence rational, and a direction vector rotated a quarter turn is obtained by swapping the two coordinates and negating one, an operation that never leaves the rationals. Every coefficient of the bisector's line equation is a finite combination of such operations on rational inputs, so the whole equation is rational. **Link to the next block.** With two rational bisector lines available, from P, Q and from Q, R, the remaining question is a 2×2 linear algebra one: do they meet, and where.

Block 3. Cue. two linear equations, two unknowns, and the noncollinearity of P, Q, R. **Why natural.** P, Q, R lie on a genuine circle, and a line meets a circle in at most two points, so the three are automatically noncollinear; that forces the chord directions PQ and QR to be linearly independent, and a quarter turn rotation, being an invertible linear map, carries that independence over to the two bisector normal vectors. Two lines with independent normals meet at exactly one point, and Cramer's rule writes that point as a ratio of rational determinants over a nonzero rational determinant, hence rational. **Link to the next block.** That rational point is forced to be the centre, since both bisectors pass through it by definition, contradicting an irrational centre; the last job is to check the bound of two is achievable.

Block 4. Cue. the word "maximum" in the question. **Why natural.** "maximum" is a claim of equality, not merely an inequality; an upper bound argument alone only proves "at most", so the problem is not finished until a circle with an irrational centre carrying two rational points is exhibited explicitly. **Link to the next block.** none, this closes the problem.

Layer 3: WHY

The phenomenon underneath this problem has nothing to do with circles specifically. The rational numbers form a field, closed under the four operations except division by zero, and any point in the plane describable purely by midpoints, differences, quarter turn rotations, and the solution of a nonsingular linear system with rational coefficients, must itself be rational, with no reference to geometry required. A circle happens to give you a natural source of such a description, through the equidistance property, since equidistance from two points is exactly a perpendicular bisector, and a bisector of a rational segment is one of those purely field theoretic objects.

The discernment mechanism to train is this: whenever a problem asks you to prove that some derived point, a centre, an intersection, a foot of a perpendicular, inherits rationality (or integrality, or membership of some other closed algebraic system) from given data, search for the smallest set of field operations that produces that point. If you can write the point down using only those operations, the rationality is automatic and you never need the full geometric picture, including quantities like a radius the target point does not actually depend on. Examiners nudge you toward this by phrasing the hypothesis on the centre rather than on the points ("centre is not a rational point"), an invitation to argue by contrapositive: assume enough rational points and force the centre.

Revolving concepts

Rational point and rational line. A point is rational if both coordinates lie in $\mathbb{Q}$; a line is rational if it can be written $ax + by = c$ with $a, b, c \in \mathbb{Q}$. Two distinct rational points determine a rational line, and two nonparallel rational lines meet at a rational point, by Cramer's rule on a 2×2 system with

rational, nonzero determinant coefficients. This is genuinely a statement about any field K in place of $\mathbb{Q}$: two K-points determine a K-line, and two nonparallel K-lines meet at a K-point. See Book 1, field: Geometry, topic: affine geometry over a field; also field: Linear algebra, topic: Cramer's rule for 2×2 systems.

Perpendicular bisector as the centre's locus. The centre of any circle through two given points lies on the perpendicular bisector of the segment joining them, since the centre is by definition equidistant from both. See Book 1, field: Geometry, topic: loci of equidistant points.

The general circle through three points. The equation of the circle through (x_1, y_1), (x_2, y_2), (x_3, y_3) can be written as a 4×4 determinant in $x^2 + y^2$, x, y, 1, giving $a(x^2 + y^2) + dx + ey + f = 0$ with centre $(-d/2a, -e/2a)$; rational input points give rational a, d, e, f, a second, heavier proof of the same fact. See Book 1, field: Matrices and determinants, topic: determinant formula for a conic through given points.

Generalisation to higher dimensions. The same bisector argument shows that at most n rational points can lie on an $(n-1)$-dimensional sphere in $\mathbb{R}^n$ whose centre is irrational: n rational points give $n-1$ independent rational bisector hyperplanes, which pin down the centre only once you have n of them, i.e. $n+1$ points. Setting $n = 2$ recovers the ceiling of 2 for a circle. See Book 1, field: Geometry, topic: spheres and hyperplane bisectors in $\mathbb{R}^n$.

Circuit Bank

- **WHEN** a problem ties rationality (or integrality) of a derived geometric point, a centre, an intersection, a foot of a perpendicular, to rationality of the given points **DO** search for a description of that point using only midpoints, differences and the solution of a nonsingular linear system with rational coefficients **BECAUSE** those operations never leave the field the given data lives in, so the derived point is rational for free, with no case work on the geometric picture needed
- **WHEN** a circle problem states a constraint on the centre and asks about points on the circle **DO** route through perpendicular bisectors of chords rather than the circle's algebraic equation **BECAUSE** a bisector uses only two points and never introduces the radius as an extra unknown, the minimum amount of algebra the question actually requires
- **WHEN** the question asks for a "maximum number" **DO** treat it as two separate deliverables: an upper bound argument, then an explicit construction attaining it **BECAUSE** "maximum" asserts equality, and an upper bound alone only proves "at most", never "exactly"
- **WHEN** you need to certify two chord directions are linearly independent **DO** invoke that a genuine circle meets a line in at most two points, so any three points on it are automatically noncollinear **BECAUSE** this removes the need for a separate non-degeneracy check on the three points

Backtesting

1. First useful cue: "centre is not a rational point" set against "maximum number of rational points", inviting the contrapositive: assume enough rational points and force the centre to be rational.
2. Method the expression suggested: equidistance from the centre routes directly to perpendicular bisectors of chords, not to the circle's algebraic equation.
3. Where you would hesitate, and the thought to reinforce: whether the two bisector lines are guaranteed to meet; reinforce with "three points on a genuine circle are automatically noncollinear, so their chord directions are independent, and a quarter turn rotation preserves that independence".

4. The trap the question sets: reaching for $(x-a)^2+(y-b)^2=r^2$ and solving for a, b, r together, dragging in the radius, an unknown the target point never depends on.
5. What makes you recognise it faster next time: whenever the proof target is "this derived point lies in the same field as the given data", check first whether it is built from midpoints, differences, and one linear solve.
6. The reusable decision rule: when the rationality of a derived point is the target, express that point purely through field operations on the given rational data, never through a full geometric equation carrying an extra unknown you do not need.

Mirror drills

1. Show that if three rational points lie on a circle, the square of its radius is also rational. *Answer:* the centre is rational by the argument above, so the squared radius, the squared distance from that centre to any one of the three rational points, is a sum of squares of rational differences, hence rational.
2. A parabola has no distinguished centre point equidistant from all of its points. Why does that mean the bisector circuit above cannot be used to bound the number of rational points on a parabola? *Answer:* the circuit is special to loci defined by equidistance from a single centre; a parabola carries no such centre, so there is no perpendicular bisector to route the argument through, and no contradiction of the same shape can be built.
3. Three rational points lie on a sphere in $\mathbb{R}^3$ whose centre is irrational. What is the true ceiling on rational points on such a sphere, and why does it differ from the circle's answer of 2? *Answer:* the ceiling is 3: two rational bisector planes from three points generically meet in a line, not a point, in $\mathbb{R}^3$, so a fourth rational point, giving three independent bisector planes, is needed before the centre is forced rational; the ceiling in $\mathbb{R}^n$ is always n.

18.8 B2. guess the closed form, then induct

Let $F_0(x) = \ln x$. *For* $n \geq 0$ *and* $x > 0$, *let* $F_{n+1}(x) = \int_0^x F_n(t)\,dt$. *Evaluate*

$$\lim_{n\to\infty} \frac{n!\,F_n(1)}{\ln n}.$$

Read it out loud

This is a recursively defined sequence of functions, each the antiderivative of the previous one, taken from 0: $F_0 = \ln x$, F_1 its integral from 0, F_2 the integral of that, and so on. The command term is "evaluate", not "find" or "prove": you are not asked to characterise F_n in general, only to extract a single numeric limit built from $F_n(1)$. That is itself a cue: whenever a problem defines an n-fold iterated operator and then asks only for the value at one fixed point, times a normalising factor, the intended route is almost always to find a closed form for $F_n(x)$ and read off the asymptotics, rather than to analyse $F_n(1)$ directly as n grows through some limiting trick.

The fields embedded are calculus, the Fundamental Theorem of Calculus and repeated integration, alongside the asymptotics of the harmonic numbers. The presence of $n!$ in the numerator, next to a limit the question asserts exists and is finite, is itself informative: $F_n(1)$ must decay roughly like $1/n!$ for the product $n!\,F_n(1)$ to stay controlled, a strong hint that $F_n(x)$ carries an $x^n/n!$ factor built in from n-fold integration. The $\ln n$ in the denominator similarly hints that whatever multiplies $x^n/n!$ grows like $\ln n$, and the most natural quantity that grows like $\ln n$ while being built from n discrete steps is the harmonic number $H_n = \sum_{k=1}^{n} 1/k$.

The first useful move, within the first minute, is to compute $F_1(x)$ explicitly: $\int_0^x \ln t\,dt = x\ln x - x$, a standard integral. Laying that next to the guess "$F_n(x)$ should look like $(\ln x - a_n)x^n/n!$ for some sequence a_n" already matches at $n = 1$ with $a_1 = 1 = H_1$. That match is the entire content of the problem: once you have the ansatz and one confirmed base case, the rest is a mechanical induction using the Fundamental Theorem of Calculus, followed by a well known bound on harmonic numbers.

Approaches, ranked worst to best

1. **direct estimation.** *Cue:* "evaluate a limit" read as an invitation to bound $F_n(1)$ recursively without ever writing a closed form. *Circuit:* ad hoc integral bounding, no real circuit.: ranks worst since every step needs a fresh, uncertain estimate on a function you have not pinned down, with no guarantee the bounds are tight enough to identify the exact value -1 rather than merely show boundedness.
2. **grind and pattern spot.** *Cue:* computing F_1, F_2, F_3 explicitly by repeated integration by parts and eyeballing the pattern. *Circuit:* pattern recognition from small cases.: genuinely useful as the discovery phase, and exactly how you should find the ansatz, but stopping here leaves the claim unproven for general n; it needs the formal induction that the ideal route supplies.
3. **ansatz plus induction.** *Cue:* an operator applied n times to a fixed seed, asked about only at one point. *Circuit:* guess a closed form from the first case, confirm it for all n by induction via the Fundamental Theorem of Calculus, then apply the harmonic number asymptotic.: the ideal route. It produces the exact closed form of $F_n(1)$ with a short, fully justified induction, after which the target limit is one line.

The ideal solution

Claim. For every $n \geq 1$ and $x > 0$, $F_n(x) = (\ln x - H_n)\dfrac{x^n}{n!}$, where $H_n = \sum_{k=1}^{n} \dfrac{1}{k}$, extended continuously by $F_n(0) = 0$.

Base case $n = 1$**.** Integrating by parts, $\int \ln t\, dt = t \ln t - t + C$, and $\lim_{t\to 0^+}(t \ln t - t) = 0$, so

$$F_1(x) = \int_0^x \ln t\, dt = x \ln x - x = (\ln x - 1)x,$$

matching the claim with $H_1 = 1$.

Inductive step. Suppose the claim holds for some $n \geq 1$, with F_n continuous on $(0, \infty)$ and $F_n(x) \to 0$ as $x \to 0^+$. Set

$$G(x) = (\ln x - H_{n+1}) \frac{x^{n+1}}{(n+1)!}, \qquad H_{n+1} = H_n + \frac{1}{n+1}.$$

Differentiating,

$$G'(x) = \frac{1}{x} \cdot \frac{x^{n+1}}{(n+1)!} + (\ln x - H_{n+1}) \cdot \frac{x^n}{n!} = \frac{x^n}{n!}\left[\frac{1}{n+1} + \ln x - H_{n+1}\right] = \frac{x^n}{n!}\big[\ln x - H_n\big] = F_n(x),$$

using $H_{n+1} - \frac{1}{n+1} = H_n$. Also $G(x) \to 0$ as $x \to 0^+$, since $x^{n+1} \ln x \to 0$ and $x^{n+1} \to 0$.

Both F_{n+1} and G are continuous on $(0, \infty)$ with the same derivative F_n there, so their difference has zero derivative on the connected interval $(0, \infty)$ and is constant; that constant is forced to 0 by the shared limit 0 as $x \to 0^+$. Hence $F_{n+1} = G$ on $(0, \infty)$, which is exactly the claim for $n + 1$. This completes the induction.

Evaluating the limit. By the claim, $F_n(1) = (\ln 1 - H_n)/n! = -H_n/n!$, so $n!\, F_n(1) = -H_n$ and

$$\lim_{n\to\infty} \frac{n!\, F_n(1)}{\ln n} = -\lim_{n\to\infty} \frac{H_n}{\ln n}.$$

Since $1/x$ is decreasing, for $k \geq 2$, $\int_k^{k+1} \frac{dx}{x} \leq \frac{1}{k} \leq \int_{k-1}^{k} \frac{dx}{x}$. Summing over $k = 2, \ldots, n$ gives

$$\ln(n+1) - \ln 2 \leq H_n - 1 \leq \ln n,$$

so $\ln(n+1) - \ln 2 + 1 \leq H_n \leq 1 + \ln n$. Dividing by $\ln n$ and letting $n \to \infty$, both bounds tend to 1, so $H_n / \ln n \to 1$ by the squeeze theorem.

Hence the required limit is $\boxed{-1}$.

Layer 1: WHAT

Block 1: recognise an n-fold operator applied to a fixed seed, asked about only at one point with a normalising factor, and route to finding a closed form, not to direct asymptotic estimation. Block 2: compute F_1 explicitly by a standard integral, matching the guessed ansatz $(\ln x - a_n)x^n/n!$. Block 3: prove the ansatz for all n by induction via the Fundamental Theorem of Calculus, differentiating the proposed F_{n+1}, matching the previous ansatz, and pinning the constant with the boundary limit at $x = 0^+$. Block 4: evaluate at $x = 1$ to get $F_n(1) = -H_n/n!$ in closed form, reducing the whole limit to $H_n/\ln n$. Block 5: apply the standard integral comparison squeeze for harmonic numbers to conclude $H_n/\ln n \to 1$, giving the answer -1.

The nominal route bifurcates immediately after computing $F_1(x) = x \ln x - x$: the zombie writer moves on to bound $F_2, F_3, \ldots$ one at a time with fresh inequalities, hoping a pattern of bounds emerges. The ideal writer instead proposes the general ansatz $(\ln x - a_n)x^n/n!$ at once and commits to proving it for every n in a single induction. That single choice, right after the first computation, is the bifurcation point.

Layer 2: HOW

Block 1. **Cue.** $n!$ sitting next to $F_n(1)$ in the same limit as $\ln n$, immediately after a definition built by n-fold integration. **Why natural.** a limit built from $n!$ times the n-th iterate value, asked to be evaluated, hence finite, signals $F_n(x)$ must carry an $x^n/n!$ factor, or $n!\,F_n(1)$ could not stay controlled; the $\ln n$ in the denominator signals whatever remains should grow like a harmonic sum, the canonical n-step-built quantity with logarithmic growth. The productive move is to conjecture a closed form of exactly that shape rather than estimate $F_n(1)$ term by term. **Link to the next block.** finding the precise closed form needs at least one explicit computation to calibrate it.

Block 2. **Cue.** $F_1(x) = \int_0^x \ln t\,dt$, a standard integral you already know. **Why natural.** computing the first iterate explicitly is the cheapest way to calibrate an ansatz; matching $(\ln x - a_1)x/1!$ against $x\ln x - x$ forces $a_1 = 1$, which is also H_1, the clue that a_n should be the n-th harmonic number rather than some other logarithm adjacent sequence. **Link to the next block.** with a calibrated ansatz for $n = 1$ you need a mechanism propagating it to every n, exactly what the defining recursion supplies via the Fundamental Theorem of Calculus.

Block 3. **Cue.** the recursive definition $F_{n+1}(x) = \int_0^x F_n(t)\,dt$ itself. **Why natural.** the Fundamental Theorem of Calculus converts "F_{n+1} is defined as an integral of F_n" into "$F'_{n+1} = F_n$ plus a boundary condition at 0", a differential equation with an initial condition, the natural tool for confirming a guessed closed form: differentiate the candidate for F_{n+1}, check it equals the already trusted formula for F_n, and match the one remaining degree of freedom using the value at the singular endpoint $x = 0$, where $x^k \ln x \to 0$ rescues you from the blow up of $\ln x$ itself. **Link to the next block.** with the closed form nailed down for every n, the actual question, a numeric limit, needs only the single value $x = 1$.

Block 4. **Cue.** the specific evaluation point $x = 1$ named in the question. **Why natural.** plugging in $x = 1$ kills the $\ln x$ term entirely, collapsing the two term closed form to the one term $-H_n/n!$, and the $n!$ in the question's numerator exists precisely to cancel the $n!$ in that denominator, a confirmation that the derived formula matches the problem's own normalising constants. **Link to the next block.** what remains is a single classical asymptotic fact about H_n.

Block 5. **Cue.** the ratio $H_n/\ln n$. **Why natural.** the harmonic numbers are sandwiched between consecutive integrals of $1/x$ by monotonicity of $1/x$, a standard integral test comparison, giving $H_n = \ln n + O(1)$ and hence $H_n/\ln n \to 1$, with no need for the sharper Euler-Mascheroni constant. **Link to the next block.** none, this closes the problem, the answer -1 picking up its sign from the ansatz's $-H_n$.

Layer 3: WHY

The underlying truth is that repeated application of a linear operator, here "integrate from 0", to a function decomposing into a log part and a polynomial part preserves that decomposition, since integration is linear and $\int x^n \ln x\,dx$ and $\int x^n\,dx$ both stay inside the span of $\{x^{n+1}\ln x,\ x^{n+1}\}$ after one integration by parts. The correct discernment mechanism whenever an n-fold iterated linear operator appears, integration, differentiation, a linear recurrence, a matrix power, is: identify the smallest finite dimensional family of functions or vectors the operator maps into itself, guess the coefficients from one or two small cases, and confirm by induction using the operator's own defining property, the Fundamental Theorem of Calculus for an integral operator, the recurrence itself for a sequence, an eigen decomposition for a matrix.

The examiners nudged you here by choosing $F_0 = \ln x$ specifically, not a polynomial, not e^x, a seed whose repeated antiderivatives are famous for staying inside the two dimensional family

$\{x^n \ln x,\ x^n\}$, and by asking for a ratio built exactly from the two natural scales in that family: $n!$ from the $x^n/n!$ coefficient that n-fold integration of a bounded function always produces, $\ln n$ from the harmonic number coefficient that using $\ln x$ as the seed always produces.

Revolving concepts

Repeated antiderivative of $\ln x$. $F_0 = \ln x$, $F_{n+1} = \int_0^x F_n$, gives the closed form $F_n(x) = (\ln x - H_n)x^n/n!$; the family $\{x^n \ln x,\ x^n\}$ is preserved by integration from 0. See Book 1, field: Calculus and integrals, topic: repeated integration by parts and reduction formulae.

Harmonic numbers and their logarithmic growth. $H_n = \sum_{k=1}^{n} 1/k$ satisfies $\ln(n+1) \le H_n \le 1 + \ln n$, proved by sandwiching $1/k$ between consecutive integrals of $1/x$; hence $H_n/\ln n \to 1$, and more precisely $H_n = \ln n + \gamma + o(1)$ for the Euler-Mascheroni constant γ, though only the coarser bound is needed here. See Book 1, field: Sequences and recurrences, topic: harmonic number asymptotics via the integral test.

Pinning a constant of integration at a singular endpoint. Two functions continuous on $(0,\infty)$, differentiable with equal derivatives there, and sharing a limit as $x \to 0^+$, are identical on $(0,\infty)$: their difference has zero derivative on a connected interval, hence is constant, and the shared limit forces that constant to 0. See Book 1, field: Real analysis, topic: uniqueness of an antiderivative via the mean value theorem.

Recursion generalisation. Any sequence of functions $F_{n+1} = T[F_n]$ for a linear operator T can be attacked by guessing a closed form ansatz from the first one or two terms, then confirming by induction that the ansatz for F_{n+1} satisfies T's own defining identity when applied to the ansatz for F_n, matching any remaining constant at a convenient boundary value. This is the standard playbook for any "apply an operator n times" Putnam problem: iterated derivatives with a Leibniz rule closed form, iterated finite differences, and linear recurrences solved by an ansatz spanned by the characteristic roots. See Book 1, field: Sequences and recurrences, topic: ansatz plus induction for linear iterated operators.

Circuit Bank

- **WHEN** a problem defines a sequence of functions by repeated integration, or differentiation, of a fixed seed and asks for the value at one point **DO** guess a closed form ansatz from the first explicit computation and prove it for all n by induction via the Fundamental Theorem of Calculus **BECAUSE** repeated application of a linear operator to a function inside a small closed family stays inside that family, so one calibration plus one inductive step pins down every term at once
- **WHEN** a limit involves $n!$ in the numerator alongside the n-th term of a sequence built by n-fold integration **DO** expect the sequence's closed form to carry a matching $x^n/n!$ factor and let the two cancel **BECAUSE** n-fold integration of a function that does not grow factorially always produces exactly that $1/n!$ scale
- **WHEN** you need to pin the constant of integration for an antiderivative defined via a definite integral from a fixed lower limit **DO** use the value, or limit, of both sides at that lower limit rather than an arbitrary point **BECAUSE** the defining integral automatically vanishes there, giving a free, exact data point instead of an unknown constant to solve for
- **WHEN** a limit reduces to a sum resembling $1 + \frac{1}{2} + \cdots + \frac{1}{n}$ divided by $\ln n$ **DO** reach immediately for the integral test comparison sandwiching the harmonic sum between consecutive integrals of $1/x$ **BECAUSE** this gives $H_n = \ln n + O(1)$ in two lines with no deeper machinery needed

Backtesting

1. First useful cue: $n!$ sitting next to $F_n(1)$ in the same limit as $\ln n$, immediately after a definition built by n-fold integration.
2. Method the expression suggested: guess a closed form $(\ln x - a_n)x^n/n!$, calibrated by computing F_1 explicitly.
3. Where you would hesitate, and the thought to reinforce: whether the constant of integration in the inductive step is really pinned down; reinforce with "match both sides at the lower limit of the defining integral, where the defining integral is automatically zero".
4. The trap the question sets: trying to estimate $F_n(1)$ via ad hoc bounds one iteration at a time instead of committing to an exact closed form, which either stalls or silently reconstructs the same formula with far more work.
5. What makes you recognise it faster next time: any "operator applied n times to a fixed seed, evaluate at one point, normalised by $n!$ or a similar scale" problem should trigger closed form ansatz plus induction before any estimation is attempted.
6. The reusable decision rule: when a recursively defined sequence of functions is built by a linear operator from a simple seed, find the finite dimensional family the operator preserves, calibrate on the first case, and induct via the operator's own defining identity.

Mirror drills

1. Let $G_0(x) = e^x$ and $G_{n+1}(x) = \int_0^x G_n(t)\,dt$. Find a closed form for $G_n(x)$. *Answer:* $G_n(x) = e^x - \sum_{k=0}^{n-1} x^k/k!$, the tail of the exponential series after subtracting its first n Taylor terms; found the same way, calibrate on $G_1(x) = e^x - 1$, guess the tail ansatz, confirm by induction via the Fundamental Theorem of Calculus.
2. Before computing anything, explain why the mere appearance of $n!\,F_n(1)$ as an evaluable quantity already predicts that F_n's closed form must contain an $x^n/n!$ factor. *Answer:* for $n!\,F_n(1)$ to stay controlled as $n \to \infty$, $F_n(1)$ itself must shrink at a rate comparable to $1/n!$, and the only natural way n-fold integration produces a $1/n!$ scale is through the $x^n/n!$ term arising from integrating a degree $n-1$ monomial exactly n times.
3. Evaluate $\lim_{n\to\infty} \dfrac{1 + \frac{1}{2} + \cdots + \frac{1}{n}}{\ln(n+1)}$. *Answer:* 1, by the same integral test squeeze used above, since $\ln(n+1)$ and $\ln n$ are asymptotically equivalent.

18.9 B3. pair up coordinates, bound each pair

What is the largest possible radius of a circle contained in a four dimensional hypercube of side length 1*?*

Read it out loud

Strip away the fear of "four dimensional": a circle is still a one parameter family of points $v_1 \cos\theta + v_2 \sin\theta$ traced about its centre, for two orthogonal vectors v_1, v_2 of equal length r, no matter what the ambient dimension is. That parametrisation is the entire content of the word "circle" and is the correct starting point regardless of dimension. "Contained in a hypercube" translates coordinate by coordinate into a bound on each coordinate function, once the hypercube is centred at the origin, which costs nothing since both the hypercube and the circle are shapes, unaffected by where you choose to place your axes. So the whole question becomes: how large can r be while every coordinate of $v_1 \cos\theta + v_2 \sin\theta$ stays inside the hypercube's range, for every θ at once.

The guiding light is "contained in", a per coordinate, for all θ constraint, not a single point constraint. The natural tool for "what is the maximum over θ of $a\cos\theta + b\sin\theta$" is the amplitude formula $\sqrt{a^2+b^2}$, an identity most people meet purely as a trigonometric rewriting trick but rarely recognise as a norm bound in disguise. Applying it once per coordinate converts a continuum of containment constraints, one for every θ, into finitely many algebraic constraints, one per coordinate, exactly the reduction turning a seemingly continuous, four dimensional geometry problem into a finite sum of squares optimisation.

There is a second, quieter idea hiding before any of that: whether the circle's centre can be assumed to sit at the hypercube's own centre at all. It can, for free, because both a circle and a hypercube are symmetric under reflection through their own centre, and that symmetry alone forces the recentred circle back inside the hypercube. Once that is granted, the whole difficulty of the problem lives entirely in the coordinate by coordinate amplitude bound and the arithmetic that follows it, not in any hard piece of four dimensional visualisation.

Approaches, ranked worst to best

1. **coupled optimisation.** *Cue:* "largest circle contained in a region" read as a joint calculus problem over the centre's position and the circle's orientation together. *Circuit:* brute multivariable optimisation, no real circuit.: ranks worst because it conflates two logically separate freedoms, where the centre sits and how the circle is oriented once centred, into one large coupled optimisation, when the centre can be fixed for free by a short symmetry lemma.
2. **low dimension analogy.** *Cue:* noticing the largest circle in a unit square or a unit cube face has radius $1/2$ and guessing the pattern continues. *Circuit:* pattern matching across dimensions, not proof.: a useful sanity check, and it correctly flags that a flat, coordinate plane circle is only a candidate, not necessarily optimal, but by itself it produces no bound and no rigorous maximum.
3. **recentre then bound per coordinate.** *Cue:* a "for all θ" containment condition on a linear combination of $\cos\theta$ and $\sin\theta$. *Circuit:* free recentring via central symmetry, then the amplitude identity applied coordinate by coordinate.: the ideal route. It splits the problem into two independent, individually easy steps, a symmetry lemma and a sum of squares bound, and produces a fully rigorous bound plus a matching construction in a handful of lines.

The ideal solution

It is convenient to solve the rescaled problem for a hypercube of side length 2, namely $H = [-1,1]^4$, and divide the resulting radius by 2 at the end; this trades the awkward bound $1/2$ per coordinate

for the cleaner bound 1, at the cost of one final division.

Recentring the circle for free. Let C be a circle of radius r contained in H, with centre P. Since C is a circle centred at P, it is invariant under the point reflection $\rho_P(x) = 2P - x$. Because $C \subseteq H$, applying ρ_P to both sides gives $C = \rho_P(C) \subseteq \rho_P(H)$, so $C \subseteq H \cap \rho_P(H)$. In coordinate i, H occupies $[-1, 1]$ and $\rho_P(H)$ occupies $[2P_i - 1, 2P_i + 1]$; these two length 2 intervals intersect in an interval centred at P_i with half width $1 - |P_i| \leq 1$ (the reflection $x \mapsto 2P_i - x$ swaps the two intervals and fixes only P_i, forcing the intersection to be symmetric about P_i, and the intersection cannot extend past 1 on either raw interval, giving half width exactly $1 - |P_i|$). Consequently $C - P$, the translate of C to be centred at the origin, lies inside $[-1, 1]^4 = H$ as well. So there is no loss of generality in assuming P is the origin O.

Reducing containment to a per coordinate bound. Write the recentred circle as $\{v_1 \cos\theta + v_2 \sin\theta : 0 \leq \theta < 2\pi\}$, where $v_1, v_2 \in \mathbb{R}^4$ are orthogonal with $|v_1| = |v_2| = r$. Containment in H means that for every coordinate i and every θ,

$$|v_{1i} \cos\theta + v_{2i} \sin\theta| \leq 1.$$

By the amplitude identity, $v_{1i} \cos\theta + v_{2i} \sin\theta = \sqrt{v_{1i}^2 + v_{2i}^2} \cos(\theta - \varphi_i)$ for a suitable phase φ_i, whose maximum over θ is exactly $\sqrt{v_{1i}^2 + v_{2i}^2}$. So the bound holds for every θ if and only if

$$v_{1i}^2 + v_{2i}^2 \leq 1 \qquad (i = 1, 2, 3, 4).$$

Bounding the radius. Since $v_1 \perp v_2$ and $|v_1| = |v_2| = r$,

$$2r^2 = |v_1|^2 + |v_2|^2 = \sum_{i=1}^{4} (v_{1i}^2 + v_{2i}^2) \leq \sum_{i=1}^{4} 1 = 4,$$

so $r \leq \sqrt{2}$.

Achieving equality. Take $v_1 = (1, 1, 0, 0)$ and $v_2 = (0, 0, 1, 1)$. Then $v_1 \cdot v_2 = 0$, $|v_1| = |v_2| = \sqrt{2}$, and coordinate by coordinate, $v_{1i}^2 + v_{2i}^2$ equals 1 for every $i = 1, 2, 3, 4$, so the circle traced by $v_1 \cos\theta + v_2 \sin\theta$ lies in H and has radius $\sqrt{2}$. This matches the bound, so the maximum radius for the side length 2 hypercube is exactly $\sqrt{2}$.

Undoing the rescaling. A hypercube of side length 1 is a hypercube of side length 2 scaled down by a factor of $\frac{1}{2}$, and scaling multiplies every length, including the radius of an inscribed circle, by the same factor. Hence the largest radius for side length 1 is $\frac{1}{2}\sqrt{2} = \boxed{\frac{\sqrt{2}}{2}}$.

Layer 1: WHAT

Block 1: recognise "circle" as the parametrisation $P + v_1 \cos\theta + v_2 \sin\theta$ with $v_1 \perp v_2$, a description valid in any dimension, and recognise "contained in the hypercube" as a per coordinate, for all θ constraint rather than a picture to visualise directly. Block 2: recentre the circle at the hypercube's own centre for free, via a short symmetry lemma, removing the centre's location from the optimisation before any inequality is written. Block 3: convert the for all θ containment constraint into a finite algebraic constraint per coordinate, using the amplitude identity for $a \cos\theta + b \sin\theta$. Block 4: sum the per coordinate bounds to get $2r^2 \leq 4$, the universal bound $r \leq \sqrt{2}$ for the rescaled hypercube. Block 5: exhibit an explicit orthogonal pair achieving equality, and undo the side length rescaling to answer the original question.

The nominal route bifurcates right after reading "circle contained in a four dimensional hypercube": the zombie writer tries to visualise, or directly optimise, over both the centre's position and the circle's orientation as one coupled multivariable problem. The ideal writer's very first move is the symmetry observation that the centre can be fixed at the hypercube's own centre without loss of generality, decoupling the problem into a pure "best orientation through a fixed point" question. That decoupling, before any inequality is written, is the bifurcation point.

Layer 2: HOW

Block 1. **Cue.** "circle" together with "contained in", a shape defined parametrically meeting a region defined by many linear constraints. **Why natural.** any circle of radius r, wherever it sits and however it is tilted, is exactly $\{P + v_1\cos\theta + v_2\sin\theta\}$ for its centre P and an orthogonal pair v_1, v_2 with $|v_1| = |v_2| = r$ spanning its plane; writing this down immediately turns a geometric containment question into a statement that must hold for every θ, an infinite family of linear inequalities in the free parameters P, v_1, v_2. **Link to the next block.** an infinite family indexed by θ is unwieldy directly, so the productive next move is removing as many of the other free parameters, in particular P, as possible before dealing with the θ dependence.

Block 2. **Cue.** the hypercube is centrally symmetric about its own centre, and so is a circle about its own centre. **Why natural.** since a circle is invariant under a 180 degree rotation about its own centre P, if $C \subseteq H$ then reflecting through P gives $C = \rho_P(C) \subseteq \rho_P(H)$, so C lies in $H \cap \rho_P(H)$, and a short coordinate computation shows this intersection, translated so P becomes the origin, still fits inside H; this costs a few lines and removes the centre entirely from the optimisation at zero cost to the bound. **Link to the next block.** with P fixed at the hypercube's own centre, the only remaining freedom is the orientation, encoded by v_1 and v_2, and containment becomes purely about how $v_1\cos\theta + v_2\sin\theta$ behaves coordinate by coordinate.

Block 3. **Cue.** a linear combination $a\cos\theta + b\sin\theta$, required to stay bounded for every θ. **Why natural.** this is exactly the amplitude, or $R\cos(\theta - \varphi)$, identity from Further Pure trigonometry, restated as a norm: the maximum over θ of $a\cos\theta + b\sin\theta$ is $\sqrt{a^2 + b^2}$, so "bounded by 1 for all θ" is equivalent to the single algebraic inequality $\sqrt{a^2 + b^2} \leq 1$; applying this once per coordinate converts the continuum of constraints from block 1 into four ordinary inequalities. **Link to the next block.** four inequalities of the same shape, each bounding a sum of two squares by 1, are begging to be added together.

Block 4. **Cue.** four separate inequalities $v_{1i}^2 + v_{2i}^2 \leq 1$, and the fact that $|v_1|^2 + |v_2|^2 = 2r^2$ by orthogonality and equal length. **Why natural.** summing the four per coordinate inequalities gives $\sum_i (v_{1i}^2 + v_{2i}^2) \leq 4$, and the left side is exactly $|v_1|^2 + |v_2|^2$ by definition of the Euclidean norm, which equals $2r^2$ since v_1, v_2 are orthogonal with common length r; this single line of arithmetic is the entire quantitative content of the bound. **Link to the next block.** an inequality alone only proves an upper bound, so the last job is checking it is achieved.

Block 5. **Cue.** equality in $\sum_i (v_{1i}^2 + v_{2i}^2) \leq 4$ needs every one of the four coordinate inequalities to be tight simultaneously. **Why natural.** pairing up the four coordinates two at a time, putting all of v_1's length into the first pair and all of v_2's length into the second, is the natural way to make each per coordinate sum equal exactly 1 while keeping $v_1 \perp v_2$; the vectors $(1, 1, 0, 0)$ and $(0, 0, 1, 1)$ do exactly this. **Link to the next block.** none, this closes the problem after undoing the side length rescaling.

Layer 3: WHY

The underlying truth is that a constraint required to hold "for all θ" against a two term trigonometric expression is always secretly a single Euclidean norm bound: the amplitude identity is nothing but the statement that the vector (a, b) has length at most the bound, viewed through the rotating unit vector $(\cos\theta, \sin\theta)$. Once you see that, any shape defined by a two parameter rotation, a circle, the principal axes of an ellipse, a rotating vector in a plane, meeting a box defined by independent coordinate bounds, reduces coordinate by coordinate to that same norm inequality, and the resulting optimisation is a plain sum of squares bound rather than a genuine multivariable calculus problem.

The discernment mechanism to train: whenever a "for all θ" or "for all directions" containment condition appears, ask whether it can be rewritten as a fixed norm bound via the amplitude identity, in two real dimensions, or the more general fact that the maximum of a linear functional over a sphere is achieved in the direction of the functional itself, in any number of dimensions, which is Cauchy-Schwarz. What the examiners were nudging toward: the number 4 in "four dimensional hypercube" is not there to make you visualise a four dimensional picture, it is there so the sum of four coordinate bounds gives a clean total of 4 and hence a clean $r = \sqrt{2}$ before rescaling, the same way "regular n-gon" problems nudge you toward roots of unity rather than literal polygon drawing.

Revolving concepts

Amplitude identity as a hidden norm bound. $a\cos\theta + b\sin\theta = \sqrt{a^2+b^2}\cos(\theta-\varphi)$ for a suitable φ; consequently $|a\cos\theta + b\sin\theta| \le M$ for every θ if and only if $\sqrt{a^2+b^2} \le M$. See Book 1, field: Inequalities, topic: amplitude form as a norm bound in disguise.

Parametrisation of a circle in $\mathbb{R}^n$. $P + v_1\cos\theta + v_2\sin\theta$, with $v_1 \perp v_2$ and $|v_1| = |v_2| = r$: the dimension free way to describe a circle once you leave the plane. See Book 1, field: Geometry, topic: circles as two dimensional slices of $\mathbb{R}^n$.

Central symmetry recentring lemma. If a centrally symmetric shape, any circle, or more generally any convex body symmetric about its own centre such as a ball or an ellipse, lies inside a centrally symmetric region, a cube or more generally any box, translating both so their centres coincide keeps the containment, proved by reflecting through the shape's centre. See Book 1, field: Geometry, topic: central symmetry and recentring arguments in containment problems.

Generalisation to k dimensional balls in n dimensional hypercubes. The largest radius of a k dimensional ball inside an n dimensional hypercube of side length 1 is $\frac{1}{2}\sqrt{n/k}$, provable via the trace identity $\operatorname{trace}(MM^{\mathsf{T}}) = \operatorname{trace}(M^{\mathsf{T}}M)$ for the matrix of radius vectors, with equality achieved by an equal norm tight frame construction; $k = 2$, $n = 4$ recovers $\sqrt{2}/2$, and $k = 2$, $n = 2$ recovers the trivial square answer $1/2$. See Book 1, field: Linear algebra, topic: trace identities and orthogonal projections bounding embedded balls.

Circuit Bank

- **WHEN** a shape's containment in a box must hold for all θ against an expression of the form $a\cos\theta + b\sin\theta$ **DO** rewrite it as the single inequality $\sqrt{a^2+b^2} \le$ the bound **BECAUSE** the amplitude identity turns an infinite family of constraints indexed by θ into one algebraic inequality, for free
- **WHEN** a centrally symmetric shape, a circle, a ball, an ellipse, is contained in a centrally symmetric region and its centre's position is one of the unknowns **DO** fix the shape's centre at the region's own centre by a short reflection argument before optimising anything else **BECAUSE** this removes an entire block of free parameters from the optimisation at no cost to the bound
- **WHEN** a maximisation is stated in n independent coordinates each individually bounded **DO** sum the per coordinate bounds and relate the total to the quantity you actually want, here $2r^2$ via orthogonality and equal length **BECAUSE** independent coordinate bounds add, turning a geometric optimisation into ordinary arithmetic
- **WHEN** you need to certify a bound is achieved **DO** construct the extremal object by spreading the available budget as unevenly as the equality case demands, here giving each of v_1, v_2 its full length in a disjoint pair of coordinates **BECAUSE** equality in a sum of independent bounds is easiest to arrange by making each bound tight on a separate, non overlapping part of the data

Backtesting

1. First useful cue: "circle", a two parameter rotation, meeting "contained in a hypercube", coordinate wise box constraints, a for all θ condition against independent per coordinate caps.
2. Method the expression suggested: rewrite the circle as $P + v_1 \cos\theta + v_2 \sin\theta$ and turn each coordinate's for all θ constraint into the amplitude bound $\sqrt{v_{1i}^2 + v_{2i}^2} \leq 1$.
3. Where you would hesitate, and the thought to reinforce: whether the circle's centre can really be assumed to sit at the hypercube's own centre without loss of generality; reinforce with "a circle is symmetric under reflection through its own centre, and a cube is symmetric under reflection through its own centre, so intersecting the cube with its own reflection through the circle's centre shows the translate to the origin still fits".
4. The trap the question sets: trying to visualise or directly parametrise the optimal circle in four dimensions by geometric intuition alone, rather than reducing containment to algebra coordinate by coordinate.
5. What makes you recognise it faster next time: any "for all θ, a linear combination of $\cos\theta$ and $\sin\theta$ stays bounded" condition should immediately become a single norm inequality via the amplitude identity.
6. The reusable decision rule: split a containment optimisation into a free recentring step, using central symmetry, followed by a per coordinate algebraic bound, and only then sum the bounds using whatever identity, orthogonality, equal length, relates the summed quantity to the thing being optimised.

Mirror drills

1. What is the largest radius of a circle contained in a two dimensional square of side length 1? *Answer:* $1/2$, by the identical per coordinate reduction with only two coordinates: for the side length 2 square, $2r^2 = \sum_{i=1}^{2}(v_{1i}^2 + v_{2i}^2) \leq 2$, so $r \leq 1$, achieved by $v_1 = (1, 0)$, $v_2 = (0, 1)$; halving for side length 1 gives $1/2$.
2. What is the largest radius of a circle contained in a three dimensional cube of side length 1, and does the flat, inscribed in a face circle of radius $1/2$ actually win? *Answer:* no: the true maximum is $\sqrt{6}/4 \approx 0.612$, strictly larger than $1/2$. For the side length 2 cube, take $v_1 = (1, -\frac{1}{2}, -\frac{1}{2})$, $v_2 = (0, \frac{\sqrt{3}}{2}, -\frac{\sqrt{3}}{2})$; these are orthogonal, $|v_1| = |v_2| = \sqrt{3/2}$, and $v_{1i}^2 + v_{2i}^2 = 1$ in every one of the three coordinates simultaneously, so $r = \sqrt{3/2}$ is achieved, giving $\sqrt{3/2}/2 = \sqrt{6}/4$ for side length 1, beating the naive coordinate plane circle by tilting.
3. Using the general formula $\frac{1}{2}\sqrt{n/k}$ for the largest k dimensional ball in an n dimensional unit hypercube, confirm it reproduces the answer to the main problem. *Answer:* $k = 2$, $n = 4$ gives $\frac{1}{2}\sqrt{4/2} = \frac{1}{2}\sqrt{2} = \sqrt{2}/2$, matching exactly.

18.10 B4. Lifting a polynomial bijection modulo prime powers

Let p be a prime number. Let $h(x)$ be a polynomial with integer coefficients such that $h(0), h(1), \ldots, h(p^2-1)$ are distinct modulo p^2. Show that $h(0), h(1), \ldots, h(p^3-1)$ are distinct modulo p^3.

Read it out loud

Say the structure to yourself in words before you write a single symbol. An integer polynomial is being evaluated on a complete residue system, and the hypothesis is that it separates every pair of residues at level two. You are asked to show that it then separates every pair at level three. So this is not a question about h; it is a question about the ladder p, p^2, p^3 and about how a polynomial map behaves when you climb one rung of that ladder. The fields wired into it are elementary number theory (congruences, the factor theorem over $\mathbb{Z}$) and the local structure of the rings $\mathbb{Z}/p^k\mathbb{Z}$.

The first cue is the word *distinct* attached to the list $h(0), \ldots, h(p^2-1)$. That list has exactly p^2 entries and there are exactly p^2 residue classes modulo p^2, so "distinct" is the same statement as "h induces a bijection of $\mathbb{Z}/p^2\mathbb{Z}$". Within thirty seconds you should have written that translation down, because injectivity is what the hypothesis hands you and surjectivity is what you will want to use. A finite set is the only place where you get both for free, and the examiners chose p^2 inputs deliberately so that you would notice.

The second cue is that the exponent moves up by exactly one. When a modulus climbs from p^k to p^{k+1}, the extra information sits in the kernel of the reduction map $\mathbb{Z}/p^{k+1}\mathbb{Z} \to \mathbb{Z}/p^k\mathbb{Z}$, which is a cyclic group of order p generated by p^k. So the real question is: what does h do to that kernel? A polynomial evaluated at $a+t$ where t is small in the p-adic sense is governed by its linear term in t, and the coefficient of that linear term is $h'(a)$. That is the whole problem. The derivative is not imported from calculus as a decoration; it is the induced map on the kernel, and it is forced on you the moment you write the increment as a separate variable.

Approaches, ranked worst to best

1. **Exhaust the residues.** *Cue:* "distinct modulo p^3" reads like a finite check. *Circuit:* compute when stuck. You cannot check anything: h has unspecified degree and coefficients, and p is arbitrary. Its only value is to force you to say out loud that the hypothesis must become a structural statement about h.

2. **Expand h coefficient by coefficient.** *Cue:* h is a polynomial, so write $h(x) = \sum c_i x^i$ and bash. *Circuit:* binomial theorem on each monomial. This works, since $(x+yp^2)^i \equiv x^i + ix^{i-1}yp^2 \pmod{p^3}$ is exactly what you need, but you rediscover the same expansion twice and carry sums over i through every line. Identical mathematics, worse bookkeeping.

3. **Rewrite h in the binomial basis.** *Cue:* integer valued polynomials are $\mathbb{Z}$ combinations of $\binom{x}{i}$. *Circuit:* change of basis to make integrality visible. Writing $h(x) = \sum_i d_i \binom{x}{i}$ with $d_i \in \mathbb{Z}$ makes the integrality of the higher Taylor coefficients transparent, and it is the right basis for many integrality problems, but here it buys nothing over one Taylor identity and you must prove the basis change first.

4. **Quote the classification of polynomial permutations of $\mathbb{Z}/p^k\mathbb{Z}$.** *Cue:* the problem is a corollary of a named theorem. *Circuit:* recognise the standard result and cite it. For $k \geq 2$, h permutes $\mathbb{Z}/p^k\mathbb{Z}$ if and only if it permutes $\mathbb{Z}/p\mathbb{Z}$ and $h'(a) \not\equiv 0 \pmod{p}$ for every integer a; the condition is independent of k, which settles the problem instantly. But the proof of that criterion is precisely the argument below, so cite it as a remark, never as the argument.

5. **Hensel style lifting of solutions.** *Cue:* lifting information from p^2 to p^3 is root lifting. *Circuit:* Hensel's lemma with a simple root. Recast the claim as: for each residue c the congruence

$h(z) \equiv c \pmod{p^3}$ has at most one solution. That is the correct conceptual home, and $p \nmid h'$ is Hensel's simple root condition, but the write up keeps translating between "solve $h(z) = c$" and "the values are distinct".

6. **First order Taylor with an integer remainder, then fibre by fibre.** *Cue:* increment plus base point; the modulus is a prime power. *Circuit:* $P(a+t) = P(a) + tP'(a) + t^2 \cdot \text{integer}$. One identity, proved once, used twice. Increment $t = p$ converts the hypothesis into $p \nmid h'$; increment $t = yp^2$ makes the map on each fibre of $\mathbb{Z}/p^3\mathbb{Z} \to \mathbb{Z}/p^2\mathbb{Z}$ affine with slope $h'(x)$, and an affine map with unit slope on a cyclic group of order p is a bijection. This is the route to write.

The ideal solution

Two lemmas first, both about a general polynomial P with integer coefficients.

Lemma 18.10.1. *For $P \in \mathbb{Z}[x]$ and integers u, v we have $(u - v) \mid P(u) - P(v)$. Consequently, if $u \equiv v \pmod{m}$ then $P(u) \equiv P(v) \pmod{m}$.*

Proof. By linearity it suffices to treat $P(x) = x^i$ with $i \geq 1$, and $u^i - v^i = (u - v)(u^{i-1} + u^{i-2}v + \cdots + v^{i-1})$ with the second factor an integer. The constant term contributes 0. □

Lemma 18.10.2. *For $P \in \mathbb{Z}[x]$ and integers a, t there is an integer R (depending on P, a, t) with*

$$P(a+t) = P(a) + t\,P'(a) + t^2 R.$$

Proof. Again by linearity take $P(x) = x^m$. The binomial theorem gives

$$(a+t)^m = a^m + ma^{m-1}t + t^2 \sum_{i=2}^{m} \binom{m}{i} a^{m-i} t^{i-2},$$

and the displayed sum is an integer (it is empty when $m \leq 1$). Summing over the monomials of P with their integer coefficients gives the claim, since the formal derivative is linear and sends x^m to mx^{m-1}. □

Step 1: the hypothesis says $p \nmid h'(x)$ for every integer x. Fix x with $0 \leq x \leq p-1$. Then $x + p \leq 2p - 1 \leq p^2 - 1$, using $p \geq 2$, so both x and $x + p$ lie in $\{0, 1, \ldots, p^2 - 1\}$ and the hypothesis gives

$$h(x+p) \not\equiv h(x) \pmod{p^2}.$$

Lemma 2 with $t = p$ gives $h(x+p) = h(x) + p\,h'(x) + p^2 R$, so $h(x+p) \equiv h(x) + p\,h'(x) \pmod{p^2}$. Combining, $p\,h'(x) \not\equiv 0 \pmod{p^2}$, that is, $h'(x) \not\equiv 0 \pmod{p}$. Now $h' \in \mathbb{Z}[x]$, so by Lemma 1 the residue $h'(x) \bmod p$ depends only on $x \bmod p$. Since every integer is congruent modulo p to one of $0, 1, \ldots, p-1$, we conclude

$$h'(x) \not\equiv 0 \pmod{p} \qquad \text{for every integer } x.$$

Step 2: the shape of h on a fibre. Let $0 \leq x \leq p^2 - 1$ and $0 \leq y \leq p - 1$. Lemma 2 with $t = yp^2$ gives

$$h(x + yp^2) = h(x) + yp^2 h'(x) + y^2 p^4 R.$$

Since $p^3 \mid y^2 p^4$, this yields the two congruences

$$h(x + yp^2) \equiv h(x) + yp^2 h'(x) \pmod{p^3}, \qquad h(x + yp^2) \equiv h(x) \pmod{p^2}.$$

Step 3: injectivity modulo p^3. Every $z \in \{0, 1, \ldots, p^3 - 1\}$ is written uniquely as $z = x + yp^2$ with $0 \leq x \leq p^2 - 1$ and $0 \leq y \leq p - 1$ (divide z by p^2 with remainder). Suppose

$$h(x_1 + y_1 p^2) \equiv h(x_2 + y_2 p^2) \pmod{p^3}$$

for two such pairs. Reducing modulo p^2 and using the second congruence of Step 2 gives $h(x_1) \equiv h(x_2) \pmod{p^2}$, and since $x_1, x_2 \in \{0, \dots, p^2-1\}$ the hypothesis forces $x_1 = x_2 =: x$. Feeding this back into the first congruence of Step 2 gives

$$y_1 p^2 h'(x) \equiv y_2 p^2 h'(x) \pmod{p^3},$$

hence $p^3 \mid p^2(y_1 - y_2)h'(x)$, hence $p \mid (y_1 - y_2)h'(x)$. By Step 1 the prime p does not divide $h'(x)$, so Euclid's lemma gives $p \mid y_1 - y_2$; and $|y_1 - y_2| \le p - 1 < p$ forces $y_1 = y_2$. Therefore $z_1 = z_2$, so the p^3 numbers $h(0), h(1), \dots, h(p^3-1)$ are pairwise distinct modulo p^3. ■

Layer 1: WHAT

Block 1. Translate the hypothesis: h induces a bijection of $\mathbb{Z}/p^2\mathbb{Z}$, and the target statement is that h induces a bijection of $\mathbb{Z}/p^3\mathbb{Z}$.

Block 2. Build one tool: for integer polynomials, $P(a+t) = P(a) + tP'(a) + t^2 \cdot (\text{integer})$, proved once by the binomial theorem.

Block 3. Set $t = p$ and read the hypothesis through the tool: $h(x+p) \not\equiv h(x) \pmod{p^2}$ becomes $p \nmid h'(x)$.

Block 4. Propagate: h' mod p depends only on x mod p, so $p \nmid h'(x)$ for all integers x, not just for $x < p$.

Block 5. Set $t = yp^2$ in the same tool. The quadratic remainder now carries p^4 and dies modulo p^3, leaving an affine formula with slope $h'(x)$.

Block 6. Split $z = x + yp^2$ and separate coordinates: the x coordinate is pinned by the hypothesis at level two, the y coordinate is pinned by $p \nmid h'(x)$.

The nominal route starts by fixing two integers $z_1 \not\equiv z_2 \pmod{p^3}$ and trying to prove $h(z_1) \not\equiv h(z_2)$ directly, splitting into the case $z_1 \equiv z_2 \pmod{p^2}$ and the case not. That is not wrong, but it invites you to expand $h(z_1) - h(z_2)$ as a difference of two unrelated evaluations and to lose the base point. The bifurcation point is Block 2: the decision to stop treating x as the variable and instead treat the *increment* t as the variable, with the base point a frozen. Once the increment is the variable, the linear coefficient in it is $h'(a)$ and the derivative appears without being invented.

Layer 2: HOW

Block 1. **Cue.** The list $h(0), \dots, h(p^2-1)$ has exactly as many entries as there are residues modulo p^2, and the word is "distinct". **Why natural.** On a finite set of matched size, injective, surjective and bijective are the same word. The hypothesis is therefore not a list of inequalities but a single structural fact, and structural facts survive being pushed through maps whereas lists of inequalities do not. **Link to the next block.** Once you say "bijection of $\mathbb{Z}/p^2\mathbb{Z}$" you immediately want the corresponding statement for $\mathbb{Z}/p^3\mathbb{Z}$, so you look at the map connecting the two rings.

Block 2. **Cue.** A prime power modulus and a step of exactly one in the exponent; equivalently, inputs that differ by something divisible by a high power of p. **Why natural.** If z and a differ by t with t divisible by a large power of p, then t^2 is divisible by twice that power, so every term of the expansion beyond the linear one is invisible modulo the relevant power. The only expansion of a polynomial in powers of an increment is the Taylor expansion, and for an integer polynomial its coefficients $P^{(i)}(a)/i!$ are integers. Two of the three facts you need are already in the binomial theorem, so prove exactly the piece you need and no more. **Link to the next block.** You now have a machine with a free parameter t. The natural thing is to try the two values of t that the problem itself names: p and p^2.

Block 3. **Cue.** The hypothesis contains pairs of inputs that differ by exactly p, namely x and $x+p$ for $x < p$, and both are inside the range $\{0, \dots, p^2-1\}$ where the hypothesis lives. **Why natural.** The hypothesis is a statement about the whole range, but you only need it on the pairs where the tool has bite. Feeding $t = p$ makes the remainder divisible by p^2, so it vanishes modulo p^2, and the hypothesis converts with no loss into a statement about h'. Notice what you did not do: you did not use the hypothesis on all $\binom{p^2}{2}$ pairs. You used it on p pairs and threw the rest away, because the rest carry no information about the derivative. **Link to the next block.** You have $p \nmid h'(x)$ for $x < p$ and you need it for $x < p^2$, so you look for a periodicity statement.

Block 4. **Cue.** h' has integer coefficients. **Why natural.** An integer polynomial evaluated modulo m depends only on the input modulo m, because $(u-v)$ divides $P(u) - P(v)$. This is the single most reused fact in congruence problems and it must be quoted, not rediscovered. **Link to the next block.** The invariant $p \nmid h'$ is now global, so it can be applied at every base point in Step 2 without further work.

Block 5. **Cue.** You must reach modulus p^3, and the inputs $0, \dots, p^3-1$ naturally split by their residue modulo p^2. **Why natural.** The set $\{0, \dots, p^3-1\}$ fibres over $\{0, \dots, p^2-1\}$ with fibres of size p, and a point of the fibre over x is $x + yp^2$. Feed that increment into the tool: $t^2 = y^2p^4$ is divisible by p^3, so the remainder dies and the restriction of h to the fibre is the affine map $y \mapsto h(x) + yp^2h'(x)$. Affine maps on a cyclic group of order p are bijections exactly when the slope is a unit, and Step 1 says the slope is a unit. **Link to the next block.** Bijectivity on the base plus bijectivity on each fibre is bijectivity upstairs, so the two coordinates can be pinned one after the other.

Block 6. **Cue.** Two inputs with equal images; you want to force them equal. **Why natural.** Reduce the equality modulo p^2 first, because that projection kills the y coordinate entirely and leaves a statement to which the hypothesis applies verbatim. Only after $x_1 = x_2$ is fixed does the y coordinate become a clean linear congruence. Order matters: doing it the other way round leaves two unknowns in one congruence. **Link to the next block.** Nothing follows; the proof is complete, and the final line should say the count, p^3 distinct values, so the grader sees the conclusion in the form the problem asked for.

Layer 3: WHY

The mathematical truth underneath is that $\mathbb{Z}/p^3\mathbb{Z} \to \mathbb{Z}/p^2\mathbb{Z}$ is a surjective ring homomorphism whose kernel is $p^2\mathbb{Z}/p^3\mathbb{Z}$, a cyclic group of order p. Any integer polynomial descends along this map, so h on $\mathbb{Z}/p^3\mathbb{Z}$ sits over h on $\mathbb{Z}/p^2\mathbb{Z}$, and for a map of finite sets respecting a fibration, bijectivity is equivalent to bijectivity of the base map together with bijectivity of every fibre map. The base map is a bijection by hypothesis, and on the fibre over x Lemma 2 says h acts by multiplication by $h'(x)$. So h is a bijection modulo p^3 precisely when $h'(x)$ is a unit modulo p for every x. The derivative is the induced map on the kernel of a reduction map; that is why calculus notation appears in a problem with no limits in it.

The discernment mechanism: when a problem hands you data at one level of a filtration and asks for data at the next, identify the map between levels, identify its kernel, and compute the induced action on the kernel. That action is always a linearisation, and for polynomial maps the linearisation is the formal derivative. Train it by rehearsing the same picture in Hensel's lemma, in Newton's method (the same formula over $\mathbb{R}$ with the same slope condition), and in the structure of $1 + p\mathbb{Z}_p$.

The examiners nudged you twice: they chose the exponents 2 and 3 rather than 1 and 2, so the base case modulo p is handed to you, and they wrote the hypothesis as a list of p^2 values rather

than as "permutes the residues", so that the translation is yours to perform.

Revolving concepts

Complete residue system and the pigeonhole equivalence. A list of m integers that are pairwise distinct modulo m is a complete residue system modulo m. For any map of finite sets of equal size, injective, surjective and bijective coincide. See Book 1, field: Number theory, topic: complete residue systems.

The factor theorem over $\mathbb{Z}$. For $P \in \mathbb{Z}[x]$ and integers u, v, $(u - v) \mid P(u) - P(v)$. Corollaries used constantly: P modulo m is a well defined function on $\mathbb{Z}/m\mathbb{Z}$; if P has a root modulo m then it has one in $\{0, \ldots, m-1\}$; and the order of $P(a) \bmod p$ is unchanged when a is shifted by a multiple of p. See Book 1, field: Number theory, topic: polynomial congruences.

Formal derivative and integer Taylor coefficients. For $P \in \mathbb{Z}[x]$ the formal derivative P' is again in $\mathbb{Z}[x]$, the product and chain rules hold formally, and $P(a+t) = \sum_{i \geq 0} \frac{P^{(i)}(a)}{i!} t^i$ with every coefficient an integer. Those coefficients are the Hasse derivatives, and their integrality is equivalent to $\binom{x}{i}$ being an integer valued polynomial. You need only the first two terms plus an integer remainder, which the binomial theorem supplies directly. See Book 1, field: Polynomials, topic: formal derivative and Hasse derivatives.

Hensel's lemma. If $P \in \mathbb{Z}[x]$, $k \geq 1$, $P(a) \equiv 0 \pmod{p^k}$ and $p \nmid P'(a)$, then there is exactly one b modulo p^{k+1} with $b \equiv a \pmod{p^k}$ and $P(b) \equiv 0 \pmod{p^{k+1}}$: write $b = a + tp^k$, expand with Lemma 2, and solve $P(a)/p^k + tP'(a) \equiv 0 \pmod{p}$, uniquely soluble because $P'(a)$ is invertible modulo p. The present problem is the same computation with the root replaced by a general value. See Book 1, field: Number theory, topic: Hensel lifting.

Polynomial permutations of $\mathbb{Z}/p^k\mathbb{Z}$ (Kempner's criterion). For $k \geq 2$, $h \in \mathbb{Z}[x]$ permutes $\mathbb{Z}/p^k\mathbb{Z}$ if and only if h permutes $\mathbb{Z}/p\mathbb{Z}$ and $h'(a) \not\equiv 0 \pmod{p}$ for every integer a. The criterion does not mention k, so permuting for one $k \geq 2$ is the same as permuting for all; the problem is the case $k = 2 \Rightarrow k = 3$. See Book 1, field: Number theory, topic: polynomial permutations of finite rings.

The recursion, generalised. Let B_k be "h induces a bijection of $\mathbb{Z}/p^k\mathbb{Z}$". The argument proves $B_k \Rightarrow B_{k+1}$ for every $k \geq 2$ with the single invariant $p \nmid h'(a)$ for all a, extracted once from B_2 and never recomputed. The general pattern is a recursion over the levels of a filtration $R_{k+1} \twoheadrightarrow R_k$ with kernel K, in which h acts on K by its linearisation. The method for every such level to level recursion is identical: find the kernel, linearise, demand invertibility of the linearisation, and check that the invertibility condition is level independent so a single induction closes all levels. The same skeleton runs the lifting of solutions of $x^2 \equiv a \pmod{2^k}$, the Teichmuller lifts, and the p-adic logarithm.

Circuit Bank

- **WHEN** a polynomial with integer coefficients is evaluated at $0, 1, \ldots, m-1$ and you are told the values are distinct modulo m **DO** rewrite it at once as "h induces a bijection of $\mathbb{Z}/m\mathbb{Z}$" **BECAUSE** the counts match, so you gain surjectivity for free, and a bijection is a structural object you can push through ring maps while a list of inequalities is not
- **WHEN** a modulus climbs from p^k to p^{k+1}, or two inputs differ by a multiple of a high power of p **DO** freeze a base point a, call the difference t, and write $P(a+t) = P(a) + tP'(a) + t^2 \cdot (\text{integer})$ **BECAUSE** the square of the increment carries twice the p power, so every term past the linear one vanishes and the map on each fibre is affine with slope $P'(a)$
- **WHEN** you have proved a congruence statement about $P(x) \bmod p$ for x in some short range and need it for all x **DO** quote $(u - v) \mid P(u) - P(v)$ and conclude that $P \bmod p$ is a function of $x \bmod p$ **BECAUSE** an integer polynomial descends to every quotient ring, so any statement about it modulo p is automatically p periodic in the input

- **WHEN** a bijection is wanted on a set that fibres over another set **DO** prove the base map is a bijection, then prove each fibre map is a bijection, and order your equalities so that the projection is taken first **BECAUSE** projecting kills one coordinate entirely, leaving a one variable congruence instead of a two variable one, which is the difference between three lines and half a page
- **WHEN** an affine congruence $y \mapsto c + \lambda y$ is to be a bijection modulo p **DO** check only that $p \nmid \lambda$ **BECAUSE** multiplication by a unit permutes a cyclic group, and Euclid's lemma is exactly the statement that non zero residues modulo a prime are units

Backtesting

1. First useful cue: p^2 values, distinct modulo p^2. Equal counts on both sides means the hypothesis is "bijection of $\mathbb{Z}/p^2\mathbb{Z}$", and that is the form in which it can be used twice.
2. Method the expression suggested: an increment expansion $P(a+t) = P(a) + tP'(a) + t^2R$ with integer R, because the modulus is a prime power and the two natural increments, p and yp^2, are exactly the ones that kill the remainder.
3. Where you would hesitate, and the thought to reinforce: at the moment the derivative appears, since the problem has no calculus in it. Reinforce that h' here is the induced map on the kernel of $\mathbb{Z}/p^{k+1}\mathbb{Z} \to \mathbb{Z}/p^k\mathbb{Z}$, not a limit; the formal derivative of an integer polynomial is an algebraic object and it is the correct linearisation over any ring.
4. The trap the question sets: proving $p \nmid h'(x)$ only for $0 \le x \le p-1$ and then silently using it for $0 \le x \le p^2 - 1$ in Step 2. The gap is real and it is cheap to close with the factor theorem, but a grader who is counting will notice it missing.
5. What makes you recognise it faster next time: the phrase "prime power modulus, exponent increases by one" should fire the increment expansion before you have finished reading the sentence. Any problem in which p^k and p^{k+1} both appear is a linearisation problem.
6. The reusable decision rule: to lift a bijectivity, injectivity or solvability statement one level up a p adic ladder, linearise on the kernel and demand the slope be a unit. Extract the slope condition once at the lowest available level, then never recompute it.

Mirror drills

1. Show that if $h \in \mathbb{Z}[x]$ satisfies the hypothesis of the problem, then $h(0), h(1), \ldots, h(p^k - 1)$ are distinct modulo p^k for every $k \ge 1$. *Answer:* the invariant $p \nmid h'$ from Step 1 is independent of k; run Steps 2 and 3 with $t = yp^k$ to get the inductive step $B_k \Rightarrow B_{k+1}$ for $k \ge 2$, and for $k = 1$ note that a bijection modulo p^2 forces a bijection modulo p by counting the p elements in each fibre.
2. Let p be an odd prime. Show that $h(x) = x^2$ does not satisfy the hypothesis, and find every h of the form $x^2 + cx$ that does. *Answer:* $h'(x) = 2x$ vanishes modulo p at $x = 0$, so $h(0) \equiv h(p) \pmod{p^2}$; for $x^2 + cx$ we need $2x + c \not\equiv 0 \pmod{p}$ for all x, which is impossible for odd p, so no such h exists (for $p = 2$ the condition becomes c odd).
3. Prove Hensel's lemma in the simple root case: if $P \in \mathbb{Z}[x]$, $P(a) \equiv 0 \pmod{p^k}$ and $p \nmid P'(a)$, then there is exactly one b modulo p^{k+1} with $b \equiv a \pmod{p^k}$ and $P(b) \equiv 0 \pmod{p^{k+1}}$. *Answer:* put $b = a + tp^k$, expand with Lemma 2 so that the remainder carries p^{2k}, and solve the linear congruence $P(a)/p^k + tP'(a) \equiv 0 \pmod{p}$, unique because $P'(a)$ is invertible modulo p.

18.11 B5. Denominator preserving smooth maps are signed integer translations

Find all continuously differentiable functions $f : \mathbb{R} \to \mathbb{R}$ such that for every rational number q, the number $f(q)$ is rational and has the same denominator as q. (The denominator of a rational number q is the unique positive integer b such that $q = a/b$ for some integer a with $\gcd(a, b) = 1$.)

Read it out loud

Say the shape of it: a hypothesis that is purely arithmetic, imposed only on the rationals, together with a hypothesis that is purely analytic, imposed on all of $\mathbb{R}$. Nothing in the problem connects the two. Your entire job is to build the connector. The fields wired in are real analysis (the mean value theorem, continuity, connectedness) and elementary number theory (denominators, coprimality, lattices in $\mathbb{R}$).

The first cue is the phrase "the same denominator". That is a very rigid demand, and the honest first move is to notice that it is really two demands stacked on top of each other: the denominator of $f(q)$ *divides* the denominator of q, and no factor is lost. The divisibility half is the useful half, because "denominator divides m" is the same as "lies in $\frac{1}{m}\mathbb{Z}$", and $\frac{1}{m}\mathbb{Z}$ is a subgroup of $(\mathbb{R}, +)$. The set of rationals with denominator exactly b is not closed under subtraction; the lattice $\frac{1}{m}\mathbb{Z}$ is. Difference quotients subtract. So within thirty seconds you should have written

$$f\left(\tfrac{1}{m}\mathbb{Z}\right) \subseteq \tfrac{1}{m}\mathbb{Z} \qquad \text{for every positive integer } m,$$

because that is the version of the hypothesis a difference quotient can eat.

The second cue is "continuously differentiable" rather than "differentiable". The setters are telling you that the final step is a continuity argument on f', not on f. Once you know that f' takes integer values somewhere dense, continuity plus the fact that $\mathbb{Z}$ is closed and totally disconnected does everything else: a continuous function from a connected set into a discrete set is constant. So the target of the whole middle of the proof is the single sentence "f' takes integer values on a dense set".

Approaches, ranked worst to best

1. **Answer with examples only.** *Cue:* "find all" invites a list. *Circuit:* pattern match and stop. Writing $f(x) = \pm x + n$ with a verification is worth almost nothing on its own, because the entire content of the problem is that there is nothing else. It is still the correct first line of the write up; it is a catastrophic last line.

2. **Build f denominator by denominator.** *Cue:* the hypothesis is indexed by b, so induct on b. *Circuit:* induction on a discrete parameter. Pin f on the integers, then on the halves, then on the thirds. It dies immediately: continuity alone does not pin f on $\mathbb{Z}$, an uncountable set of continuous functions maps $\mathbb{Z}$ into $\mathbb{Z}$, and you never obtain any statement about irrational inputs, which is where "find all" actually lives.

3. **Treat it as a functional equation.** *Cue:* "find all f" is Cauchy shaped. *Circuit:* additivity, then $\mathbb{Q}$ linearity, then continuity. There is no additive hypothesis. The condition relates f to the arithmetic of its argument, not to f at other points, so no substitution generates a relation between $f(x + y)$, $f(x)$ and $f(y)$. This is the classic misfire: matching the words "find all continuous f" rather than the structure.

4. **Difference quotient at a fixed rational.** *Cue:* differentiability is a limit of quotients, so choose a sequence of rationals converging to q. *Circuit:* make the quotient an integer by design. Fix $q = a/b$ and set $q_j = q + \frac{1}{bj}$; when $b \mid j$ this is in lowest terms with denominator bj, so

$bj(f(q_j) - f(q))$ is an integer, and its limit $f'(q)$ is therefore an integer that the sequence attains for large j. Then squeeze $f'(q)$ to ± 1 by inspecting the denominator of $f(q_j)$. This works, it is the route the official solution takes, and it is entirely correct; it costs an extra coprimality lemma and an extra "a convergent integer sequence is eventually constant" lemma, and it forces the value ± 1 at every rational separately when one global argument would do.

5. **Lattice preservation plus the mean value theorem.** *Cue:* "denominator divides m" is "lies in $\frac{1}{m}\mathbb{Z}$"; two lattice points sit a distance $\frac{1}{m}$ apart. *Circuit:* turn a secant slope into a derivative value. For $t \in \frac{1}{m}\mathbb{Z}$ the numbers $f(t)$ and $f(t + \frac{1}{m})$ both lie in $\frac{1}{m}\mathbb{Z}$, so the secant slope over that step is an integer, and the mean value theorem converts it into a value of f'. Those points are dense as m grows, so f' is integer valued on a dense set, hence everywhere, hence constant. Now f is affine with integer slope, and the exactness clause tested at $q = 1/b$ forces the slope to be a unit. No coprimality lemma, no sequence lemma, and the arithmetic is used exactly twice. This is the route to write.

The ideal solution

The answer is

$$\boxed{f(x) = x + n \quad \text{and} \quad f(x) = -x + n, \qquad n \in \mathbb{Z}.}$$

These work. Let $q = a/b$ in lowest terms with $b \geq 1$, and let $n \in \mathbb{Z}$. Then $\pm q + n = (\pm a + nb)/b$ and $\gcd(\pm a + nb, b) = \gcd(\pm a, b) = 1$, so $\pm q + n$ is in lowest terms with denominator b. Both functions are continuously differentiable. So every member of the stated family satisfies the condition.

Now suppose f satisfies the condition; we show it lies in the family.

Lemma 18.11.1. *For every positive integer m, the set of rationals whose denominator divides m is exactly $\frac{1}{m}\mathbb{Z}$, and $f\left(\frac{1}{m}\mathbb{Z}\right) \subseteq \frac{1}{m}\mathbb{Z}$.*

Proof. If $q = a/b$ in lowest terms and $b \mid m$ then $q = a(m/b)/m \in \frac{1}{m}\mathbb{Z}$; conversely if $q = u/m$ then cancelling gives $q = (u/g)/(m/g)$ with $g = \gcd(u, m)$, in lowest terms, and m/g divides m. For the second statement, let $q \in \frac{1}{m}\mathbb{Z}$ and let b be its denominator, so $b \mid m$. By hypothesis $f(q)$ is rational with denominator b, so $b \mid m$ again, so $f(q) \in \frac{1}{m}\mathbb{Z}$. □

Step 1: f' takes integer values on a dense set. Fix a positive integer m and an element $t \in \frac{1}{m}\mathbb{Z}$. Both t and $t + \frac{1}{m}$ lie in $\frac{1}{m}\mathbb{Z}$, so by the Lemma both $f(t)$ and $f(t + \frac{1}{m})$ lie in $\frac{1}{m}\mathbb{Z}$, whence

$$m\left(f\left(t + \tfrac{1}{m}\right) - f(t)\right) \in \mathbb{Z}.$$

Since f is differentiable on $\mathbb{R}$ it is continuous on $[t, t + \frac{1}{m}]$ and differentiable on $(t, t + \frac{1}{m})$, so the mean value theorem supplies $\xi \in (t, t + \frac{1}{m})$ with

$$f'(\xi) = \frac{f(t + \frac{1}{m}) - f(t)}{\frac{1}{m}} = m\left(f\left(t + \tfrac{1}{m}\right) - f(t)\right) \in \mathbb{Z}.$$

Let $Z = \{x \in \mathbb{R} : f'(x) \in \mathbb{Z}\}$. We have just shown that Z meets the open interval $(t, t + \frac{1}{m})$ for every $m \geq 1$ and every $t \in \frac{1}{m}\mathbb{Z}$. Given any nonempty open interval I, choose m with $\frac{1}{m} < \frac{1}{2}\,\text{length}(I)$; the points of $\frac{1}{m}\mathbb{Z}$ are spaced $\frac{1}{m}$ apart, so some $t \in \frac{1}{m}\mathbb{Z}$ has $[t, t + \frac{1}{m}] \subseteq I$, and then $Z \cap I \neq \emptyset$. Hence Z is dense in $\mathbb{R}$.

Step 2: f' is a constant integer. The function f' is continuous by hypothesis, $\mathbb{Z}$ is a closed subset of $\mathbb{R}$, and $f'(Z) \subseteq \mathbb{Z}$ with Z dense. For any $x \in \mathbb{R}$ pick $x_i \in Z$ with $x_i \to x$; then $f'(x) = \lim_i f'(x_i)$ lies in $\mathbb{Z}$ because $\mathbb{Z}$ is closed. So $f'(\mathbb{R}) \subseteq \mathbb{Z}$. Now $\mathbb{R}$ is connected and f' is continuous, so $f'(\mathbb{R})$ is a connected subset of $\mathbb{Z}$, that is, a single point. (Equivalently: if f' took two distinct integer values it

would take a non integer value in between, by the intermediate value theorem.) Write $f'(x) = k$ for all x, with $k \in \mathbb{Z}$.

Step 3: identify f. From $f' \equiv k$ we get $f(x) = kx + n$ where $n = f(0)$. The number 0 has denominator 1, so $f(0)$ has denominator 1, that is, $n \in \mathbb{Z}$.

Step 4: the exactness clause forces $k = \pm 1$. Let $b \geq 1$ be any integer. The rational $1/b$ has denominator b, and

$$f\left(\tfrac{1}{b}\right) = \frac{k + nb}{b}.$$

For a fraction u/v with $v \geq 1$ the denominator is $v/\gcd(u, v)$, so the denominator of $f(1/b)$ is $b/\gcd(k + nb, b) = b/\gcd(k, b)$, using $k + nb \equiv k \pmod{b}$. The hypothesis says this equals b, so $\gcd(k, b) = 1$ for every positive integer b. Taking $b = 2$ rules out $k = 0$, since $\gcd(0, 2) = 2$. If $|k| \geq 2$, taking $b = |k|$ gives $\gcd(k, |k|) = |k| \geq 2$, a contradiction. Hence $k = \pm 1$.

Therefore $f(x) = x + n$ or $f(x) = -x + n$ with $n \in \mathbb{Z}$, and both were verified above. ■

Remark 18.11.2. The hypothesis "continuously differentiable" is used exactly once, in Step 2, to spread integrality of f' from a dense set to all of $\mathbb{R}$. If you only assume f differentiable, Darboux's theorem still gives f' the intermediate value property, which replaces the connectedness argument at the end but not the density argument in the middle; so the weakening is not free, and you should not claim it.

Layer 1: WHAT

Block 1. Write down the family $f(x) = \pm x + n$ and verify it, so that "find all" has an answer on the page before the hard half begins.

Block 2. Split the hypothesis: "denominator of $f(q)$ divides m whenever the denominator of q divides m" is the group theoretic half, $f(\frac{1}{m}\mathbb{Z}) \subseteq \frac{1}{m}\mathbb{Z}$; "no cancellation" is the rigidity half, saved for the very end.

Block 3. Two neighbouring lattice points are a distance $\frac{1}{m}$ apart and their images lie in the same lattice, so the secant slope over that step is an integer.

Block 4. The mean value theorem converts that integer secant slope into an actual value of f', at an unknown point inside a known short interval.

Block 5. Shrink the mesh: the set where f' is an integer meets every interval, so it is dense.

Block 6. Continuity of f' plus closedness of $\mathbb{Z}$ gives $f'(\mathbb{R}) \subseteq \mathbb{Z}$; connectedness of $\mathbb{R}$ collapses that to a single integer k, so $f(x) = kx + n$ with $n = f(0) \in \mathbb{Z}$.

Block 7. Reintroduce the rigidity half at the single test point $q = 1/b$: the denominator of $(k + nb)/b$ is $b/\gcd(k, b)$, so $\gcd(k, b) = 1$ for all b, so $k = \pm 1$.

The nominal route fixes a rational q, builds a sequence $q_j \to q$ of rationals with growing denominators, and grinds out $f'(q) = \pm 1$ one rational at a time, paying for a coprimality lemma about $\frac{aj+1}{bj}$ and for the observation that a convergent integer sequence is eventually constant. It reaches the same place. The bifurcation point is Block 2: whether you keep the hypothesis in its "exactly b" form, which is not stable under subtraction and therefore forces you to fight coprimality every time you form a difference, or weaken it at once to the lattice form, which is a subgroup and therefore free under subtraction. Weakening a hypothesis to make it structural, and paying the lost strength back once at the end, is the move.

Layer 2: HOW

Block 1. Cue. "Find all" with a condition that is obviously satisfied by something. **Why natural.** You cannot aim a uniqueness argument at an unknown target. Verifying the family first also tells you what the final constant must be constrained to, so Step 4 has a known destination. **Link to the next block.** Having the answer $\pm x + n$ tells you that the derivative is the object to pin down, since the family is exactly "derivative ± 1, intercept an integer".

Block 2. **Cue.** The word "denominator", which is a divisibility statement in disguise, sitting next to a differentiability hypothesis, which needs subtraction. **Why natural.** Difference quotients subtract two values of f. The set of rationals with denominator exactly b is not closed under subtraction ($\frac{1}{3}$ and $\frac{2}{3}$ both have denominator 3, their difference has denominator 3, but $\frac{1}{3}$ and $\frac{4}{3}$ differ by an integer). The set with denominator dividing m is exactly $\frac{1}{m}\mathbb{Z}$, a subgroup, so it is closed under subtraction by construction. Weaken the hypothesis to the closed version; you can always spend the strength you dropped at the end, at one convenient point. **Link to the next block.** A subgroup of $\mathbb{R}$ preserved by f has a mesh, and a mesh gives you a canonical increment to differentiate along.

Block 3. **Cue.** The lattice $\frac{1}{m}\mathbb{Z}$ has mesh $\frac{1}{m}$, and $\frac{1}{m}$ is exactly the reciprocal of the number multiplying the values. **Why natural.** If $f(t)$ and $f(t+\frac{1}{m})$ both lie in $\frac{1}{m}\mathbb{Z}$, their difference lies in $\frac{1}{m}\mathbb{Z}$ and dividing by $\frac{1}{m}$ lands in $\mathbb{Z}$. The increment and the lattice mesh are the same number, so the quotient is an integer with nothing to check. This is the whole bridge between the arithmetic and the analysis, and it is one line. **Link to the next block.** You now hold an integer that looks like a derivative but is only a secant slope, so you need a theorem that turns secant slopes into derivative values.

Block 4. **Cue.** A secant slope over a short interval, with f differentiable everywhere. **Why natural.** The mean value theorem is precisely the statement "every secant slope is attained by the derivative". You do not care where it is attained, only that it is attained somewhere inside a controlled interval, because the next step only needs density. **Link to the next block.** The interval has length $\frac{1}{m}$ and m is yours to choose, so the located points can be made to fall inside any prescribed interval.

Block 5. **Cue.** m is a free parameter and $\frac{1}{m} \to 0$. **Why natural.** Density is not a limit statement, it is a "meets every interval" statement, and the lattice $\frac{1}{m}\mathbb{Z}$ with fine mesh has a full step inside any given interval. Notice what you did not do: you did not try to compute $\lim_m$ of anything, because you do not need the value of f' at any named point, only its membership of $\mathbb{Z}$ on a dense set. **Link to the next block.** Membership of a closed set on a dense set is exactly the input that continuity upgrades.

Block 6. **Cue.** "Continuously differentiable" in the hypothesis; $\mathbb{Z}$ is closed and has no interval in it. **Why natural.** Continuity carries a dense set to a dense subset of the image, and a closed set absorbs limits, so $f'(\mathbb{R}) \subseteq \overline{\mathbb{Z}} = \mathbb{Z}$. Then a continuous map from the connected set $\mathbb{R}$ into $\mathbb{Z}$ must be constant, since a connected subset of $\mathbb{Z}$ is a point; the intermediate value theorem is the same statement in exam language. **Link to the next block.** A constant derivative integrates to an affine function, so f is now completely known up to two integers.

Block 7. **Cue.** You have not yet used "the same denominator", only "the denominator divides". **Why natural.** The lost strength has to be repaid, and it is cheapest to repay it at a single family of test points. Choosing $q = 1/b$ is optimal because its numerator is 1, so the greatest common divisor computation collapses to $\gcd(k, b)$ with nothing else in the way. Requiring $\gcd(k, b) = 1$ for every b says k is a unit of $\mathbb{Z}$. **Link to the next block.** Nothing follows; state both families and confirm the verification from Block 1 closes the "if and only if".

Layer 3: WHY

The mathematical truth underneath is that $\frac{1}{m}\mathbb{Z}$ is a discrete subgroup of $(\mathbb{R}, +)$ and that these subgroups, ordered by divisibility of m, refine to a dense set. A C^1 map preserving a discrete subgroup of mesh $\frac{1}{m}$ has all of its mesh step secant slopes in $\mathbb{Z}$, and the mean value theorem transfers that integrality from secants to the derivative itself. Because $\mathbb{Z}$ is closed and totally disconnected while $\mathbb{R}$ is connected, integrality of a continuous function forces constancy. So the

analysis does not determine the constant; it only proves that a constant exists. The arithmetic then determines which constant: multiplication by k preserves denominators exactly if and only if $\gcd(k, b) = 1$ for every b, that is, if and only if k is a unit of $\mathbb{Z}$.

That two part anatomy is the general shape of rigidity theorems: discreteness plus connectedness produces a constant, and a separate arithmetic condition cuts the constant down to a unit. The same skeleton runs elsewhere: a continuous homomorphism $\mathbb{R} \to \mathbb{R}$ carrying $\mathbb{Z}$ into $\mathbb{Z}$ is $x \mapsto kx$ with $k \in \mathbb{Z}$, a continuous self map of the circle $\mathbb{R}/\mathbb{Z}$ has an integer degree that cannot change under deformation, and a C^1 map of $\mathbb{R}^n$ carrying every $\frac{1}{m}\mathbb{Z}^n$ into itself is affine with matrix in $M_n(\mathbb{Z})$.

The discernment mechanism: when a hypothesis is arithmetic and pointwise but the conclusion is analytic and global, the connector is always a difference quotient, and the design question is which increment makes the quotient arithmetically clean. Do not pick the increment for analytic convenience; pick it so that the numerator and the denominator have the same arithmetic scale. Ask, before writing any limit, what the common denominator of the two values you are subtracting is, and whether your increment can be its reciprocal.

The setters nudged you twice. They defined "denominator" explicitly, a hint that the split between "divides" and "is exactly" is where the work is, and they wrote "continuously differentiable" rather than "differentiable", which tells you the endgame is a topological argument about f' rather than a further computation with f.

Revolving concepts

Denominator of a rational. For integers u and $v \geq 1$, the fraction u/v in lowest terms has denominator $v/\gcd(u, v)$. Two facts get used constantly: $u \equiv u' \pmod{v}$ implies $\gcd(u, v) = \gcd(u', v)$, and the denominator of a sum divides the least common multiple of the denominators. See Book 1, field: Number theory, topic: fractions and coprimality.

The lattice $\frac{1}{m}\mathbb{Z}$. It equals $\{q \in \mathbb{Q} : \text{denominator of } q \text{ divides } m\}$, is a subgroup of $(\mathbb{R}, +)$ of mesh $\frac{1}{m}$, and satisfies $\frac{1}{m}\mathbb{Z} \subseteq \frac{1}{jm}\mathbb{Z}$ for every $j \geq 1$. The union over all m is $\mathbb{Q}$. By contrast the set of rationals with denominator exactly m is not a group. See Book 1, field: Number theory, topic: lattices and discrete subgroups.

Discrete subgroups of $\mathbb{R}$. Every subgroup of $(\mathbb{R}, +)$ is either dense or equal to $\alpha\mathbb{Z}$ for a unique $\alpha \geq 0$. This dichotomy is the reason "preserves a discrete subgroup" is a rigid condition, and it is the standard tool behind equidistribution and three distance arguments. See Book 1, field: Real analysis, topic: subgroups of the reals.

Mean value theorem and Darboux's theorem. The mean value theorem says every secant slope of a differentiable function is attained as a derivative value; Darboux's theorem says the derivative of a differentiable function has the intermediate value property even when it is not continuous. The first is the bridge in Step 1; the second is what would replace Step 2's connectedness argument if the hypothesis were weakened, though the density argument would still need repair. See Book 1, field: Real analysis, topic: mean value theorems.

Continuous maps into discrete sets. If g is continuous on a connected set and g takes values in a set with no interior points and no interval, such as $\mathbb{Z}$, then g is constant, because the image is connected. Combined with "closed sets absorb limits along dense sets", this is the standard two line rigidity engine. See Book 1, field: Real analysis, topic: connectedness.

The recursion, generalised. The recursion here is refinement of lattices: $G_1 \subseteq G_2 \subseteq G_6 \subseteq \cdots$ along $G_m = \frac{1}{m}\mathbb{Z}$ ordered by divisibility, with mesh $\frac{1}{m} \to 0$. The general statement is this. Let f be differentiable on $\mathbb{R}$ and let (G_m) be subgroups $\frac{1}{m}\mathbb{Z}$ with $f(G_m) \subseteq G_m$ for all m in an infinite set S of positive integers with unbounded elements. Then every mesh step secant slope over G_m lies in $\mathbb{Z}$, so by the mean value theorem f' is integer valued on a set meeting every interval of length greater than $2/\min\{m \in S, m > M\}$ for each M, hence on a dense set. If moreover f' is continuous then f' is a constant integer and f is affine. The same statement holds verbatim in $\mathbb{R}^n$ with $\frac{1}{m}\mathbb{Z}^n$ and the derivative matrix, giving an affine map with matrix in $M_n(\mathbb{Z})$; and if the map is required

to preserve denominators exactly in each coordinate, the matrix must be invertible over $\mathbb{Z}$. This is the general method for any "preserves all refinements of a lattice" hypothesis: read the mesh step quotient, apply the mean value theorem, use density and continuity, then repay the strength you dropped.

Circuit Bank

- **WHEN** a hypothesis constrains f only on $\mathbb{Q}$ or another countable set, while the regularity hypothesis is on all of $\mathbb{R}$ **DO** force the arithmetic into a difference quotient, and choose the increment so that the quotient is an integer by construction **BECAUSE** the derivative is the only object that reads both the local analysis and the discrete values, and an integer valued quantity in a limit is far more rigid than a real one
- **WHEN** a set defined by "denominator equals b" appears **DO** split it into "denominator divides m", which is the subgroup $\frac{1}{m}\mathbb{Z}$, and the leftover coprimality, and use only the subgroup until the last step **BECAUSE** difference quotients subtract, subgroups are closed under subtraction, and sets defined by exact denominators are not
- **WHEN** you hold an integer secant slope and want a statement about f' **DO** apply the mean value theorem and keep only the fact that the point exists inside a short interval **BECAUSE** you never need to know where the derivative attains the slope, only that the attaining points are dense, and density is cheaper than convergence
- **WHEN** a continuous function takes values in $\mathbb{Z}$, or in any closed set with empty interior, on a dense set **DO** conclude the values lie in that set everywhere, and if the set is discrete and the domain is an interval, conclude the function is constant **BECAUSE** closed sets absorb limits along dense sequences, and a continuous image of a connected set is connected, so it cannot straddle two isolated points
- **WHEN** you deliberately weakened a hypothesis to make it structural **DO** repay the lost strength at the very end, at one cheap test point chosen so the algebra collapses **BECAUSE** a single well chosen evaluation, here $q = 1/b$ with numerator 1, replaces a general coprimality argument that would otherwise be repeated at every rational

Backtesting

1. First useful cue: "the same denominator" read as "denominator divides m", giving $f(\frac{1}{m}\mathbb{Z}) \subseteq \frac{1}{m}\mathbb{Z}$, the only form of the hypothesis stable under the subtraction inside a difference quotient.
2. Method the expression suggested: an integer secant slope over one lattice step, upgraded to a derivative value by the mean value theorem, then density plus continuity.
3. Where you would hesitate, and the thought to reinforce: at the choice to weaken the hypothesis, since throwing away strength feels wrong. Reinforce that you are not throwing it away, you are deferring it; the exactness clause is spent once, at $q = 1/b$, and buys the last constraint $\gcd(k, b) = 1$.
4. The trap the question sets: assuming $q + \frac{1}{n}$ has denominator bn for every n. It does not; $\frac{1}{3} + \frac{1}{2} = \frac{5}{6}$ is fine but $\frac{1}{3} + \frac{1}{6} = \frac{1}{2}$ is not. Anyone who takes the difference quotient route without the divisibility condition $b \mid j$ has a false lemma in the middle of the proof.
5. What makes you recognise it faster next time: the pairing of a countable arithmetic hypothesis with a differentiability hypothesis. That pairing has exactly one bridge, the difference quotient, and the only design decision is the increment.

6. The reusable decision rule: to exploit an arithmetic constraint analytically, choose the increment equal to the reciprocal of the common denominator of the two values being subtracted, so the quotient is an integer with nothing to verify.

Mirror drills

1. Find all continuously differentiable $f : \mathbb{R} \to \mathbb{R}$ such that for every rational q, the denominator of $f(q)$ divides the denominator of q. *Answer:* $f(x) = kx + n$ with $k, n \in \mathbb{Z}$. The proof is Steps 1 to 3 unchanged; Step 4 disappears because the exactness clause is exactly what was dropped.
2. Suppose $g : \mathbb{R} \to \mathbb{R}$ is differentiable at 0, $g(0) = 0$, and $g(1/m) \in \frac{1}{m}\mathbb{Z}$ for every positive integer m. Show $g'(0) \in \mathbb{Z}$. *Answer:* $m\,g(1/m)$ is an integer for every m and converges to $g'(0)$; a convergent sequence of integers is eventually constant, so the limit is that integer.
3. Let f be differentiable on $\mathbb{R}$ with $f'(x) \in \{-1, 1\}$ for every x, with no continuity assumed on f'. Show $f(x) = x + c$ or $f(x) = -x + c$. *Answer:* Darboux's theorem gives f' the intermediate value property, so it cannot take both values without taking 0; hence f' is constant, and integrate. This isolates precisely which use of "continuously" in the original problem is removable and which is not.

18.12 B6. Parity of banded permutations via determinants modulo two

Let n and k be positive integers. Say that a permutation σ of $\{1, 2, \ldots, n\}$ is k-limited *if $|\sigma(i) - i| \leq k$ for all i. Prove that the number of k-limited permutations of $\{1, 2, \ldots, n\}$ is odd if and only if $n \equiv 0$ or $1 \pmod{2k+1}$.*

Read it out loud

Read the structure: you are asked for the parity of a count of permutations, where the constraint on the permutation is local and of the form "$\sigma(i)$ must lie in a prescribed set of columns". Two words in that sentence are doing everything. *Parity* means you never have to compute the count, and it means that all signs are free, because $+1$ and -1 are the same element of $\mathbb{F}_2$. *Prescribed set of columns* means the constraint is exactly a $0/1$ matrix: put a 1 in position (i, j) when $\sigma(i) = j$ is allowed. The fields wired in are combinatorics (counting permutations with restricted positions), linear algebra over a finite field (determinant, rank, kernel), and linear recurrences (because the matrix is translation invariant along the diagonal).

The first cue is "odd if and only if". That is never an invitation to compute the number; it is an invitation to find a structure that pairs almost all the objects up. Two structures do that here: an involution on the set of permutations, or a determinant, since a determinant is a signed count and modulo 2 the signs disappear. Within thirty seconds you should have written

$$\#\{k\text{-limited } \sigma\} \;=\; \operatorname{per}(M_{n,k}) \;\equiv\; \det(M_{n,k}) \pmod 2, \qquad (M_{n,k})_{ij} = \begin{cases} 1 & |i-j| \leq k \\ 0 & \text{otherwise,} \end{cases}$$

because the Leibniz formula for the determinant is a sum over permutations of $\pm$ the same products, and $\pm 1 \equiv 1$ modulo 2. The combinatorial problem has become a linear algebra problem: for which n is the banded matrix $M_{n,k}$ invertible over $\mathbb{F}_2$?

The second cue is the modulus in the answer. Each interior row of $M_{n,k}$ has exactly $2k+1$ ones. The answer's modulus is $2k+1$. That is not a coincidence you discover at the end; it is a prediction you make at the start, and it tells you what kind of object controls the problem: something with period equal to the band width. In a translation invariant matrix, the object with that period is a linear recurrence satisfied by the vectors in the kernel.

Approaches, ranked worst to best

1. **Count k-limited permutations exactly.** *Cue:* "the number of permutations" sounds like a counting problem. *Circuit:* find a closed form, then reduce it. There is no usable closed form. For fixed k the count satisfies a linear recurrence in n, but its order grows exponentially in k, and k is a free parameter here. This route does not stall late; it stalls at the first attempt to write a formula, and its only use is to make you realise that you must never evaluate the count.

2. **Transfer matrix dynamic programming modulo two.** *Cue:* the constraint is local, so scan left to right and carry a state. *Circuit:* profile dynamic programming. Correct in principle: the state is which of the last $2k$ positions have been used, so the transfer matrix has size $2^{O(k)}$ over $\mathbb{F}_2$, and the claim becomes a statement about the order of that matrix. You would still have to discover the period $2k+1$ by hand, and the bookkeeping is unmanageable for general k in the arena. Rank it above brute force only because it is at least finite.

3. **Pair permutations with their inverses.** *Cue:* "odd if and only if" asks for an involution with few fixed points. *Circuit:* parity via a fixed point free pairing. Since σ is k-limited if and only if σ^{-1} is, inversion is an involution on the set, so the count has the parity of the number of k-limited *involutions*. Grouping those by their restriction to $\{k+2, \ldots, n\}$ makes every class even except the ones in bijection with the k-limited involutions of $\{2k+2, \ldots, n\}$, which

gives $F_{n,k} \equiv F_{n-2k-1,k}$ directly. It is the most purely combinatorial solution, but the class size bookkeeping is fiddly to make airtight under time pressure.

4. **Determinant modulo two, then block row reduction.** *Cue:* parity of a restricted permanent is a determinant modulo two. *Circuit:* permanent equals determinant over $\mathbb{F}_2$, then Gaussian elimination. Clear the top left $(2k+1) \times (2k+1)$ corner of $M_{n,k}$ to a permutation matrix using only rows and columns 1 to $k+1$, which never touches the bottom right block; what remains is block diagonal with an invertible corner and $M_{n-2k-1,k}$, giving $\det M_{n,k} \equiv \det M_{n-2k-1,k}$, and then you check the base cases $n = 0, \ldots, 2k$. Correct, and the standard write up, but the elimination is much easier to draw than to write.

5. **Determinant modulo two, then a direct kernel analysis.** *Cue:* a banded translation invariant matrix has a kernel governed by a linear recurrence. *Circuit:* Toeplitz band gives $a_{m-k} = a_{m+k+1}$, so kernel vectors have period $2k+1$. Same first move, but instead of eliminating, ask directly when $M_{n,k}$ has a nonzero kernel over $\mathbb{F}_2$. Subtracting consecutive coordinate equations turns the band into the recurrence $a_{m+k+1} = a_{m-k}$; the boundary zeros then force k coordinates to vanish at each end, and where those two blocks land modulo $2k+1$ is decided entirely by n modulo $2k+1$. In the bad residues you can even write the dependency down explicitly, as two different arithmetic progressions of rows that both sum to the all ones vector. No pictures, no elimination, and the modulus $2k+1$ appears as the degree of a recurrence rather than as the outcome of a computation. This is the route to write.

The ideal solution

Write $F_{n,k}$ for the number of k-limited permutations of $\{1, \ldots, n\}$, and let $M = M_{n,k}$ be the $n \times n$ matrix over the field $\mathbb{F}_2$ with

$$M_{ij} = \begin{cases} 1 & |i-j| \le k, \\ 0 & \text{otherwise.} \end{cases}$$

Denote its rows by $r_1, \ldots, r_n$; the support of r_t is $[t-k, t+k] \cap [1, n]$.

Step 1: the parity is a determinant. By the Leibniz formula, computed in $\mathbb{F}_2$,

$$\det M = \sum_{\sigma} \operatorname{sgn}(\sigma)\, M_{1\sigma(1)} \cdots M_{n\sigma(n)}.$$

The product $M_{1\sigma(1)} \cdots M_{n\sigma(n)}$ equals 1 precisely when $|\sigma(i) - i| \le k$ for every i, and equals 0 otherwise; and $\operatorname{sgn}(\sigma) = \pm 1 \equiv 1$ in $\mathbb{F}_2$. Hence $\det M$ is the image of $F_{n,k}$ in $\mathbb{F}_2$, that is,

$$F_{n,k} \equiv \det M_{n,k} \pmod 2.$$

So it suffices to prove that $M_{n,k}$ is invertible over $\mathbb{F}_2$ if and only if $n \equiv 0$ or $1 \pmod{2k+1}$.

Step 2: if $n \equiv 2, 3, \ldots, 2k \pmod{2k+1}$ then M is singular. Let s be the residue of n modulo $2k+1$, with $2 \le s \le 2k$ (so $k \ge 1$ and $n \ge 2$). Put $t = 2k+1-s$, so $1 \le t \le 2k-1$. Choose integers a, b with $0 \le a < b \le k$ and $a + b = t$: take $(a, b) = (0, t)$ if $t \le k$, and $(a, b) = (t-k, k)$ if $k < t \le 2k-1$. Set

$$j = \frac{n+a+b}{2k+1} \in \mathbb{Z}_{\ge 1},$$

which is an integer because $n + a + b \equiv s + t \equiv 0 \pmod{2k+1}$, and is positive because $n + a + b \ge 3$.

Consider the row indices $u_i = (k+1-a) + (2k+1)i$ for $i = 0, 1, \ldots, j-1$. They lie in $[1, n]$: the smallest is $k + 1 - a \ge 1$ because $a \le k$, and the largest is

$$(k+1-a) + (2k+1)(j-1) = (k+1-a) + (n+a+b) - (2k+1) = n + b - k \le n$$

because $b \le k$. Consecutive indices differ by $2k+1$, so the supports $[u_i - k, u_i + k]$ are pairwise disjoint abutting blocks of length $2k+1$, and their union is the interval $[u_0 - k, u_{j-1} + k] = [1 - a,\ n + b]$, which contains $[1, n]$. Every column of $[1, n]$ therefore lies in exactly one of these supports, so over $\mathbb{F}_2$

$$\sum_{i=0}^{j-1} r_{u_i} = (1, 1, \ldots, 1).$$

Running the identical computation with b in place of a, the indices $v_i = (k + 1 - b) + (2k+1)i$ for $i = 0, \ldots, j-1$ also lie in $[1, n]$ (smallest $k + 1 - b \ge 1$, largest $n + a - k \le n$) and likewise

$$\sum_{i=0}^{j-1} r_{v_i} = (1, 1, \ldots, 1).$$

Since $0 < b - a \le k < 2k+1$, the indices u_i and $v_{i'}$ are never equal (they lie in different residue classes modulo $2k+1$), so $\{u_i\}$ and $\{v_i\}$ are disjoint nonempty sets of row indices. Adding the two identities gives

$$\sum_{i=0}^{j-1} r_{u_i} + \sum_{i=0}^{j-1} r_{v_i} = 0$$

with all $2j \ge 2$ coefficients equal to 1, a nontrivial linear dependence among the rows. Hence $\det M = 0$ and $F_{n,k}$ is even.

Step 3: if $n \equiv 0$ or $1 \pmod{2k+1}$ then M is invertible. If $n = 1$ then $M = (1)$ and there is nothing to prove, so assume $n \ge 2$; then $n \ge 2k+1$ in the first case and $n \ge 2k+2$ in the second, so in both cases $n \ge 2k + 1 > k$.

Let $a_1, \ldots, a_n \in \mathbb{F}_2$ satisfy $a_1 r_1 + \cdots + a_n r_n = 0$, and extend the notation by setting $a_i = 0$ for $i \le 0$ and for $i > n$. Column m of r_i is 1 exactly when $|i - m| \le k$, so the m-th coordinate of the dependence reads

$$c_m := a_{m-k} + a_{m-k+1} + \cdots + a_{m+k} = 0 \qquad (1 \le m \le n).$$

Adding c_m and c_{m+1} (in $\mathbb{F}_2$, adding is subtracting) cancels the $2k$ common terms and leaves

$$a_{m-k} = a_{m+k+1} \qquad (1 \le m \le n-1). \tag{$*$}$$

Equivalently, $a_i = a_{i+2k+1}$ whenever $1 - k \le i \le n - 1 - k$: the coefficients are periodic with period $2k+1$ throughout the range where the relation applies.

Take $m = 1, 2, \ldots, k$ in $(*)$. The left hand sides are $a_{1-k}, a_{2-k}, \ldots, a_0$, all with index at most 0 and therefore zero, so

$$a_{k+2} = a_{k+3} = \cdots = a_{2k+1} = 0. \tag{L}$$

Take $m = n - k, n - k + 1, \ldots, n - 1$ in $(*)$. The right hand sides are $a_{n+1}, \ldots, a_{n+k}$, all with index greater than n and therefore zero, so

$$a_{n-2k} = a_{n-2k+1} = \cdots = a_{n-1-k} = 0. \tag{R}$$

Suppose first $n = (2k+1)N$ with $N \ge 1$. For $1 \le i \le k$, repeated use of $(*)$ gives

$$a_i = a_{i+(2k+1)} = \cdots = a_{i+(2k+1)(N-1)} = a_{n-2k-1+i},$$

every step being legitimate because the indices used never exceed $n - 1 - k$. As i runs over $1, \ldots, k$ the right hand index runs over $n - 2k, \ldots, n - k - 1$, which is exactly the block (R). Hence $a_1 = \cdots = a_k = 0$. Together with (L) this gives $a_i = 0$ for every $i \in \{1, \ldots, 2k+1\}$ except possibly $i = k+1$. Now the equation $c_{k+1} = 0$ reads $a_1 + a_2 + \cdots + a_{2k+1} = 0$, which forces $a_{k+1} = 0$ as well. Finally $(*)$ propagates $a_i = a_{i+2k+1}$ from the block $\{1, \ldots, 2k+1\}$ rightwards through all of $[1, n]$, so every a_i is zero.

Suppose instead $n = (2k+1)N+1$ with $N \geq 1$. Now $n-2k = 2+(2k+1)(N-1)$ and $n-1-k = (k+1)+(2k+1)(N-1)$, so for $2 \leq i \leq k+1$ the same chain

$$a_i = a_{i+(2k+1)} = \cdots = a_{i+(2k+1)(N-1)}$$

identifies $a_2, \ldots, a_{k+1}$ with the entries of the block (R), so $a_2 = \cdots = a_{k+1} = 0$. Together with (L) this gives $a_2 = a_3 = \cdots = a_{2k+1} = 0$, and $c_{k+1} = 0$ then forces $a_1 = 0$. Propagating by $(*)$ as before, every a_i is zero.

In both cases the rows of M are linearly independent over $\mathbb{F}_2$, so $\det M \neq 0$ in $\mathbb{F}_2$ and $F_{n,k}$ is odd.

Steps 2 and 3 together say $\det M_{n,k} \neq 0$ in $\mathbb{F}_2$ exactly when $n \equiv 0$ or $1 \pmod{2k+1}$, which by Step 1 is exactly the assertion that $F_{n,k}$ is odd in those cases and even otherwise. ■

Layer 1: WHAT

Block 1. Parity of a count of permutations with prescribed allowed positions is the permanent of a $0/1$ matrix modulo 2, and modulo 2 the permanent is the determinant.

Block 2. So the question is: for which n is the banded matrix $M_{n,k}$ invertible over $\mathbb{F}_2$? The combinatorics is finished at this line.

Block 3. The matrix is translation invariant along the diagonal, so consecutive coordinate equations for a kernel vector differ in only two places, giving the recurrence $a_{m-k} = a_{m+k+1}$.

Block 4. That recurrence says kernel vectors are periodic with period $2k+1$. The modulus of the answer is now explained, before any case analysis.

Block 5. The zero padding outside $[1, n]$ turns the first k and last k equations into two blocks of k forced zeros, one anchored at each end.

Block 6. Where the right hand block lands modulo $2k+1$ depends only on n modulo $2k+1$. For $n \equiv 0$ or 1 the two blocks are complementary and cover all residues except one, which the central equation kills; the kernel is trivial.

Block 7. For the other residues, do not argue by counting; exhibit the dependency. Two arithmetic progressions of rows with common difference $2k+1$, offset by $k+1-a$ and $k+1-b$, each tile $[1, n]$ and so each sum to the all ones vector; their sum is a nontrivial relation.

The nominal route conditions on $\sigma(1)$, obtains a recurrence for $F_{n,k}$ in n, and tries to track it modulo 2. For $k = 1$ that route works and gives the Fibonacci numbers; for general k the recurrence has order growing with k and the parity pattern is invisible. The bifurcation point is Block 1: the moment you decide to stop counting permutations and start computing a determinant. The licence for that decision is the single observation that signs are invisible modulo 2, and the payoff is that all of linear algebra becomes available for a combinatorial question.

Layer 2: HOW

Block 1. Cue. "Odd if and only if", attached to a count of permutations whose constraint is "$\sigma(i)$ lies in an allowed set". **Why natural.** A count of permutations with prescribed allowed positions is by definition a permanent of a $0/1$ matrix. The permanent is intractable in general, but modulo 2 it coincides with the determinant, because the only difference between the two is the factor $\operatorname{sgn}(\sigma) = \pm 1$ and $-1 = 1$ in $\mathbb{F}_2$. Parity is therefore the one regime in which a permanent is computable, and the problem asks for exactly that regime. Notice what you did not do: you did not look for a bijection or a formula, because the words "odd if and only if" told you signs would be free. **Link to the next block.** You now own a concrete matrix, so every question about it is a question about its rank.

Block 2. Cue. The matrix is defined by $|i-j| \leq k$, which depends only on $i-j$. **Why natural.** A matrix whose entries depend only on $i-j$ is Toeplitz, and Toeplitz means translation invariant, and translation invariant systems have solution spaces described by linear recurrences. Reading

the band width off the definition, an interior row has $2k+1$ ones, which is already the modulus in the answer; that match tells you that you are looking for a recurrence of order $2k+1$ and that you should stop looking for anything else. **Link to the next block.** To find the recurrence, look at what changes between one coordinate equation and the next.

Block 3. **Cue.** Consecutive coordinate equations of the dependence share $2k$ of their $2k+1$ terms. **Why natural.** Adding two equations that overlap in $2k$ terms annihilates the overlap and leaves the two extreme terms, a_{m-k} and a_{m+k+1}. Over $\mathbb{F}_2$ that says they are equal. This is the standard way to extract a recurrence from a band: differencing kills the interior. The distance between the two survivors is $(m+k+1)-(m-k)=2k+1$, the band width, which is why the period is what it is. **Link to the next block.** A recurrence needs boundary conditions, and the boundary conditions are hiding in the convention that $a_i = 0$ outside $[1,n]$.

Block 4. **Cue.** The padding $a_i = 0$ for $i \le 0$ and $i > n$ was introduced only to write the coordinate equation uniformly, and free information usually hides in such conventions. **Why natural.** Apply the recurrence at the k leftmost usable values of m: each one equates a padded zero with a genuine coefficient, so it kills a genuine coefficient. The same at the right end. Two blocks of k zeros, at positions $k+2,\dots,2k+1$ from the left and $n-2k,\dots,n-1-k$ from the right. That is $2k$ of the $2k+1$ residues. **Link to the next block.** The one remaining residue must be killed by an equation that has not yet been used, and the only untouched equation is a coordinate equation rather than a difference of two.

Block 5. **Cue.** Exactly one residue class is still alive, and $c_{k+1} = 0$ is a single equation involving precisely $a_1,\dots,a_{2k+1}$. **Why natural.** The middle equation $c_{k+1} = 0$ is the unique coordinate equation whose window is exactly $[1,2k+1]$, with no padding on either side. If all but one of its terms are already known to vanish, it determines the last one. The arithmetic of which residues the two blocks occupy is the entire content of the case split, and it depends only on n modulo $2k+1$, because periodicity transports the right hand block leftwards by multiples of $2k+1$. **Link to the next block.** If the two blocks overlap or fail to cover $2k$ residues, the argument breaks, and where it breaks you should expect an actual dependency, so switch from proving independence to constructing a relation.

Block 6. **Cue.** For the bad residues you need to *produce* a dependency, not to fail to rule one out. **Why natural.** Rows whose indices form an arithmetic progression of common difference $2k+1$ have supports that are consecutive disjoint blocks of length $2k+1$, so they tile an interval and sum to its indicator. Choose the starting offset so that the tiled interval contains $[1,n]$; the sum is then the all ones vector. Two different offsets $k+1-a$ and $k+1-b$ both work exactly when $a+b \equiv -n \pmod{2k+1}$ with $0 \le a < b \le k$, and such a pair exists precisely for the residues $2,\dots,2k$. The condition $a,b \le k$ is what keeps the progressions inside $[1,n]$; the condition $a \ne b$ is what makes the relation nontrivial. **Link to the next block.** The two directions together are an if and only if, so state the conclusion in the language of the question, in terms of parity, not of determinants.

Layer 3: WHY

The first truth is that the permanent and the determinant of a matrix coincide over any ring of characteristic 2. Every question of the form "is the number of systems of distinct representatives odd" is therefore a rank question over $\mathbb{F}_2$, and rank questions have tools: elimination, kernels, explicit dependencies. This single transfer is the most valuable thing in the problem, and it is worth memorising in the general form: parity of a restricted permutation count equals the $\mathbb{F}_2$ determinant of the allowed positions matrix.

The second truth is that a banded Toeplitz matrix is a difference operator in disguise. Its kernel over a field is the solution set of a constant coefficient linear recurrence, subject to boundary conditions imposed by the finite window. The recurrence is read off the *symbol* of the band: here the interior row is $1 + x + \cdots + x^{2k}$, and over $\mathbb{F}_2$

$$(1 + x)\left(1 + x + \cdots + x^{2k}\right) = 1 + x^{2k+1},$$

so multiplying the coordinate equations by $(1 + x)$, which is exactly what adding consecutive equations does, produces the relation $a_i = a_{i+2k+1}$ with characteristic polynomial $x^{2k+1} + 1$. Every solution of that recurrence is periodic with period dividing $2k + 1$, and the solution space over the whole line is $(2k + 1)$ dimensional. The finite window imposes k conditions from the left edge, k from the right edge, and one from the central equation; $k + k + 1 = 2k + 1$, so generically the kernel is cut down to zero. It is only when the two edge blocks collide, that is when they occupy overlapping residues, that fewer than $2k + 1$ independent conditions are imposed and a kernel vector survives. That collision is governed by n modulo $2k + 1$, and the arithmetic works out to the residues $2, \ldots, 2k$.

The discernment mechanism is a two step reflex. First, when the parity of a combinatorial count is requested, ask whether the objects are matchings, permutations or systems of distinct representatives; if they are, the count is a permanent and modulo 2 it is a determinant. Second, when the resulting matrix is translation invariant, do not eliminate; difference the equations to get a recurrence, read the period off the band width, and let the boundary conditions do the case analysis. Train it on $k = 1$, $n = 1, \ldots, 5$, watching the Fibonacci parity pattern $1, 0, 1, 1, 0, \ldots$ emerge with period $3 = 2k + 1$.

The examiners nudged you by asking for parity rather than the count, by making the constraint symmetric in i and j so that inversion is available as a second route, and above all by putting $2k + 1$ in the answer, which is the number of ones in an interior row. When the modulus in the conclusion equals a structural width in the hypothesis, the proof runs through an object of that period.

Revolving concepts

Permanent, determinant and characteristic 2. For an $n \times n$ matrix, $\operatorname{per}(M) = \sum_\sigma \prod_i M_{i\sigma(i)}$ and $\det(M) = \sum_\sigma \operatorname{sgn}(\sigma) \prod_i M_{i\sigma(i)}$. Over any commutative ring in which $2 = 0$ these agree. For a $0/1$ matrix, $\operatorname{per}(M)$ counts the permutations σ with $M_{i\sigma(i)} = 1$ for all i, that is, the systems of distinct representatives. See Book 1, field: Matrices and determinants, topic: permanent and Leibniz formula.

Linear algebra over $\mathbb{F}_2$. A square matrix over a field is invertible if and only if its rows are linearly independent, if and only if its determinant is nonzero. Over $\mathbb{F}_2$, addition is subtraction, a linear combination is a subset sum of rows, and a linear dependency is precisely a pair of distinct subsets of rows with equal sums. That last reformulation is what makes Step 2 constructive. See Book 1, field: Linear algebra, topic: rank over finite fields.

Banded and Toeplitz matrices. A matrix with M_{ij} depending only on $i - j$ is Toeplitz; if additionally $M_{ij} = 0$ for $|i - j| > k$ it is banded with band width $2k + 1$. Its symbol is the Laurent polynomial $\sum_{|d| \le k} c_d x^d$. Kernel vectors of the finite section satisfy the linear recurrence whose characteristic polynomial is the symbol, with the finite window supplying boundary conditions. See Book 1, field: Matrices and determinants, topic: Toeplitz and banded matrices.

Involutions and parity. The map $\sigma \mapsto \sigma^{-1}$ is an involution on S_n whose fixed points are the involutions, so the number of involutions of an m-element set has the same parity as $m!$, hence is odd exactly when $m \le 1$. Since σ is k-limited if and only if σ^{-1} is, $F_{n,k}$ has the same parity as the number of k-limited involutions; this is the engine of the third ranked approach. See Book 1, field: Combinatorics, topic: sign reversing and fixed point involutions.

The recursion, generalised. Two recursions live here and both should be stated in general form. The first is the recurrence satisfied by a kernel vector: $a_{i+2k+1} = a_i$, a constant coefficient linear

recurrence of order $2k+1$ over $\mathbb{F}_2$ with characteristic polynomial $x^{2k+1}+1$; its general solution is "periodic with period dividing $2k+1$", the solution space has dimension $2k+1$, and the general method for any banded Toeplitz kernel is identical: multiply the symbol by whatever polynomial clears the interior (here $1+x$), read off the characteristic polynomial, and count boundary conditions against the dimension. The second is the congruence $F_{n,k} \equiv F_{n-2k-1,k} \pmod 2$ with the convention $F_{0,k}=1$, a recursion in n of step $2k+1$; its general solution is $F_{n,k} \equiv F_{r,k}$ where r is the residue of n modulo $2k+1$, so the whole problem reduces to the base window $r = 0, 1, \ldots, 2k$, in which $F_{r,k} = r!$ for $r \le k+1$ and the matrix has two equal rows for $k+2 \le r \le 2k$. Any "period P in the answer" problem should be attacked by looking for exactly this kind of step P congruence. See Book 1, field: Sequences and recurrences, topic: linear recurrences over finite fields.

Knowledge gaps, stated plainly. This problem is not approachable from A level Further Mathematics alone, and the missing theory is specific. You must have installed: the Leibniz formula for a determinant of arbitrary size (Edexcel FP1 gives you 2×2 and 3×3 determinants as formulas, not as sums over permutations); the notion of a field with two elements and of linear algebra over it, including "determinant nonzero if and only if rows independent" over a field that is not $\mathbb{R}$; the permanent, and the fact that it counts systems of distinct representatives; the sign of a permutation as a group homomorphism to $\{\pm 1\}$; and the vocabulary of Toeplitz and banded matrices. None of these is hard, but each is a genuine prerequisite, and without the first three the problem cannot even be started.

Circuit Bank

- **WHEN** you are asked for the parity of the number of permutations, matchings or systems of distinct representatives satisfying "$\sigma(i)$ lies in an allowed set S_i" **DO** build the $0/1$ matrix $M_{ij} = [\, j \in S_i \,]$ and replace the count by $\det M$ computed over $\mathbb{F}_2$ **BECAUSE** the permanent and the determinant differ only by $\operatorname{sgn}(\sigma) = \pm 1$, which is 1 modulo 2, so a parity question is the one case where a permanent is computable, and it hands you the whole of linear algebra

- **WHEN** a matrix over a field has entries depending only on $i-j$, especially with a narrow band **DO** add consecutive coordinate equations of a kernel vector to cancel the overlap, producing a linear recurrence whose order is the band width **BECAUSE** a Toeplitz band is a difference operator, its kernel is the solution set of a constant coefficient recurrence, and differencing is exactly multiplication of the symbol by $(1+x)$, which collapses $1+x+\cdots+x^{2k}$ to $1+x^{2k+1}$

- **WHEN** the answer to a problem contains a modulus, and some structural width in the hypothesis equals that modulus **DO** stop searching for a formula and search for an object of that period, typically a recurrence or a tiling **BECAUSE** a modulus in a conclusion is almost always the order of a translation symmetry present in the setup, and identifying it early tells you the shape of the whole proof

- **WHEN** you must show a matrix over $\mathbb{F}_2$ is singular **DO** exhibit two distinct sets of rows with the same sum, most conveniently the all ones vector **BECAUSE** over $\mathbb{F}_2$ the sum of the two sets is a nontrivial dependency, and rows of a banded matrix whose indices form an arithmetic progression of common difference equal to the band width have disjoint abutting supports, so they tile an interval and sum to its indicator

- **WHEN** a parity question concerns a set carrying a natural involution **DO** pair each object with its image and count only the fixed points **BECAUSE** everything not fixed cancels in pairs, so the parity of the whole set equals the parity of the fixed point set, and for permutations the involution $\sigma \mapsto \sigma^{-1}$ preserves almost every natural positional constraint

Backtesting

1. First useful cue: "odd if and only if" attached to a permutation count with prescribed allowed positions, which converts at once into $\det M_{n,k}$ over $\mathbb{F}_2$ with $M_{ij} = [\,|i-j| \le k\,]$.
2. Method the expression suggested: the band condition $|i-j| \le k$ is Toeplitz, so difference the coordinate equations of a kernel vector to obtain $a_{i+2k+1} = a_i$ and let the boundary padding supply the conditions.
3. Where you would hesitate, and the thought to reinforce: at the case analysis, because it looks like arithmetic bookkeeping with no idea in it. Reinforce the dimension count: the recurrence space has dimension $2k+1$, the left edge imposes k conditions, the right edge k, the central equation 1, and $k+k+1 = 2k+1$; the kernel survives exactly when the two edge blocks fail to be complementary, and that failure is a statement about n modulo $2k+1$.
4. The trap the question sets: proving only one direction. Showing that the residues 0 and 1 give an invertible matrix is the easier half and reads like a complete solution; without the explicit dependency in the other residues you have proved "odd if" and not "only if", which is half the marks. The second trap is forgetting that $n = 1$ falls outside the index ranges used in Step 3 and must be checked directly.
5. What makes you recognise it faster next time: the modulus of the conclusion equalling the number of ones in an interior row. Any problem in which a stated period matches a structural width is a translation symmetry problem, and the proof will run through a recurrence of exactly that order.
6. The reusable decision rule: parity of a restricted permutation count is an $\mathbb{F}_2$ determinant; an $\mathbb{F}_2$ determinant of a banded Toeplitz matrix is decided by a linear recurrence of order equal to the band width together with the boundary conditions from the finite window.

Mirror drills

1. Show that the number of involutions of an m-element set is odd if and only if $m \le 1$. *Answer:* pair each non involution σ with σ^{-1}; the pairs cancel, so the count has the same parity as $|S_m| = m!$, which is odd exactly for $m = 0, 1$.
2. Take $k = 1$. Show that $F_{n,1}$ is the Fibonacci number with $F_{1,1} = 1$, $F_{2,1} = 2$, and confirm directly that it is odd exactly when $n \equiv 0$ or $1 \pmod 3$. *Answer:* condition on $\sigma(1)$, which is 1 or 2; if $\sigma(1) = 1$ the rest is a 1-limited permutation of $\{2, \dots, n\}$, and if $\sigma(1) = 2$ then $\sigma(2) = 1$ forcibly, leaving $\{3, \dots, n\}$. So $F_{n,1} = F_{n-1,1} + F_{n-2,1}$, and the Fibonacci parity sequence is $1, 0, 1, 1, 0, 1, \dots$ with period $3 = 2k+1$.
3. Prove $F_{n,k} \equiv F_{n-2k-1,k} \pmod 2$ for $n \ge 2k+1$, with $F_{0,k} = 1$, using only the determinant. *Answer:* in $M_{n,k}$ add row 1 to rows $2, \dots, k+1$ and then column 1 to columns $2, \dots, k+1$; repeat with rows $2, \dots, k+1$ to clear the remaining entries of the top left $(2k+1) \times (2k+1)$ corner, using only the first $k+1$ rows and columns so that the unseen part of the matrix is never disturbed. The result is block diagonal with an invertible corner and $M_{n-2k-1,k}$, so the determinants agree modulo 2.

Chapter 19

Putnam 2007

19.1 A1. Two mirrored parabolas and their tangency values

Find all values of α for which the curves

$$y = \alpha x^2 + \alpha x + \frac{1}{24}, \qquad x = \alpha y^2 + \alpha y + \frac{1}{24}$$

are tangent to each other.

Read it out loud

Say the two equations aloud and you hear the same sentence twice, once with x and y swapped. That is the loudest cue on the page, louder than "tangent", louder than the $\frac{1}{24}$. Before you write a single derivative, register it explicitly: the second curve is not some unrelated parabola, it is the first curve reflected in the line $L : y = x$. Everything downstream should be organised around that fact, because a pair of curves exchanged by a reflection has its whole intersection structure constrained by the reflection: any point where they meet is either fixed by the reflection (sits on L) or comes paired with its mirror image, which also has to be a meeting point.

The fields wired into the problem are coordinate geometry (two conics), the calculus of tangency (matching derivatives, or a repeated root under elimination) and, once you commit to eliminating a variable, ordinary quadratic theory (discriminants, Vieta). What is really being tested is whether you let the symmetry do the organising work before you touch a derivative, or whether you set up three unknowns (x_0, y_0, α) and a slope matching condition and grind.

The first useful line to have written within thirty seconds is $f(x) = \alpha x^2 + \alpha x + \frac{1}{24}$, so that the two curves read $y = f(x)$ and $x = f(y)$: the same function, twice.

Approaches, ranked worst to best

1. **raw simultaneous system.** *Cue:* "tangent" reads as "matching slopes". *Circuit:* set $f'(x_0) = 1/g'(y_0)$ alongside both curve equations. (three unknowns tied by three separate conditions, no organising principle, and the count of four solutions never becomes visible until very late, if at all; easy to lose a sign carrying x_0, y_0, α together)
2. **geometric case split on L.** *Cue:* C_2 is C_1 reflected in L. *Circuit:* case on whether C_1 meets L, is tangent to L, or crosses it, then rule out tangency off L by counting intersections. (close to a standard route, and rigorous once finished, but its last step borrows the fact that two conics meet in four points counted with multiplicity, machinery you would be reaching for mid exam if asked to defend it from scratch)
3. **eliminate to one quartic, force a factorisation.** *Cue:* the same function f appears on both sides. *Circuit:* substitute $y = f(x)$ into C_2, get a single quartic in x, and hunt for a repeated root. (turns two curves and a slope condition into one polynomial and one derivative; the four solutions and their split into on L and off L both fall out of the algebra rather than needing to be argued geometrically)

The ideal solution

Write $f(x) = \alpha x^2 + \alpha x + \frac{1}{24}$, so C_1 is $y = f(x)$ and C_2 is $x = f(y)$: C_2 is exactly C_1 reflected in $L : y = x$. If $\alpha = 0$, C_1 is the line $y = \frac{1}{24}$ and C_2 is the line $x = \frac{1}{24}$, which cross once, transversally; discard this case and take $\alpha \neq 0$ from here on.

A point $(x, f(x))$ of C_1 lies on C_2 exactly when $x = f(f(x))$. Subtracting the two defining equations gives

$$f(x) - f(y) = \alpha(x^2 - y^2) + \alpha(x - y) = \alpha(x - y)(x + y + 1),$$

so every common point of C_1, C_2 satisfies $(x-y)\big[\alpha(x+y+1)+1\big] = 0$. A subtraction alone always loses information, but keeping $y = f(x)$ alongside it does not: substituting throughout and expanding, the quartic that records every intersection factorises cleanly (check it by direct expansion) as

$$x - f(f(x)) \;=\; -\alpha\, Q_1(x)\, Q_2(x),$$

$$Q_1(x) = \alpha x^2 + (\alpha-1)x + \tfrac{1}{24}, \qquad Q_2(x) = \alpha x^2 + (\alpha+1)x + \tfrac{1}{24} + \tfrac{\alpha+1}{\alpha}.$$

$Q_1(x) = 0$ says exactly $f(x) = x$: the point sits on L. $Q_2(x) = 0$ says exactly $x + f(x) = -\frac{\alpha+1}{\alpha}$: the point sits on the mirror partner line. Both are genuine quadratics since $\alpha \neq 0$, so between them they carry all four intersections of the two conics, with nothing left over.

Now convert "tangent" into algebra. C_1 never has a vertical tangent, since it is a graph in x throughout, so at any point where C_1, C_2 touch, C_2's tangent there is non vertical too, i.e. $f'(y_0) \neq 0$. Writing $H(x) = x - f(f(x))$ and differentiating by the chain rule, $H'(x_0) = 1 - f'(f(x_0))f'(x_0) = 1 - f'(y_0)f'(x_0)$, which vanishes exactly when C_1's slope $f'(x_0)$ equals C_2's slope $1/f'(y_0)$: exactly tangency, not a transversal crossing. So you want the α for which $H(x) = -\alpha Q_1(x)Q_2(x)$ has a repeated root.

From the factorisation, $H' = -\alpha(Q_1'Q_2 + Q_1Q_2')$. At a root of Q_1 this vanishes iff Q_1 itself has a double root there, or Q_2 also vanishes there (symmetrically for a root of Q_2). So you need Q_1 to have a repeated root, or Q_2 to.

Q_1 repeated root: $\operatorname{disc}(Q_1) = (\alpha-1)^2 - \frac{\alpha}{6} = 0$, i.e. $6\alpha^2 - 13\alpha + 6 = 0$, giving $\alpha = \frac{2}{3}, \frac{3}{2}$.

Q_2 repeated root: its vertex sits at $x = -\frac{\alpha+1}{2\alpha}$, which is also the unique root of the linear equation $Q_1(x) - Q_2(x) \;=\; -2x - 1 - \frac{1}{\alpha} \;=\; 0$: so Q_1 and Q_2 agree there for every α, and the value of a quadratic at its own vertex is $-\operatorname{disc}/(4\cdot\text{leading coefficient})$, so both equal $-\operatorname{disc}(Q_2)/(4\alpha)$ at that point. In other words, whenever Q_2 has a repeated root, Q_1 vanishes there too, and that point is forced onto L: there is no off axis tangency left to hunt for separately, this case collapses back onto the first. Computing, $\operatorname{disc}(Q_2) = (\alpha+1)^2 - 4\alpha\left(\frac{1}{24} + \frac{\alpha+1}{\alpha}\right) = 0$ reduces to $6\alpha^2 - 13\alpha - 18 = 0$, giving $\alpha = \frac{13 \pm \sqrt{601}}{12}$.

No other value of α makes H have a repeated root, so

$$\boxed{\alpha \in \left\{ \frac{2}{3},\ \frac{3}{2},\ \frac{13+\sqrt{601}}{12},\ \frac{13-\sqrt{601}}{12} \right\}}.$$

Layer 1: WHAT

Block 1: rewrite both curves through one function f, see C_2 as C_1 reflected in L. Block 2: dismiss $\alpha = 0$ directly. Block 3: subtract the defining equations to get a necessary factorised condition. Block 4: keep $y = f(x)$ alive alongside that condition so nothing is lost, and land on the full quartic factorisation $H = -\alpha Q_1 Q_2$. Block 5: convert "tangent" to "H has a repeated root" via one chain rule differentiation, having ruled out the vertical tangent edge case. Block 6: read off, from the product rule shape of H', that a repeated root of the quartic means a repeated root of Q_1 or of Q_2. Block 7: two discriminant computations, two quadratics in α, four roots.

The nominal route sets up x_0, y_0, α as three tied unknowns from the start, with a slope matching condition alongside both curve memberships, and grinds an elimination without ever building a single polynomial to ask a repeated root question of. It can be made to work, but it hides the "exactly four, in two clean pairs" structure until very late. The bifurcation point is Block 1: the moment you decide whether to carry three unknowns through a slope condition, or eliminate immediately to one variable and ask when a polynomial has a repeated root. Every unknown you keep past that moment is ink you did not need to spend.

Layer 2: HOW

Block 1. **Cue.** The second equation is the first with x, y swapped. **Why natural.** A reflection exchanging two curves forces their combined intersection set to be invariant under that reflection: points are fixed by it (on L) or come in mirror pairs. **Link to the next block.** This tells you, before any algebra, that the eventual case split is on L versus off L, not some undifferentiated four points.

Block 2. **Cue.** α multiplies every squared term; at $\alpha = 0$ both curves degenerate to a line. **Why natural.** A degenerate case sitting right at the edge of your later working (you will shortly divide by α) has to be checked and dismissed explicitly, not left implicit. **Link to the next block.** Clears the way to treat $\alpha \neq 0$ as a standing hypothesis for the rest of the solution.

Block 3. **Cue.** Two equations, same shape, only x, y swapped: subtracting cancels the shared structure and factors. **Why natural.** $f(x) - f(y)$ is a difference of the same quadratic at two points, which always factors through $(x - y)$. **Link to the next block.** The factor $(x - y)[\ldots] = 0$ names the two lines, L and its mirror partner, that every intersection point has to sit on.

Block 4. **Cue.** Subtracting throws away information; you still have $y = f(x)$ sitting unused. **Why natural.** Keeping one original equation alongside the subtracted one is exactly enough to recover the other original equation, since it can be written as the first equation minus the subtracted one. **Link to the next block.** Substituting $y = f(x)$ throughout turns the necessary condition into the full quartic, no solutions missing.

Block 5. **Cue.** C_1 is a graph in x, so it never carries a vertical tangent. **Why natural.** Tangency of two curves, one of which is a graph, is exactly a repeated root of the polynomial obtained by substituting that graph into the other curve's equation, away from any point where the second curve itself goes vertical. **Link to the next block.** Differentiating $H(x) = x - f(f(x))$ by the chain rule gives a slope matching condition directly.

Block 6. **Cue.** H is presented as a product $-\alpha Q_1 Q_2$. **Why natural.** The product rule sends "H has a repeated root at a root of Q_1" to "Q_1 has a double root there, or Q_2 also vanishes there": nothing else can make the derivative vanish. **Link to the next block.** Two clean quadratic discriminant conditions replace what looked like a four way case split.

Layer 3: WHY

The mathematical truth underneath the route is that eliminating one variable between two conics, when one of them is literally a graph, never loses solutions and always produces a polynomial whose repeated roots are exactly the tangencies: this is a general fact about intersection multiplicity, not a coincidence of this problem, and it is worth installing permanently rather than re deriving under pressure. The discernment mechanism for choosing this route over the raw three unknown system is recognising, in the first thirty seconds, that the problem handed you a symmetry (the same function, two variables swapped) for free; a symmetry between two objects always constrains where their interaction can live, and the fastest working habit is to ask immediately what is fixed by it and what it pairs up. The examiners nudged you toward this by using the identical formula on both sides rather than two unrelated quadratics: had they wanted you to grind three unknowns, they would not have handed you a reflection.

Training this discernment is a matter of repetition on the cue itself: any time a problem presents the same expression with two variables exchanged, write down the reflection sentence before anything else, and ask what lies on the fixed line before asking what lies off it. The second half of the mechanism, checking whether the off axis branch secretly collapses back onto the fixed line, is

the part most solvers skip; here it is exactly what happens, a repeated root of Q_2 forces Q_1 to vanish at the same point too, and missing that check would have cost you a search for a second, and here nonexistent, family of solutions.

Revolving concepts

Tangency via elimination: if C_1 is $y = f(x)$ and C_2 another curve, substitute to get $H(x)$; common points are roots of H, and where C_2 has no vertical tangent, C_1, C_2 touch at x_0 exactly when $H(x_0) = H'(x_0) = 0$. See Book 1, field: Geometry, topic: intersection multiplicity of plane curves.

Two conics meet in four points: neither curve here is degenerate, so elimination gives degree four exactly, and Q_1, Q_2, two genuine quadratics, divide that four between them with nothing lost. See Book 1, field: Geometry, topic: Bezout style counting for two conics.

Discriminant and vertex value: a quadratic $Ax^2 + Bx + C$ has a repeated root iff $B^2 - 4AC = 0$, and its value at its own vertex $x = -B/(2A)$ equals $-(B^2 - 4AC)/(4A)$. See Book 1, field: Polynomials, topic: quadratic discriminant.

Reflection symmetry of an intersection set: if two curves are exchanged by a reflection, or any involution, of the plane, their mutual intersection set is invariant under it; points are fixed (on the mirror line) or paired. See Book 1, field: Geometry, topic: symmetry arguments in coordinate geometry.

No recursive structure appears in this problem: there is no sequence or self referential definition to generalise here, the repeated root idea is algebraic rather than recursive.

Circuit Bank

- **WHEN** two curves are literally each other under $x \leftrightarrow y$ or another visible swap **DO** rewrite both through one shared function and ask what the swap fixes **BECAUSE** the mutual intersection set of two curves exchanged by an involution is invariant under it, so the fixed locus organises the whole search
- **WHEN** you need "curve tangent to curve" and one of them is a graph $y = f(x)$ **DO** eliminate to a single polynomial $H(x)$ and ask for a repeated root, rather than matching two separately computed slopes by hand **BECAUSE** substitution into a graph never drops solutions, and repeated roots of one polynomial are far easier to characterise via discriminants than a three unknown slope matching system
- **WHEN** a case split produces two branches that came from the same underlying symmetry **DO** check whether the second branch collapses back into the first via Vieta or a vertex value comparison before treating it as independent **BECAUSE** branches born from one symmetry are rarely independent, and skipping the check risks either a missed rigour gap or an imaginary extra family of solutions
- **WHEN** you need the value of a quadratic at its own vertex **DO** use $-\text{disc}/(4A)$ directly **BECAUSE** it is a fact you should have installed, and it replaces a resubstitution with a one line read off

Backtesting

1. First useful cue: the right hand side of C_2's equation is the identical function f used for C_1, only with x, y exchanged.
2. Method the expression suggested: eliminate y via $y = f(x)$ and hunt for a repeated root of the resulting quartic in x, rather than juggle two curve equations and a separate slope condition.

3. Where you would hesitate, and the thought to reinforce: proving a repeated root of Q_2 cannot give a new tangency of its own; reinforce it by comparing Q_1 and Q_2 directly at the unique root of their difference, not by a vague symmetry appeal.
4. The trap the question sets: $\alpha = 0$ looks like it needs the general machinery, but it is simply two perpendicular lines and should be dismissed by one direct line, kept separate from the rest of the argument.
5. What makes you recognise it faster next time: a problem handing you two equations built from the identical formula with variables swapped is a reflection problem before it is a calculus problem, and the reflection sentence should be the first line you write.
6. The reusable decision rule: when two curves are exchanged by reflection in a line, organise the intersection search around that line, points on it and points off it in mirror pairs, before any derivative is taken.

Mirror drills

1. For which α are $y = \alpha x^2 + 1$ and $x = \alpha y^2 + 1$ tangent? *Hint:* same reflection, same $f(x) = \alpha x^2 + 1$; work out Q_1 (on L) first and its discriminant.
2. Two conics are given by equations of identical shape with x, y swapped; without computing anything, state which two conditions between them, on the line of symmetry and on its mirror partner line, account for every intersection point, and explain why nothing is left over.
3. For $y = x^2 + c$ and $x = y^2 + c$, find every c for which the two parabolas are tangent, and check directly whether each tangency you find sits on $y = x$.

19.2 A2. Minimal convex area meeting both hyperbolas

Find the least possible area of a convex set in the plane that intersects both branches of the hyperbola $xy = 1$ and both branches of the hyperbola $xy = -1$. (A set S in the plane is called convex if for any two points in S the line segment connecting them is contained in S.)

Read it out loud

Four disconnected pieces of curve, one in each quadrant, and a single convex set asked to touch all four. Convexity is the loud word here, and it is doing real work: the moment your set contains one point from each of the four branches, it contains the segment between any two of them, so it contains the whole convex hull of those four points. That turns a question about an unknown shaped region into a question about four numbers, one parameter per branch, which is exactly the kind of reduction that should happen in your head before any calculus.

The fields wired in are coordinate geometry (parametrising a point on $xy = 1$ or $xy = -1$), elementary convexity (segments and hulls), and an inequality: the area, once written out, is a cyclic sum of ratios, and a cyclic sum of a positive quantity and its reciprocal is an AM-GM signal you should react to immediately. What is really being tested is whether you trust the reduction "area of S is at least the area of some hull of four chosen points" enough to make it precise, rather than reaching for calculus of variations on a general convex region, which is far more machinery than a four branch touching condition needs.

The first useful line to have within thirty seconds: name one point per branch, $A = (a, \frac{1}{a})$ in the first quadrant with $a > 0$, and the analogous B, C, D in the other three, and ask what shape they force.

Approaches, ranked worst to best

1. **calculus of variations on a general region.** *Cue:* "least area" sounds like an optimisation over shapes. *Circuit:* parametrise the boundary of S and impose the four touching constraints with Lagrange multipliers. (treats S as an unknown curve rather than using convexity to collapse the problem to four numbers; heavy machinery, and the positivity and quadrant constraints are awkward to encode)

2. **assert the hull quadrilateral bound outright.** *Cue:* four points, one per branch, so their hull is a quadrilateral. *Circuit:* write $\text{area}(S) \geq \text{area}(ABCD)$ and apply AM-GM to the shoelace sum. (reaches the right number fast, but the assertion is not automatic: for some choices of a, b, c, d the point C, say, sits inside the triangle formed by A, B, D, so A, B, C, D in that cyclic order is not even a convex quadrilateral, and the bound needs a proof that survives that case, not just a picture of the nice case)

3. **fan out from the origin.** *Cue:* the four points sweep once around the origin, one quadrant each. *Circuit:* decompose the region into the four triangles from O to consecutive points, show each lies in S, sum areas, then AM-GM. (fully general, no hidden assumption about convex position, and the origin's membership in S falls out of the same convexity already used)

The ideal solution

Let S be convex and suppose it meets both branches of $xy = 1$ and both branches of $xy = -1$. Choose one point of S on each branch: $A = (a, \frac{1}{a})$ with $a > 0$ (quadrant 1, on $xy = 1$), $B = (-b, \frac{1}{b})$ with $b > 0$ (quadrant 2, on $xy = -1$), $C = (-c, -\frac{1}{c})$ with $c > 0$ (quadrant 3, on $xy = 1$), $D = (d, -\frac{1}{d})$ with $d > 0$ (quadrant 4, on $xy = -1$).

As seen from the origin O, A, B, C, D sit at angles increasing through quadrants 1, 2, 3, 4 in turn, each consecutive gap well under a straight angle, and the four gaps sweep out the full circle once.

A point surrounded this way, with every consecutive angular gap less than 180°, always lies inside the convex hull of the four points: no line through O can have all four points strictly on one side of it. Since S is convex and contains A, B, C, D, it contains their hull, so it contains O too.

Now O, A, B, C, D are all in S, so each of the four triangles OAB, OBC, OCD, ODA lies in S (three points of a convex set, hence the whole triangle). These four triangles have disjoint interiors, again because the angular order around O increases strictly through a full turn with no gap wider than a straight angle, so they tile without overlap, and

$$\text{area}(S) \;\geq\; [OAB] + [OBC] + [OCD] + [ODA].$$

This sidesteps entirely the question of whether $ABCD$ is itself a convex quadrilateral (it need not be: for some a, b, c, d one of the four points sits inside the triangle of the other three), because the fan decomposition from O is valid regardless.

Each triangle area is half a cross product: $[OAB] = \frac{1}{2}\left(a \cdot \frac{1}{b} + b \cdot \frac{1}{a}\right) = \frac{1}{2}\left(\frac{a}{b} + \frac{b}{a}\right)$, and similarly for the other three. Summing,

$$\text{area}(S) \;\geq\; \frac{1}{2}\left(\frac{a}{b} + \frac{b}{a} + \frac{b}{c} + \frac{c}{b} + \frac{c}{d} + \frac{d}{c} + \frac{d}{a} + \frac{a}{d}\right).$$

By AM-GM, each of the four bracketed pairs is at least 2 (equality iff the two entries are equal), so the right side is at least $\frac{1}{2} \cdot 8 = 4$, with equality iff $a = b = c = d$.

This bound is achieved: take $a = b = c = d = 1$, so $A = (1, 1)$, $B = (-1, 1)$, $C = (-1, -1)$, $D = (1, -1)$, and let S be the square with these vertices. S is convex, touches $xy = 1$ exactly at A and C and $xy = -1$ exactly at B and D, and has area $2 \times 2 = 4$. So the minimum is exactly

$$\boxed{4}.$$

Layer 1: WHAT

Block 1: convexity plus "touches all four branches" reduces the problem to choosing one point per branch. Block 2: the four points sweep once around the origin, so the origin sits inside their hull, hence inside S. Block 3: fan S into four triangles from O, argue they tile without overlap, giving a clean lower bound that needs no assumption about the four points being in convex position. Block 4: write each triangle's area as a cross product, sum to the eight term cyclic expression. Block 5: AM-GM on the four reciprocal pairs. Block 6: exhibit the square as the exact extremiser.

The nominal route treats "least area convex set" as an optimisation over an unknown region and reaches for calculus of variations or a general parametrisation of the boundary, which buries the actual content of the problem, the fact that convexity alone forces a hull containment, under machinery that was never necessary. A softer nominal route asserts that S contains the quadrilateral $ABCD$ without checking that $ABCD$ is actually convex in that cyclic order, which happens to fail for some parameter choices; the bound then needs a repair it was never given. The bifurcation point is the choice between naming the hull quadrilateral by picture and hoping it behaves, against building the lower bound from a decomposition, the origin fan, that is valid for every choice of a, b, c, d without exception.

Layer 2: HOW

Block 1. **Cue.** Convex, and required to meet four disjoint pieces of curve. **Why natural.** Convexity turns "contains one point from each piece" into "contains the segment, and hence the hull, of any four such points": this is the definition of convex applied to each pair. **Link to the next block.** Reduces an optimisation over shapes to an optimisation over four positive parameters.

Block 2. **Cue.** The four chosen points lie one in each quadrant. **Why natural.** Four points whose consecutive angular gaps, viewed from a common centre, are all under a straight angle always have that centre inside their hull: no line through the centre can miss all four on one side. **Link to the next block.** Since the hull is inside S, the origin is inside S too, which is exactly the extra point needed to build a triangle fan.

Block 3. **Cue.** A lower bound on area(S) is wanted without knowing whether A, B, C, D are in convex position. **Why natural.** Triangles fanning out from a point that is itself surrounded, angularly, by the other points always tile without overlapping interiors, regardless of whether the outer points happen to form a convex quadrilateral. **Link to the next block.** The sum of the four triangle areas is a valid lower bound for area(S) unconditionally, sidestepping the convex quadrilateral question entirely.

Block 4. **Cue.** Each triangle has one vertex at the origin. **Why natural.** The area of a triangle with one vertex at the origin is half the absolute cross product of the other two position vectors, a one line formula rather than a full shoelace computation. **Link to the next block.** Summing four such halves produces exactly the eight term cyclic sum of ratios.

Block 5. **Cue.** Four pairs, each of the shape $t + \frac{1}{t}$. **Why natural.** Any positive number and its reciprocal sum to at least 2, with equality only when the number is 1; this is the single inequality fact doing all the work here. **Link to the next block.** Gives the numeric bound 4 and identifies the equality condition $a = b = c = d$ in the same step.

Block 6. **Cue.** Equality in AM-GM demands $a = b = c = d$. **Why natural.** An inequality without a matching example is only half a solution; the equality condition tells you exactly which convex set to write down. **Link to the next block.** Setting $a = b = c = d = 1$ produces the square, confirmed directly to touch all four branches, closing the proof.

Layer 3: WHY

The truth underneath the bound is that convexity is a containment statement, hull inside the set, long before it is an optimisation statement, and any problem that says "convex set meeting several disjoint pieces" should immediately be read as "convex set containing the hull of one representative point per piece". The discernment mechanism for choosing the origin fan over the naive quadrilateral is a standing habit: whenever a bound is built from a picture of a nice configuration, four points forming an obviously convex quadrilateral, check the claim against a skewed choice of parameters before trusting it, because a Putnam problem rarely hands you a symmetric picture as the general case for free. Here the skew is real: unequal a, b, c, d can push one of the four points inside the triangle of the other three, and the fan argument is the version of the proof that does not care.

The examiners nudged you toward the hull reduction by making S merely convex, not any more specific shape, and by asking for the branches to be merely intersected, not passed through in a particular way: both choices are invitations to reduce to finitely many points and use convexity as containment. Training this discernment is a matter of always asking, when a convex set is required to meet several separate pieces, what the cheapest witnessing points would give you for free, before reaching for any calculus.

Revolving concepts

Convexity as hull containment: a convex set containing finitely many points contains their convex hull; this converts "touches several regions" into "contains a specific finite configuration". See Book 1, field: Geometry, topic: convex sets and hulls.

A point surrounded angularly lies in the hull: if finitely many points, viewed from a centre, have every consecutive angular gap under a straight angle, the centre lies inside their convex hull. See Book 1, field: Geometry, topic: angular criteria for hull membership.

Triangle area via a vertex at the origin: for points P, Q, the triangle OPQ has area $\frac{1}{2}|x_P y_Q - x_Q y_P|$, the half cross product; summing such triangles sharing vertex O, taken in angularly increasing order with no two overlapping, reproduces the shoelace formula for a polygon. See Book 1, field: Geometry, topic: shoelace and cross product area.

AM-GM on a reciprocal pair: for $t > 0$, $t + \frac{1}{t} \geq 2$ with equality iff $t = 1$; a cyclic sum of ratios around a loop of positive variables is the standard trigger for applying this pair by pair. See Book 1, field: Inequalities, topic: AM-GM and equality analysis.

No recursive or self referential structure appears in this problem; the extremal configuration is found in one inequality step, not by an iterative or inductive construction.

Circuit Bank

- **WHEN** a convex set is required to meet several disjoint pieces of curve **DO** name one point of the set on each piece and work with the hull of those points instead of the set itself **BECAUSE** convexity means the hull of any finite subset of the set is contained in the set, so any bound proved for the hull transfers immediately
- **WHEN** several points, viewed from a common centre, sweep around it with every gap under a straight angle **DO** conclude the centre lies inside their hull, and use it as an extra point **BECAUSE** no line through such a centre can separate it from all the points, which is exactly the condition for hull membership
- **WHEN** a picture of the nice, symmetric configuration seems to justify a bound **DO** stress test it against a skewed, unequal choice of the free parameters before trusting the bound **BECAUSE** a Putnam configuration is rarely handed to you already in the convenient special case, and the proof needs to survive the general one
- **WHEN** an area splits into a cyclic sum of a variable and its reciprocal **DO** apply AM-GM pair by pair and read off both the bound and the equality condition together **BECAUSE** a cyclic sum of $t + \frac{1}{t}$ terms is the single most common inequality shape linking area or perimeter extremal problems to AM-GM

Backtesting

1. First useful cue: four disjoint branches, one convex set, immediately suggests naming one witnessing point per branch and working with their hull.
2. Method the expression suggested: a cyclic sum of ratios, t and $1/t$ terms around a loop, is an AM-GM trigger, applied one pair at a time.
3. Where you would hesitate, and the thought to reinforce: whether the four witnessing points form a convex quadrilateral in the obvious cyclic order; reinforce it by checking a skewed numeric example rather than trusting the picture, and fall back on the origin fan, which needs no such assumption.
4. The trap the question sets: the tempting shortcut, that area(S) is at least the area of quadrilateral $ABCD$, is not automatically true as stated, since that quadrilateral can fail to be convex for unequal parameters; the safe route bounds area(S) by a sum of triangles from a point already known to lie in S.
5. What makes you recognise it faster next time: any time a convex region must meet several separated pieces, reach for hull containment before reaching for calculus on the boundary.

6. The reusable decision rule: when a lower bound depends on several chosen points being in convex position, replace it with a decomposition, such as a fan from an interior point, that is valid regardless of their configuration.

Mirror drills

1. Find the least area of a convex set meeting both branches of $xy = 1$ alone, with no second hyperbola. *Answer:* zero; the segment joining one point on each branch is itself a convex set of area zero meeting both, which is exactly why the original problem needed a second hyperbola, to force four separated pieces instead of two.
2. Four points lie one in each quadrant, at angles $10^\circ, 100^\circ, 190^\circ, 280^\circ$ from the origin. Explain in one line why the origin must lie inside their convex hull. *Answer:* every consecutive angular gap is under 180°, so no line through the origin can put all four points on one side, which is exactly hull membership.
3. Minimise $\frac{a}{b} + \frac{b}{c} + \frac{c}{d} + \frac{d}{a}$ over positive reals with no other constraint, and state the equality case. *Answer:* the minimum is 4, at $a = b = c = d$, by AM-GM on the four terms directly; notice this is only half the shape of the eight term sum used in the main problem, and think about why the solution there needed both cyclic directions.

19.3 A3. Random order avoiding sums divisible by three

Let k be a positive integer. Suppose that the integers $1, 2, 3, \ldots, 3k+1$ are written down in random order. What is the probability that at no time during this process, the sum of the integers that have been written up to that time is a positive integer divisible by 3? Your answer should be in closed form, but may include factorials.

Read it out loud

Every partial sum of a nonempty set of positive integers is itself positive, so "positive integer divisible by 3" is just "divisible by 3": the condition is that no prefix sum is ever a multiple of 3. That is a walk. Read each number by its residue mod 3 and you are walking on $\{0, 1, 2\}$, starting at 0 before anything is written, and asking never to return to 0 once you have moved. The actual sizes of the numbers are irrelevant to the constraint; only their residues mod 3 matter, so the first move is to stop thinking about the numbers $1, \ldots, 3k+1$ and start thinking about a multiset of residues.

The fields wired in are elementary number theory (residues mod 3), a counting argument disguised as a random walk (which residue sequences avoid returning to 0), and probability computed as a ratio of counts rather than by conditioning. What is really being tested is whether you notice that a residue 0 term never changes the running residue, so it can never itself cause a violation except by inheriting whatever residue was already there, and whether you correctly count how the specific residue counts, k zeros, $k+1$ ones, k twos, forced by the range 1 to $3k+1$, constrain which walks are even possible.

The first useful line within thirty seconds: count the residues. Among $1, \ldots, 3k+1$, exactly k are $\equiv 0$, $k+1$ are $\equiv 1$, and k are $\equiv 2 \pmod 3$; the extra 1 over 2 is the single fact the whole answer turns on.

Approaches, ranked worst to best

1. **direct casework or generating functions on k.** *Cue:* a probability in closed form, for general k. *Circuit:* try to build up the count of valid orderings by recursion on the length of the prefix without first reducing to residues. (the actual numbers never matter, only residues do, so working with them directly multiplies the state space for no benefit and rarely lands on a clean closed form)

2. **residue walk without separating the two constraints.** *Cue:* no prefix sum is 0 mod 3. *Circuit:* track the full walk, including the zero residue steps, and attempt to compute the probability of avoidance in one pass. (correct in principle, but conflates "where do the multiples of 3 sit" with "what order are the 1's and 2's in", which is exactly the decomposition that makes the count of valid arrangements land in one line)

3. **count valid residue arrangements directly, as multiset placements.** *Cue:* residue 0 steps never move the walk, and residues 1, 2 force an alternating pattern once you leave 0. *Circuit:* place the k zeros anywhere except first, and show the 1's and 2's have exactly one valid relative order among themselves, then take a ratio of two counts. (the entire probability drops out as a single ratio, with the numerator a clean product of two simple counts)

The ideal solution

Only the residues mod 3 of $1, 2, \ldots, 3k+1$ matter to the condition. Among these $3k+1$ numbers, exactly k are $\equiv 0$, $k+1$ are $\equiv 1$, and k are $\equiv 2 \pmod 3$ (the residue 1 class starts the list and is one longer). Treat 1 as a step of $+1$ and 2 as a step of -1 on $\mathbb{Z}/3\mathbb{Z}$; a 0 residue leaves the running total unchanged. The condition "no prefix sum is 0 mod 3" is exactly "the running residue never returns to 0 after the walk starts".

A residue 0 number placed first immediately gives a prefix sum of 0: forbidden. So the first number written must be a 1 or a 2. After that, the walk moves only when a 1 or 2 is written; a 0 never changes the current residue, so it is safe to insert wherever the current residue already avoids 0, and it can only be dangerous where the residue is already 0, which cannot happen once the walk has started correctly. So the whole condition reduces to two requirements: the first number is not a multiple of 3, and reading only the 1's and 2's in the order they appear, deleting every multiple of 3 from the sequence, the running total mod 3 of this subsequence never hits 0.

For that subsequence: starting at residue 1, the only safe next step is $+1$ again (since -1 would return to 0), landing on residue 2; from residue 2 the only safe step is -1, back to 1. So once the first step is fixed, every subsequent step is forced, and the subsequence of steps is forced into one of exactly two patterns: $+1, +1, -1, +1, -1, \ldots$ or $-1, -1, +1, -1, +1, \ldots$. The first pattern uses one more $+1$ than -1 (it is $k+1$ pluses and k minuses over $2k+1$ terms), the second uses one more -1 than $+1$. Since the residue 1's outnumber the residue 2's by exactly one ($k+1$ against k), only the first pattern is numerically possible: the relative order of the 1's and 2's among themselves is forced into this single pattern, with no freedom left.

Now count. A valid arrangement of $1, \ldots, 3k+1$ is one where the k multiples of 3 occupy any k of the $3k$ positions other than position 1, and the remaining $2k+1$ positions carry the 1's and 2's in the one forced relative pattern (which numbers exactly, among the $k+1$ numbers $\equiv 1$ and k numbers $\equiv 2$, go where among themselves is then free, since only the residue pattern is constrained). The number of ways to choose which positions the multiples of 3 occupy, avoiding position 1, is $\binom{3k}{k}$, choosing k positions from the $3k$ positions other than the first; given that choice, there is exactly one valid residue pattern to lay into the rest, and then $k!$ ways to arrange the actual multiples of 3 among their positions, $(k+1)!$ ways to arrange the actual 1 mod 3 numbers among theirs, and $k!$ ways to arrange the actual 2 mod 3 numbers among theirs, each residue class arranged freely within its own forced slots. So

$$\#\{\text{valid orderings}\} = \binom{3k}{k} k!\,(k+1)!\,k!.$$

The total number of orderings is $(3k+1)!$. Hence

$$\Pr = \frac{\binom{3k}{k} k!\,(k+1)!\,k!}{(3k+1)!} = \frac{\frac{(3k)!}{k!(2k)!}\,k!\,(k+1)!\,k!}{(3k+1)!} = \frac{(3k)!\,(k+1)!\,k!}{(2k)!\,(3k+1)!} = \frac{k!\,(k+1)!}{(3k+1)\,(2k)!},$$

using $(3k+1)! = (3k+1)(3k)!$ in the last step. So

$$\boxed{\Pr = \frac{k!\,(k+1)!}{(3k+1)(2k)!}}.$$

(For $k=1$ this gives $\frac{1!\cdot 2!}{4\cdot 2!} = \frac{1}{4}$, which a direct check of all 24 orderings of $1,2,3,4$ confirms exactly: 6 of them survive.)

Layer 1: WHAT

Block 1: strip the problem down to residues mod 3, count how many of each residue class appear $(k, k+1, k)$. Block 2: read the condition as a walk that must never return to 0, and note a 0 residue step never moves the walk. Block 3: separate the condition into "first term is not 0" and "the subsequence of 1's and 2's never returns to 0 on its own". Block 4: show the second condition forces a single, unique relative pattern of 1's and 2's, using that 1's outnumber 2's by exactly one. Block 5: count valid orderings directly as positions for the zeros, avoiding first, times arrangements within each residue class, and divide by the total $(3k+1)!$.

The nominal route tries to compute this probability by conditioning step by step on the actual numbers written, or by setting up a recursion in k without first discarding everything except

residues; both approaches drown in bookkeeping that the residue reduction throws away for free. The bifurcation point is Block 1: the instant you notice that only residues matter, the problem shrinks from a question about $3k+1$ distinct labelled integers to a question about a multiset of three symbols, and every count becomes a binomial or multinomial count rather than a sum over cases.

Layer 2: HOW

Block 1. **Cue.** The condition only mentions divisibility by 3. **Why natural.** Only the residue mod 3 of each number, not its actual value, can affect whether a partial sum is divisible by 3. **Link to the next block.** Replace the $3k+1$ distinct integers by a multiset of residues, k zeros, $k+1$ ones, k twos, and work with that instead.

Block 2. **Cue.** "No prefix sum divisible by 3" is a statement about a running total returning to a value. **Why natural.** This is exactly a walk on $\mathbb{Z}/3\mathbb{Z}$ starting at 0, and a 0 residue step adds nothing, so it can never be blamed for a violation except by inheriting the current, already illegal, residue. **Link to the next block.** The whole condition splits cleanly into where the zeros are allowed to sit, anywhere but first, and how the nonzero steps must behave on their own.

Block 3. **Cue.** From residue 1 only $+1$ avoids 0; from residue 2 only -1 avoids 0. **Why natural.** Once you leave 0, at each step exactly one of the two possible moves returns to 0 and is therefore forbidden, so the rest of the subsequence is forced, not merely constrained. **Link to the next block.** The subsequence of 1's and 2's has only two candidate patterns, and one is eliminated purely by counting how many of each residue exist.

Block 4. **Cue.** $k+1$ ones against k twos, a specific known imbalance. **Why natural.** Each of the two candidate patterns produces a different, computable count of pluses versus minuses, and the actual imbalance in the problem matches exactly one of them. **Link to the next block.** Only one relative order of the nonzero residues survives, collapsing what looked like a count over many orders into a single arrangement.

Block 5. **Cue.** A count of favourable orderings is wanted, divided by $(3k+1)!$. **Why natural.** Once positions for the zero residues are chosen, avoiding the first slot, and the single valid nonzero pattern is laid into the rest, the specific numbers within each residue class can be permuted freely among their own slots without affecting validity. **Link to the next block.** The favourable count factors as a binomial coefficient times three factorials, and the ratio to $(3k+1)!$ collapses to the closed form.

Layer 3: WHY

The truth underneath the solution is that a constraint stated purely through divisibility only ever depends on residues, and residues live in a much smaller space than the original labelled objects; recognising this is what turns a permutation counting problem with $(3k+1)!$ raw outcomes into one with three symbols and a binomial count. The discernment mechanism is to ask, the moment a problem's condition is phrased with a modulus, whether the actual values are doing any work beyond determining that residue, and to discard them the instant the answer is no.

The second mechanism, spotting that the nonzero steps are forced rather than merely constrained, comes from treating "avoid returning to a state" on a three element cycle as a walk with, at every step, exactly one safe move and one unsafe move once you are off the starting state: this is the general shape of any avoid a residue walk on $\mathbb{Z}/n\mathbb{Z}$ with steps in $\{\pm 1\}$, not a coincidence of $n=3$. Training this discernment means, whenever a walk has few enough states and steps, checking

directly whether some moves are simply illegal at each state rather than assuming a genuine branching count is needed.

The examiners nudged you toward the residue reduction by choosing $1, 2, \ldots, 3k+1$, a range whose residue counts, $k, k+1, k$, are forced and asymmetric by exactly one, rather than a range with equal residue classes; that one count imbalance is the single fact that collapses two candidate walk patterns to one, and is worth reading as a deliberate signal, not an incidental detail of the range chosen.

Revolving concepts

Reduction to residues: a condition phrased entirely through divisibility by m depends only on residues mod m; replace the original labelled objects by a multiset of residues before counting. See Book 1, field: Number theory, topic: working modulo m.

Walks that must avoid a state: a walk on $\mathbb{Z}/n\mathbb{Z}$ with steps restricted to $\{\pm 1\}$, required never to return to its starting state after leaving it, has at each nonstarting state exactly one safe continuation, so the entire step sequence is forced once the first step is chosen; only the actual counts of each step type then decide which, if any, of the finitely many candidate sequences is achievable. Generalisation: for general n and a walk avoiding state 0, the same forcing argument gives at most $n-1$ candidate cyclic patterns, one per starting step, each ruled in or out purely by matching the count of each step type against what the pattern demands over a walk of any given length. See Book 1, field: Sequences and recurrences, topic: forced walks and residue avoidance.

Counting via placement: the number of ways to interleave several labelled groups into ordered positions, once the pattern of which group occupies which position is fixed, is the product of the factorials of the group sizes; combined with a binomial count for where one group's positions sit among the rest, this gives exact favourable counts without casework. See Book 1, field: Combinatorics, topic: multinomial coefficients and placement counting.

Probability as a ratio of two counts built the same way: writing both the favourable count and the total count as products of factorials and binomial coefficients lets heavy cancellation happen before any numbers are substituted, which is how the closed form here collapses to $k!(k+1)!/[(3k+1)(2k)!]$ rather than staying as a raw ratio of large factorials. See Book 1, field: Combinatorics, topic: probability by counting.

Circuit Bank

- **WHEN** a condition is phrased purely through divisibility by a fixed modulus **DO** replace the objects by their residues and discard their actual values **BECAUSE** only residues can affect a divisibility condition, so carrying the actual values further is wasted information
- **WHEN** a walk on finitely many residue states must avoid returning to its start **DO** check, at each state other than the start, how many of the available steps are actually safe **BECAUSE** on a small cycle with steps of size 1, typically only one step out of each state avoids the forbidden state, turning what looks like a branching count into a forced sequence
- **WHEN** a forced pattern's step counts can be compared against the known composition of the underlying multiset **DO** use the mismatch, or match, in counts to eliminate all but one candidate pattern **BECAUSE** a forced pattern predicts an exact count of each step type, and the problem's own numbers either satisfy that prediction or rule the pattern out entirely, with no need for further casework
- **WHEN** a probability is wanted for a random permutation subject to a placement style condition **DO** write both favourable and total counts as products of binomial coefficients and factorials before dividing **BECAUSE** heavy algebraic cancellation between the two counts is the difference between a clean closed form and an unreduced mess

Backtesting

1. First useful cue: the condition only ever mentions divisibility by 3, so only residues mod 3 can matter.
2. Method the expression suggested: a walk on three residue states that must never return to 0, with the zero residue elements acting as steps that do not move the walk at all.
3. Where you would hesitate, and the thought to reinforce: convincing yourself that the nonzero steps are entirely forced rather than merely restricted; reinforce it by checking, at each non-starting residue, exactly which of the two possible steps returns to 0 and eliminating only that one.
4. The trap the question sets: it is tempting to treat "first element not a multiple of 3" and "the 1,2 pattern is valid" as needing a joint, harder argument, when in fact the residue 0 elements can be inserted into any of the non first slots independently once the nonzero pattern is fixed, so a clean product of counts suffices.
5. What makes you recognise it faster next time: any random ordering question whose condition is stated purely via a modulus should immediately be rewritten as a question about a multiset of residues, before any probability machinery is set up.
6. The reusable decision rule: when a walk on a small cyclic residue system must avoid a state, check state by state which steps are forced, then match the resulting pattern's step counts against the actual composition of the multiset to see which patterns survive.

Mirror drills

1. The integers $1, 2, \ldots, 2k+1$ are written in random order. Find the probability that no prefix sum is ever even. *Hint:* the forbidden state is residue 0 mod 2, and the odd numbers outnumber the even ones by exactly one among $1, \ldots, 2k+1$; the same forcing argument applies with one fewer residue class to track.
2. The integers $1, \ldots, 3k$, with no extra term, are written in random order. Explain in one line why the probability that no prefix sum is divisible by 3 must be 0. *Answer:* residues $0, 1, 2$ each occur exactly k times, so the forced pattern would need one more of one nonzero residue than the other, which the equal counts can never supply.
3. For a walk on $\mathbb{Z}/4\mathbb{Z}$ with steps in $\{\pm 1\}$ starting at 0 and never allowed to return to 0, list the candidate step patterns and state, for each, the exact count of $+1$ steps versus -1 steps it requires over a walk of length n.

19.4 A4. Polynomials sending repunits to repunits

A repunit *is a positive integer whose digits in base 10 are all ones. Find all polynomials f with real coefficients such that if n is a repunit, then so is $f(n)$.*

Read it out loud

Read the structure, not the words: this is a rigidity question about a polynomial that is constrained only on a sparse infinite set. The set is $1, 11, 111, \dots$, and the constraint is that f maps that set into itself. Two fields are wired in, and only two. The first is elementary number theory, and you need exactly one fact from it, namely how to write a repunit in closed form. The second is the identity theorem for polynomials: a polynomial is a rigid object, so the instant you pin its values at infinitely many points you have pinned the polynomial itself. Everything between those two facts is bookkeeping.

The first cue is the word repunit. A named set is an invitation to re-describe it, because a name is a hint that the natural description is not the one printed. Write R_m for the repunit with m ones. Then $R_{m+1} = 10R_m + 1$, which is an affine recursion, not a multiplicative one, and that single observation explains why the set looks awkward: polynomials interact beautifully with multiplicative sets and badly with additive ones. So inside thirty seconds you should have on the page

$$R_m = \frac{10^m - 1}{9}, \qquad 9R_m + 1 = 10^m.$$

The second line is the entire problem. The affine map $x \mapsto 9x + 1$ carries the repunits bijectively onto $\{10, 100, 1000, \dots\}$, a geometric set, and a geometric set has the property you are about to exploit: on a multiplicative scale it is discrete, so nothing can sneak up on it.

What is really being tested is whether you will change coordinates. A polynomial condition stated on an ugly set is almost never attacked on that set. It is attacked by conjugating f with the affine map that cleans the set up, because affine substitution is an automorphism of the polynomial ring: it preserves degree, it is invertible, and it carries polynomials to polynomials in both directions. If you do not conjugate, you will spend the whole session pushing 9s around.

One more cue, planted deliberately: *real* coefficients. The examiners are stopping you from assuming f has integer coefficients, and they are right to, because the answer contains a factor 10^c with c allowed to be negative. Any argument of yours that quietly uses integrality of the coefficients is a wrong argument that happens to reach a correct list.

Approaches, ranked worst to best

1. **Undetermined coefficients.** *Cue:* "find all polynomials" with a finite list of sample inputs $1, 11, 111$. *Circuit:* write $f(x) = \sum a_i x^i$ and substitute. This is the zombie route. You get equations $f(R_m) = R_{h(m)}$ in which the outputs carry unknown *exponents* $h(m)$, so the system is not linear, not finite, and not closed. It stalls immediately, and it stalls for a reason worth naming: the unknown is not just the coefficient vector, it is also the index function h.

2. **Guess the family and stop.** *Cue:* the examples $f(x) = x$, $f(x) = 10x + 1$ are easy to find. *Circuit:* pattern-match, verify, hope. You will find $f(n) = 10n + 1$ within a minute because it appends a one. Verification is instant. But "find all" is a uniqueness demand, and a list with no proof of completeness is worth very little on the ten point scale. Rank it here only because it produces the correct family, which is a real asset for the next two routes.

3. **Growth comparison in the original coordinate.** *Cue:* f has a degree, repunits grow geometrically. *Circuit:* compare $\log_{10}$ of both sides. Write $f(R_m) = R_{h(m)}$. Since $R_m \sim 10^m/9$ and $f(x) \sim ax^d$, you get $h(m) = dm + \log_{10}(9^{1-d}a) + o(1)$, and since $h(m)$ is an integer it is

eventually exactly $dm + c$. That is a complete proof, and it is essentially the ideal one; it just carries a cloud of 9s and a nuisance $o(1)$ through every line. It costs more ink for no gain.

4. **Rationality plus a valuation argument.** *Cue:* polynomials taking rational values at rational points. *Circuit:* Lagrange interpolation, then induct on the degree. By interpolation f has rational coefficients, so $g(0)$ is rational; if $g(0) \neq 0$ its numerator has a bounded power of 10 in it, while $g(10^k)$ eventually has an unbounded one, so $g(0) = 0$ and you induct on $g(n)/n$. This is the 2-adic and 5-adic mirror of the ideal route and it is genuinely instructive, but it needs interpolation and a divisibility estimate you must set up carefully.

5. **Affine conjugation plus discreteness of the powers of ten.** *Cue:* $9n + 1 = 10^m$, and a polynomial is determined by infinitely many values. *Circuit:* move to the fixed point of the recursion, then read off the leading coefficient as a limit. Best. One substitution turns the problem into "polynomials mapping powers of 10 to powers of 10", after which a single limit finishes it. The whole solution is six lines.

The ideal solution

Write R_m for the repunit with m ones, $m \geq 1$, so that

$$R_m = \frac{10^m - 1}{9}, \qquad 9R_m + 1 = 10^m.$$

Thus a positive integer n is a repunit if and only if $9n + 1$ is a power of 10 exceeding 1.

Given a real polynomial f, define the real polynomial

$$g(x) = 9f\left(\frac{x-1}{9}\right) + 1, \qquad \text{equivalently} \qquad f(x) = \frac{g(9x+1) - 1}{9}.$$

The two formulas are mutually inverse, so $f \mapsto g$ is a bijection from real polynomials to real polynomials, and it preserves degree. For every $m \geq 1$,

$$g(10^m) = 9f(R_m) + 1,$$

so $f(R_m)$ is a repunit if and only if $g(10^m)$ is a power of 10 exceeding 1. Hence the problem is equivalent to the following.

Claim 19.4.1. *The real polynomials g such that $g(10^m) \in \{10^k : k \geq 1\}$ for every integer $m \geq 1$ are exactly the polynomials*

$$g(x) = 10^c x^d, \qquad d \in \mathbb{Z}_{\geq 0},\ c \in \mathbb{Z},\ c + d \geq 1.$$

Proof. Sufficiency first. If $g(x) = 10^c x^d$ with $d \geq 0$ and $c + d \geq 1$, then $g(10^m) = 10^{c+md}$, and for $m \geq 1$ and $d \geq 0$ we have $c + md \geq c + d \geq 1$, so $g(10^m)$ is a power of 10 exceeding 1.

Now necessity. The zero polynomial fails at $x = 10$, so $g \neq 0$; let $d = \deg g \geq 0$ and let a be its leading coefficient. Since $g(10^m) \geq 10 > 0$ for all $m \geq 1$ and $10^m \to \infty$, the leading coefficient satisfies $a > 0$.

For each $m \geq 1$ write $g(10^m) = 10^{k_m}$ with k_m a positive integer. Then

$$\frac{g(10^m)}{10^{md}} = 10^{k_m - md},$$

and the left hand side tends to a as $m \to \infty$, because $g(x)/x^d \to a$ as $x \to \infty$.

The set of integer powers of 10 has no accumulation point in $(0, \infty)$, and here is the quantitative form of that, which is all we need. Choose M so that $|10^{k_m - md} - a| < a/2$ for all $m \geq M$. Then every such value lies in the interval $(a/2, 3a/2)$, whose endpoints have ratio $3 < 10$, so the interval

contains at most one integer power of 10. Therefore the numbers $10^{k_m - md}$ for $m \geq M$ are all equal to a single value 10^c with $c \in \mathbb{Z}$, and passing to the limit gives $a = 10^c$.

Consequently $g(10^m) = 10^c \cdot 10^{md} = 10^c (10^m)^d$ for every $m \geq M$. The polynomial $g(x) - 10^c x^d$ therefore vanishes at the infinitely many distinct points $10^M, 10^{M+1}, \ldots$, so it is the zero polynomial and $g(x) = 10^c x^d$.

Finally, evaluating at $m = 1$ shows $10^{c+d} = g(10)$ is a power of 10 exceeding 1, so $c + d \geq 1$. $\square$

Translating back through $f(x) = (g(9x+1) - 1)/9$, the answer is

$$\boxed{f(x) = \frac{10^c (9x+1)^d - 1}{9}, \qquad d \in \mathbb{Z}_{\geq 0},\ c \in \mathbb{Z},\ c \geq 1 - d}$$

and every such f works, sending R_m to R_{c+md}.

Three sanity checks, each of which you should do in the arena. Taking $d = 0$ and $c \geq 1$ gives the constant polynomials $f \equiv R_c$. Taking $d = 1, c = 0$ gives $f(x) = x$. Taking $d = 1, c = 1$ gives $f(x) = 10x + 1$, which appends a one. Taking $d = 2, c = -1$ gives $f(x) = \frac{1}{10}\left(\frac{(9x+1)^2}{9}\right) - \frac{1}{9}$, and indeed $f(R_m) = R_{2m-1}$, for example $f(11) = 111$.

Layer 1: WHAT

Block 1. Name the set properly: n is a repunit if and only if $9n + 1 = 10^m$ with $m \geq 1$.

Block 2. Conjugate: define $g(x) = 9f((x-1)/9) + 1$, note the correspondence is a degree preserving bijection, and restate the problem as "polynomials carrying powers of 10 to powers of 10".

Block 3. Normalise by the degree: look at $g(10^m)/10^{md}$, which tends to the leading coefficient a.

Block 4. Use discreteness: a sequence of integer powers of 10 converging to a positive limit is eventually constant, so $a = 10^c$ and $g(10^m) = 10^c 10^{md}$ for all large m.

Block 5. Close with rigidity: $g(x) - 10^c x^d$ has infinitely many roots, hence is identically zero.

Block 6. Translate back and pin the inequality by testing $m = 1$: $c + d \geq 1$. Verify sufficiency.

The nominal route branches at Block 1. The zombie writes $f(x) = \sum a_i x^i$ and substitutes $1, 11, 111$, generating equations whose right hand sides contain unknown exponents; that route never closes. The bifurcation point is the decision of *which coordinate to work in*: the moment you choose to study $9n + 1$ rather than n, every later step is forced. Notice what the ideal route does not do: it never expands a single power of $(9x + 1)$, and it never uses a divisibility argument. Ink saved is error avoided.

Layer 2: HOW

Block 1. Cue. The problem names a set instead of describing it algebraically, and the set is defined by a base 10 digit pattern. **Why natural.** A digit pattern of all ones is a geometric series: $R_m = 1 + 10 + \cdots + 10^{m-1}$. Summing it is not a manoeuvre, it is the definition of what the digits mean. The sum is $(10^m - 1)/9$, and clearing the 9 gives the clean characterisation $9R_m + 1 = 10^m$. There is also a deeper reason to look for exactly this shift, which I give in Layer 3: $x \mapsto 10x + 1$ generates the repunits, and its fixed point is $-1/9$. **Link to the next block.** You now hold a set that is awkward for polynomials, and beside it a set that is not. That asymmetry is the instruction to move.

Block 2. Cue. Two descriptions of the same set differ by the affine map $x \mapsto 9x + 1$. **Why natural.** Substituting an affine function into a polynomial is an automorphism of $\mathbb{R}[x]$: it is invertible with affine inverse, it preserves degree, and it sends real polynomials to real polynomials. So there is no loss whatever in transporting the whole problem across it. This is the same reasoning that lets you

complete the square without worrying that you have changed the problem. The condition "f maps repunits into repunits" becomes "g maps $\{10^m : m \geq 1\}$ into the powers of 10". **Link to the next block.** The new condition is multiplicative in shape, and the objects it constrains, powers of 10, are spread out geometrically. Both features are exactly what a polynomial growth argument wants.

Block 3. **Cue.** A polynomial constrained at points going to infinity, with the constraint set thinning out. **Why natural.** The only thing about a polynomial that survives at infinity is its degree and leading coefficient: $g(x)/x^d \to a$. Dividing by x^d is therefore not an idea, it is the unique way to extract information from "large x" behaviour. And on the constraint set the quotient is $10^{k_m - md}$, that is, it stays inside a fixed discrete set. **Link to the next block.** A convergent sequence trapped in a discrete set is eventually constant. That is the pin.

Block 4. **Cue.** A convergent sequence whose terms all lie in $\{10^j : j \in \mathbb{Z}\}$. **Why natural.** Distinct integer powers of 10 differ by a factor of at least 10, so any interval whose endpoints have ratio less than 10 contains at most one of them. Take the interval $(a/2, 3a/2)$, which the terms eventually enter because they converge to $a > 0$. Hence the terms are eventually a single value, and that value is the limit a. **Link to the next block.** You have upgraded an asymptotic statement into an exact one: $g(10^m) = 10^c(10^m)^d$ for all large m, an identity at infinitely many points.

Block 5. **Cue.** An identity between two polynomials holding at infinitely many points. **Why natural.** A nonzero polynomial of degree D has at most D roots. So a polynomial with infinitely many roots is the zero polynomial, and two polynomials agreeing at infinitely many points are equal as polynomials. This is the rigidity that makes the whole problem finite. **Link to the next block.** The form of g is now known exactly; only the admissible range of c and d remains.

Block 6. **Cue.** The problem said "repunit", which means a positive integer, so $m \geq 1$ everywhere, not $m \geq 0$. **Why natural.** Boundary conditions of this kind are decided by testing the extreme case, here $m = 1$. The requirement $c + md \geq 1$ for all $m \geq 1$ with $d \geq 0$ is monotone in m, so it is equivalent to its value at $m = 1$. Then transport back through the inverse substitution and check sufficiency, since "find all" demands both directions. **Link to the next block.** Nothing follows; the answer is complete, and you have written the verification that a grader will look for.

Layer 3: WHY

Three mathematical truths make this route the right one.

First, rigidity. Polynomials form a vector space of finite dimension in each degree, so agreement at enough points is equality. Every "find all polynomials such that" problem is ultimately solved by manufacturing infinitely many agreements with a guessed candidate. Once you internalise that the endgame is always "two polynomials agree infinitely often, hence coincide", you stop trying to force the coefficients and start trying to force the values.

Second, coordinates. The ring $\mathbb{R}[x]$ has affine substitutions as automorphisms. Any hypothesis phrased on a set S can be rephrased on $\varphi(S)$ for affine φ at zero cost. So the correct first question for a constrained polynomial problem is not "what is f" but "what is the cleanest affine image of the constraint set". Here the repunits are the orbit of 1 under $T(x) = 10x + 1$. An affine map with slope $b \neq 1$ has a unique fixed point $p = c/(1 - b)$, and conjugating by the translation $x \mapsto x - p$ turns T into pure multiplication by b. For $T(x) = 10x + 1$ the fixed point is $p = -1/9$, and indeed $R_m + \frac{1}{9} = \frac{10^m}{9}$. The substitution $9n + 1$ is nothing more than "translate to the fixed point, then rescale". That is the honest reason the substitution is natural, and it is a reason you can reuse on any problem whose special set is an orbit of an affine map.

Third, discreteness. The powers of 10 have no positive accumulation point, so any convergent sequence inside them is eventually constant. This is the mechanism that upgrades "asymptotically" to "exactly", and it is the same mechanism behind "an integer sequence that converges is eventually constant", which you will meet constantly. Whenever a limit lands in a discrete set, you have not merely estimated, you have determined.

The discernment mechanism, stated as a rule you can run in the arena: when a polynomial is constrained on a set S, ask three questions in order. Is S an orbit of a simple map, so that a change of coordinate makes it geometric or arithmetic? Is the target set discrete on some scale? Does the constraint survive to infinity, so that the leading coefficient can be read off as a limit? If all three answer yes, the problem is a six line problem. Training it is a matter of collecting constrained sets and rehearsing the normalising substitution for each: repunits become powers of 10; numbers of the form $2^k - 1$ become powers of 2; an arithmetic progression $a + bk$ becomes $\mathbb{Z}$; the Fibonacci numbers do not normalise, which is precisely why Fibonacci constraints are hard.

What the examiners nudged with: the word repunit (a named set, therefore re-describe it), the phrase "real coefficients" (a warning not to use integrality), and the phrase "find all" (a demand for both necessity and sufficiency).

Revolving concepts

Repunits and base representation. $R_m = (10^m - 1)/9$ is the base 10 instance of $\frac{b^m-1}{b-1}$, the repunit in base b. All the structure used here is base free: n is a base b repunit if and only if $(b-1)n + 1 = b^m$. See Book 1, field: Number theory, topic: base representation and digit sums.

Affine maps and their fixed points, with the recursion generalised. The repunits satisfy $R_1 = 1$, $R_{m+1} = 10R_m + 1$. The general first order affine recursion is $u_{m+1} = bu_m + c$ with $b \neq 1$; its unique fixed point is $p = c/(1-b)$, and $u_m - p = b^m(u_0 - p)$, so

$$u_m = b^m \left(u_0 - \frac{c}{1-b} \right) + \frac{c}{1-b}.$$

The general solution method is exactly the conjugation used above: translate the fixed point to the origin, whereupon the recursion becomes multiplication by b. If $b = 1$ the map is a translation and $u_m = u_0 + mc$ is arithmetic instead. Every time a problem hands you a set generated by an affine rule, run this normalisation before anything else. See Book 1, field: Sequences and recurrences, topic: first order linear recurrences.

The identity theorem for polynomials. A nonzero polynomial of degree D over a field has at most D roots; hence two polynomials agreeing at more than $\max(\deg)$ points are identical, and agreement at infinitely many points is equality. This is the finiteness engine of every "find all polynomials" problem. See Book 1, field: Polynomials, topic: roots and the identity theorem.

Behaviour at infinity. For $\deg g = d$ with leading coefficient a, $g(x)/x^d \to a$. This is the only asymptotic fact about polynomials you need, and it converts leading coefficient questions into limit questions. See Book 1, field: Polynomials, topic: growth and leading terms.

Discrete sets and eventual constancy. A subset $S \subseteq \mathbb{R}$ with no accumulation point at L has the property that any sequence in S converging to L is eventually constant. Instances: the integers, the powers of a fixed base, any set with a uniform multiplicative gap. See Book 1, field: Real analysis, topic: discreteness and eventual constancy.

Rational values and interpolation. If a polynomial of degree d takes rational values at $d + 1$ distinct rational points, then Lagrange interpolation writes it with rational coefficients. That is the entry point to the p-adic route, and it is worth installing separately because it converts an arithmetic hypothesis into a structural one. See Book 1, field: Polynomials, topic: Lagrange interpolation.

Circuit Bank

- **WHEN** a polynomial is constrained on a named set such as repunits, Mersenne numbers, or an orbit of an affine map **DO** first rewrite the set in closed form, then conjugate the polynomial by the affine map that turns the set into a geometric or arithmetic one **BECAUSE** affine substitution is a degree preserving automorphism of the polynomial ring, so it costs nothing and it can turn an additive description into a multiplicative one
- **WHEN** a set is generated by $u_{m+1} = bu_m + c$ with $b \neq 1$ **DO** translate by the fixed point $c/(1-b)$ **BECAUSE** the conjugated map is multiplication by b, so the orbit becomes a geometric progression
- **WHEN** you must prove a polynomial equals a specific polynomial **DO** prove they agree at infinitely many points and invoke the identity theorem **BECAUSE** a nonzero polynomial has finitely many roots, so agreement at infinitely many points is equality
- **WHEN** the constraint places $g(x_m)$ inside a discrete set while $x_m \to \infty$ **DO** divide by $x_m^{\deg g}$ and take the limit **BECAUSE** the quotient converges to the leading coefficient and is trapped in a discrete set, so it is eventually constant, which pins the leading coefficient exactly
- **WHEN** the question says "find all" **DO** write the verification of your family as a separate paragraph **BECAUSE** necessity and sufficiency are two different proofs and the grader marks both

Backtesting

1. First useful cue: the word repunit, a named set, which forces the closed form $R_m = (10^m - 1)/9$ and then the clean statement $9R_m + 1 = 10^m$.
2. Method the expression suggested: conjugate f by $x \mapsto 9x + 1$ to reduce to polynomials sending powers of 10 to powers of 10, then read the leading coefficient off as a limit and close with the identity theorem.
3. Where you would hesitate, and the thought to reinforce: at the limit step, because it feels like an analysis move inside a number theory problem. Reinforce this sentence: a convergent sequence inside a discrete set is eventually constant, so the limit is not an estimate, it is an exact determination.
4. The trap the question sets: assuming integer coefficients. The correct family contains 10^c with c possibly negative, so f need not have integer coefficients, and any argument built on integrality is invalid even though it may reach the right list. A second, smaller trap: forgetting that repunits are positive integers, so the exponent must be at least 1, which is exactly what produces the constraint $c \geq 1 - d$.
5. What makes you recognise it faster next time: the shape "polynomial constrained on a sparse set closed under a simple map". That shape always means normalise the set first.
6. The reusable decision rule: when the hypothesis lives on an ugly set, change coordinates until the set is an orbit under multiplication, then let growth plus discreteness do the work.

Mirror drills

1. Find all real polynomials f such that $f(2^m)$ is a power of 2 for every integer $m \geq 1$. *Answer:* identical mechanism with no conjugation needed: $f(x) = 2^c x^d$ with $d \geq 0$ and $c + d \geq 0$, the last inequality coming from whichever positivity convention the powers are required to satisfy.

2. Find all real polynomials f mapping the set $\{2^m - 1 : m \geq 1\}$ into itself. *Answer:* the generating map is $x \mapsto 2x + 1$ with fixed point -1, so put $g(x) = f(x-1) + 1$; then g must send powers of 2 to powers of 2, giving $f(x) = 2^c(x+1)^d - 1$ with $d \geq 0$ and $c + d \geq 1$.

3. Let $u_0 = 2$ and $u_{m+1} = 3u_m - 1$. Find all real polynomials f with $f(u_m) \in \{u_j : j \geq 0\}$ for every m. *Answer:* the fixed point is $1/2$, so $u_m = \frac{1}{2}(3^{m+1} + 1)$ and $2u_m - 1 = 3^{m+1}$; conjugating by $x \mapsto 2x - 1$ reduces to polynomials sending powers of 3 to powers of 3, so $f(x) = \frac{1}{2}\left(3^c(2x-1)^d + 1\right)$ with $d \geq 0$ and $c + d \geq 1$.

19.5 A5. Counting elements of prime order in a group

Suppose that a finite group has exactly n elements of order p, where p is a prime. Prove that either $n = 0$ or p divides $n + 1$.

Read it out loud

Read it as a congruence about a cardinality: a finite group G is given, a prime p is given, and the count of elements of a prescribed order is claimed to satisfy $n \equiv -1 \pmod{p}$ whenever it is nonzero. The field is finite group theory, and the sub-field is the one that matters most on the Putnam: group actions used as counting devices. Nothing here is about the structure of G; G is arbitrary, so no classification, no commutativity, no Sylow structure is available to you as a hypothesis. Only counting can survive that level of generality.

The decisive cue is the $+1$. Why would a natural theorem be about $n + 1$ rather than n? Because $n + 1$ is the count of a set that is closed and canonical, whereas n is the count of a set that is not. Write

$$S_p = \{x \in G : x^p = e\}.$$

Since p is prime, $x^p = e$ forces $\operatorname{ord}(x) \in \{1, p\}$, so S_p consists of the identity together with the n elements of order exactly p, giving $|S_p| = n + 1$. The claim to be proved is therefore the cleaner statement: if G contains any element of order p, then p divides $|S_p|$. Within thirty seconds of reading the problem those two lines should be on your page, and with them the problem has already changed shape.

The second cue is the word "exactly" attached to a count with no structure attached to G. When a count in an arbitrary finite group must be divisible by p, there is one reliable mechanism: exhibit a set whose cardinality you know outright, let a group of order p act on it, and observe that every orbit has size 1 or p. The count you want then appears as the number of fixed points. Everything in the ideal solution is an implementation of that sentence.

A caveat worth clarifying before you start. The set $\{x : \operatorname{ord}(x) = p\}$ itself is a poor object: it is not closed under any natural operation, and it is not the fixed point set of anything. That is why the theorem is stated with the $+1$. Adding the identity back is not cosmetic; it is what creates the structure.

Approaches, ranked worst to best

1. **Case analysis on the structure of G.** *Cue:* finite group, small prime, familiar examples. *Circuit:* classify, then count in each case. Hopeless and not even a route. G is an arbitrary finite group; there is no list to run through. Rank it last because attempting it consumes the session.

2. **Pair each element with its inverse.** *Cue:* elements of order p come in pairs x, x^{-1}. *Circuit:* involution counting. This proves only that n is even when $p > 2$, since $x \neq x^{-1}$ for $\operatorname{ord}(x) = p > 2$. It is the reflex of someone who knows one counting circuit and applies it regardless of the modulus. It gives divisibility by 2, not by p.

3. **Partition into cyclic subgroups.** *Cue:* every element of order p generates a subgroup of order p, and distinct such subgroups meet trivially. *Circuit:* orbit style partition. This is genuinely useful and you should always write it down: if k is the number of subgroups of order p, then $n = k(p - 1)$, because each such subgroup is cyclic of prime order and contains exactly $p - 1$ elements of order p, and each element of order p lies in exactly one of them, namely the one it generates. Hence $n + 1 = k(p - 1) + 1 \equiv 1 - k \pmod{p}$, so the problem is *equivalent* to proving $k \equiv 1 \pmod{p}$. That is a Sylow flavoured statement, and it is not easier than the original. Rank it here: it reformulates cleanly but does not close.

4. **Quote Frobenius.** *Cue:* a divisibility statement about solutions of $x^m = e$. *Circuit:* cite a named theorem. Frobenius proved that if m divides $|G|$ then m divides $|\{x : x^m = e\}|$. That instantly gives $p \mid n+1$. The trouble is that you cannot prove Frobenius in the arena, and a proof consisting of a citation of a theorem strictly stronger than the problem invites a grader to ask whether you have any idea why it is true. Use it as a sanity check, never as the write-up.

5. **Conjugation action of a fixed subgroup of order p.** *Cue:* if $n > 0$ there is a subgroup H of order p to act with. *Circuit:* let H act on the other elements of order dividing p by conjugation. Take $H = \langle a \rangle$ where a has order p, and let H act by conjugation on the set of elements of order p outside H. Orbits have size 1 or p; the fixed points are elements commuting with H, and each such element together with H generates an elementary abelian group of order p^2, which lets you group the fixed points into blocks of size $p^2 - p$. Correct and self contained, but the block counting needs care and the write-up is twice as long as it needs to be.

6. **Cyclic rotation of p tuples with product the identity.** *Cue:* a count divisible by p is wanted, and $|G|^{p-1}$ is a count you know exactly. *Circuit:* build a set of known size, act by $\mathbb{Z}/p$, count fixed points. Best. One set, one action, three lines of verification. It uses no Sylow theory, no Cauchy theorem, and no structure of G beyond Lagrange, and as a bonus it proves Cauchy's theorem on the way past.

The ideal solution

Let G be a finite group with identity e and $|G| = m$, and let p be prime. Put

$$S_p = \{x \in G : x^p = e\}.$$

Since p is prime, an element with $x^p = e$ has order dividing p, hence order 1 or p; the only element of order 1 is e. Therefore

$$|S_p| = n + 1.$$

Lemma 19.5.1. *If p divides m, then p divides $|S_p|$.*

Proof. Let

$$S = \{(a_0, a_1, \ldots, a_{p-1}) \in G^p : a_0 a_1 \cdots a_{p-1} = e\}.$$

The first $p-1$ coordinates may be chosen arbitrarily, and then $a_{p-1} = (a_0 \cdots a_{p-2})^{-1}$ is uniquely determined and lies in G. Hence the map $(a_0, \ldots, a_{p-1}) \mapsto (a_0, \ldots, a_{p-2})$ is a bijection $S \to G^{p-1}$, so

$$|S| = m^{p-1}.$$

Because $p \mid m$ and $p - 1 \geq 1$, we get $p \mid |S|$.

Define $\sigma : G^p \to G^p$ by the cyclic shift $\sigma(a_0, a_1, \ldots, a_{p-1}) = (a_1, \ldots, a_{p-1}, a_0)$. This maps S into S: if $a_0 a_1 \cdots a_{p-1} = e$ then $a_1 \cdots a_{p-1} = a_0^{-1}$, so

$$a_1 \cdots a_{p-1} a_0 = a_0^{-1} a_0 = e.$$

Since σ^p is the identity map, $j \mapsto \sigma^j$ is a homomorphism from the cyclic group $\mathbb{Z}/p$ to the symmetric group on S, that is, an action of $\mathbb{Z}/p$ on S.

By the orbit stabiliser theorem every orbit of this action has size dividing $|\mathbb{Z}/p| = p$, so every orbit has size 1 or p. Consequently, writing F for the number of orbits of size 1,

$$|S| \equiv F \pmod{p},$$

since S is the disjoint union of its orbits and the orbits of size p contribute 0 modulo p.

Finally identify the fixed points. A tuple $(a_0, \ldots, a_{p-1}) \in S$ is fixed by σ if and only if $a_i = a_{i+1}$ for all indices i read modulo p, that is, if and only if all coordinates are equal, say to a. Such a tuple

lies in S exactly when $a^p = e$. So the fixed points of σ on S are precisely the constant tuples $(a, \ldots, a)$ with $a \in S_p$, and a tuple fixed by σ is fixed by every power of σ, so these are exactly the orbits of size 1. Hence $F = |S_p|$, and

$$0 \equiv |S| \equiv |S_p| \pmod{p},$$

which is the assertion. □

Now the problem itself. If $n = 0$ there is nothing to prove. If $n > 0$, pick $x \in G$ of order p. Then $\langle x \rangle$ is a subgroup of G of order p, so by Lagrange's theorem p divides m. The lemma applies and gives

$$p \mid |S_p| = n + 1,$$

as required. ■

Remark 19.5.2. The lemma is stronger than the problem, and it contains Cauchy's theorem. If $p \mid m$ then $p \mid |S_p|$, and S_p is nonempty because it contains e, so $|S_p| \geq p \geq 2$; hence G has an element of order p. In other words, the same three lines that solve this problem prove that a prime dividing the order of a group is realised as the order of an element.

Layer 1: WHAT

Block 1. Replace the set of elements of order exactly p by $S_p = \{x : x^p = e\}$, and record $|S_p| = n + 1$ using primality of p.

Block 2. Dispose of the trivial case: if $n > 0$ then Lagrange forces $p \mid |G|$, so the only case needing work is $p \mid m$.

Block 3. Manufacture a set of exactly known size: p tuples with product e, of size m^{p-1}, which is divisible by p.

Block 4. Find a $\mathbb{Z}/p$ action on that set: cyclic shift, well defined because the product condition is rotation invariant.

Block 5. Orbits have size 1 or p, so the cardinality is congruent to the number of fixed points modulo p.

Block 6. Identify the fixed points as the constant tuples $(a, \ldots, a)$ with $a^p = e$, hence with S_p, and conclude.

The nominal route branches at Block 1. The zombie tries to count elements of order p directly, pairing x with x^{-1} or with its conjugates, and gets divisibility by 2 or a statement about conjugacy class sizes that does not assemble. The bifurcation point is the decision *which set to count*: elements of order exactly p, or elements of order dividing p. The second set is a fixed point set; the first is not. Everything downstream follows from that choice, and the $+1$ in the problem statement was the examiners telling you which one they had in mind.

Layer 2: HOW

Block 1. **Cue.** The conclusion is about $n + 1$, not n, and p is prime. **Why natural.** A shifted count almost always means the natural object is the count before the shift. Primality of p says the divisors of p are 1 and p, so "order divides p" splits as "identity, or order exactly p". Therefore $|S_p| = n + 1$ exactly, and the problem becomes a statement about S_p. The reason to want S_p rather than the order p elements is structural: S_p is defined by an equation, and sets defined by equations are the ones that arise as fixed point sets. **Link to the next block.** You now need to prove that p divides the size of a set defined by the equation $x^p = e$, in a group with no other hypotheses.

Block 2. **Cue.** The dichotomy "$n = 0$ or" in the statement is an invitation to split immediately. **Why natural.** If an element of order p exists, it generates a cyclic subgroup of order p, and Lagrange's theorem says the order of a subgroup divides the order of the group. So $n > 0$ forces $p \mid m$. Clearing

this in one sentence is not tidiness; it buys the hypothesis $p \mid m$, which is the only hypothesis the main argument needs. **Link to the next block.** The remaining task is a clean lemma: $p \mid m$ implies $p \mid |S_p|$.

Block 3. **Cue.** You must produce a multiple of p out of thin air, and the only number you know exactly is m. **Why natural.** If you want to count something modulo p, you need a set whose size you can compute without any knowledge of G. The one such construction available in an arbitrary group is a tuple space with one constraint: fix $p-1$ coordinates freely and let the last one absorb the constraint. That gives $|S| = m^{p-1}$, computed by an explicit bijection, and the exponent $p-1$ is at least 1, so $p \mid m$ transfers to $p \mid |S|$. The choice of p as the tuple length is not arbitrary: you want the acting group to have order p so that orbit sizes are only 1 or p. **Link to the next block.** A set of size divisible by p is useless without an action; the constraint $a_0 \cdots a_{p-1} = e$ was chosen because it is rotation invariant.

Block 4. **Cue.** The defining condition of S is a product around a cycle, and a product around a cycle does not change when you rotate the starting point. **Why natural.** In a group, $a_0 a_1 \cdots a_{p-1} = e$ is equivalent to $a_1 \cdots a_{p-1} = a_0^{-1}$, and multiplying on the right by a_0 gives the rotated product. This is conjugation of the identity by a_0 in disguise, and it is exactly why one writes the condition as a product equal to e rather than, say, $a_0 a_1 \cdots a_{p-1} = g$ for some other g. The tuple length being p makes σ^p the identity, so $\mathbb{Z}/p$ acts. **Link to the next block.** You have a group of prime order acting on a set of size divisible by p, which is the exact input to the fixed point lemma.

Block 5. **Cue.** A group of prime order acting on a finite set. **Why natural.** Orbits partition the set, and by orbit stabiliser each orbit size divides the order of the acting group, here p. Since p is prime, the only orbit sizes are 1 and p. Summing over orbits, $|S|$ is congruent modulo p to the number of singleton orbits, that is, the number of fixed points. This is the single most reusable lemma in this circle of ideas. **Link to the next block.** Everything now hinges on identifying the fixed points, which is a purely mechanical step.

Block 6. **Cue.** A tuple fixed by a cyclic shift. **Why natural.** Being fixed by the shift means each coordinate equals the next, and running the index around the cycle forces all coordinates equal. The membership condition then reads $a^p = e$. So the fixed points biject with S_p, which was arranged in Block 1 to have size $n+1$. **Link to the next block.** Nothing follows; combining Blocks 3, 5 and 6 gives $0 \equiv |S| \equiv n+1 \pmod{p}$.

Layer 3: WHY

The mathematical truth beneath the route is a single lemma, and it deserves to be installed as a named object rather than rediscovered each time.

Lemma 19.5.3 (Fixed point lemma)**.** *Let H be a group of order p^k with p prime, acting on a finite set X. Then $|X| \equiv |X^H| \pmod{p}$, where X^H is the set of points fixed by every element of H.*

The proof is the same in general as it was above: orbits partition X, orbit sizes divide p^k, so every orbit has size 1 or a positive power of p, and only the singletons survive modulo p. This lemma is the engine behind Cauchy's theorem, behind the class equation and the nontriviality of the centre of a p group, behind Sylow's theorems, behind the necklace proof of Fermat's little theorem, and behind Lucas's theorem on binomial coefficients modulo p. If you install one piece of group theory for the Putnam, install this.

The discernment mechanism, stated so you can run it under time pressure. When a problem asks you to prove that p divides some count N in a setting with no structural hypotheses, run three

questions. Can I find a set X whose size I can compute exactly and which is visibly divisible by p? Is there a natural action of a group of order p on X? Do the fixed points of that action biject with the N objects I care about? The construction in this problem answers all three in one stroke, and the reason it can is that the tuple space G^p has size m^p computable with no knowledge of G at all, while the single relation $a_0 \cdots a_{p-1} = e$ cuts it down by exactly one factor of m and is simultaneously rotation invariant.

Training it is a matter of collecting the pattern "count modulo p equals fixed point count". The drills below are the standard three. After them, the reflex to build the tuple set becomes automatic, and you will recognise the same shape when a competition asks for the number of necklaces, the number of solutions of $x^m = e$, or the number of Sylow subgroups.

There is also a structural truth worth naming. The problem is equivalent to the statement that the number k of subgroups of order p satisfies $k \equiv 1 \pmod{p}$, because $n = k(p-1)$ and so $n + 1 \equiv 1 - k$. That equivalence is worth writing on your page even though it is not the route you take, because it tells you the theorem is a Sylow type congruence and therefore that a Sylow type mechanism, namely a p group acting with few fixed points, will be involved. Recognising the family a statement belongs to is how you choose the tool before you know the proof.

What the examiners nudged with: the $+1$ (switch to elements of order dividing p), the phrase "exactly n elements" with no structure on the group (only counting can work), and the dichotomy "$n = 0$ or" (dispose of the degenerate case with Lagrange in one sentence).

Revolving concepts

Order of an element. $\operatorname{ord}(x)$ is the least positive integer t with $x^t = e$, and $x^s = e$ if and only if $\operatorname{ord}(x) \mid s$. Hence for prime p, $x^p = e$ forces $\operatorname{ord}(x) \in \{1, p\}$. See Book 1, field: Abstract algebra, topic: cyclic groups and element orders.

Lagrange's theorem. The order of a subgroup divides the order of the group, so $\operatorname{ord}(x) \mid |G|$ for every x. Used once, to show $n > 0$ implies $p \mid |G|$. See Book 1, field: Abstract algebra, topic: cosets and Lagrange.

Group actions and orbit stabiliser. An action of H on X partitions X into orbits, and $|Hx| = [H : H_x]$ divides $|H|$. The corollaries used here are that a prime order group has only orbits of size 1 and p, and that the fixed points are exactly the singleton orbits. See Book 1, field: Abstract algebra, topic: group actions.

Fixed point lemma and its induction, the recursion generalised. For $|H| = p^k$ acting on finite X we have $|X| \equiv |X^H| \pmod{p}$. The general form is an induction on k in disguise: orbit sizes are p^j for $0 \le j \le k$, so the congruence holds modulo p for every k, and modulo p^k if all nonsingleton orbits are free. The general recursion pattern here is "a group of order p^k acting on a set decomposes it into blocks whose sizes form a geometric family $1, p, p^2, \ldots, p^k$, so counting modulo p discards all but the blocks of size 1". Applied with $H = \mathbb{Z}/p$ acting on strings of length p over an alphabet of size a, it yields $a^p \equiv a \pmod{p}$, which is Fermat's little theorem. See Book 1, field: Abstract algebra, topic: p groups and fixed points.

Cauchy's theorem. If a prime p divides $|G|$ then G has an element of order p. The lemma proved above contains it. The tuple proof is due to McKay and is the shortest known. See Book 1, field: Abstract algebra, topic: Cauchy and Sylow.

Sylow's theorems. A Sylow p subgroup exists, they are all conjugate, and their number n_p satisfies $n_p \equiv 1 \pmod{p}$ and $n_p \mid |G|$. The congruence $n_p \equiv 1$ is the same phenomenon as this problem, and indeed for $|G|$ divisible by p but not p^2 the two statements coincide. See Book 1, field: Abstract algebra, topic: Cauchy and Sylow.

Frobenius's theorem. If t divides $|G|$ then t divides the number of $x \in G$ with $x^t = e$. This is the correct general statement of which the problem is the case $t = p$. Its proof is much harder and is a knowledge gap you should record, not disguise. See Book 1, field: Abstract algebra, topic: counting solutions of $x^t = e$.

Elementary abelian p groups. A group in which every nonidentity element has order p and which is abelian is a vector space over $\mathbb{F}_p$, hence of order p^k. Needed only for the alternative solutions, but it is the standard object when several commuting elements of order p appear. See Book 1, field: Abstract algebra, topic: abelian p groups.

Circuit Bank

- **WHEN** you must prove that a prime p divides some count in an arbitrary finite structure **DO** build a set whose size you can compute exactly, let a group of order p act on it, and identify the fixed points with the objects you are counting **BECAUSE** orbits have size 1 or p, so the size of the set is congruent to the number of fixed points modulo p
- **WHEN** a statement counts elements of order exactly p and contains a $+1$ **DO** switch to the set $\{x : x^p = e\}$ of elements of order dividing p **BECAUSE** that set is defined by an equation and therefore arises as a fixed point set, whereas the set of elements of order exactly p has no closure property at all
- **WHEN** you need a set of known size inside an arbitrary group of order m **DO** take tuples with one product constraint, so that the last coordinate is determined by the others **BECAUSE** the count is $m^{\text{length}-1}$ regardless of the structure of the group
- **WHEN** a defining condition is invariant under cyclic rotation of its arguments **DO** let $\mathbb{Z}/p$ act by rotation and look at the constant tuples **BECAUSE** rotation invariance is precisely the statement that the shift preserves the set, and the fixed points of a rotation are the constant configurations
- **WHEN** a problem offers the dichotomy "either the count is zero or a congruence holds" **DO** dispose of the zero case first with Lagrange or an existence argument, and use the surviving hypothesis in the main argument **BECAUSE** the nonzero case supplies the divisibility hypothesis $p \mid |G|$ that the main count requires

Backtesting

1. First useful cue: the $+1$ in the conclusion, which says the natural object is $\{x : x^p = e\}$ of size $n + 1$ rather than the set of elements of order exactly p.
2. Method the expression suggested: prove that p divides the number of solutions of $x^p = e$ by letting $\mathbb{Z}/p$ act by cyclic shift on the set of p tuples with product e, whose size $|G|^{p-1}$ is divisible by p.
3. Where you would hesitate, and the thought to reinforce: at the moment of inventing the tuple set, because it looks unmotivated. Reinforce this sentence: I need a set of known size with a $\mathbb{Z}/p$ action, and in an arbitrary group the only sizes I know are powers of $|G|$, so I take tuples with exactly one relation, and I choose the relation to be rotation invariant.
4. The trap the question sets: attempting to count elements of order exactly p directly, which yields only $n = k(p-1)$ and pushes the difficulty into proving $k \equiv 1 \pmod{p}$. A second trap is forgetting to verify that the cyclic shift really preserves the set S, which is the one line where the group axioms are used.
5. What makes you recognise it faster next time: the shape "divisibility by a prime of a count, in a structure with no hypotheses". That shape always means a group of order p acting, and fixed points.
6. The reusable decision rule: to prove $p \mid N$, realise N as the number of fixed points of a $\mathbb{Z}/p$ action on a set whose cardinality is visibly divisible by p.

Mirror drills

1. Prove Fermat's little theorem, $a^p \equiv a \pmod{p}$ for a prime p and a positive integer a, by a counting argument. *Answer:* let $\mathbb{Z}/p$ act by cyclic rotation on the a^p strings of length p over an alphabet of size a; orbits have size 1 or p, and the fixed points are the a constant strings.
2. Let G be a finite group of even order. Prove that G contains an element of order 2. *Answer:* take $p = 2$ in the lemma: $|\{x : x^2 = e\}|$ is even and contains e, so it contains a second element, which has order 2. Equivalently, pair each x with x^{-1} and note the unpaired elements are exactly the solutions of $x^2 = e$.
3. Let p be a prime dividing the order of a finite group G. Prove that the number k of subgroups of G of order p satisfies $k \equiv 1 \pmod{p}$. *Answer:* distinct subgroups of order p meet only in e and each contains $p - 1$ elements of order p, so $n = k(p-1)$; the theorem just proved gives $p \mid n + 1 = k(p-1) + 1 \equiv 1 - k \pmod{p}$.

19.6 A6. Bounding admissible triangulations of a polygon

A triangulation $\mathcal{T}$ *of a polygon P is a finite collection of triangles whose union is P, and such that the intersection of any two triangles is either empty, or a shared vertex, or a shared side. Moreover, each side is a side of exactly one triangle in $\mathcal{T}$. Say that $\mathcal{T}$ is* admissible *if every internal vertex is shared by 6 or more triangles. Prove that there is an integer M_n, depending only on n, such that any admissible triangulation of a polygon P with n sides has at most M_n triangles.*

Read it out loud

Say what the hypothesis is really about. "Every internal vertex is shared by 6 or more triangles" is a curvature condition. In the flat triangular lattice every internal vertex is shared by exactly six triangles, so six is the flat value; anything above six is negatively curved, and anything below six is positively curved. The problem is therefore: a triangulated disc with no positive curvature in the interior has its size controlled by its boundary. That is a discrete Gauss and Bonnet statement, and once you say it in those words you know which identity to write.

The fields wired in are planar graph theory (Euler's formula and edge counting), a small amount of plane topology (a triangulated polygon is a disc, so the link of an interior vertex is a cycle and of a boundary vertex a path), and induction with a peeling argument. There is no metric content at all: the triangles may be long, thin, wildly unequal, and the statement is unaffected, which is the signal that the whole problem is combinatorial and that any attempt to argue by areas or angles alone will fail.

Within thirty seconds you should have written three lines. Write T for the number of triangles, I for the number of internal vertices, E for the number of edges, and a_i for the number of edges at the boundary vertex v_i. Then

$$T = n + 2I - 2, \qquad 3T = 2E - n, \qquad \sum_{i=1}^{n} a_i \leq 4n - 6.$$

The first says bounding T is the same as bounding I. The third is the whole content of the hypothesis, and here is the sting you must notice immediately: it does not involve I at all. So the global count, which is the obvious weapon, cannot bound the number of triangles. That is the discovery on which the problem turns, and it is the reason the problem is an A6 rather than an A2.

Since the global identity will not bound I, it must be used locally: to force the existence of a special configuration on the boundary that you can shave off, reducing n by one at a bounded cost. That is the plan, and everything after this point is executing it while refusing to wave your hands at the degenerate cases.

Approaches, ranked worst to best

1. **Metric estimates.** *Cue:* triangles, angles, areas. *Circuit:* bound the area or the angles of each triangle. There is no scale in the problem and no lower bound on any angle, so a triangulation may contain arbitrarily thin triangles of arbitrarily small area. Every metric quantity you might bound is unbounded here. Rank last: it is not merely hard, it is false as a strategy.

2. **Use the degree inequality globally.** *Cue:* $\sum a_i \leq 4n - 6$ looks like a strong constraint. *Circuit:* combine with $T = n + 2I - 2$. This is the zombie route and it is worth doing once so you feel it fail. The inequality is completely free of I, so no amount of algebra will extract a bound on I from it. What it does tell you is the curvature budget on the boundary, which is what the ideal route uses. Rank it here because it is a necessary experiment, not a solution.

3. **Peel by distance to the boundary.** *Cue:* layers of a lattice patch shrink as you go inward. *Circuit:* induct on the layer index. The instinct is right, and in the hexagonal lattice patch the layers shrink by exactly six each step. But the set of vertices at distance r from the boundary need not bound a connected region, so you must carry a case analysis over components and prove the layer really does shrink. More bookkeeping than the collar removal for the same conclusion.
4. **Remove an ear when some boundary vertex has degree** 2. *Cue:* a degree 2 boundary vertex sits in exactly one triangle. *Circuit:* induct on n. Correct and easy, but it handles only one case; a triangulation may have every boundary degree at least 3. It is a component of the ideal route, not a route.
5. **Curvature budget, chord splitting, then collar removal.** *Cue:* six or more triangles per internal vertex is a nonpositive curvature condition. *Circuit:* prove the global identity, then use it to force a boundary run of degrees $3, 4, \dots, 4, 3$ and peel that run. Best. The identity is proved once, used once globally to reduce "bound T" to "bound I", and then used repeatedly as a local existence device. The induction on n has exactly two cases and both are clean.

The ideal solution

Throughout, P is a simple polygon with vertices $v_1, \dots, v_n$ in cyclic order and $\mathcal{T}$ is an admissible triangulation of P. Indices on boundary vertices are read cyclically.

Step 0: the triangulation is a simplicial disc.

No vertex of $\mathcal{T}$ lies in the relative interior of an edge of $\mathcal{T}$. Indeed, suppose the triangle Δ' has a vertex q in the relative interior of a side s of the triangle Δ. Then $\Delta \cap \Delta'$ is nonempty, so by hypothesis it is a shared vertex or a shared side; either way q would have to be a vertex of Δ, which it is not. Hence $\mathcal{T}$ is edge to edge, and since each side of P is a side of exactly one triangle, every side of P is an edge of $\mathcal{T}$. Consequently the vertices of $\mathcal{T}$ lying on ∂P are exactly $v_1, \dots, v_n$.

Since the union of the triangles is P, a closed disc, the complex is a triangulated closed disc. We use the standard consequences: the link of an internal vertex is a cycle, so an internal vertex of degree a lies in exactly a triangles; the link of a boundary vertex is a path, so a boundary vertex of degree a lies in exactly $a - 1$ triangles, and its neighbours in the fan order begin and end with its two neighbours on ∂P. In particular $a_i \geq 2$ for each i, and admissibility says exactly that every internal vertex has degree at least 6.

Step 1: the counting lemma.

Lemma 19.6.1. *Let T be the number of triangles, I the number of internal vertices, E the number of edges, and a_i the degree of v_i. Then*

$$\text{(i)} \quad T = n + 2I - 2, \qquad \text{(ii)} \quad \sum_{i=1}^{n} a_i \leq 4n - 6.$$

Proof. Sum all the angles of all the triangles. Each triangle contributes π, so the total is $T\pi$. Grouping instead by vertex, each internal vertex contributes 2π and each boundary vertex v_i contributes the interior angle of P at v_i; those interior angles sum to $(n-2)\pi$. Hence $T\pi = 2\pi I + (n-2)\pi$, which is (i).

Count incidences between triangles and edges. Each triangle has three edges, and each edge lies in two triangles except the n edges of ∂P, which lie in one. Hence $3T = 2(E - n) + n = 2E - n$, so

$$2E = 3T + n = 3(n + 2I - 2) + n = 6I + 4n - 6.$$

Now sum degrees over all vertices: $2E = \sum_{i=1}^{n} a_i + \sum_{u \text{ internal}} \deg u \geq \sum_{i=1}^{n} a_i + 6I$, using admissibility. Combining the two displays gives $\sum a_i + 6I \leq 6I + 4n - 6$, which is (ii). □

It is worth recording (ii) in its curvature form. Define $\kappa(u) = 6 - \deg u$ for an internal vertex and $\kappa(v_i) = 4 - a_i$ for a boundary vertex. The computation above says

$$\sum_{u \text{ internal}} \kappa(u) + \sum_{i=1}^{n} \kappa(v_i) = 6I + 4n - 2E = 6,$$

always, for any triangulated disc. Admissibility says the internal curvature is nonpositive, so the boundary must carry at least the whole of the total curvature:

$$\sum_{i=1}^{n} (4 - a_i) \geq 6.$$

Step 2: the bound.

Theorem 19.6.2. *Every admissible triangulation of a polygon with $n \geq 3$ sides has at most $M_n = n^2 - 2n - 2$ triangles.*

Proof. Strong induction on n. Note $M_3 = 1$ and $M_n = M_{n-1} + 2n - 3$ for $n \geq 4$.

Base case $n = 3$. By (ii), $a_1 + a_2 + a_3 \leq 6$, and each $a_i \geq 2$, so $a_1 = a_2 = a_3 = 2$. Then v_1 has only the two sides of P as edges, so the unique triangle of $\mathcal{T}$ at v_1 has vertices v_1, v_2, v_3 and equals P. Its interior is all of the interior of P, and distinct triangles of $\mathcal{T}$ have disjoint interiors, so $\mathcal{T} = \{P\}$ and $T = 1 = M_3$.

Induction step. Let $n \geq 4$ and assume the theorem for all polygons with at least 3 and at most $n - 1$ sides. Call an edge of $\mathcal{T}$ a *chord* if it is not a side of P but both its endpoints are boundary vertices.

Case A: $\mathcal{T}$ has a chord $v_i v_j$. The relative interior of the chord lies in the interior of P, so the chord divides P into two simple polygons P_1 and P_2 with n_1 and n_2 sides, where $n_1 + n_2 = n + 2$ and $n_1, n_2 \geq 3$; hence $3 \leq n_1, n_2 \leq n - 1$. Every triangle of $\mathcal{T}$ lies in P_1 or in P_2, since its interior is connected and disjoint from the chord, so $\mathcal{T}$ splits as $\mathcal{T}_1 \cup \mathcal{T}_2$. A vertex internal to P_k is internal to P, because no vertex lies in the relative interior of the chord, and all of its triangles lie in P_k; so its degree is unchanged and $\mathcal{T}_k$ is admissible. By the induction hypothesis and the convexity of $x \mapsto x^2$,

$$T = T_1 + T_2 \leq M_{n_1} + M_{n_2} = n_1^2 + n_2^2 - 2(n+2) - 4 \leq 9 + (n-1)^2 - 2n - 8 = n^2 - 4n + 2 \leq M_n,$$

the last inequality because $M_n - (n^2 - 4n + 2) = 2n - 4 \geq 0$.

Case B: $\mathcal{T}$ has no chord. First, every $a_i \geq 3$. For if $a_i = 2$ then the only edges at v_i are the two sides of P at v_i, so the unique triangle at v_i is $v_{i-1}v_iv_{i+1}$ and $v_{i-1}v_{i+1}$ is an edge of $\mathcal{T}$ joining two boundary vertices; as $n \geq 4$ it is not a side of P, so it is a chord, contrary to assumption.

Write $t_j = |\{i : a_i = j\}|$ for $j \geq 3$. From $\sum_i (4 - a_i) \geq 6$ we get

$$t_3 \geq 6 + \sum_{j \geq 5} (j - 4) t_j \geq 6 + \sum_{j \geq 5} t_j.$$

In particular $t_3 \geq 6$, so $n \geq 6$. Mark each boundary vertex with $a_i = 3$ as low and each with $a_i \geq 5$ as high, and read the marks in cyclic order. There are strictly more low marks than high marks, so two low marks are adjacent in this cyclic reading, that is, there are boundary vertices w_1 and w_k with $a = 3$ such that the vertices strictly between them along one arc, namely $w_2, \ldots, w_{k-1}$, all have $a = 4$. Here $k \geq 2$, and since at least $t_3 - 2 \geq 4$ vertices of degree 3 lie outside the run, $k \leq n - 4$.

Write w_0 and w_{k+1} for the boundary neighbours of w_1 and w_k outside the run. Read the fan at each run vertex. At w_1, of degree 3, the neighbours in fan order are w_0, z_1, w_2 for some vertex z_1, and the two triangles at w_1 are $w_0w_1z_1$ and $z_1w_1w_2$. At w_j for $1 < j < k$, of degree 4, the neighbours in fan order are $w_{j-1}, z_{j-1}, z_j, w_{j+1}$, and the three triangles are $w_{j-1}w_jz_{j-1}$, $z_{j-1}w_jz_j$ and $z_jw_jw_{j+1}$. At w_k, of degree 3, the neighbours are $w_{k-1}, z_{k-1}, w_{k+1}$ and the triangles are $w_{k-1}w_kz_{k-1}$ and $z_{k-1}w_kw_{k+1}$. Since consecutive entries of a fan are distinct, $z_{j-1} \neq z_j$.

Each z_j is an internal vertex: it is a neighbour of w_j other than w_{j-1} and w_{j+1}, so the edge z_jw_j is not a side of P, and if z_j were a boundary vertex that edge would be a chord.

The z_j are pairwise distinct. Adjacent ones are distinct as noted. Suppose $z_i = z_j = z$ with $j \geq i+2$. The closed walk $z \to w_{i+1} \to w_{i+2} \to \cdots \to w_j \to z$ is a simple cycle C of edges of $\mathcal{T}$, since the w's are distinct boundary vertices and z is internal. As $C \subseteq P$ and the complement of P is connected and unbounded, the bounded component R of the complement of C satisfies $R \subseteq P$; and C is a union of edges, so $\overline{R}$ is the union of the triangles of $\mathcal{T}$ inside it and is an admissible triangulation of a polygon Q with $m = j - i + 1$ sides. Near each w_l with $i+2 \leq l \leq j-1$ the cycle C coincides with ∂P, so all four edges at w_l lie in Q and its degree there is 4; at w_{i+1} the region Q is the sector from z round to w_{i+2}, containing the three edges to z, z_{i+1} and w_{i+2}; symmetrically the degree of w_j in Q is 3; and the degree of z in Q is at least 2. Summing, the boundary degrees of Q total at least $2 + 3 + 4(j-i-2) + 3 = 4(j-i)$, while (ii) applied to Q gives at most $4m - 6 = 4(j-i) - 2$. That is a contradiction, so no such coincidence occurs.

Now delete from $\mathcal{T}$ the triangles containing at least one of $w_1, \ldots, w_k$. By the fan description these are $w_0w_1z_1$, the $k-1$ triangles $w_jw_{j+1}z_j$ for $1 \leq j \leq k-1$, the $k-2$ triangles $z_{j-1}w_jz_j$ for $2 \leq j \leq k-1$, and $z_{k-1}w_kw_{k+1}$; that is $2k-1$ triangles in all. The edges deleted are exactly those incident to some w_l, so no other vertex loses an edge. The remaining triangles form a triangulation $\mathcal{T}'$ of the region P' obtained from P by replacing the boundary path $w_0, w_1, \ldots, w_k, w_{k+1}$ by the path $w_0, z_1, \ldots, z_{k-1}, w_{k+1}$, whose edges w_0z_1, z_jz_{j+1} and $z_{k-1}w_{k+1}$ are all edges of deleted triangles. Since the z_j are distinct internal vertices of P, this new boundary is a simple closed polygon, and P' has $n - (k+1) + k = n - 1$ sides. Every vertex internal to P' is internal to P with the same degree, so $\mathcal{T}'$ is admissible. By the induction hypothesis and $k \leq n-4$,

$$T \leq M_{n-1} + 2k - 1 \leq M_{n-1} + 2n - 3 = M_n.$$

This completes the induction. □

So $M_n = n^2 - 2n - 2$ works, and in particular such a bound exists, which is what was asked. The truth is smaller: a refinement of the run argument gives $M_n = M_{n-1} + 2\lfloor n/6 \rfloor + 1$, of order $\frac{1}{6}n^2$, and a hexagonal patch of the triangular lattice with a boundary as close to regular as possible attains order $\frac{1}{6}n^2$, so the exponent and the constant $\frac{1}{6}$ are both sharp.

Layer 1: WHAT

Block 1. Read "6 or more triangles at each internal vertex" as nonpositive combinatorial curvature, with the triangular lattice as the flat model.

Block 2. Convert triangles into internal vertices: $T = n + 2I - 2$, by summing angles or by Euler.

Block 3. Derive the boundary budget $\sum a_i \leq 4n - 6$, equivalently $\sum \kappa(v_i) \geq 6$, and notice it contains no I.

Block 4. Conclude that the identity must be used locally, not globally, and set up an induction on n.

Block 5. Dispose of chords first: a chord splits P into two polygons with strictly fewer sides each.

Block 6. In the chordless case the budget forces at least six boundary vertices of degree 3, hence a run of consecutive degrees $3, 4, \ldots, 4, 3$.

Block 7. Read the fans along the run, check the collar vertices z_j are internal and distinct, and delete the $2k-1$ triangles of the collar to obtain an admissible triangulation of an $(n-1)$ gon.

Block 8. Bookkeeping: $M_3 = 1$, $M_n = M_{n-1} + 2n - 3$, so $M_n = n^2 - 2n - 2$.

The nominal route runs Blocks 2 and 3 and then tries to solve for I. The bifurcation point is the sentence "the inequality $\sum a_i \leq 4n - 6$ does not involve I". The zombie treats that as a dead end and goes back to the drawing board; the ideal solver treats it as a specification of what the inequality is for. It is an existence device for special boundary vertices, and existence devices feed inductions.

Layer 2: HOW

Block 1. **Cue.** The number six, attached to internal vertices of a triangulation. **Why natural.** Six is the number of triangles at a vertex of the regular triangular lattice, which is the unique flat triangulated model of the plane. A condition "at least six" is therefore "curvature at most zero". Recognising the flat model is what tells you the extremal object will be a lattice patch and that no metric argument can help, since the lattice can be sheared arbitrarily without changing any incidence. **Link to the next block.** If curvature is the language, the next thing to write is the curvature identity, that is, Euler's formula.

Block 2. **Cue.** A triangulated polygon with a triangle count to be bounded. **Why natural.** Both quantities you can compute for a triangulated disc are forced: the total angle is $T\pi$ by triangles and $2\pi I + (n-2)\pi$ by vertices. Equating them gives $T = n + 2I - 2$ with no inequality wasted. The same identity follows from $V - E + F = 2$ with $F = T + 1$ and $3T = 2E - n$; use whichever you can write faster, but write one of them before anything else, because it replaces the object you must bound by a simpler one. **Link to the next block.** Bounding T is now the same as bounding I, the number of internal vertices.

Block 3. **Cue.** Admissibility is a statement about degrees, and $\sum_v \deg v = 2E$ is the only identity that consumes degrees. **Why natural.** Every internal vertex contributes at least 6 to $2E$, and $2E$ has already been computed as $6I + 4n - 6$. The $6I$ cancels exactly, which is not a coincidence: it is the statement that the hypothesis is exactly critical, that six is the flat value. What survives is a pure boundary statement, $\sum(4 - a_i) \geq 6$. **Link to the next block.** The cancellation of I is the crucial discovery, and it must be interpreted rather than lamented.

Block 4. **Cue.** A global inequality that is silent about the quantity you want. **Why natural.** If a constraint cannot bound a quantity globally, its role is to force structure locally. Here it says the boundary carries at least six units of positive curvature, and since each boundary vertex carries at most 1 unit when $a_i \geq 3$, at least six boundary vertices are special. Special boundary vertices are exactly what a peeling induction needs, because they are the places where the triangulation is thin enough to shave. **Link to the next block.** Induct on n; the task becomes finding a piece of boundary whose removal drops n by one at bounded cost.

Block 5. **Cue.** The removal will replace a piece of boundary by another piece, so anything that already cuts the polygon should be handled first. **Why natural.** A chord is an edge of the triangulation joining two boundary vertices without being a side. It splits P into two polygons with $n_1 + n_2 = n + 2$, and since each part has at least three sides, each has at most $n - 1$; so both parts are covered by a strong induction hypothesis. Clearing this case also buys a strong structural fact for the other case: with no chord, no boundary vertex can have degree 2. **Link to the next block.** The surviving case has all $a_i \geq 3$, which is what makes the curvature accounting sharp.

Block 6. **Cue.** A budget of at least six units distributed over the boundary, with each degree 3 vertex worth $+1$ and each degree $j \geq 5$ vertex worth $-(j-4)$. **Why natural.** Since $t_3 \geq 6 + \sum_{j \geq 5} t_j$, the low marks strictly outnumber the high marks. In a cyclic sequence of low and high marks in which lows outnumber highs, two lows must be adjacent, since otherwise each low would be followed by at least one high. The vertices between those two lows all have degree 4. This is pigeonhole applied to a cyclic arrangement, and the cyclic form is what makes it work. **Link to the next block.** You now have a run $3, 4, \ldots, 4, 3$ of consecutive boundary vertices, which is precisely a strip of triangulation two triangles deep.

Block 7. **Cue.** Degrees 3 and 4 at boundary vertices mean fans of exactly two and exactly three triangles. **Why natural.** Write the fans out. The degree 3 vertex at each end contributes one new interior neighbour, and each degree 4 vertex in the middle passes the collar along by one. The result is a chain $z_1, \dots, z_{k-1}$ forming the new boundary. Then check the two ways this can degenerate, because a marker will look for exactly those checks: a z_j on the boundary would be a chord, which the case assumption forbids; a repetition $z_i = z_j$ would pinch the region, and it is killed by applying the counting lemma to the sub polygon cut off by the repeated vertex, where the boundary degrees are forced to be $3, 4, \dots, 4, 3$ again and overshoot the budget by exactly 2. **Link to the next block.** Deleting the collar removes $2k - 1$ triangles and exactly one boundary vertex, so the induction closes.

Block 8. **Cue.** A recursion $M_n = M_{n-1} + 2n - 3$ with $M_3 = 1$. **Why natural.** First order recursions with polynomial forcing are solved by summing. Here $\sum_{m=4}^{n}(2m-3) = n^2 - 2n - 3$, so $M_n = n^2 - 2n - 2$. Then check that this closed form also dominates the chord case, which is the only place the two branches could conflict. **Link to the next block.** The proof is complete; the sharp constant is a separate question and is worth knowing but not worth attempting in the arena.

Layer 3: WHY

The mathematical truth underneath is discrete Gauss and Bonnet. For any triangulated closed disc,

$$\sum_{u \text{ internal}} (6 - \deg u) \;+\; \sum_{v \in \partial} (4 - \deg v) \;=\; 6,$$

which is Euler's formula in curvature clothing: the total curvature is 6χ with $\chi = 1$. The two normalisations, $6 - \deg$ inside and $4 - \deg$ on the boundary, are not arbitrary. Inside, six is the degree in the flat lattice. On the boundary, four is the degree of a vertex of a straight boundary in the flat lattice, where the boundary vertex sees three triangles and four edges. So the identity says: total curvature is a topological invariant, and any hypothesis that suppresses curvature in the interior forces it out to the boundary.

That is the discernment mechanism. Whenever a combinatorial problem imposes a degree condition at interior points of a planar structure, first convert it to a curvature statement and write the total curvature identity, then ask whether the hypothesis makes the interior curvature negative, zero, or merely nonpositive. If it is strictly negative, you get an immediate global bound: for instance, if every internal vertex had degree at least 7, then the internal curvature would be at most $-I$, so $I \le \sum_{\partial}(4 - a_i) - 6 \le 2n - 6$, and the number of triangles would be at most $5n - 14$ with no induction at all. The reason the problem is hard at exactly six is that the interior curvature is allowed to be exactly zero, so the budget yields nothing globally and must be spent locally. Recognising which of those two regimes you are in is the whole discernment, and it is trainable in one sitting by doing the degree 7 version first and then the degree 6 version.

The second truth is that a nonpositively curved disc has thin boundary somewhere. The budget of six positive units cannot be carried by long stretches of degree 4 boundary, because those are curvature neutral; it has to be carried by degree 3 corners, and there must be at least six of them. Two of them must be adjacent in the reduced cyclic order, which means there is a stretch of boundary bounded by two corners with only neutral vertices between: a collar two triangles deep. Peeling that collar is the discrete analogue of pushing in a boundary curve in a nonpositively curved surface, which shortens it. This is the same phenomenon that makes the extremal example a hexagon: a hexagonal patch is the shape whose six corners of degree 3 exactly carry the six units, with all other boundary vertices neutral, and it is the largest patch for a given boundary length.

What the examiners nudged with: the number six (find the flat model), the phrase "depending only on n" (they want existence, not the sharp constant, so an induction with a crude recursion is fully sufficient), and the definition insisting the intersection of two triangles is empty, a shared

vertex, or a shared side, which is there to make the complex simplicial so that Euler's formula applies without caveat.

Revolving concepts

Euler's formula. For a connected plane graph, $V - E + F = 2$ with F counting the unbounded face. For a triangulated polygon, $F = T + 1$ and $3T = 2E - n$, giving $T = n + 2I - 2$ and $2E = 6I + 4n - 6$. See Book 1, field: Discrete structures and graphs, topic: planar graphs and Euler's formula.

Triangulated discs and links. In a simplicial triangulation of a closed disc, the link of an internal vertex is a cycle and the link of a boundary vertex is a path whose endpoints are its two boundary neighbours. Hence an internal vertex of degree a lies in a triangles and a boundary vertex of degree a lies in $a - 1$ triangles. This is a plane topology fact, and it is a genuine knowledge gap for a reader coming from A level; it rests on the Jordan curve theorem and should be quoted as a standard fact, not smuggled in. See Book 1, field: Geometry, topic: simplicial complexes and links.

Combinatorial curvature and discrete Gauss and Bonnet. With $\kappa = 6 - \deg$ inside and $4 - \deg$ on the boundary, $\sum \kappa = 6\chi$. For a disc that is 6; for a sphere it is 12, which is why a triangulated sphere with all degrees at most 6 has at least twelve vertices of degree 5, the reason a football has twelve pentagons. See Book 1, field: Discrete structures and graphs, topic: discrete curvature.

The triangular lattice as the flat model, and sharpness. Every internal vertex of the triangular lattice has degree 6; a hexagonal patch with n boundary vertices has about $n^2/6$ triangles, so the true growth rate is quadratic with constant $1/6$. The refinement $M_n = M_{n-1} + 2\lfloor n/6 \rfloor + 1$ follows from the sharper statement that six disjoint runs of total length at most $n + 6$ exist. See Book 1, field: Geometry, topic: lattices and extremal configurations.

Chord splitting. A diagonal of a simple polygon that lies inside it divides it into two simple polygons with $n_1 + n_2 = n + 2$; both have at least three and at most $n - 1$ sides, which is exactly what a strong induction on the number of sides needs. See Book 1, field: Geometry, topic: polygon dissection.

The recursion, generalised. The bound obeys $M_3 = 1$ and $M_n = M_{n-1} + q(n)$ with $q(n) = 2n - 3$. Any recursion $u_n = u_{n-1} + q(n)$ with q a polynomial of degree d has $u_n = u_{n_0} + \sum_{m=n_0+1}^{n} q(m)$, a polynomial in n of degree $d + 1$, computable either by Faulhaber sums or by writing an undetermined polynomial of degree $d + 1$ and matching the difference. Here $\sum_{m=4}^{n}(2m - 3) = n^2 - 2n - 3$ and $M_n = n^2 - 2n - 2$. The same template covers every peeling induction whose per step cost is polynomial in the parameter. See Book 1, field: Sequences and recurrences, topic: telescoping and summation of polynomial forcing.

Pigeonhole on a cyclic sequence. If a cyclic sequence contains more marks of type L than of type H, then two L marks are adjacent among the marks. This tiny lemma is the whole of Block 6 and is worth banking on its own. See Book 1, field: Combinatorics, topic: cyclic pigeonhole.

Circuit Bank

- **WHEN** a planar subdivision problem imposes a degree condition at interior vertices **DO** write Euler's formula and the edge to face count, and convert the hypothesis into the curvature identity with $6 - \deg$ inside and $4 - \deg$ on the boundary **BECAUSE** degree conditions are curvature conditions, and total curvature is a topological invariant, so this is the only global constraint available

- **WHEN** the global identity you derived turns out not to involve the quantity you want to bound **DO** stop using it globally and use it as an existence device to force a special local configuration, then induct **BECAUSE** a constraint that cancels the unknown is telling you the hypothesis is exactly critical, which means the content is local

- **WHEN** you want to reduce a triangulated polygon from n sides to $n-1$ **DO** find a boundary vertex or a run of consecutive boundary vertices of low degree and delete the triangles meeting it **BECAUSE** low boundary degree means a shallow fan, so the deletion is a collar whose removal leaves a polygon with one fewer vertex
- **WHEN** an edge joins two boundary vertices without being a side **DO** split the polygon along it and apply strong induction to both parts **BECAUSE** both parts have at least three and at most $n-1$ sides, so the induction hypothesis applies to each
- **WHEN** a peeling argument produces new boundary vertices **DO** prove they are internal and pairwise distinct before claiming the result is a polygon **BECAUSE** a repeated or boundary vertex pinches the region and destroys the induction, and applying the counting lemma to the cut off sub polygon is what rules it out

Backtesting

1. First useful cue: the number six, which is the degree of an internal vertex in the flat triangular lattice, so the hypothesis is nonpositive curvature.
2. Method the expression suggested: write $T = n + 2I - 2$ and the degree sum, obtain $\sum(4 - a_i) \geq 6$, notice it is free of I, and therefore use it locally to force a run $3, 4, \ldots, 4, 3$ on the boundary which can be peeled to reduce n by one.
3. Where you would hesitate, and the thought to reinforce: at the discovery that the inequality does not bound I. Reinforce this sentence: an inequality in which the unknown cancels is not a failure, it is a statement that the hypothesis is critical, and critical hypotheses are used to produce local structure for an induction.
4. The trap the question sets: believing that the Euler count alone finishes the problem, and, at the other end, writing "clearly the remaining region is an $(n-1)$ gon" without checking that the new boundary vertices are internal and distinct. That check is where the marks are, because it is the only place the hypothesis is used twice.
5. What makes you recognise it faster next time: the shape "interior degree bound in a planar complex, bound the size". That shape always means total curvature identity first, then decide whether the interior curvature is strictly negative, in which case you are done globally, or exactly critical, in which case you peel.
6. The reusable decision rule: convert degree hypotheses to curvature, write the invariant, and if the unknown cancels, hunt for the forced local configuration.

Mirror drills

1. Show that any triangulation of a polygon with n sides and no internal vertices has exactly $n-2$ triangles. *Answer:* put $I = 0$ in $T = n + 2I - 2$; equivalently sum angles, $T\pi = (n-2)\pi$.
2. Prove that for any triangulated closed disc, $\sum_{\text{internal}}(6 - \deg u) + \sum_{\partial}(4 - \deg v) = 6$. *Answer:* $2E = 3T + n$ and $T = n + 2I - 2$ give $2E = 6I + 4n - 6$, and the left side of the claim is $6I + 4n - 2E$.
3. Suppose every internal vertex of a triangulation of an n gon has degree at least 7. Prove directly, with no induction, that there are at most $2n - 6$ internal vertices and at most $5n - 14$ triangles. *Answer:* internal curvature is at most $-I$, so $\sum_{\partial}(4 - a_i) \geq 6 + I$; each boundary term is at most 2 since $a_i \geq 2$, so $2n \geq 6 + I$, and then $T = n + 2I - 2 \leq 5n - 14$. Compare with the problem to see why degree six is the critical threshold.

19.7 B1. the congruence that collapses to f(1)

Let f be a polynomial with positive integer coefficients. Prove that if n is a positive integer, then $f(n)$ divides $f(f(n)+1)$ if and only if $n = 1$. (Editor's note: assume f is nonconstant; a constant f makes the claim false, since then $f(n)$ divides $f(f(n)+1)$ for every n.)

Read it out loud

This is a divisibility statement about the value of a polynomial at a shifted argument, not a statement about the polynomial as an algebraic object with roots. The object under scrutiny is a specific positive integer, $f(n)$, and the question is whether another specific integer, $f(f(n)+1)$, sits on its multiple ladder. The instant you see "does A divide B" for two integers built from the same underlying function, the honest first move is not calculus and not factoring the polynomial; it is modular arithmetic. Treat $f(n)$ as a modulus and ask what $f(f(n)+1)$ reduces to modulo it.

The phrase $f(n)+1$ is a guiding light, written as the modulus plus one. Modulo $f(n)$, the modulus itself is 0, so $f(n)+1 \equiv 1$, trivially. Once you see that, pushing this reduced argument through the outer f is a one-line computation term by term, because a polynomial with integer coefficients respects congruence: if $a \equiv b \pmod{m}$ then $f(a) \equiv f(b) \pmod{m}$ for any such f. That fact is the whole engine of the problem; everything after it is bookkeeping.

Positivity of the coefficients does two separate jobs, and you should clock both before writing a line of proof. First, it forces f to be strictly increasing on the positive integers, since every monomial $a_i n^i$ with $a_i > 0$ only grows as n grows; so $f(n) > f(1)$ for every $n > 1$. Second, it forces $f(1) > 0$, so once you reduce the target to $f(1)$ modulo $f(n)$, that residue is a genuine positive integer strictly between 0 and $f(n)$ whenever $n > 1$, and a number strictly between 0 and a modulus can never be the zero residue. The "if and only if" in the statement is two one-line arguments glued together by a single congruence.

Fields embedded: elementary number theory, specifically congruences, and the congruence-respecting property of polynomial evaluation with integer coefficients. The knowledge base needed sits just past A level Further Mathematics; the manipulation itself is short. Within thirty seconds you should have written $f(n)+1 \equiv 1 \pmod{f(n)}$ and be reaching to apply f to both sides.

Approaches, ranked worst to best

1. **expand $f(f(n)+1)$ in full as a polynomial in n.** *Cue:* the composite argument tempts you to substitute and multiply out. *Circuit:* none, this is the zombie route. (why it ranks here: for a degree d polynomial you must expand $(f(n)+1)^i$ for every i up to d and collect terms in n; the algebra grows fast, the modulus structure never becomes visible, and a single slipped term wrecks the whole proof.)

2. **test small explicit f (linear, then quadratic) and conjecture the pattern.** *Cue:* the problem lets you pick any f, so instinct reaches for examples. *Circuit:* inductive conjecture, not a proof circuit. (why it ranks here: taking $f(n) = n$ gives the correct guess, $n \mid n+1$ only when $n = 1$, but a handful of examples never discharges "for every polynomial with positive integer coefficients"; you still owe the general argument afterward, so the examples cost ink without closing the proof.)

3. **rewrite $f(f(n)+1)$ modulo $f(n)$ using $f(n)+1 \equiv 1$.** *Cue:* the divisor $f(n)$ sits literally inside the argument of the dividend, one more than itself. *Circuit:* congruence is preserved by polynomial evaluation. (why it ranks here: this is the ideal route; one substitution, one general lemma about polynomials modulo m, and one strict inequality finish the problem, and every step is forced rather than guessed.)

The ideal solution

Assume, per the editor's note, that f is nonconstant, of degree $d \geq 1$, and write

$$f(x) = \sum_{i=0}^{d} a_i x^i, \qquad a_i \in \mathbb{Z}_{>0} \text{ for every } i = 0, 1, \ldots, d.$$

Since $f(n) \equiv 0 \pmod{f(n)}$ trivially, adding 1 to both sides gives

$$f(n) + 1 \equiv 1 \pmod{f(n)}.$$

A polynomial with integer coefficients preserves congruences under evaluation: if $a \equiv b \pmod{m}$, then $a^i \equiv b^i \pmod{m}$ for every $i \geq 0$ (multiply the congruence by itself i times), so $a_i a^i \equiv a_i b^i \pmod{m}$ term by term, and summing over i gives $f(a) \equiv f(b) \pmod{m}$. Applying this with $a = f(n) + 1$, $b = 1$, $m = f(n)$,

$$f(f(n) + 1) \equiv f(1) \pmod{f(n)}.$$

This single congruence settles both directions at once.

If $n = 1$: here $f(n) = f(1)$, so the congruence reads $f(f(1) + 1) \equiv f(1) \equiv 0 \pmod{f(1)}$, i.e. $f(1) \mid f(f(1) + 1)$. This holds for every nonconstant f with positive integer coefficients, no further hypothesis needed.

If $n > 1$: since every coefficient is a positive integer and f is nonconstant, f is strictly increasing on the positive integers, because

$$f(n) - f(1) = \sum_{i=1}^{d} a_i \left(n^i - 1 \right) > 0$$

for $n > 1$, as every term $n^i - 1$ is positive for $i \geq 1$. So $0 < f(1) < f(n)$. A positive integer strictly between 0 and $f(n)$ is never a multiple of $f(n)$, since the least positive multiple of $f(n)$ is $f(n)$ itself; hence $f(1) \not\equiv 0 \pmod{f(n)}$, and by the congruence above, $f(f(n) + 1) \equiv f(1) \not\equiv 0 \pmod{f(n)}$, so $f(n) \nmid f(f(n) + 1)$.

Combining both cases, $f(n) \mid f(f(n) + 1)$ exactly when $n = 1$, as required. ■

For the editor's note: if f is allowed to be constant, $f(n) = c$ for a positive integer c, then $f(f(n) + 1) = c = f(n)$ for every n, so the divisibility holds always, not only at $n = 1$. This is why nonconstant f is an essential hypothesis, not a cosmetic one.

Layer 1: WHAT

Block 1: rewrite the divisibility question as a congruence, $f(n) \mid f(f(n) + 1) \iff f(f(n) + 1) \equiv 0 \pmod{f(n)}$. Block 2: reduce the inner argument, $f(n) + 1 \equiv 1 \pmod{f(n)}$, and apply the congruence-preserving property of polynomial evaluation to collapse the target to $f(1) \pmod{f(n)}$. Block 3: split on $n = 1$ (where $f(n) = f(1)$, congruence trivially 0) versus $n > 1$ (where strict monotonicity gives $0 < f(1) < f(n)$, so the congruence can never be 0).

The nominal route expands $f(f(n) + 1)$ term by term in powers of n and hunts for a common factor by staring at the expanded polynomial; the ideal route never expands anything, it reduces the argument modulo $f(n)$ first and evaluates second. The bifurcation point is the instant you decide whether to expand $f(f(n) + 1)$ as a polynomial in n, or to treat $f(n)$ as a modulus and reduce before touching f at all. Reducing first is the reroute; expanding first is the zombie route, and it is the entire reason this problem looks harder than it is.

Layer 2: HOW

Block 1. **Cue.** The phrase "does A divide B" for two specific integers is, by definition, a congruence statement, $B \equiv 0 \pmod{A}$. Whenever a problem hands you a divisibility claim about the values of a function rather than its coefficients, translate immediately into modular language. **Why natural.** This is forced by the definition of divisibility, not a choice. Writing $f(n) \mid f(f(n)+1)$ as $f(f(n)+1) \equiv 0 \pmod{f(n)}$ costs no ink and names the correct tool at once, modular reduction, not factoring, not calculus. **Link to the next block.** Once the target is a congruence, the next question answers itself: what is the argument $f(n)+1$ congruent to modulo the same $f(n)$? That is a far smaller question than the divisibility itself.

Block 2. **Cue.** $f(n)+1$ is written as the modulus plus one. Modulo $f(n)$, the modulus itself is 0, so $f(n)+1 \equiv 1$. **Why natural.** $f(n) \equiv 0 \pmod{f(n)}$ is a tautology; adding 1 to both sides of a congruence is always legal. Recognising this costs one line and replaces a messy argument by the constant 1. **Link to the next block.** The reduced argument still has to pass through the outer f, and that only works cleanly if evaluating a polynomial respects congruence, which it does; stating that fact explicitly is the next block.

Block 3. **Cue.** The outer function is the same polynomial f, applied to something already known modulo $f(n)$. A polynomial with integer coefficients is built only from addition and multiplication, both of which respect congruence. **Why natural.** If $a \equiv b \pmod{m}$ then $a^i \equiv b^i \pmod{m}$ for every $i \geq 0$, so $a_i a^i \equiv a_i b^i \pmod{m}$ term by term, and summing gives $f(a) \equiv f(b) \pmod{m}$. With $a = f(n)+1$, $b = 1$, $m = f(n)$, the entire right-hand side collapses to $f(1) \pmod{f(n)}$, regardless of the degree of f. **Link to the next block.** The problem has now reduced to one comparison: is $f(1) \equiv 0 \pmod{f(n)}$? That is answered by comparing the size of $f(1)$ to $f(n)$, which is the final block.

Block 4. **Cue.** The coefficients of f are positive integers; that hypothesis has done nothing yet in the argument, so it is needed here, for a size comparison. **Why natural.** Positive coefficients make f strictly increasing on the positive integers, so $f(n) > f(1)$ whenever $n > 1$, hence $0 < f(1) < f(n)$; a number strictly between 0 and a modulus is never congruent to 0 modulo it. At $n = 1$ the two values coincide and $f(1) \equiv f(1) \equiv 0 \pmod{f(1)}$ trivially. **Why natural.** This is the second use of the same hypothesis in two roles: it made $f(1) > 0$ earlier and now makes $f(1) < f(n)$; both uses are the same monotonicity fact read at different scales. **Link to the next block.** This closes both directions of the iff at once; there is no fifth block.

Layer 3: WHY

The deep fact underneath the whole problem is that reduction modulo m is a ring homomorphism $\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$, and evaluating a fixed integer-coefficient polynomial is built only from addition and multiplication inside that target ring; consequently it makes no difference whether you reduce the argument first or evaluate first and reduce afterward, the two routes land on the same residue. That is the honest reason $f(f(n)+1) \equiv f(1) \pmod{f(n)}$ is not a coincidence of this particular f, it is true for every integer-coefficient polynomial.

The discernment mechanism for spotting this route: whenever a divisibility claim reads "$g(n)$ divides $h(n)$" and h is built by feeding a shifted or transformed version of $g(n)$ back into a polynomial, check first whether that shift is congruent to something simple modulo $g(n)$. It almost always is, because problem setters choose the shift precisely so it collapses; the shift $+1$ here is the cheapest possible choice, landing on the residue 1, which is the multiplicative identity and therefore the simplest input f can receive. Training this discernment is a matter of repetition: every time the

divisor reappears literally inside the dividend's own argument, resist evaluating and reduce first instead.

The examiners nudged you toward this by writing $f(n)+1$ rather than, say, $2f(n)+3$; the simplicity of the shift is itself a signal that a one-line congruence closes the problem, and the "if and only if" phrasing is there to test whether you notice that positivity of the coefficients is load-bearing twice over, once to make $f(1)$ positive and once to make it strictly less than $f(n)$.

Revolving concepts

Congruence. $a \equiv b \pmod{m}$ means $m \mid (a-b)$. The set of residues modulo m forms a ring $\mathbb{Z}/m\mathbb{Z}$ under the induced addition and multiplication; see Book 1, field: Number theory, topic: modular arithmetic.

Congruence-compatibility of polynomial evaluation. For any polynomial f with integer coefficients, $a \equiv b \pmod{m} \Rightarrow f(a) \equiv f(b) \pmod{m}$. This is the reduction-modulo-m homomorphism composed with evaluation; it generalises to any ring homomorphism composed with any polynomial map, not merely $\mathbb{Z} \to \mathbb{Z}/m\mathbb{Z}$. See Book 1, field: Abstract algebra, topic: ring homomorphisms and polynomial evaluation.

Strict monotonicity of positive-coefficient polynomials. If every coefficient of f is a positive real number, f is strictly increasing on the positive reals, since every monomial strictly increases and no term can cancel another. See Book 1, field: Polynomials, topic: monotonicity from sign-definite coefficients.

Residues strictly between 0 and the modulus. If $0 < r < m$ then $r \not\equiv 0 \pmod{m}$, since the least positive integer congruent to 0 modulo m is m itself. This is the closing move in any proof that a specific integer is not divisible by another, once you have bounded it strictly between consecutive multiples.

No recursive or self-referential structure appears in this problem; there is nothing to generalise along that axis here.

Circuit Bank

- **WHEN** you must decide whether $g(n)$ divides $h(n)$ for two integer expressions built from a common function **DO** translate to a congruence, $h(n) \equiv 0 \pmod{g(n)}$, and try to reduce the argument of h modulo $g(n)$ before evaluating anything **BECAUSE** divisibility is congruence by definition, and reducing before evaluating avoids expanding a polynomial in full.

- **WHEN** a polynomial with integer coefficients is evaluated at an expression you already know modulo some m **DO** replace the argument by its residue and evaluate the same polynomial at that smaller residue **BECAUSE** polynomial evaluation is built only from addition and multiplication, both of which respect congruence.

- **WHEN** a polynomial has positive coefficients and you need to compare its values at two positive integers **DO** use strict monotonicity directly, $f(n) > f(k)$ whenever $n > k \geq 1$, term by term **BECAUSE** every term $a_i n^i$ only grows with n when $a_i > 0$, so no cancellation between terms can occur.

- **WHEN** a divisibility claim carries an if and only if between one specific value and a general range **DO** dispose of the specific value in one trivial line, then rule out every other value uniformly via a strict inequality on residues **BECAUSE** a residue strictly between 0 and the modulus can never be the zero class, the cheapest possible way to rule out divisibility.

Backtesting

1. First useful cue: the argument $f(n)+1$ sitting exactly one integer above the modulus $f(n)$ itself.
2. Method the expression suggested: reduce the inner argument modulo $f(n)$ before applying the outer f.
3. Where you would hesitate, and the thought to reinforce: hesitating over whether evaluating a polynomial commutes with reduction; the reinforcement is that evaluation is only addition and multiplication, both congruence-preserving.
4. The trap the question sets: expanding $f(f(n)+1)$ in full powers of n before noticing the modulus structure, burying a one-line argument in algebra.
5. What makes you recognise it faster next time: seeing the modulus written literally inside the argument of the same function, "f of something plus one" where the something is the divisor itself.
6. The reusable decision rule: whenever the divisor reappears inside the dividend's own argument, reduce first and evaluate second, never the reverse order.

Mirror drills

1. (drill) Let $f(x) = x^3 + 2x + 5$. For which positive integers n does $f(n)$ divide $f(f(n)+2)$? *Answer:* $f(n)+2 \equiv 2 \pmod{f(n)}$, so the condition becomes $f(n) \mid f(2)$; since f is strictly increasing, this forces $f(n) \le f(2)$, and direct evaluation gives $f(1) = 8$, $f(2) = 17$, so $n = 2$ always works, and $n = 1$ works only if $f(1) \mid f(2)$, which fails here since 17 is not a multiple of 8; the shift $+1$ in the original problem is exactly what made the answer universal instead of f-dependent.
2. (drill) With f as in the original problem, determine for which n, $f(n)$ divides $f(f(n)-1)$. *Answer:* $f(n)-1 \equiv -1 \pmod{f(n)}$, so the condition becomes $f(n) \mid f(-1)$; since $f(-1)$ is a fixed integer independent of n while $f(n) \to \infty$, this can hold for at most finitely many n, in sharp contrast to the original shift.
3. (drill) Why does the proof break if f is allowed a coefficient equal to 0 or negative? *Answer:* the congruence collapse $f(f(n)+1) \equiv f(1) \pmod{f(n)}$ still holds, since it only needs integer coefficients, but the closing step, $0 < f(1) < f(n)$ for $n > 1$, needs every coefficient strictly positive to guarantee strict monotonicity; a mixed-sign f can have $f(n) \le f(1)$ or even $f(n) \le 0$ for some $n > 1$, and the size comparison collapses.

19.8 B2. the antiderivative's extremum forces f to vanish

Suppose that $f : [0,1] \to \mathbb{R}$ has a continuous derivative and that $\int_0^1 f(x)\,dx = 0$. Prove that for every $\alpha \in (0,1)$,

$$\left| \int_0^\alpha f(x)\,dx \right| \leq \frac{1}{8} \max_{0 \leq x \leq 1} |f'(x)|.$$

Read it out loud

Strip away the specific letters and look at the shape underneath. You are given an object built by integrating f, and asked how far that object can wander from zero at an interior point. That is a question about a function pinned to zero at two boundary points with limited curvature, not a question about f itself. Introduce $g(x) = \int_0^x f(t)\,dt$ so that $g(0) = 0$ automatically and $g(1) = \int_0^1 f = 0$ is precisely the hypothesis you were handed; the quantity you must bound, $\int_0^\alpha f$, is nothing but $g(\alpha)$. Rewriting the problem in terms of g is not decoration, it is the whole translation step: the real hypotheses (vanishing at both ends of an interval, a bounded second derivative) are hypotheses about g, dressed up in the statement as hypotheses about f.

Once the problem reads "g vanishes at two points, bound g strictly between them using a bound on g''", the guiding light is the two-point remainder identity that also governs the error of linear interpolation: if g agreed with some polynomial of degree at most one at both 0 and 1, that polynomial would be exactly determined, and the gap between g and that polynomial anywhere in between is controlled by g'' at some unseen point. Here the polynomial that already matches g at 0 and 1 is the zero function, since both values are zero, so the remainder identity hands you a bound on $g(\alpha)$ directly, with no extra structure to invent.

The role of $\max_{0 \leq x \leq 1} |f'(x)|$ is now obvious once you notice $f' = g''$: bounding the curvature of g is exactly the hypothesis on hand, written in terms of f because f is the function the problem starts from, not because f' is the natural object to work with. Reading the problem as a statement about g rather than about f is the single decision that separates every clean solution from every laboured one.

Fields embedded: real analysis, specifically Rolle's theorem and the mean value theorem, plus the fundamental theorem of calculus needed to introduce g in the first place. The knowledge base needed sits past A level Further Mathematics: Rolle's theorem applied twice in sequence, to pin down a point where a second derivative vanishes, is standard analysis but not standard sixth form material. Within thirty seconds of reading you should have written $g(x) = \int_0^x f(t)\,dt$ and restated the two hypotheses as $g(0) = g(1) = 0$.

Approaches, ranked worst to best

1. **crude pointwise bound on f from the Lipschitz condition.** *Cue:* $|f'| \leq B$ makes f Lipschitz, so bounding $f(x)$ in terms of a nearby known value feels natural. *Circuit:* none, this stalls. (why it ranks here: without first locating a point where f is known to vanish, you have no base point to anchor the Lipschitz estimate; you either need the mean-zero hypothesis to locate such a point, which this approach never does explicitly, or you integrate a bound that still carries an unknown constant, and the argument never closes.)

2. **locate the interior critical point of g, use symmetry $\alpha \leftrightarrow 1 - \alpha$, then bound f linearly from that point.** *Cue:* $g(0) = g(1) = 0$ forces the maximum of $|g|$ to occur at an interior critical point of g, i.e. a zero of f. *Circuit:* extremal point plus symmetry reduction. (why it ranks here: this is correct and closes the problem, but it needs three separate reductions chained together, first to the extremal α, then by symmetry to $\alpha \leq \frac{1}{2}$, then a sign normalisation, each requiring its own justification.)

3. **two-point remainder via a double application of Rolle's theorem.** *Cue:* g vanishes at both endpoints of $[0,1]$, which is exactly the setup for the interpolation-error identity. *Circuit:* Rolle's theorem applied twice to an auxiliary function. (why it ranks here: this is the ideal route; one auxiliary function and two uses of Rolle's theorem give the bound at every α simultaneously, with no case split, no symmetry argument, and a constant that is visibly sharp.)

The ideal solution

Define $g : [0,1] \to \mathbb{R}$ by $g(x) = \int_0^x f(t)\,dt$. By the fundamental theorem of calculus, since f is continuous, g is differentiable with $g' = f$; since f itself has a continuous derivative, g is twice differentiable with $g'' = f'$, and g'' is continuous. Put

$$B = \max_{0\le x\le 1} |f'(x)| = \max_{0\le x\le 1} |g''(x)|,$$

which is attained since f' is continuous on the closed interval $[0,1]$. By hypothesis $g(0) = 0$ and $g(1) = \int_0^1 f(t)\,dt = 0$.

Lemma 19.8.1 (Two-point remainder)**.** *Let $g \in C^2[0,1]$ with $g(0) = g(1) = 0$. For every $x \in (0,1)$ there is some $\xi \in (0,1)$, depending on x, such that*

$$g(x) = \frac{g''(\xi)}{2}\, x(x-1).$$

Proof. Fix $x \in (0,1)$ and define

$$\phi(t) = g(t) - \frac{g(x)}{x(x-1)}\, t(t-1), \qquad t \in [0,1].$$

Then $\phi(0) = g(0) = 0$, $\phi(1) = g(1) = 0$, and $\phi(x) = g(x) - g(x) = 0$, so ϕ vanishes at the three points $0 < x < 1$. By Rolle's theorem applied on $[0,x]$ and again on $[x,1]$, ϕ' vanishes at some $\eta_1 \in (0,x)$ and some $\eta_2 \in (x,1)$. Applying Rolle's theorem once more, to ϕ' on $[\eta_1,\eta_2]$, ϕ'' vanishes at some $\xi \in (\eta_1,\eta_2) \subset (0,1)$. Since $t(t-1) = t^2 - t$ has second derivative 2,

$$\phi''(t) = g''(t) - \frac{2g(x)}{x(x-1)},$$

so $\phi''(\xi) = 0$ gives $g''(\xi) = \dfrac{2g(x)}{x(x-1)}$, which rearranges to the claim. □

Apply the lemma with $x = \alpha$: there is $\xi \in (0,1)$ with $g(\alpha) = \frac{g''(\xi)}{2}\alpha(\alpha-1)$, so

$$|g(\alpha)| = \frac{|g''(\xi)|}{2}\,\alpha(1-\alpha) \le \frac{B}{2}\,\alpha(1-\alpha).$$

Completing the square, $\alpha(1-\alpha) = \frac{1}{4} - \left(\alpha - \frac{1}{2}\right)^2 \le \frac{1}{4}$ for every $\alpha \in [0,1]$, with equality only at $\alpha = \frac{1}{2}$. Hence

$$\left|\int_0^\alpha f(x)\,dx\right| = |g(\alpha)| \le \frac{B}{2}\cdot\frac{1}{4} = \frac{B}{8} = \frac{1}{8}\max_{0\le x\le 1}|f'(x)|,$$

for every $\alpha \in (0,1)$, which is the required inequality. ■

Sharpness. The constant $\frac{1}{8}$ is best possible. Take $f(x) = \frac{1}{2} - x$; then $f'(x) = -1$ everywhere, so $B = 1$, and $\int_0^1 f(x)\,dx = \left[\frac{x}{2} - \frac{x^2}{2}\right]_0^1 = 0$, satisfying the hypothesis. Here $g(\alpha) = \int_0^\alpha \left(\frac{1}{2} - x\right)dx = \frac{\alpha(1-\alpha)}{2}$, which attains the value $\frac{1}{8}$ exactly at $\alpha = \frac{1}{2}$. The affine function that has constant $|f'|$ throughout saturates every inequality used in the proof at once, so the bound cannot be improved.

Layer 1: WHAT

Block 1: introduce $g(x) = \int_0^x f$ and translate the two hypotheses, $g(0) = g(1) = 0$, $g' = f$, $g'' = f'$; restate the goal as bounding $|g(\alpha)|$ by $\frac{1}{8}\max|g''|$. Block 2: recognise the shape, a function vanishing at two points with bounded curvature, as the two-point remainder identity, and build the auxiliary function $\phi(t) = g(t) - \frac{g(x)}{x(x-1)}t(t-1)$ that shares three zeros with g. Block 3: apply Rolle's theorem twice in sequence (once on each half, once on the pair of resulting critical points) to extract a point ξ where $\phi'' = 0$, giving the exact remainder formula for $g(x)$. Block 4: bound $|g''(\xi)|$ by B and $x(1-x)$ by its maximum $\frac{1}{4}$ at $x = \frac{1}{2}$, closing the inequality, then confirm sharpness with an explicit affine example.

The nominal route works with f directly, trying to bound $\int_0^\alpha f$ from a Lipschitz estimate on f anchored at some point that has to be found by other means; it never promotes g to a first-class object, so the mean-zero hypothesis and the derivative bound never combine cleanly. The ideal route renames the problem in terms of g before estimating anything. The bifurcation point is exactly that renaming: the decision to write down $g = \int_0^x f$ and restate both hypotheses in its language before touching an inequality sign is what separates the two routes, and skipping it is what sends every other approach into case analysis.

Layer 2: HOW

Block 1. **Cue.** The quantity to be bounded, $\int_0^\alpha f$, and the hypothesis, $\int_0^1 f = 0$, are both integrals of f starting at 0. Two integrals of the same function starting at the same point are values of a single new function. **Why natural.** Naming that function, $g(x) = \int_0^x f(t)\,dt$, turns two separate statements about f into two boundary values of one function, $g(0) = 0$ and $g(1) = 0$, which is a far more usable shape than two disconnected integral conditions. **Link to the next block.** With g vanishing at both ends of $[0,1]$, the natural next question is how far g can stray from zero in between, which is a curvature question once you notice $g'' = f'$.

Block 2. **Cue.** A twice-differentiable function that vanishes at two points and needs a bound in between, in terms of a bound on its second derivative, is the exact setup of the two-point interpolation remainder: subtract off the unique low-degree polynomial that already agrees with g at those two points. **Why natural.** Since $g(0) = g(1) = 0$, the matching polynomial is the zero function, and the auxiliary function $\phi(t) = g(t) - \frac{g(x)}{x(x-1)}t(t-1)$ is engineered to also vanish at the third point $t = x$, giving three zeros to feed into Rolle's theorem. **Link to the next block.** Three zeros of ϕ hand you a zero of ϕ' between each adjacent pair, and a zero of ϕ'' between those, which is the mechanism that produces ξ.

Block 3. **Cue.** ϕ has zeros at $0 < x < 1$; Rolle's theorem converts n zeros of a differentiable function into $n-1$ zeros of its derivative. **Why natural.** Applying Rolle's theorem on $[0,x]$ and on $[x,1]$ gives zeros $\eta_1 \in (0,x)$ and $\eta_2 \in (x,1)$ of ϕ'; applying it once more on $[\eta_1,\eta_2]$ gives a zero ξ of ϕ''. Since $\phi''(t) = g''(t) - \frac{2g(x)}{x(x-1)}$, the equation $\phi''(\xi) = 0$ is precisely $g(x) = \frac{g''(\xi)}{2}x(x-1)$. **Link to the next block.** This closed-form remainder is now a genuine equation, not an inequality, and every subsequent step is pure estimation.

Block 4. **Cue.** The remainder formula expresses $g(\alpha)$ as $\frac{g''(\xi)}{2}\alpha(\alpha-1)$, a product of a bounded quantity and an explicit, fully computable quadratic in α. **Why natural.** Bounding $|g''(\xi)|$ by B and maximising $\alpha(1-\alpha)$ by completing the square, both one-line facts, multiply to give exactly $\frac{B}{8}$; no case split on the sign of $\alpha - \frac{1}{2}$ or on the sign of $g(\alpha)$ is ever needed, because the remainder formula already carries the correct sign built in. **Why natural.** Checking sharpness with $f(x) = \frac{1}{2} - x$ confirms every inequality in the chain, $|g''(\xi)| \le B$ and $\alpha(1-\alpha) \le \frac{1}{4}$, is attained simultaneously

by the same simplest possible example, an affine function, which is the honest reason $\frac{1}{8}$ is the best constant and not merely a convenient one. **Link to the next block.** There is no fifth block; the bound holds for every $\alpha \in (0,1)$ at once.

Layer 3: WHY

The truth underneath this problem is the classical bound on how far a twice-differentiable function can be pushed away from zero once it is pinned to zero at both ends of an interval and its curvature is controlled: for $g \in C^2[0,L]$ with $g(0) = g(L) = 0$ and $|g''| \le B$, the same double-Rolle argument gives $|g(x)| \le \frac{B}{2}x(L-x) \le \frac{BL^2}{8}$ for $x \in (0,L)$. The exponent 2 on the interval length is not an accident, it is the same exponent that appears in every bound relating a function to its second derivative through a boundary condition, from the error term of linear interpolation to the discretisation error of the trapezoidal rule to the Green's function of the operator $-\frac{d^2}{dx^2}$ on $[0,L]$ with zero boundary values, whose maximum value is exactly $\frac{L}{4}$ at the midpoint, matching the $\frac{1}{4}$ that appears here.

The discernment mechanism for spotting this route: whenever a problem's hypotheses describe an integral of a given function vanishing at two points, and asks for a bound on that integral in terms of a bound on the given function's derivative, the hypotheses are secretly about the antiderivative, not the original function. Promote the antiderivative to a named object immediately and restate the problem in its language; the bound you need is then always governed by the two-point remainder identity, never by ad hoc estimation of the original function. Training this discernment is a matter of drilling the auxiliary function $\phi(t) = g(t) - \frac{g(x)}{x(x-1)}t(t-1)$ until writing it down becomes as automatic as writing $g(x) = \int_0^x f$ in the first place.

The examiners nudged you toward exactly this shape by restricting α to the open interval $(0,1)$, which is precisely where the auxiliary function's denominator $x(x-1)$ is nonzero, and by phrasing the bound using $\max|f'|$ rather than an integral norm of f', which only makes sense as a bound on a pointwise remainder term, not as a bound obtained from an averaging argument.

Revolving concepts

Fundamental theorem of calculus. For f continuous on $[0,1]$, $g(x) = \int_0^x f(t)\,dt$ is differentiable with $g' = f$; this is the move that manufactures the object the whole solution is built around. See Book 1, field: Calculus and integrals, topic: antiderivatives from integral hypotheses.

Rolle's theorem. If h is continuous on $[a,b]$, differentiable on (a,b), and $h(a) = h(b)$, then $h'(c) = 0$ for some $c \in (a,b)$. Applied twice in sequence, first to ϕ then to ϕ', it is the engine that produces the point ξ where the second derivative is pinned down. See Book 1, field: Real analysis, topic: Rolle's theorem and iterated applications.

Two-point Lagrange remainder. The generalisation is the remainder formula for polynomial interpolation at $k+1$ nodes $x_0 < x_1 < \cdots < x_k$ in $[a,b]$: if P is the unique polynomial of degree at most k agreeing with a C^{k+1} function h at those nodes, then for any $x \in (a,b)$ there is $\xi \in (a,b)$ with

$$h(x) - P(x) = \frac{h^{(k+1)}(\xi)}{(k+1)!}\prod_{i=0}^{k}(x - x_i).$$

This problem is the case $k = 1$, two nodes 0 and 1, with $P \equiv 0$ since g vanishes at both nodes; the same identity with $k = 2$ underlies Simpson's rule error bounds, and with general k underlies polynomial interpolation error throughout numerical analysis. See Book 1, field: Real analysis, topic: interpolation remainder formulas.

Maximising a quadratic by completing the square. $\alpha(1-\alpha) = \frac{1}{4} - \left(\alpha - \frac{1}{2}\right)^2$, maximised at $\alpha = \frac{1}{2}$; a one-line alternative to differentiating and solving for a stationary point. See Book 1, field: Inequalities, topic: completing the square for extremal problems.

No sequence-type recursion appears in this problem; the natural generalisation here runs along interval length and interpolation order instead, as stated above.

Circuit Bank

- **WHEN** a bound is requested on $\int_a^x f$ and the hypotheses force this integral to vanish at both endpoints of an interval **DO** introduce $g(x) = \int_a^x f$ explicitly and restate every hypothesis in the language of g before estimating anything **BECAUSE** the true structure, zero boundary values and a bounded higher derivative, belongs to g, not to the original f, and estimating f directly throws that structure away.
- **WHEN** a twice-differentiable function vanishes at two points and you need a bound on it strictly between them from a bound on its second derivative **DO** construct $\phi(t) = g(t) - \frac{g(x)}{x(x-1)} t(t-1)$ and apply Rolle's theorem twice **BECAUSE** this is the two-point Lagrange remainder in explicit form, and it produces the sharpest possible pointwise bound with an equation, not merely an inequality.
- **WHEN** you need the maximum of a quadratic such as $\alpha(1-\alpha)$ on a fixed interval **DO** complete the square rather than differentiate and solve for a stationary point **BECAUSE** completing the square exposes the location and the value of the maximum in the same line, with no need to check it is a maximum rather than a minimum.
- **WHEN** you have just proved a sharp-looking constant in an inequality **DO** exhibit the simplest function that saturates every inequality used in the proof, typically an affine or piecewise-constant one **BECAUSE** affine functions make every derivative-based inequality in a proof of this shape an equality simultaneously, confirming the constant is genuinely best possible rather than an artefact of the method.

Backtesting

1. First useful cue: the two conditions $\int_0^\alpha f$ and $\int_0^1 f = 0$ are both integrals of f from 0, begging to be named as one function.
2. Method the expression suggested: introduce $g(x) = \int_0^x f$ and treat the problem as bounding a function that vanishes at two points via a bound on its second derivative.
3. Where you would hesitate, and the thought to reinforce: hesitating over how to produce a point where the second derivative is pinned down exactly; the reinforcement is that Rolle's theorem, applied twice to an auxiliary function built to share three zeros with g, always produces exactly such a point.
4. The trap the question sets: trying to bound f directly from its derivative bound without first locating a base point where f vanishes, which leaves an unresolved constant in the estimate.
5. What makes you recognise it faster next time: seeing a function forced to equal zero at both ends of an interval alongside a bound on a derivative one order higher than that function's own order.
6. The reusable decision rule: whenever boundary values vanish and a higher derivative is bounded, reach for the two-point remainder identity via a double application of Rolle's theorem before reaching for any direct estimate.

Mirror drills

1. (drill) With g as in the lemma, show $\int_0^1 |g(x)|\,dx \le \frac{B}{12}$. *Answer:* integrate the pointwise bound $|g(x)| \le \frac{B}{2} x(1-x)$ over $[0,1]$; since $\int_0^1 x(1-x)\,dx = \frac{1}{6}$, the bound follows directly.
2. (drill) Suppose instead $g(0) = 0$ and $g(1) = c \neq 0$, with $|g''| \le B$ still. Bound $g(\alpha) - \ell(\alpha)$, where ℓ is the line through $(0,0)$ and $(1,c)$. *Answer:* subtracting a linear function does not

change g'', so the same double-Rolle argument applies to $g - \ell$, which vanishes at both endpoints, giving $|g(\alpha) - \ell(\alpha)| \le \frac{B}{8}$ unchanged.

3. (drill) Redo the lemma on $[0, L]$ instead of $[0, 1]$: if $g \in C^2[0, L]$ with $g(0) = g(L) = 0$ and $|g''| \le B$, find the sharp bound on $|g(x)|$. *Answer:* the identical construction gives $g(x) = \frac{g''(\xi)}{2} x(x - L)$, so $|g(x)| \le \frac{B}{2} x(L - x) \le \frac{BL^2}{8}$, showing the constant $\frac{1}{8}$ scales with the square of the interval length.

19.9 B3. eliminating a floor via its algebraic conjugate

Let $x_0 = 1$ and for $n \geq 0$, let $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5}\rfloor$. In particular, $x_1 = 5$, $x_2 = 26$, $x_3 = 136$, $x_4 = 712$. Find a closed-form expression for x_{2007}. ($\lfloor a\rfloor$ means the largest integer $\leq a$.)

Read it out loud

The defining recursion mixes an ordinary integer multiple, $3x_n$, with the floor of an irrational multiple of the same variable, $\lfloor x_n\sqrt{5}\rfloor$. The instant a floor function wraps around an irrational multiple of a variable that is itself forced to stay an integer, the guiding light is to bound the floor between the irrational quantity and one less than it, then hunt for an integer combination of consecutive terms that makes the irrational part cancel exactly. That is the generic move for any sequence of the form $x_{n+1} = ax_n + \lfloor x_n c\rfloor$ with c irrational: eliminate the floor algebraically, do not try to reason about it term by term.

The problem hands you four terms explicitly, x_1 through x_4. That is a strong cue that a pattern is meant to be noticed, but noticing is not proving, and it depends on recognising the numbers 13, 34, 89 hiding inside $x_2/2$, $x_3/4$, $x_4/8$ as Fibonacci numbers, a recognition that either fires instantly or not at all. A systematic route that never needs this recognition is strictly more valuable to install as a circuit, because it produces the answer whether or not the Fibonacci connection ever crosses your mind, and it produces the SAME answer, since both routes describe the same sequence.

Once the floor is eliminated algebraically, what remains is a genuine linear recurrence with constant coefficients, and everything downstream is standard machinery: a characteristic equation, two roots, and a 2×2 linear system for two undetermined coefficients using the two earliest terms you have. Nothing about the index 2007 is special once a closed form exists; raising a fixed real number to the 2007th power already is the closed form, the problem is not asking for a decimal.

Fields embedded: sequences and recurrences, specifically floor-function recursions and constant-coefficient linear recurrences solved via characteristic equations, with a touch of number theory, the irrationality of $\sqrt{5}$ and the use of its algebraic conjugate. The specific combination, bounding a floor via an irrational argument and then eliminating that irrationality with a conjugate, sits past A level Further Mathematics, though each separate ingredient, solving a linear recurrence and using a conjugate surd identity like $(a + \sqrt{d})(a - \sqrt{d}) = a^2 - d$, is individually familiar. Within thirty seconds of reading you should have written $x_n\sqrt{5} - 1 < \lfloor x_n\sqrt{5}\rfloor < x_n\sqrt{5}$.

Approaches, ranked worst to best

1. **compute terms one by one up to $n = 2007$.** *Cue:* the recursion is explicit, so direct computation always technically works. *Circuit:* none, this is not a closed form at all. (why it ranks here: this never answers the question, since the problem explicitly demands a closed-form expression, and even as a warm-up it gives no route to the general term.)
2. **spot the Fibonacci numbers inside $x_1, \ldots, x_4$, conjecture $x_n = 2^{n-1}F_{2n+3}$, prove by induction.** *Cue:* the given values $13, 34, 89$ are recognisable Fibonacci numbers once divided out by powers of 2. *Circuit:* pattern recognition plus inductive verification. (why it ranks here: this works, but it depends on recognising specific numbers by sight, a recognition that is not guaranteed to fire, and the inductive step still has to reconstruct essentially the same conjugate bound derived below in order to verify the floor identity, so the recognition buys you the answer in advance without buying you less work.)
3. **eliminate the floor using the conjugate $3 - \sqrt{5}$, derive a floor-free linear recurrence, solve it directly.** *Cue:* the recursion has the shape $x_{n+1} = ax_n + \lfloor x_n c\rfloor$ with $a = 3$, $c = \sqrt{5}$, and $a - c = 3 - \sqrt{5}$ lies strictly between 0 and 1. *Circuit:* conjugate elimination of a floor. (why it ranks here: this is the ideal route; it needs no guessing, produces the recurrence and hence the closed form mechanically, and generalises immediately to any recursion of this shape.)

The ideal solution

Step 1: positivity and irrationality. By induction, x_n is a positive integer for every $n \geq 0$: $x_0 = 1$, and if x_n is a positive integer then $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5} \rfloor$ is an integer, and positive since $x_n\sqrt{5} > 0$ gives $\lfloor x_n\sqrt{5} \rfloor \geq 0$. Since $\sqrt{5}$ is irrational and $x_n \neq 0$, the number $x_n\sqrt{5}$ is irrational, hence never an integer, so its fractional part lies strictly between 0 and 1:

$$x_n\sqrt{5} - 1 < \lfloor x_n\sqrt{5} \rfloor < x_n\sqrt{5}.$$

Step 2: a bracket for x_{n+1}. Since $\lfloor x_n\sqrt{5} \rfloor = x_{n+1} - 3x_n$, substituting into the bound above and rearranging gives

$$(3+\sqrt{5})x_n - 1 \;<\; x_{n+1} \;<\; (3+\sqrt{5})x_n. \tag{$\star$}$$

Step 3: eliminate $\sqrt{5}$ via the conjugate. Multiply $(\star)$ by $3-\sqrt{5}$, which is positive since $\sqrt{5} < 3$, so the inequalities are preserved. Using $(3+\sqrt{5})(3-\sqrt{5}) = 9 - 5 = 4$,

$$4x_n - (3-\sqrt{5}) \;<\; (3-\sqrt{5})x_{n+1} \;<\; 4x_n.$$

Rewrite $x_{n+1}\sqrt{5} = 3x_{n+1} - (3-\sqrt{5})x_{n+1}$, so negating the bracket above and adding $3x_{n+1}$ throughout gives

$$3x_{n+1} - 4x_n \;<\; x_{n+1}\sqrt{5} \;<\; 3x_{n+1} - 4x_n + (3-\sqrt{5}).$$

Since x_n, x_{n+1} are integers, $M := 3x_{n+1} - 4x_n$ is an integer, and $0 < 3 - \sqrt{5} < 1$, so $x_{n+1}\sqrt{5}$ lies strictly between M and $M + (3-\sqrt{5})$, an interval of length less than 1 contained in $(M, M+1)$. Hence

$$\lfloor x_{n+1}\sqrt{5} \rfloor = M = 3x_{n+1} - 4x_n.$$

Step 4: the floor-free recurrence. Substituting into the definition at index $n+1$,

$$x_{n+2} = 3x_{n+1} + \lfloor x_{n+1}\sqrt{5} \rfloor = 3x_{n+1} + (3x_{n+1} - 4x_n) = 6x_{n+1} - 4x_n,$$

valid for every $n \geq 0$. A check against the given values confirms it: $6 \cdot 5 - 4 \cdot 1 = 26 = x_2$, $6 \cdot 26 - 4 \cdot 5 = 136 = x_3$, $6 \cdot 136 - 4 \cdot 26 = 712 = x_4$.

Step 5: solve the linear recurrence. The characteristic equation $t^2 - 6t + 4 = 0$ has roots $t = 3 \pm \sqrt{5}$, both real and distinct, so

$$x_n = A(3+\sqrt{5})^n + B(3-\sqrt{5})^n$$

for constants A, B determined by $x_0 = 1$ and $x_1 = 5$. From x_0: $A + B = 1$. From x_1: $A(3+\sqrt{5}) + B(3-\sqrt{5}) = 5$, i.e. $3(A+B) + \sqrt{5}(A-B) = 5$, so $\sqrt{5}(A-B) = 2$ and $A - B = \frac{2}{\sqrt{5}}$. Solving the pair $A + B = 1$, $A - B = \frac{2}{\sqrt{5}}$ gives

$$A = \frac{1}{2} + \frac{1}{\sqrt{5}} = \frac{5+2\sqrt{5}}{10}, \qquad B = \frac{1}{2} - \frac{1}{\sqrt{5}} = \frac{5-2\sqrt{5}}{10}.$$

Therefore, for every $n \geq 0$,

$$x_n = \frac{5+2\sqrt{5}}{10}(3+\sqrt{5})^n + \frac{5-2\sqrt{5}}{10}(3-\sqrt{5})^n.$$

At $n = 2007$,

$$\boxed{x_{2007} = \frac{5+2\sqrt{5}}{10}(3+\sqrt{5})^{2007} + \frac{5-2\sqrt{5}}{10}(3-\sqrt{5})^{2007}.}$$

Remark (agreement with the Fibonacci-flavoured form). Writing $\varphi = \frac{1+\sqrt{5}}{2}$, the identity $\varphi^2 = \varphi + 1 = \frac{3+\sqrt{5}}{2}$ gives $3 + \sqrt{5} = 2\varphi^2$, and since $\varphi\left(\frac{1-\sqrt{5}}{2}\right) = -1$ also gives $3 - \sqrt{5} = 2\varphi^{-2}$. A short computation with $\varphi^3 = 2\varphi + 1 = \sqrt{5} + 2$ converts the boxed answer into

$$x_n = \frac{2^{n-1}}{\sqrt{5}}\left(\varphi^{2n+3} + \varphi^{-(2n+3)}\right) = 2^{n-1}F_{2n+3},$$

where F_k is the kth Fibonacci number, matching $x_1 = 2^0F_5 = 5$, $x_2 = 2^1F_7 = 26$, $x_3 = 2^2F_9 = 136$, $x_4 = 2^3F_{11} = 712$ exactly against the given data. This second form is equivalent to the boxed one; it is included as a cross check, not as a replacement, since the boxed form was reached without ever needing to notice a Fibonacci number.

Layer 1: WHAT

Block 1: establish x_n is a positive integer and $x_n\sqrt{5}$ is irrational, giving strict bounds $x_n\sqrt{5} - 1 < \lfloor x_n\sqrt{5}\rfloor < x_n\sqrt{5}$. Block 2: translate the definition into a bracket for x_{n+1} itself, $(3+\sqrt{5})x_n - 1 < x_{n+1} < (3+\sqrt{5})x_n$. Block 3: multiply by the conjugate $3 - \sqrt{5}$ to collapse the width of the bracket to less than 1, pinning $\lfloor x_{n+1}\sqrt{5}\rfloor$ down exactly as $3x_{n+1} - 4x_n$. Block 4: substitute back to obtain the floor-free recurrence $x_{n+2} = 6x_{n+1} - 4x_n$. Block 5: solve via the characteristic equation, fit the two constants to x_0, x_1, and evaluate at $n = 2007$.

The nominal route tries to recognise the given numbers as a known sequence and prove the guess correct afterward; the ideal route never guesses, it eliminates the floor first and the closed form falls out mechanically. The bifurcation point is the decision, on first reading the four given values, between reaching for pattern recognition (does this remind me of a sequence I know) and reaching for algebraic elimination (can I multiply by a conjugate to kill the irrational part). The second choice never depends on what you happen to remember.

Layer 2: HOW

Block 1. **Cue.** The recursion produces x_n as an integer at every step by construction, and involves $\sqrt{5}$, an irrational number, multiplied by that integer. **Why natural.** An irrational number times a nonzero integer is irrational, so it can never land exactly on an integer, which means its fractional part is strictly between 0 and 1, never touching either endpoint. That strictness is what will later guarantee a floor is pinned down uniquely rather than merely bounded. **Link to the next block.** Rewriting the floor as a strict double inequality converts the recursion from an equation involving an unfriendly floor function into an inequality you can manipulate freely.

Block 2. **Cue.** The floor in the definition equals $x_{n+1} - 3x_n$ by rearranging the recursion itself. **Why natural.** Substituting this into the bound from Block 1 turns a statement about $\lfloor x_n\sqrt{5}\rfloor$ into a direct two-sided bound on x_{n+1}, namely $(3+\sqrt{5})x_n - 1 < x_{n+1} < (3+\sqrt{5})x_n$, with no floor symbol left in it. **Link to the next block.** This bracket still contains $\sqrt{5}$ on both sides in an unavoidable way; the next move must remove it if a rational, integer linear recurrence is the goal.

Block 3. **Cue.** The bracket contains $3 + \sqrt{5}$, and its algebraic conjugate $3 - \sqrt{5}$ satisfies $(3 + \sqrt{5})(3 - \sqrt{5}) = 4$, a rational integer, and also satisfies $0 < 3 - \sqrt{5} < 1$. **Why natural.** Multiplying the whole bracket by $3 - \sqrt{5}$ turns the irrational coefficient $3 + \sqrt{5}$ into the rational integer 4 on one side, while the smallness of $3 - \sqrt{5}$ shrinks the width of the resulting window below 1; a window of width less than 1 around an irrational number determines its floor uniquely, exactly the fact needed to compute $\lfloor x_{n+1}\sqrt{5}\rfloor$. **Link to the next block.** With $\lfloor x_{n+1}\sqrt{5}\rfloor$ now known exactly as $3x_{n+1} - 4x_n$, feeding it back into the original definition at the next index produces a recurrence with no floor function at all.

Block 4. **Cue.** The defining recursion applies just as well at index $n+1$ as at index n. **Why natural.** Substituting the newly pinned-down floor, $x_{n+2} = 3x_{n+1} + (3x_{n+1} - 4x_n) = 6x_{n+1} - 4x_n$, is now a constant-coefficient linear recurrence, and every tool for that class of recurrence applies without further justification. **Link to the next block.** A linear recurrence with constant coefficients is solved by its characteristic equation; there is nothing sequence-specific left to discover.

Block 5. **Cue.** The recurrence $x_{n+2} = 6x_{n+1} - 4x_n$ has characteristic polynomial $t^2 - 6t + 4$. **Why natural.** Its two distinct real roots $3 \pm \sqrt{5}$ give a general solution $A(3+\sqrt{5})^n + B(3-\sqrt{5})^n$, and the two given values $x_0 = 1$, $x_1 = 5$ determine A, B uniquely via a 2×2 linear system, which is routine algebra. **Link to the next block.** Plugging $n = 2007$ into the resulting closed form finishes the problem; no further block is needed.

Layer 3: WHY

The mechanism underneath this problem is the algebra of the ring $\mathbb{Z}[\sqrt{5}]$: the number $3+\sqrt{5}$ has norm $(3+\sqrt{5})(3-\sqrt{5}) = 4$, a rational integer, and its conjugate $3-\sqrt{5}$ has absolute value strictly less than 1. Whenever a recursion is built from an algebraic integer whose conjugate is smaller than 1 in absolute value, bounding a floor and then multiplying by that conjugate always shrinks the resulting window below width 1, which is exactly the condition needed to read off a floor exactly rather than merely bound it. This is the same phenomenon, restricted to a single step, that underlies Pisot numbers: an algebraic integer greater than 1 all of whose other conjugates have absolute value less than 1 produces sequences, such as $\lfloor \beta^n \rfloor$ or the sequence in this problem, whose fractional parts behave in a controlled, eventually predictable way, because the conjugate's contribution to the exact value shrinks geometrically.

The discernment mechanism for spotting this route: whenever a recursion has the shape $x_{n+1} = ax_n + \lfloor x_n\sqrt{d} \rfloor$ for integers a, d, immediately compute the norm $a^2 - d$ and check whether $|a - \sqrt{d}| < 1$. If so, the three-step pattern, bound the floor, multiply by the conjugate, re-derive the next floor exactly, will always produce a floor-free integer linear recurrence with characteristic roots $a + \sqrt{d}$ and $a - \sqrt{d}$. Training this discernment means drilling that three-step pattern on several values of a, d until reaching for the conjugate is as automatic as reaching for the quadratic formula.

The examiners nudged you toward this by supplying four explicit terms, more than the two needed to seed a second-order recurrence; that surplus is there so you can either notice the Fibonacci pattern by trivia recognition, or, if you derive the recurrence systematically, use the extra terms as a numerical check that your derivation is correct, exactly as done in Step 4 above. The specific choice of $\sqrt{5}$ is a second-layer nudge toward the golden ratio for those who spot it, but the argument above never depends on recognising that connection at all.

Revolving concepts

Conjugate surds and norms. For a, d integers with d not a perfect square, $(a+\sqrt{d})(a-\sqrt{d}) = a^2 - d \in \mathbb{Z}$. This is the algebraic fact that converts an irrational coefficient into a rational one after multiplication. See Book 1, field: Number theory, topic: quadratic irrationals and their conjugates.

Floor bounds from irrationality. If y is irrational, $y - 1 < \lfloor y \rfloor < y$ strictly, with no equality possible at either end; this strictness is what lets a window of width less than 1 pin a floor down exactly. See Book 1, field: Number theory, topic: fractional parts of irrational multiples.

Linear recurrences with constant coefficients (the recursion here, generalised). For a recurrence $x_{n+k} = c_{k-1}x_{n+k-1} + \cdots + c_0 x_n$, form the characteristic polynomial $t^k - c_{k-1}t^{k-1} - \cdots - c_0$; if it has k distinct roots $\lambda_1, \ldots, \lambda_k$, every solution has the form $x_n = \sum_i A_i \lambda_i^n$, with the A_i determined by the first k terms via a linear system, and repeated roots contribute terms of the form $n^j A\lambda^n$ instead. The present problem is the case $k = 2$, $\lambda_{1,2} = 3 \pm \sqrt{5}$. The same solving method applies verbatim to any floor-eliminated recursion produced by the conjugate technique above, for any integers a, d with

$|a - \sqrt{d}| < 1$: the resulting recurrence is always $x_{n+2} = 2a\,x_{n+1} - (a^2 - d)x_n$, with roots $a \pm \sqrt{d}$. See Book 1, field: Sequences and recurrences, topic: characteristic equations for linear recurrences.

Fibonacci numbers and Binet's formula. $F_k = \frac{\varphi^k - \psi^k}{\sqrt{5}}$ with $\varphi = \frac{1+\sqrt{5}}{2}$, $\psi = \frac{1-\sqrt{5}}{2}$, itself the $k = 2$ instance of the general linear recurrence solving method above, applied to $F_{k+2} = F_{k+1} + F_k$; the connection to this problem is that $3 \pm \sqrt{5} = 2\varphi^{\pm 2}$, so the two closed forms are the same recurrence viewed through two different but algebraically equivalent root labellings. See Book 1, field: Sequences and recurrences, topic: Binet's formula and the golden ratio.

Circuit Bank

- **WHEN** a recursion mixes an integer coefficient with the floor of an irrational multiple of the same variable **DO** bound the floor between the irrational quantity and one less than it, then multiply by the algebraic conjugate of that irrational coefficient **BECAUSE** the conjugate's norm is a rational integer and its size below 1 shrinks the resulting window enough to pin the next floor down exactly, eliminating the floor from the recursion for good.
- **WHEN** a floor has just been eliminated and a genuine constant-coefficient linear recurrence remains **DO** solve it via its characteristic equation and fit the free constants using the two earliest available terms **BECAUSE** the resulting closed form is valid at every index simultaneously, so an index as large as 2007 costs no more work than an index of 2.
- **WHEN** a problem supplies more explicit terms of a sequence than the order of its recurrence requires **DO** treat the surplus terms as a numerical check on a systematically derived recurrence, not as the primary route to it **BECAUSE** a derivation confirmed against real data is stronger evidence than a pattern spotted from the data alone, and does not depend on recognising any particular named sequence.
- **WHEN** a quadratic irrational $a + \sqrt{d}$ appears inside a floor-elimination argument **DO** compute its norm $a^2 - d$ and its conjugate $a - \sqrt{d}$ immediately **BECAUSE** the conjugate's size relative to 1 tells you at a glance whether the elimination technique will succeed in pinning the next floor down exactly.

Backtesting

1. First useful cue: an integer-valued recursion built from $3x_n$ plus the floor of $x_n\sqrt{5}$, mixing a rational and an irrational coefficient.
2. Method the expression suggested: bound the floor strictly, then multiply by the conjugate $3 - \sqrt{5}$ to eliminate the irrational part from the next term.
3. Where you would hesitate, and the thought to reinforce: hesitating over whether the resulting window after multiplying by the conjugate is really short enough to pin a floor down; the reinforcement is to check explicitly that the conjugate's absolute value, here $3 - \sqrt{5} \approx 0.764$, is less than 1.
4. The trap the question sets: the four explicit values invite you to recognise Fibonacci numbers and stop there, without ever proving that the guessed formula satisfies the defining floor recursion in general.
5. What makes you recognise it faster next time: seeing a floor wrapped around an irrational multiple, $\lfloor x_n c \rfloor$, sitting inside an otherwise integer-valued recursion.
6. The reusable decision rule: whenever a floor and a conjugable irrational appear together in a recursion, eliminate the floor via the conjugate before attempting any pattern recognition.

Mirror drills

1. (drill) Let $y_0 = 1$ and $y_{n+1} = 2y_n + \lfloor y_n\sqrt{3}\rfloor$. Find the floor-free linear recurrence satisfied by y_n. *Answer:* here $a = 2$, $d = 3$, conjugate $2 - \sqrt{3} \approx 0.268 \in (0, 1)$, norm $(2 + \sqrt{3})(2 - \sqrt{3}) = 1$; the same three-step argument gives $y_{n+2} = 4y_{n+1} - y_n$.
2. (drill) Why does the elimination argument fail if the conjugate $a - \sqrt{d}$ is not strictly between -1 and 1? *Answer:* the window bounding the next term, of width $|a - \sqrt{d}|$ after scaling, would then have width at least 1, and could straddle two consecutive integers, so the floor would no longer be determined uniquely by the argument.
3. (drill) Confirm that the boxed closed form and the Fibonacci-flavoured form agree at $n = 2$ without recomputing x_2 from the original recursion. *Answer:* using $3 + \sqrt{5} = 2\varphi^2$ and $3 - \sqrt{5} = 2\varphi^{-2}$ with $\varphi = \frac{1+\sqrt{5}}{2}$, both forms reduce to the same rational-plus-surd expression once expanded, since they were shown above to be algebraically identical for every n, not merely checked numerically at small cases.

19.10 B4. Counting sums of two polynomial squares

Let n be a positive integer. Find the number of pairs P, Q of polynomials with real coefficients such that

$$(P(X))^2 + (Q(X))^2 = X^{2n} + 1$$

and $\deg P > \deg Q$.

Read it out loud

Say the structure aloud before you touch a pen. This is a counting question about representations of one fixed real polynomial as a sum of two squares of real polynomials, with a degree condition attached to break the symmetry between the two slots. Three words carry the whole problem: *sum of two squares*, *real coefficients*, *count*. Nothing else in the statement does any work; $X^{2n}+1$ is just a convenient target that will turn out to be squarefree with no real roots.

A sum of two squares is never to be read as a sum. It is a norm. Over the integers $a^2 + b^2 = (a + bi)(a - bi)$ is the norm of a Gaussian integer, and everything you know about which integers are sums of two squares, and in how many ways, comes from factorisation in $\mathbb{Z}[i]$. The same sentence transplants verbatim to polynomials: $P^2 + Q^2 = (P + iQ)(P - iQ)$ in $\mathbb{C}[X]$, and $\mathbb{C}[X]$ is a Euclidean domain, so it has unique factorisation and you know every factorisation of the target explicitly, because over $\mathbb{C}$ every polynomial splits into linear factors and you can write down the roots of $X^{2n}+1$ in your sleep. The fields wired into this problem are therefore: factorisation in $\mathbb{C}[X]$, the conjugation automorphism that pins down which factorisations come from *real* P and Q, roots of unity, and at the very end a two line count.

What is actually being tested is whether you convert an additive statement about real polynomials into a multiplicative statement about complex ones, and then whether you can say precisely which multiplicative data descends back to the reals. The second half is where marks are lost. Splitting $X^{2n}+1$ into $2n$ linear factors is free; deciding which subsets of those factors can be the factors of $P + iQ$ is the whole problem, and there is a published write-up of this problem that gets that step wrong (it gets the right count from the wrong subsets; see Layer 3).

Within thirty seconds you should have on the page: $P^2 + Q^2 = (P+iQ)(P-iQ)$; $\deg P = n$ with leading coefficient ± 1 and $\deg Q \le n-1$; and $X^{2n}+1 = \prod_{m=0}^{2n-1}(X - \zeta_m)$ with $\zeta_m = \exp\big(i\pi(2m+1)/(2n)\big)$. Everything after that is bookkeeping about conjugate pairs.

Approaches, ranked worst to best

1. **Undetermined coefficients.** *Cue:* a polynomial identity with unknown coefficients. *Circuit:* expand and compare coefficients. You set $P = \sum_{j=0}^{n} p_j X^j$ and $Q = \sum_{j=0}^{n-1} q_j X^j$ and match $2n+1$ coefficients. The equations are quadratic in $2n+1$ unknowns, and they are not decoupled. For $n = 1$ this is over in a line; for general n there is no route from here to a count, because you cannot even see the solution set. Worst, and it is worst for a structural reason: you have thrown away the multiplicative structure of the identity and replaced it by an unstructured nonlinear system.

2. **Small cases and a guessed pattern.** *Cue:* the answer is a number depending on n. *Circuit:* compute $n = 1, 2$, guess, stop. You get 4 and 8, guess 2^{n+1}, and write it down. The answer is right and the write-up is worth almost nothing, because the command term is "find the number" on a proof paper: the graders want the bijection. Worth doing in the arena as a sanity check on the real solution, never as the solution.

3. **Real quadratic factors plus the two squares identity.** *Cue:* $X^{2n}+1$ factors over $\mathbb{R}$ into n quadratics. *Circuit:* Brahmagupta and Fibonacci identity, multiply representations together. Write $X^{2n}+1 = \prod_{j=1}^{n}\left(X^2 - 2\cos\theta_j X + 1\right)$ with $\theta_j = \pi(2j-1)/(2n)$, notice each factor is

itself a sum of two squares, $X^2 - 2\cos\theta_j X + 1 = (X - \cos\theta_j)^2 + (\sin\theta_j)^2$, and combine them with $(a^2+b^2)(c^2+d^2) = (ac-bd)^2 + (ad+bc)^2 = (ac+bd)^2 + (ad-bc)^2$. Each of the $n-1$ multiplications offers two outcomes, so you manufacture representations in bulk. This is the correct mathematics in disguise, and it constructs plenty of pairs, but it gives you no upper bound: to prove there are no others you must show every representation arises this way, which is exactly the unique factorisation argument you were avoiding. It also costs far more ink, since the identity is just the norm being multiplicative, written out in real coordinates.

4. **Norm factorisation in $\mathbb{C}[X]$.** *Cue:* sum of two squares of reals. *Circuit:* $P^2 + Q^2 = N(P + iQ)$; factor the target, then impose reality by conjugation. Set $A = P + iQ$, so that the equation reads $A \cdot \sigma(A) = X^{2n} + 1$ where σ conjugates coefficients. Unique factorisation in $\mathbb{C}[X]$ turns the problem into a choice of a subset of the $2n$ roots, and reality turns that into "exactly one root from each conjugate pair". Best: it produces the construction and the completeness proof in the same breath, and it generalises immediately to any real target (see Revolving concepts).

The ideal solution

The number of such pairs is $\boxed{2^{n+1}}$.

Degrees. Suppose $P^2 + Q^2 = X^{2n} + 1$ with $\deg P > \deg Q$. Use the convention $\deg 0 = -\infty$, so a priori $Q = 0$ is not excluded by the degree condition and must be excluded by hand later. Let $d = \deg P$ and let p be the leading coefficient of P. Since $\deg Q^2 < \deg P^2$, the coefficient of X^{2d} in $P^2 + Q^2$ equals $p^2 \neq 0$, so $2d = 2n$ and $p^2 = 1$. Thus

$$\deg P = n, \qquad p = \pm 1, \qquad \deg Q \leq n - 1.$$

The target. The polynomial $X^{2n} + 1$ has the $2n$ roots

$$\zeta_m = \exp\Big(\frac{i\pi(2m+1)}{2n}\Big), \qquad m = 0, 1, \ldots, 2n-1,$$

which are distinct because their arguments $\pi(2m+1)/(2n)$ are distinct in $[0, 2\pi)$. Counting degrees,

$$X^{2n} + 1 = \prod_{m=0}^{2n-1} (X - \zeta_m).$$

None of these roots is real, since $x^{2n} + 1 \geq 1 > 0$ for every real x. Complex conjugation permutes the root set (the coefficients are real) and has no fixed point on it, so it partitions the $2n$ roots into exactly n conjugate pairs $\{\zeta, \bar{\zeta}\}$.

The norm form. Let $\sigma\colon \mathbb{C}[X] \to \mathbb{C}[X]$ be the map that conjugates every coefficient, $\sigma\big(\sum a_k X^k\big) = \sum \overline{a_k} X^k$. It is a ring automorphism, it satisfies $\sigma^2 = \mathrm{id}$, its fixed ring is exactly $\mathbb{R}[X]$, and $\sigma(X - \zeta) = X - \bar{\zeta}$. Put

$$A = P + iQ, \qquad \text{so that} \qquad \sigma(A) = P - iQ, \qquad A\,\sigma(A) = P^2 + Q^2.$$

Conversely, any $A \in \mathbb{C}[X]$ determines the pair

$$P = \tfrac{1}{2}\big(A + \sigma(A)\big), \qquad Q = \tfrac{1}{2i}\big(A - \sigma(A)\big),$$

and both are fixed by σ, hence have real coefficients: $\sigma(P) = \frac{1}{2}(\sigma(A) + A) = P$ and $\sigma(Q) = \frac{1}{-2i}(\sigma(A) - A) = Q$. So real pairs (P, Q) correspond bijectively to elements $A \in \mathbb{C}[X]$, and the equation to be solved becomes

$$A\,\sigma(A) = X^{2n} + 1.$$

Which A occur. $\mathbb{C}[X]$ is a Euclidean domain, hence a unique factorisation domain, its units are the nonzero constants, and the primes dividing $X^{2n}+1$ are the $2n$ distinct monic linear polynomials $X-\zeta_m$. Therefore A, being a divisor of $X^{2n}+1$, has the form

$$A = c\prod_{m\in S}(X-\zeta_m), \qquad c\in\mathbb{C}^\times, \quad S\subseteq\{0,1,\dots,2n-1\},$$

and then $\sigma(A)=\bar{c}\prod_{m\in S}(X-\bar{\zeta}_m)$. Write $Z(S)=\{\zeta_m : m\in S\}$ and let $\overline{Z(S)}$ be its image under conjugation. Comparing $A\,\sigma(A)=X^{2n}+1$ with the factorisation of the target and using uniqueness, we need $c\bar{c}=1$, that is $|c|=1$, together with

$$\overline{Z(S)} = \{\zeta_0,\dots,\zeta_{2n-1}\}\setminus Z(S).$$

This last condition says precisely that S contains exactly one member of each conjugate pair. Indeed, if $Z(S)$ contained both ζ and $\bar{\zeta}$ then $\bar{\zeta}\in\overline{Z(S)}$ would also lie outside $Z(S)$, a contradiction; if it contained neither, then ζ lies outside $Z(S)$, hence in $\overline{Z(S)}$, so $\bar{\zeta}\in Z(S)$, again a contradiction. Conversely, one from each pair clearly gives $\overline{Z(S)}=Z(S)^c$. In particular $|S|=n$ and $\deg A=n$.

The constant. The coefficient of X^n in $A=P+iQ$ is the leading coefficient of P, because $\deg Q\le n-1$; that coefficient is c, so $c=\pm1$. Conversely, take any S with one index from each conjugate pair and any $c\in\{1,-1\}$, and set $A=c\prod_{m\in S}(X-\zeta_m)$. Then $A\,\sigma(A)=c\bar{c}\prod_{m=0}^{2n-1}(X-\zeta_m)=X^{2n}+1$; the associated P and Q are real by the paragraph above; and the coefficient of X^n in A is the real number c, so the coefficient of X^n in Q is 0, giving $\deg Q\le n-1<n=\deg P$. Every such (S,c) is therefore admissible. (This also disposes of the degenerate case: $Q=0$ would force $A=\sigma(A)$, hence $Z(S)=\overline{Z(S)}=Z(S)^c$, which is impossible for a nonempty S.)

The count. The correspondence $(S,c)\mapsto A\mapsto(P,Q)$ is injective, since $A=P+iQ$ recovers c as its leading coefficient and S as its root set. There are n conjugate pairs and two choices in each, so 2^n admissible sets S, and two choices of c. Hence the number of pairs is $2\cdot 2^n=2^{n+1}$. □

Layer 1: WHAT

Block 1. Degree bookkeeping: $\deg P=n$, leading coefficient ±1, $\deg Q\le n-1$.

Block 2. Rewrite the additive equation multiplicatively: $P^2+Q^2=(P+iQ)(P-iQ)$, and name $A=P+iQ$.

Block 3. Name the reality operator: σ conjugates coefficients, its fixed ring is $\mathbb{R}[X]$, and the equation is $A\,\sigma(A)=X^{2n}+1$.

Block 4. Factor the target: $2n$ distinct roots, none real, so n conjugate pairs.

Block 5. Unique factorisation: A is a constant times a subproduct; reality forces the root set of A to meet each conjugate pair exactly once; the degree condition forces the constant to be ±1.

Block 6. Count the free choices: $2^n\cdot 2=2^{n+1}$, and check injectivity so that the choices really are the pairs.

Juxtaposition. The nominal route stays inside $\mathbb{R}[X]$: expand, compare coefficients, or factor into real quadratics and grind the two squares identity. It can manufacture examples but never closes the count, because inside $\mathbb{R}[X]$ the object P^2+Q^2 is irreducible information, whereas inside $\mathbb{C}[X]$ it is a product. The bifurcation point is Block 2, the single line where you write $P^2+Q^2=(P+iQ)(P-iQ)$ and commit to leaving the reals. Everything before that line is common to both routes; everything after it is either a two paragraph proof or a page of algebra with no ending.

Layer 2: HOW

Block 1. Cue. The problem hands you a degree inequality rather than a symmetric condition, which means it is there to remove an ambiguity, not to be used cleverly. **Why natural.** Squares of reals cannot cancel leading terms: if $\deg P>\deg Q$ then the top coefficient of P^2+Q^2 is p^2, which is positive, so no cancellation can occur and the degree of P is forced. Reading a degree condition

as "leading coefficients cannot interfere" is the standard first move with real squares. **Link to the next block.** Now you know the shape of the answer set: P monic up to sign of degree n, Q strictly smaller. That is exactly the data you will later need to pin the unit constant c, so keep it.

Block 2. Cue. The expression $P^2 + Q^2$, with real coefficients, and a question that asks how many. **Why natural.** Counting questions are answered by structure, and the only structure a sum of two squares has is that it is a norm. Over $\mathbb{R}$, $a^2 + b^2$ does not factor; over $\mathbb{C}$ it always does, into $(a + bi)(a - bi)$. Since the ring $\mathbb{C}[X]$ has unique factorisation and you can list the primes of the target explicitly, factoring converts "how many representations" into "how many ways to split a known list of primes", which is a finite, countable choice. Nothing here is a leap: the identity $a^2 + b^2 = (a + ib)(a - ib)$ is the definition of the complex norm, and you were handed a norm. **Link to the next block.** Having introduced i, you have widened the ring, so you must immediately record what distinguishes the polynomials you actually want (real ones) inside the bigger ring. That is Block 3.

Block 3. Cue. You have moved from $\mathbb{R}[X]$ to $\mathbb{C}[X]$ and the problem said "real coefficients". **Why natural.** Whenever you enlarge a ring to gain factorisation, you must carry the symmetry that characterises the smaller ring. Here that symmetry is coefficientwise conjugation σ, whose fixed ring is exactly $\mathbb{R}[X]$. Writing the equation as $A\,\sigma(A) = X^{2n} + 1$ is not a cosmetic rewrite: it records both halves of the problem at once, since the left factor is arbitrary and the right factor is then determined by it. **Link to the next block.** The equation now says that the multiset of primes of the target splits into a part S and the conjugate of that part. To use that you need to know the primes and how conjugation acts on them, which is Block 4.

Block 4. Cue. $X^{2n} + 1$, a binomial with a root of unity flavour. **Why natural.** You know the roots of $z^{2n} = -1$ from de Moivre, so the factorisation is free. Two facts must be checked because the argument depends on them: the roots are distinct (so the target is squarefree, and subsets rather than multisets suffice), and none is real (so conjugation pairs them up without fixed points). Both are one line, and both are load bearing, so both go in the write-up. **Link to the next block.** You now have n conjugate pairs. The equation $\overline{Z(S)} = Z(S)^c$ is a condition on how S meets those pairs, which is Block 5.

Block 5. Cue. A divisor of a squarefree product of known primes. **Why natural.** Unique factorisation says a divisor is a unit times a subproduct, with no other freedom. The reality condition then reads off pairwise, and the two possible failures (both members of a pair chosen, neither chosen) are each contradictory in one line. The unit c is constrained twice: $|c| = 1$ from the norm, and $c \in \mathbb{R}$ from the degree condition, which is where the hypothesis $\deg P > \deg Q$ finally pays for itself. **Link to the next block.** The construction is now reversible, so the last block is a count of choices plus a check that different choices give different pairs.

Block 6. Cue. A bijection has been set up, so count the domain. **Why natural.** Two independent choices: n binary choices for the root set, one binary choice for the sign. Injectivity is not decoration: without it you would be counting factorisations rather than pairs. Say it in one sentence and take the mark.

Layer 3: WHY

The mathematical truth under this route is that $\mathbb{R}[X] \subset \mathbb{C}[X]$ is a degree two extension of unique factorisation domains, with the nontrivial automorphism σ, and $P^2 + Q^2$ is its norm form $N(A) = A\sigma(A)$. Counting the ways to write a fixed element as a norm is counting the ways to split its prime

factorisation into a part and the conjugate of the complementary part. That is the same sentence that governs sums of two squares in $\mathbb{Z}$: the number of representations of N is controlled by how the primes of N split in $\mathbb{Z}[i]$, and the primes that split contribute a factor to the count while the inert ones contribute nothing. Here every prime splits, because $\mathbb{C}[X]$ has no inert primes over $\mathbb{R}[X]$ except the ones coming from real roots, and $X^{2n}+1$ has none.

The discernment mechanism to install is a dictionary between quadratic expressions and rings: a^2+b^2 means $\mathbb{Z}[i]$ or $\mathbb{C}$ over $\mathbb{R}$; a^2-db^2 means $\mathbb{Z}[\sqrt{d}]$; a^2+ab+b^2 means the Eisenstein integers; $a^2+b^2+c^2+d^2$ means the quaternions and Lagrange. When one of these forms is set equal to something and you are asked which values occur or how many representations there are, the answer is always obtained by factoring in the corresponding ring, never by manipulating the sum. Train it by rewriting every sum of two squares you meet as a norm for a month, even the ones where you do not need it.

There is one more thing to extract here, and it is worth more than the problem itself. The published solution factors $X^{2n}+1$ as $\prod_{j=0}^{n-1}(X-\omega_j)(X+\omega_j)$ with $\omega_j=\exp(i\pi(2j+1)/(2n))$, and then asserts that you must choose exactly one factor from each pair $\{X-\omega_j, X+\omega_j\}$ in order to make $P+iQ$ and $P-iQ$ conjugate. The count that comes out, 2^n subsets, is correct, but the criterion is not: negation and conjugation are different fixed point free involutions of the root set, and they pair the roots differently as soon as $n\geq 2$. Take $n=2$ and the pair $\{X-e^{i\pi/4},\ X+e^{i\pi/4}\}$: choosing $X-e^{i\pi/4}$ and $X+e^{3i\pi/4}=X-e^{-i\pi/4}$ is legal under that criterion and gives $A=X^2-\sqrt{2}X+1$, which is real, hence $Q=0$ and $P^2=(X^2-\sqrt{2}X+1)^2\neq X^4+1$. The correct criterion is one root from each *conjugate* pair, which for $n=2$ selects the four sets $\{e^{i\pi/4},e^{3i\pi/4}\}$, $\{e^{i\pi/4},e^{-3i\pi/4}\}$, $\{e^{-i\pi/4},e^{3i\pi/4}\}$, $\{e^{-i\pi/4},e^{-3i\pi/4}\}$ and yields the eight genuine pairs. Both involutions have n orbits of size 2, which is why the final number survives the slip. The lesson is not that the authors were careless; it is that a counting argument must name the exact equivalence it is counting orbits of, and that you should verify any such pairing on the first case where two candidate pairings differ.

What the examiners nudged toward: they chose $X^{2n}+1$ rather than a general target precisely so that the target is squarefree with no real roots, which makes the orbit structure uniform and the answer a clean power of two. The degree condition $\deg P>\deg Q$ is the nudge that kills the unit circle of constants c down to $\{1,-1\}$; without it the problem would have infinitely many solutions, since A could be multiplied by any unimodular constant.

Revolving concepts

Sum of two squares as a norm. For a ring R with $\mathbb{C}$ style extension $R[i]$, $N(a+bi)=a^2+b^2$ is multiplicative: $N(uv)=N(u)N(v)$. Written out in coordinates this multiplicativity is the Brahmagupta and Fibonacci identity $(a^2+b^2)(c^2+d^2)=(ac-bd)^2+(ad+bc)^2$, and the second expression $(ac+bd)^2+(ad-bc)^2$ comes from conjugating one factor. See Book 1, field: Complex numbers, topic: norms and the two squares identity.

Unique factorisation in $\mathbb{C}[X]$. $\mathbb{C}[X]$ is Euclidean under the degree function, hence a principal ideal domain, hence a unique factorisation domain; its units are the nonzero constants and its primes are the monic linear polynomials. A polynomial is squarefree exactly when it shares no root with its derivative, which for $X^{2n}+1$ is immediate since the derivative $2nX^{2n-1}$ vanishes only at 0. See Book 1, field: Polynomials, topic: factorisation and squarefreeness.

The conjugation automorphism. σ acting coefficientwise on $\mathbb{C}[X]$ fixes exactly $\mathbb{R}[X]$; for a real polynomial the root multiset is stable under conjugation; a real polynomial of odd degree has a real root; and a real root of P^2+Q^2 with P,Q real must be a common root of P and Q, hence a root of the sum of even multiplicity. See Book 1, field: Polynomials, topic: real versus complex factorisation.

The generalisation, which is where the recursion lives. The count $c_n=2^{n+1}$ satisfies $c_{n+1}=2c_n$, and the reason is structural rather than numerical: each new conjugate pair of roots contributes an independent binary choice. Generalise: let $T\in\mathbb{R}[X]$ be nonconstant with positive leading coefficient, and suppose every real root of T has even multiplicity (otherwise no representation

exists at all, since $T = P^2 + Q^2$ forces every real root of T to be a common root of P and Q). Write the nonreal roots as conjugate pairs $\{\zeta_j, \bar{\zeta}_j\}$, $j = 1, \ldots, r$, with ζ_j of multiplicity e_j. Then the number of pairs (P, Q) of real polynomials with $P^2 + Q^2 = T$ and $\deg P > \deg Q$ is

$$2\prod_{j=1}^{r}(e_j + 1),$$

provided the leading coefficient of T is a square in $\mathbb{R}$, which it is. The proof is the one above verbatim: A takes a_j copies of ζ_j and the condition $\overline{Z(S)} = Z(S)^c$ forces the multiplicity of $\bar{\zeta}_j$ in A to be $e_j - a_j$, leaving $a_j \in \{0, 1, \ldots, e_j\}$ free, while real roots are forced to be split evenly and contribute a factor 1. Our problem is the case $r = n$, all $e_j = 1$. This is the exact analogue of the divisor formula for $r_2(N)$ in $\mathbb{Z}[i]$, where a prime $p \equiv 1 \pmod 4$ appearing to the power e contributes $e + 1$. See Book 1, field: Number theory, topic: sums of two squares and $r_2(N)$.

Circuit Bank

- **WHEN** you see $a^2 + b^2$ set equal to something, with a counting or existence question attached **DO** write $a^2 + b^2 = (a + ib)(a - ib)$ and move to the ring where i lives **BECAUSE** a sum of two squares is a norm, and norms are studied by factorisation, never by addition
- **WHEN** you widen a ring to gain factorisation **DO** immediately name the automorphism whose fixed ring is the original one, and rewrite the equation as $A\,\sigma(A) = T$ **BECAUSE** the reality or rationality condition is the only thing keeping the count finite, and it must be carried explicitly, not remembered
- **WHEN** you must count divisors of a squarefree element in a unique factorisation domain subject to a symmetry **DO** count orbits of the symmetry on the prime list, then count one choice per orbit **BECAUSE** unique factorisation converts a divisor into a subset of the primes, so a symmetry constraint becomes a pairing constraint
- **WHEN** a problem gives a degree or size inequality that seems only to break a tie **DO** use it at the very end to pin down the unit or constant **BECAUSE** units are exactly the freedom that unique factorisation leaves behind, and inequalities of this kind are the examiners handing you the tool to remove them
- **WHEN** two candidate pairings of a set both give the same orbit count **DO** test them on the smallest case where they differ before trusting either **BECAUSE** equinumerous pairings can give different answers to everything except the count, and a counting proof must name the correct equivalence

Backtesting

1. First useful cue: the expression $P^2 + Q^2$ with real coefficients and a counting command term. That is a norm form, so the first line written is $(P + iQ)(P - iQ)$.
2. Method the expression suggested: factor the target completely in $\mathbb{C}[X]$, then impose reality through coefficientwise conjugation, converting the equation into a subset choice on the $2n$ roots.
3. Where you would hesitate, and the thought to reinforce: at the moment of deciding which subsets S are legal. Reinforce this: the condition is $\sigma(A)$ equals the complementary factor, so the root set of A conjugates onto its own complement, which is exactly one root from each conjugate pair. Do not reason about which factors "look paired"; compute the involution.

4. The trap the question sets: the degree condition. It looks like a tiebreaker to prevent double counting (P, Q) and (Q, P), and it is really the constraint that reduces the unit constant from the whole unit circle to $\{1, -1\}$. Miss that and you either report 2^n or claim infinitely many.
5. What makes you recognise it faster next time: the shape "fixed real object equals a norm form, count the representations". That shape has one answer in every ring: split the primes.
6. The reusable decision rule: WHEN a real quadratic form equals a fixed element and you are asked how many representations, DO factor in the corresponding quadratic ring and count prime splittings subject to the conjugation constraint, BECAUSE representation counts are divisor counts in disguise.

Mirror drills

1. Write down all pairs (P, Q) of real polynomials with $P^2 + Q^2 = X^4 + 1$ and $\deg P > \deg Q$. *Answer:* the eight pairs $(\pm X^2, \pm 1)$ and $\big(\pm(X^2 - 1), \pm\sqrt{2}X\big)$; check $(X^2 - 1)^2 + 2X^2 = X^4 + 1$.
2. How many pairs (P, Q) of real polynomials satisfy $P^2 + Q^2 = (X^2 + 1)^2$ with $\deg P > \deg Q$? *Answer:* six, by the generalisation with one conjugate pair of multiplicity 2, giving $2(2 + 1)$; explicitly $\pm(X^2 + 1)$ with $Q = 0$, and $\pm(X^2 - 1)$ with $Q = \pm 2X$. The lesson is that squarefreeness was doing real work in the original problem.
3. Prove that there are no real polynomials with $P^2 + Q^2 = X^{2n} - 1$. *Answer:* substituting $X = 1$ gives $P(1)^2 + Q(1)^2 = 0$, so $P(1) = Q(1) = 0$, so $(X - 1)^2$ divides $P^2 + Q^2$; but 1 is a simple root of $X^{2n} - 1$.

19.11 B5. Reducing powers of a floor function

Let k be a positive integer. Prove that there exist polynomials $P_0(n), P_1(n), \ldots, P_{k-1}(n)$ (which may depend on k) such that for any integer n,

$$\left\lfloor \frac{n}{k} \right\rfloor^k = P_0(n) + P_1(n) \left\lfloor \frac{n}{k} \right\rfloor + \cdots + P_{k-1}(n) \left\lfloor \frac{n}{k} \right\rfloor^{k-1}.$$

($\lfloor a \rfloor$ means the largest integer $\leq a$.)

Read it out loud

Read the shape, not the symbols. Something of the form u^k is to be written as a combination of $1, u, u^2, \ldots, u^{k-1}$ with coefficients that are polynomials in n. That sentence has one meaning in mathematics and only one: u satisfies a monic polynomial equation of degree k over the ring $\mathbb{Q}[n]$. This is not an existence question about mysterious polynomials P_i; it is the assertion that $u = \lfloor n/k \rfloor$ is integral of degree at most k over $\mathbb{Q}[n]$, and the only way to prove such an assertion is to write the monic degree k equation down.

Now look at the two k's. The modulus is k and the exponent is k. That coincidence is the guiding light the proposers planted. Why would the exponent be tied to the modulus? Because $\lfloor n/k \rfloor$, for integer n, is not one mysterious quantity: it is one of exactly k explicit quantities. If $n \equiv j \pmod{k}$ then $\lfloor n/k \rfloor = (n-j)/k$, and j ranges over $\{0, 1, \ldots, k-1\}$. So u always equals one of the k known rational functions $(n-0)/k, (n-1)/k, \ldots, (n-k+1)/k$. A quantity that is always one of k known values is always a root of the monic degree k polynomial whose roots are those k values. There is the equation, and the problem is finished before it starts.

The fields embedded: the division algorithm and the definition of the floor (that is all the number theory there is here), Vieta and elementary symmetric functions to expand the product, and the algebraic idea of an annihilating polynomial, which is the same idea as a minimal polynomial of an algebraic number and as Cayley and Hamilton for a matrix. What is being tested is whether you can convert a case analysis into a single algebraic identity, because a k way case split is unusable when k is a variable, while a product over k cases is one line.

Within thirty seconds you should have written: $n = kq + r$ with $q = \lfloor n/k \rfloor$ and $r \in \{0, 1, \ldots, k-1\}$, and then $\prod_{j=0}^{k-1} \left(kq - n + j\right) = 0$. Everything else is expansion.

Approaches, ranked worst to best

1. **Compute small k and hope for an induction on k.** *Cue:* the statement is indexed by k. *Circuit:* do $k = 2$, do $k = 3$, look for a pattern. You get $k = 2$ in three lines, $k = 3$ with effort, and then you stall, because the statement for k says nothing whatever about the statement for $k+1$: the objects $\lfloor n/2 \rfloor$ and $\lfloor n/3 \rfloor$ are unrelated. There is no induction available in the variable that indexes the problem. Worst, and worth recognising quickly: when the induction variable changes the object rather than extending it, induction is not a candidate.

2. **Split into residue classes and interpolate.** *Cue:* the floor is piecewise linear in n on each residue class. *Circuit:* Lagrange interpolation across the k classes. On $n = km + r$ the required identity becomes a polynomial identity in m, so you try to build P_i class by class and then interpolate in r to glue. It can be made to work, but it fights itself: the P_i must be polynomials in n alone, while the natural data is indexed by r, and r is not a polynomial function of n. You end up reintroducing the very case analysis you are trying to remove. More ink, more places to be wrong.

3. **Vandermonde system with Cramer.** *Cue:* for each fixed n there are k equations and k unknowns $P_0, \ldots, P_{k-1}$. *Circuit:* invert a Vandermonde matrix. For fixed n, the requirement is $\sum_{i=0}^{k-1} P_i(n) q_r^i = q_r^k$ for each r, where $q_r = (n-r)/k$. The q_r are distinct, so the Vandermonde

matrix is invertible and Cramer gives each $P_i(n)$ as a ratio of determinants. This actually works cleanly, and there is a pleasant detail: $q_s - q_r = (r-s)/k$ is a constant, so the Vandermonde determinant is a nonzero constant independent of n, hence Cramer returns polynomials rather than rational functions with denominators to worry about. It ranks third only because it costs a page of determinant bookkeeping and returns the coefficients in a form you cannot read.

4. **Annihilating polynomial from the division algorithm.** *Cue:* $\lfloor n/k \rfloor$ is always one of k explicit values. *Circuit:* multiply the k cases together into a single vanishing product. Write $\prod_{j=0}^{k-1} \left(q - (n-j)/k\right) = 0$, which is true because one factor is zero, expand by Vieta, and read off the P_i. Best by a distance: it is three lines, it is uniform in k, it gives the coefficients in closed form, and it proves as a bonus that the P_i are unique.

The ideal solution

Write $q = \lfloor n/k \rfloor$ throughout, and recall that polynomials here are allowed rational coefficients; the problem asks only for polynomials, and rational coefficients will turn out to be unavoidable.

Step 1: the k possible values of q. Let n be any integer. By definition of the floor, $q \le n/k < q+1$, so $0 \le n - kq < k$; since $r := n - kq$ is an integer, $r \in \{0, 1, \ldots, k-1\}$. In other words, for every integer n there exists $j \in \{0, 1, \ldots, k-1\}$ with

$$q = \frac{n-j}{k}.$$

This holds for negative n as well as positive; the floor is defined by the inequality above and nothing in the argument used a sign.

Step 2: the annihilating product. Consider the polynomial in one variable t, with coefficients in $\mathbb{Q}[n]$,

$$F(t) \;=\; \prod_{j=0}^{k-1} \left(t - \frac{n-j}{k} \right).$$

It is monic of degree k in t. By Step 1, for every integer n one of the k factors of $F(q)$ vanishes, hence

$$F(q) = \prod_{j=0}^{k-1} \left(\left\lfloor \frac{n}{k} \right\rfloor - \frac{n-j}{k} \right) = 0.$$

Step 3: expand. By Vieta, for indeterminates $v_0, \ldots, v_{k-1}$,

$$\prod_{j=0}^{k-1} (t - v_j) \;=\; \sum_{i=0}^{k} (-1)^{k-i} e_{k-i}(v_0, \ldots, v_{k-1})\, t^i,$$

where e_m denotes the mth elementary symmetric polynomial and $e_0 = 1$. Substituting $v_j = (n-j)/k$ and $t = q$, and using $e_m\left(\frac{n}{k}, \frac{n-1}{k}, \ldots, \frac{n-k+1}{k}\right) = k^{-m} e_m(n, n-1, \ldots, n-k+1)$, the identity $F(q) = 0$ reads

$$q^k + \sum_{i=0}^{k-1} \frac{(-1)^{k-i}}{k^{k-i}}\, e_{k-i}(n, n-1, \ldots, n-k+1)\, q^i \;=\; 0.$$

Therefore, defining

$$P_i(n) \;=\; \frac{(-1)^{k-i-1}}{k^{k-i}}\, e_{k-i}\big(n,\, n-1,\, \ldots,\, n-k+1\big), \qquad i = 0, 1, \ldots, k-1,$$

we have $q^k = \sum_{i=0}^{k-1} P_i(n) q^i$ for every integer n, as required. Each P_i is a polynomial in n with rational coefficients, of degree exactly $k-i$, since e_{k-i} of k monic linear expressions in n is a monic multiple of n^{k-i} up to the constant k^{i-k}. $\square$

Worked instance. For $k = 2$: $F(t) = (t - \frac{n}{2})(t - \frac{n-1}{2})$, so

$$\left\lfloor \frac{n}{2} \right\rfloor^2 = \frac{2n-1}{2} \left\lfloor \frac{n}{2} \right\rfloor - \frac{n(n-1)}{4}.$$

Check $n = -3$: the left side is $(-2)^2 = 4$, the right side is $\frac{-7}{2} \cdot (-2) - \frac{(-3)(-4)}{4} = 7 - 3 = 4$.

Proposition (the P_i are unique). Suppose $P_0, \ldots, P_{k-1} \in \mathbb{Q}[n]$ satisfy the identity for every integer n. Set $R(n,t) = t^k - \sum_{i=0}^{k-1} P_i(n) t^i$. Fix $r \in \{0, \ldots, k-1\}$ and substitute $n = km + r$, so that $q = m$ for every integer m: the identity gives $R(km + r, m) = 0$ for all integers m. A one variable polynomial vanishing at infinitely many points is the zero polynomial, so $R(km + r, m)$ is identically zero in m; since $m \mapsto km + r$ is an invertible affine substitution, this is the same as saying that $R\big(n, (n-r)/k\big)$ is identically zero in n. Thus, over the field $\mathbb{Q}(n)$, the monic degree k polynomial $R(n,t)$ has the k distinct roots $(n-r)/k$ for $r = 0, \ldots, k-1$, so $R(n,t) = F(t)$ and the P_i are exactly the ones constructed. $\square$

Corollary. The P_i cannot in general be taken with integer coefficients: for $k = 2$ uniqueness forces $P_1(n) = (2n-1)/2$. So the word "polynomials" in the statement must be read over $\mathbb{Q}$, and the denominator k^{k-i} is not an artefact of the method.

Layer 1: WHAT

Block 1. Read the target shape: u^k in terms of $1, u, \ldots, u^{k-1}$ means u is a root of a monic degree k polynomial over $\mathbb{Q}[n]$.

Block 2. Ask what $u = \lfloor n/k \rfloor$ actually is for integer n: by the division algorithm it is one of the k explicit values $(n-j)/k$, $j = 0, \ldots, k-1$.

Block 3. Convert the k way case split into one identity: the product of the k differences vanishes, since one factor is zero.

Block 4. Expand the product by Vieta; the coefficient of q^k is 1, so solving for q^k requires no division by anything that could vanish.

Block 5. Read off P_i in closed form and record the degree and the denominator.

Block 6 (the extra mark). Prove uniqueness, which turns "here is a construction" into "this is the only construction", and explains why rational coefficients are necessary.

Juxtaposition. The nominal route is: split into residue classes, do the algebra in each, then try to glue the pieces into polynomials in n. It is a route through k separate worlds, and k is a variable, so it can never terminate on the page. The ideal route never leaves the single world of $\mathbb{Q}[n][t]$. The bifurcation point is Block 3: the moment you decide that "one of these k things holds" should be written as a product equal to zero rather than as a case analysis. Everything before that line is shared; after it the nominal route is still writing and the ideal route is finished.

Layer 2: HOW

Block 1. Cue. The exponent on the left is k and the highest exponent on the right is $k-1$, with k coefficient functions. **Why natural.** Reducing the kth power to lower powers is the definition of satisfying a monic degree k equation. This is not an analogy; it is a restatement. Recognising it converts an unfamiliar existence claim into a familiar target: produce the monic degree k equation. **Link to the next block.** A monic equation must have roots, so the next question is forced: what are the possible values of $\lfloor n/k \rfloor$?

Block 2. Cue. The modulus k inside the floor and the exponent k outside agree. **Why natural.** The floor of n/k is not an opaque object for integer n. The division algorithm says $n = kq + r$ with $r \in \{0, \ldots, k-1\}$, so $q = (n-r)/k$, and there are exactly k candidates because there are exactly k residues. The matching of the two k's is the examiners telling you where the degree k comes from. **Link to the next block.** You now have k candidate values and a quantity known to equal one of them. That is the standard input to the product move.

Block 3. **Cue.** A statement of the form "$u = v_0$ or $u = v_1$ or ... or $u = v_{k-1}$". **Why natural.** A disjunction of equalities is exactly a single product equalling zero: $\prod_j (u - v_j) = 0$. This is the cleanest way to hold a variable length case analysis in one hand, and it is the reason the argument does not care that k is a variable. Note that you never need to know which case holds, only that some case holds, which is precisely what makes the move legitimate here where r depends on n in a non polynomial way. **Link to the next block.** The product is a polynomial identity, so the remaining work is expansion.

Block 4. **Cue.** A product of k linear factors in t to be expanded. **Why natural.** Vieta is the expansion, and elementary symmetric functions are its output; you do not need to know them explicitly for the existence proof, only that they are polynomials in n. The important structural check is that the leading coefficient in t is 1, so the highest power can be isolated without dividing by anything that might be zero. Had you written the product as $\prod(kq - n + j) = 0$ instead, the leading coefficient would be k^k, which is still harmless because $k \geq 1$, but writing the monic version keeps the statement in the exact shape the problem asks for. **Link to the next block.** Expansion complete, the coefficients can be named, and then the only thing left worth having is uniqueness.

Block 5. **Cue.** The problem said "polynomials", without specifying the coefficient field. **Why natural.** Your P_i carry a denominator k^{k-i}, so you should say explicitly that rational coefficients are intended, and then earn the point by showing that this is forced rather than sloppy. **Link to the next block.** That forcing is exactly the uniqueness statement.

Block 6. **Cue.** An existence proof that produced a canonical looking answer. **Why natural.** If a construction is forced, saying so is cheap and it upgrades the write-up. The mechanism is the standard one: an identity holding for infinitely many integers is a polynomial identity, and a monic degree k polynomial with k distinct roots in a field equals the product of the corresponding linear factors.

Layer 3: WHY

The mathematical truth underneath is a statement about integrality. The ring $\mathbb{Q}[n]$ acts on the set of functions from $\mathbb{Z}$ to $\mathbb{Q}$, and the function $n \mapsto \lfloor n/k \rfloor$, which is not itself a polynomial and not even continuous as a function of a real variable, nevertheless satisfies a monic polynomial equation of degree k over $\mathbb{Q}[n]$. Consequently the $\mathbb{Q}[n]$ algebra generated by $\lfloor n/k \rfloor$ is spanned as a module by $1, \lfloor n/k \rfloor, \ldots, \lfloor n/k \rfloor^{k-1}$, and every polynomial expression in n and $\lfloor n/k \rfloor$ can be reduced to that basis. The problem asks you to prove exactly the generating relation of that module. Seen this way it belongs to the same family as: an algebraic number of degree k has $\mathbb{Q}(\alpha) = \mathbb{Q} \oplus \mathbb{Q}\alpha \oplus \cdots \oplus \mathbb{Q}\alpha^{k-1}$; a square matrix satisfies its own characteristic polynomial, so every power reduces to the first n; an eigenvalue satisfies the characteristic polynomial even though you do not know which eigenvalue you have.

That last phrase is the discernment mechanism, and it is the transferable content of this problem. You do not know which residue class n lies in, and you must never find out, because the answer would not be a polynomial in n. What you do know is the finite list of possibilities. The rule is: when an object is known to belong to a finite explicit list, do not case split, multiply. The product over the list is a single algebraic identity that is valid uniformly and is insensitive to which case actually occurs. This is why the same move solves apparently unrelated problems: quantities defined by "the nearest integer", "the remainder", "the sign", "whichever root is real", "the eigenvalue in this region". Train it by rewriting every case analysis you meet as a product or a symmetric function and asking whether the case labels ever needed to be known.

How the examiners nudged. First, they matched the exponent to the modulus, which is the only structural hint and it points directly at "degree k relation". Second, they wrote "which may depend

on k", which sounds like a permission but is really a warning that the coefficients are objects of the ambient parameter, not universal constants. Third, they wrote "for any integer n", not "for any positive integer n", so negative n must be covered; the division algorithm formulation covers it without comment, whereas a solution that reasons about "how many multiples of k are below n" has to stop and think.

The honest knowledge gap to name: nothing here is beyond a well trained A level student except the vocabulary. The reason this is a Putnam B5 rather than a school exercise is that the move "disjunction becomes product" is not taught anywhere in the A level syllabus, and without it you are left staring at a case analysis of variable length. That single move is the entire problem.

Revolving concepts

Floor and the division algorithm. For integers n and $k \geq 1$ there are unique integers q, r with $n = kq + r$ and $0 \leq r < k$, and then $q = \lfloor n/k \rfloor$. The defining inequality $\lfloor x \rfloor \leq x < \lfloor x \rfloor + 1$ is the only property used, and it makes no sign assumption. Related standard facts worth having installed: $\lfloor x + m \rfloor = \lfloor x \rfloor + m$ for integer m; Hermite's identity $\sum_{i=0}^{k-1} \lfloor x + i/k \rfloor = \lfloor kx \rfloor$; and $\lfloor \lfloor x/a \rfloor / b \rfloor = \lfloor x/(ab) \rfloor$ for positive integers a, b. See Book 1, field: Number theory, topic: floor function and division.

Annihilating polynomials, and the recursion they generate. General statement: let u be any quantity depending on a parameter, and suppose that for each value of the parameter u equals one of the k expressions $v_0, \ldots, v_{k-1}$. Then $\prod_j (u - v_j) = 0$ identically, hence

$$u^k = \sum_{i=0}^{k-1} (-1)^{k-i-1} e_{k-i}(v_0, \ldots, v_{k-1})\, u^i.$$

This is a recursion in disguise and it must be generalised as such: it says the sequence of powers $u^0, u^1, u^2, \ldots$ satisfies a linear recurrence of order k with coefficients $(-1)^{k-i-1} e_{k-i}$, so every power u^N with $N \geq k$ reduces to the span of $1, u, \ldots, u^{k-1}$, and the reduction is computed by repeatedly applying that recurrence. The general solution method for such a reduction is the same as for any linear recurrence with constant coefficients relative to the base ring: the characteristic polynomial is F itself, and the closed form of the coefficients of u^N in the reduced basis is obtained from the power sums or, equivalently, from Lagrange interpolation at the v_j. The identical structure appears as Cayley and Hamilton for matrices (every power of an m by m matrix reduces to the span of $I, A, \ldots, A^{m-1}$), as reduction modulo the minimal polynomial of an algebraic number, and as the reduction of $\lfloor n/k \rfloor^N$ for any $N \geq k$ in this problem. See Book 1, field: Abstract algebra, topic: minimal polynomials and integrality.

Elementary symmetric functions and Vieta. $\prod_{j=0}^{k-1} (t - v_j) = \sum_{i=0}^{k} (-1)^{k-i} e_{k-i}(v) t^i$, with e_m the sum of all products of m distinct v_j. In this problem $e_m(n, n-1, \ldots, n-k+1)$ appears; expanding instead in powers of $kq - n$ gives the alternative closed form

$$P_a(n) = -\frac{1}{k^{k-a}} \sum_{i=a}^{k} \begin{bmatrix} k \\ i \end{bmatrix} \binom{i}{a} (-n)^{i-a},$$

where $\begin{bmatrix} k \\ i \end{bmatrix}$ are the unsigned Stirling numbers of the first kind, defined by the rising factorial expansion $\prod_{j=0}^{k-1} (x + j) = \sum_i \begin{bmatrix} k \\ i \end{bmatrix} x^i$ and satisfying the recursion $\begin{bmatrix} k+1 \\ i \end{bmatrix} = k \begin{bmatrix} k \\ i \end{bmatrix} + \begin{bmatrix} k \\ i-1 \end{bmatrix}$. Generalised, that recursion is the statement that multiplying the rising factorial by $(x + k)$ shifts and scales the coefficient list, which is how every table of Stirling numbers is generated. See Book 1, field: Combinatorics, topic: Stirling numbers and factorial bases.

Uniqueness and the vanishing principle. A polynomial over an infinite integral domain that vanishes at infinitely many points is the zero polynomial; a monic degree k polynomial over a field with k distinct roots factors completely. Both are used in the uniqueness proposition, and together they are the standard machine for upgrading a pointwise identity to a formal one. See Book 1, field: Polynomials, topic: identity theorem for polynomials.

Circuit Bank

- **WHEN** a quantity is known to be one of k explicit values but you cannot tell which **DO** write $\prod_j (u - v_j) = 0$ and expand **BECAUSE** a disjunction of equalities is a single vanishing product, and the product does not need to know which case occurred
- **WHEN** you are asked to express u^k in terms of $1, u, \ldots, u^{k-1}$ with coefficients from a ring **DO** prove u satisfies a monic degree k polynomial over that ring, then rearrange **BECAUSE** that request is the definition of integrality of degree k, so producing the monic relation is the entire task
- **WHEN** a floor, a remainder or a nearest integer appears with integer input **DO** immediately write the division algorithm form $n = kq + r$, $0 \le r < k$ **BECAUSE** it removes the floor from the problem and replaces it by a finite explicit list, which is the only usable form
- **WHEN** you have constructed the objects a problem asks to exist and the construction looks canonical **DO** spend three lines proving uniqueness **BECAUSE** uniqueness converts a construction into a classification, and it usually explains the shape of the coefficients, such as why denominators are unavoidable
- **WHEN** a proposed induction would change the object rather than extend it, such as an induction on a modulus **DO** abandon induction and look for a uniform identity in the parameter **BECAUSE** induction needs the case k to be contained in the case $k + 1$, and a change of modulus destroys that containment

Backtesting

1. First useful cue: the exponent k matches the modulus k, and the right hand side is a combination of the lower powers. That is a monic degree k relation waiting to be written.
2. Method the expression suggested: division algorithm to list the k possible values of $\lfloor n/k \rfloor$, then multiply the k differences to get a single vanishing product, then Vieta.
3. Where you would hesitate, and the thought to reinforce: at the temptation to work residue class by residue class. Reinforce this: you must never learn which class n is in, because the coefficients are required to be polynomials in n alone, and r is not a polynomial in n. The product is the only object that respects that constraint.
4. The trap the question sets: the word "polynomials" with no coefficient field named. Denominators k^{k-i} are unavoidable, so a solver who insists on integer coefficients will conclude the statement is false. State that you work over $\mathbb{Q}$ and prove uniqueness to justify it.
5. What makes you recognise it faster next time: the phrase "express the kth power in terms of the lower powers". That phrase is always integrality, and integrality is always a monic annihilating polynomial.
6. The reusable decision rule: WHEN a quantity ranges over a finite explicit list, DO multiply the case equations into one vanishing product instead of splitting into cases, BECAUSE the product is uniform in the number of cases and eliminates the unknown case label.

Mirror drills

1. Write the identity explicitly for $k = 3$ and verify it at $n = -4$. *Answer:* $F(t) = (t - \frac{n}{3})(t - \frac{n-1}{3})(t - \frac{n-2}{3})$ gives $q^3 = \frac{3n-3}{3} q^2 - \frac{e_2}{9} q + \frac{n(n-1)(n-2)}{27}$ with $e_2 = n(n-1) + n(n-2) + (n-1)(n-2)$; at $n = -4$, $q = -2$ and both sides equal -8.

2. Let $r = n - k\lfloor n/k \rfloor$ be the remainder. Show that r^k is an integer coefficient combination of $1, r, \ldots, r^{k-1}$. *Answer:* $r \in \{0, \ldots, k-1\}$, so $\prod_{j=0}^{k-1}(r-j) = 0$; expanding gives $r^k = \sum_{i<k}(-1)^{k-i-1}e_{k-i}(0,1,\ldots,k-1)r^i$, and those elementary symmetric values are integers. Contrast with the main problem, where the corresponding coefficients are genuinely non integral.
3. Let x be a real number with $x^3 = 3x + 1$. Express x^5 in the basis $1, x, x^2$. *Answer:* $x^4 = 3x^2 + x$, so $x^5 = 3x^3 + x^2 = x^2 + 9x + 3$. Same circuit: a monic relation reduces every higher power.

19.12 B6. Counting factorial coin representations asymptotically

For each positive integer n, let $f(n)$ be the number of ways to make $n!$ cents using an unordered collection of coins, each worth $k!$ cents for some k, $1 \le k \le n$. Prove that for some constant C, independent of n,

$$n^{n^2/2-Cn}e^{-n^2/4} \;\le\; f(n) \;\le\; n^{n^2/2+Cn}e^{-n^2/4}.$$

Read it out loud

There are two things in this problem and you must separate them before doing anything. The first is a counting object: $f(n)$ is the number of solutions in nonnegative integers of $a_1 \cdot 1! + a_2 \cdot 2! + \cdots + a_n \cdot n! = n!$, because an unordered collection of coins is exactly a multiset, and a multiset is exactly the vector of multiplicities. The second is an asymptotic target with a free constant C in the exponent, and the free constant is the entire strategic content of the problem.

Take logarithms of the target immediately. The claim is

$$\tfrac{1}{2}n^2 \log n - \tfrac{1}{4}n^2 - Cn \log n \;\le\; \log f(n) \;\le\; \tfrac{1}{2}n^2 \log n - \tfrac{1}{4}n^2 + Cn \log n,$$

since $n^{Cn} = e^{Cn \log n}$. So you are asked for the first two terms of $\log f(n)$, with an error term of size $n \log n$ permitted. The main term is of size $n^2 \log n$. That is a relative precision of order $1/n$, which is an enormous licence: any step that changes $f(n)$ by a factor of $n^{O(n)}$ is free. Multiplying by 2^n is free. Replacing $n!$ by $(n-1)!$ in each of n coordinates is free, since it costs a factor n per coordinate. Throwing away an entire coordinate is free. Once you have read the tolerance, the method chooses itself: bracket the solution set between two boxes, count the boxes, and never think about the actual combinatorics again.

The fields embedded: counting lattice points in a polytope (that is what a coin problem is), and asymptotics of sums of the form $\sum i \log i$, which is Stirling's method applied to the hyperfactorial rather than the factorial. Indeed, recognise the target: $n^{n^2/2}e^{-n^2/4}$ is, up to $n^{O(n)}$, the hyperfactorial $\prod_{j=1}^{n} j^j$. So the shape of the answer is telling you in advance that the count will come out as a product of consecutive integers raised to powers proportional to themselves, which is exactly what a product of factorial ratios $\prod_i n!/i!$ is. Read the answer backwards to find the method.

Within thirty seconds you should have on the page: the tuple formulation $\sum a_i\, i! = n!$; the reduction of the claim to $\left|\log f(n) - \left(\tfrac{1}{2}n^2 \log n - \tfrac{1}{4}n^2\right)\right| \le Cn \log n$; and the observation that $a_i \le n!/i!$ for every i.

Approaches, ranked worst to best

1. **Exact evaluation of $f(n)$.** *Cue:* a coin counting function. *Circuit:* find a recursion or closed form, then estimate it. There is a recursion in principle, since a solution is a factorial base object, but no usable closed form exists and you would be computing an exact quantity in order to throw almost all of it away. Worst, and worst for a reason worth naming: the problem asked for two terms of a logarithm and you would be computing every term. That is consumption, not analysis.

2. **Generating functions and a saddle point.** *Cue:* coin counting means $\prod_{i=1}^{n}(1 - x^{i!})^{-1}$. *Circuit:* extract $[x^{n!}]$ by a contour integral and estimate. $f(n) = [x^{n!}] \prod_{i=1}^{n}(1 - x^{i!})^{-1}$ is correct and is the standard machine for partition asymptotics. Here the number of parts grows with n and the parts are spaced factorially, so the saddle point analysis is a research level computation, and it would deliver an answer far sharper than the one requested. Ranks second from bottom purely on cost: right tool, wrong problem, because the requested precision is crude.

3. **Volume heuristic for the simplex.** *Cue:* the solution set is the lattice points on $\sum x_i\, i! = n!$, $x \ge 0$. *Circuit:* lattice point count is approximately the volume. For fixed weights $w_1, \ldots, w_k$

the number of nonnegative integer solutions of $\sum a_i w_i = N$ is asymptotic to $N^{k-1}/\big((k-1)!\,w_1\cdots w_k\big)$ as $N\to\infty$. Substituting $N=n!$, $k=n$, $w_i=i!$ gives $\log(\text{volume}) = \sum_j (j-2)\log j + O(n\log n)$, which is the correct answer. But the asymptotic is for fixed k and here k grows with N, so this is a compass, not a proof. Use it to know where you are going, then never write it down.

4. **Two boxes and one Stirling type estimate.** *Cue:* the tolerance is a factor $n^{O(n)}$. *Circuit:* bound the solution set above by a product of intervals and below by a sub box in which one coordinate absorbs the remainder. Upper bound: $0\le a_i\le n!/i!$ for each i, so $f(n)$ is at most the product of the interval lengths. Lower bound: fix $a_n=0$, let $a_2,\dots,a_{n-1}$ range freely over intervals chosen small enough that no overshoot is possible, and let a_1 absorb whatever is left, which is legal because the smallest coin is worth one cent. Both counts are products of factorial ratios, and one integral comparison estimates both. Best by a distance: half a page, and the two crude bounds differ by exactly the factor the problem allows.

The ideal solution

Throughout, log is the natural logarithm and

$$G(n) \;=\; \tfrac12 n^2\log n - \tfrac14 n^2.$$

Since $n^{n^2/2\pm Cn}e^{-n^2/4} = \exp\big(G(n)\pm Cn\log n\big)$, the assertion to be proved is exactly

$$\big|\log f(n) - G(n)\big| \;\le\; C\,n\log n.$$

Assume $n\ge 2$. This is necessary: $f(1)=1$ while both sides of the required inequality equal $e^{-1/4}<1$ when $n=1$, and no choice of C repairs this because $1^{Cn}=1$. For $n\ge 2$ the factor n^{Cn} is unbounded in C, so any finite set of remaining small cases can be absorbed by enlarging C.

Step 1: the counting model. An unordered collection of coins of denominations $1!,2!,\dots,n!$ is determined by the multiplicities, so

$$f(n) \;=\; \#\Big\{(a_1,\dots,a_n)\in\mathbb{Z}^n_{\ge 0} \;:\; \sum_{i=1}^{n} a_i\,i! \;=\; n!\Big\}.$$

Step 2: an estimate for the sums that will appear.

Lemma 19.12.1. *Let c be a real number and $n\ge 2$. Then*

$$\Big|\sum_{i=1}^{n}(i+c)\log i \;-\; G(n)\Big| \;\le\; (2+|c|)\,n\log n.$$

Proof. The function $g(x)=x\log x$ is nonnegative and nondecreasing on $[1,\infty)$. For a nonnegative nondecreasing g we have $\int_{i-1}^{i} g \le g(i) \le \int_i^{i+1} g$; summing the left inequality for $i=2,\dots,n$ gives $\sum_{i=1}^n g(i)\ge \int_1^n g$, and summing the right inequality for $i=1,\dots,n-1$ gives $\sum_{i=1}^n g(i)\le \int_1^n g + g(n)$. Since

$$\int_1^n x\log x\,dx \;=\; \tfrac12 n^2\log n - \tfrac14 n^2 + \tfrac14 \;=\; G(n)+\tfrac14,$$

we obtain $\big|\sum_{i=1}^n i\log i - G(n)\big| \le n\log n + \frac14$. Also $\big|c\sum_{i=1}^n \log i\big| \le |c|\,n\log n$. Adding, and using $\frac14 \le n\log n$ for $n\ge 2$, gives the claim. □

Step 3: the upper bound. If $(a_1,\dots,a_n)$ is a solution then $a_i\,i!\le n!$ for each i, so a_i lies in $\{0,1,\dots,n!/i!\}$, a set of $n!/i!+1\le 2\,(n!/i!)$ elements, the inequality because $n!/i!\ge 1$. Hence

$$f(n) \;\le\; \prod_{i=1}^{n} 2\,\frac{n!}{i!} \;=\; 2^n \prod_{i=1}^{n}\frac{n!}{i!} \;=\; 2^n\prod_{j=2}^{n} j^{j-1},$$

where the last equality holds because $n!/i! = \prod_{j=i+1}^{n} j$, and a given j occurs in this product for exactly the $j-1$ indices $i = 1, \dots, j-1$. Taking logarithms and applying the Lemma with $c = -1$,

$$\log f(n) \;\le\; n\log 2 + \sum_{j=1}^{n}(j-1)\log j \;\le\; n\log n + G(n) + 3n\log n \;=\; G(n) + 4n\log n,$$

using $\log 2 \le \log n$ for $n \ge 2$.

Step 4: the lower bound. Consider the tuples with $a_n = 0$ and

$$a_i \in \Big\{0, 1, \dots, \tfrac{(n-1)!}{i!} - 1\Big\} \qquad (i = 2, 3, \dots, n-1),$$

which is a legitimate range because $(n-1)!/i!$ is a positive integer for $i \le n-1$. For any such choice,

$$\sum_{i=2}^{n-1} a_i\, i! \;\le\; \sum_{i=2}^{n-1}\Big(\frac{(n-1)!}{i!} - 1\Big) i! \;=\; \sum_{i=2}^{n-1}\big((n-1)! - i!\big) \;\le\; (n-2)\,(n-1)! \;\le\; n!,$$

so setting

$$a_1 \;=\; n! - \sum_{i=2}^{n-1} a_i\, i! \;\ge\; 0$$

produces a solution, because the smallest coin is worth exactly one cent and therefore absorbs any nonnegative remainder. Distinct choices of $(a_2, \dots, a_{n-1})$ give distinct solutions, so

$$f(n) \;\ge\; \prod_{i=2}^{n-1} \frac{(n-1)!}{i!} \;=\; \prod_{j=3}^{n-1} j^{\,j-2},$$

by the same regrouping as before: $(n-1)!/i! = \prod_{j=i+1}^{n-1} j$ and a given j occurs for the $j-2$ indices $i = 2, \dots, j-1$. Since the terms $j = 1$ and $j = 2$ of $\sum_{j=1}^{n}(j-2)\log j$ both vanish,

$$\sum_{j=3}^{n-1}(j-2)\log j \;=\; \sum_{j=1}^{n}(j-2)\log j \;-\; (n-2)\log n,$$

so by the Lemma with $c = -2$,

$$\log f(n) \;\ge\; G(n) - 4n\log n - (n-2)\log n \;\ge\; G(n) - 5n\log n.$$

Step 5: conclusion. For every $n \ge 2$,

$$\big|\log f(n) - G(n)\big| \;\le\; 5\, n\log n,$$

that is, $n^{n^2/2 - 5n} e^{-n^2/4} \le f(n) \le n^{n^2/2 + 5n} e^{-n^2/4}$. So $C = 5$ works for all $n \ge 2$, and the statement is false for $n = 1$. $\square$

Remark 19.12.2. The two bounds differ by a factor $n^{\Theta(n)}$, so this method cannot prove anything sharper than the stated form, and it does not need to. The volume heuristic indicates that $\log f(n) - G(n)$ is genuinely of order $n\log n$ and negative, so the tolerance $n^{\pm Cn}$ in the problem is matched to the truth rather than generous: the same statement with $e^{\pm Cn}$ in place of $n^{\pm Cn}$ would not be provable, being false.

Layer 1: WHAT

Block 1. Convert the coin language into the tuple equation $\sum a_i\, i! = n!$ with $a_i \geq 0$ integers.

Block 2. Take logarithms of the target and read the tolerance: the claim is $|\log f(n) - G(n)| \leq Cn \log n$, so any factor of size $n^{O(n)}$ is free.

Block 3. Upper bound by the box: each coordinate is confined to an interval of length $n!/i!$, so $f(n) \leq 2^n \prod_i n!/i!$, and regroup into $\prod_j j^{j-1}$.

Block 4. Lower bound by a sub box: shrink each range from $n!/i!$ to $(n-1)!/i!$ so that overshoot is impossible, discard the top coordinate, and let a_1 absorb the remainder. Regroup into $\prod_j j^{j-2}$.

Block 5. Estimate $\sum_i (i+c) \log i$ once, by comparing with $\int x \log x\, dx$, and apply it to both products.

Block 6. Assemble, and dispose of $n = 1$ honestly.

Juxtaposition. The nominal route treats this as a partition asymptotics problem: generating function, saddle point, sharp constants. It is not wrong, it is disproportionate, and in the arena it does not finish. The ideal route never computes anything about $f(n)$ beyond two inclusions of sets. The bifurcation point is Block 2, the line where you take logarithms and measure the permitted error before choosing a method. Everything upstream of that line is shared; downstream, one route is doing complex analysis and the other is counting boxes.

Layer 2: HOW

Block 1. **Cue.** "Unordered collection of coins" and a fixed total. **Why natural.** An unordered collection is a multiset, and a multiset over a finite denomination set is exactly its multiplicity vector. Nothing is lost and the object becomes a lattice point set defined by one linear equation. Do this translation before reading the rest of the problem, because you cannot estimate what you have not written as a set. **Link to the next block.** You now have a lattice point count in a simplex, and you next need to know how accurately it must be counted.

Block 2. **Cue.** A free constant C sitting in the exponent, on both sides, with the same main term. **Why natural.** A two sided bound with a free constant is never a request for a value; it is a request for the first terms of an asymptotic expansion, and the position of the constant tells you exactly how many terms. Taking logarithms turns the product form into $G(n) \pm Cn \log n$, and comparing $n \log n$ with the main term $\frac{1}{2} n^2 \log n$ shows the error budget is a factor $n^{O(n)}$ on $f(n)$ itself. This measurement is the single most valuable act in the whole solution, and it takes one line. **Link to the next block.** Knowing that a factor $n^{O(n)}$ is free, you may now use the crudest bounds available, so long as each is crude by only that much.

Block 3. **Cue.** A linear equation with positive coefficients and nonnegative unknowns. **Why natural.** Each coordinate is immediately bounded: $a_i\, i! \leq n!$ because all other terms are nonnegative. So the solution set sits inside a box, and the number of lattice points in a box is the product of the side lengths. The wasteful step, replacing $n!/i! + 1$ by $2 \cdot n!/i!$, costs a factor 2^n, which is far inside the budget; noticing that you are allowed to be wasteful here is the reason this step is written in one line rather than fussed over. **Link to the next block.** An upper bound alone proves nothing about a two sided claim, so you now need a lower bound of the same shape, which means a box that consists entirely of solutions.

Block 4. **Cue.** The denomination $1! = 1$ is present. **Why natural.** A unit coin means every nonnegative remainder can be completed. So to build solutions in bulk, choose all the other coordinates freely subject only to not overshooting $n!$, and let a_1 take up the slack. To guarantee no overshoot with room to spare, shrink each range from $n!/i!$ to $(n-1)!/i!$, which costs a factor of about n per coordinate, that is $n^{O(n)}$ in total, which the budget already told you was free. Dropping

a_n costs a factor 2. The resulting map from partial tuples to solutions is injective because a_1 is determined by the others. **Link to the next block.** Both bounds are now products of factorial ratios, so one estimate serves both.

Block 5. Cue. Products of the form $\prod_j j^{j+c}$, that is, sums $\sum_j (j+c)\log j$. **Why natural.** This is Stirling's method one level up. Stirling estimates $\sum \log i$ by comparing with $\int \log x\,dx$; here you compare $\sum i \log i$ with $\int x \log x\,dx$, and the monotone comparison gives an error at most one term, namely $n \log n$, which is exactly the size of the permitted error. Do not reach for a sharp Euler and Maclaurin expansion: the crude monotone bound is already inside budget, and a sharp expansion would cost ink for nothing. **Link to the next block.** Substituting $c = -1$ and $c = -2$ into the same lemma finishes both sides.

Block 6. Cue. The problem says "for some constant C independent of n". **Why natural.** That phrasing permits finitely many small cases to be absorbed, provided the bound is nontrivial there, which it is for $n \geq 2$ since $n^{Cn} \to \infty$ as C grows. It does not permit $n = 1$, where $n^{Cn} = 1$; there the statement is false, and saying so plainly is worth more than pretending otherwise. A grader who has read the problem knows this.

Layer 3: WHY

The mathematical truth beneath the route is that counting lattice points in a polytope is, to first order, computing a volume, and that both of the crude bounds used here are really volume computations in disguise. The solution set is the set of lattice points on the simplex $\Sigma = \{x \in \mathbb{R}^n_{\geq 0} : \sum x_i\, i! = n!\}$. The upper box is the bounding box of Σ, and the lower sub box is an inscribed box. For a fixed number of weights the count is asymptotic to the volume, $N^{k-1}/\big((k-1)!\,w_1 \cdots w_k\big)$, and substituting $N = n!$, $k = n$, $w_i = i!$ reproduces $\sum_j (j-2)\log j$ after taking logarithms, which is the same main term the rigorous bounds give. That is why the two crude boxes could possibly have agreed to the required precision: in high dimension the ratio of the bounding box to the inscribed box is exponential in the dimension, $n^{O(n)}$ here, and the problem's tolerance was written to be exactly that size. The proposers chose the tolerance to make box counting sufficient. Recognising that a stated error term has been calibrated to a specific method is a skill in itself, and it is usually the fastest way into an asymptotic problem.

The discernment mechanism: before choosing a method for an estimate, compute the ratio between the permitted error and the main term, in logarithmic form. If the permitted relative error is a constant factor, you need sharp constants and probably a saddle point or a precise Stirling expansion. If the permitted error is $e^{O(n)}$ on a quantity of size $e^{\Theta(n^2)}$, you may be exponentially wasteful per dimension but not more. If, as here, the permitted error is $e^{O(n \log n)}$, you may lose a factor of order n per coordinate, which is precisely what "replace $n!$ by $(n-1)!$" does. Train this by writing the tolerance as a factor per coordinate before starting: budget first, method second.

The second transferable truth is the role of the unit coin. Lower bounds for counting problems are usually harder than upper bounds, because you must exhibit solutions rather than merely constrain them. The presence of a denomination equal to 1 converts the lower bound into a triviality: any conservative choice of the other coordinates extends uniquely to a solution. The general principle: if the weights $w_1 \mid w_2 \mid \cdots$ or merely $w_1 = 1$, then the solution set of $\sum a_i w_i = N$ contains a full box of dimension one less, obtained by bounding the other coordinates so that their total cannot exceed N. Here the factorials are even better than coprime, they are a divisibility chain, which is also why the exact count has a factorial base flavour.

Knowledge gaps to name, because this problem is not approachable without them. First, the integral comparison estimate for monotone sums, which is Euler and Maclaurin at its crudest and is the engine behind Stirling. Second, comfort with the identity $n^{Cn} = e^{Cn \log n}$ and with reading an exponent as a tolerance, which is not taught at A level at all. Third, the hyperfactorial $H(n) = \prod_{j=1}^{n} j^j$

and its asymptotic $H(n) = A\, n^{n^2/2+n/2+1/12} e^{-n^2/4}$, where A is the Glaisher and Kinkelin constant: you do not need the constant, but you must recognise the shape $n^{n^2/2} e^{-n^2/4}$ as "a product of j^j" or you will not know what kind of product to aim for. Fourth, the fact that a coin counting problem is a lattice point count and that lattice point counts approximate volumes, which is the Ehrhart circle of ideas.

Revolving concepts

Coin counting, multisets and generating functions. The number of multisets of denominations $w_1, \dots, w_k$ with total N is the number of $(a_1, \dots, a_k) \in \mathbb{Z}^k_{\geq 0}$ with $\sum a_i w_i = N$, and it equals $[x^N] \prod_{i=1}^{k} (1 - x^{w_i})^{-1}$. For fixed k this count is a quasi polynomial in N of degree $k-1$ (Ehrhart), with leading term $N^{k-1} / \big((k-1)!\, w_1 \cdots w_k\big)$. See Book 1, field: Combinatorics, topic: coin problems and Ehrhart quasi polynomials.

The recursion behind the count, generalised. Write $f_m(N)$ for the number of representations of N using only the denominations $w_1, \dots, w_m$. Then

$$f_m(N) \;=\; \sum_{a \geq 0} f_{m-1}(N - a w_m) \;=\; f_{m-1}(N) + f_m(N - w_m),$$

with $f_0(N) = [N = 0]$. That two term form is the standard dynamic programming recursion for coin counting and is the general solution method for every problem of this shape: it says the generating function factorises, $\sum_N f_m(N) x^N = \prod_{i \leq m} (1 - x^{w_i})^{-1}$. In this problem the weights satisfy the divisibility chain $i! \mid (i+1)!$, which specialises the recursion further: writing $N = n!$ in the factorial base, the unique representation with digits $c_i \leq i$ is the greedy one, and $f(n)$ counts all representations, greedy or not. The general lesson to carry: a linear Diophantine counting problem always has both a product generating function and a one step recursion, and you choose between them by whether you want exact values or asymptotics. See Book 1, field: Generating functions, topic: partitions with prescribed parts.

Stirling's method and the hyperfactorial. For a nonnegative nondecreasing g on $[1, n]$, $\int_1^n g \leq \sum_{i=1}^{n} g(i) \leq \int_1^n g + g(n)$. With $g = \log x$ this gives $\log n! = n \log n - n + O(\log n)$, and with $g = x \log x$ it gives $\sum_{i \leq n} i \log i = \frac{1}{2} n^2 \log n - \frac{1}{4} n^2 + O(n \log n)$, that is, $\prod_{j \leq n} j^j = n^{n^2/2 + O(n)} e^{-n^2/4}$. The sharp versions are Stirling's series and the Glaisher and Kinkelin asymptotic $H(n) = A n^{n^2/2+n/2+1/12} e^{-n^2/4}$ with $A \approx 1.2824$. See Book 1, field: Real analysis, topic: Euler and Maclaurin comparison and Stirling.

Reading a tolerance. If a claim is $M(n) e^{-E(n)} \leq X(n) \leq M(n) e^{+E(n)}$, then in logarithmic form it is $|\log X - \log M| \leq E$, and the useful quantity is E divided by the number of coordinates, which is the permitted loss per coordinate. Here $E = Cn \log n$ over n coordinates, so a factor n^C per coordinate is free. See Book 1, field: Real analysis, topic: asymptotic bookkeeping.

Circuit Bank

- **WHEN** a bound is stated with a free constant in an exponent, on both sides of the quantity **DO** take logarithms first and compute the permitted error divided by the number of coordinates **BECAUSE** the position of the free constant tells you how crude you are allowed to be, and that determines the method before any mathematics is done
- **WHEN** you must count solutions of a linear equation in nonnegative integers **DO** bound above by the bounding box of the simplex and below by an inscribed box, and only sharpen if the tolerance demands it **BECAUSE** lattice point counts are volumes to first order, and in high dimension the bounding box and an inscribed box differ by a factor exponential in the dimension, which is often exactly the permitted error
- **WHEN** one of the weights in a linear Diophantine problem equals 1 or divides all the others **DO** build a lower bound by choosing every other coordinate conservatively and letting that

coordinate absorb the remainder **BECAUSE** a unit weight makes completion automatic, which converts the hard direction of a counting bound into an injection from a box

- **WHEN** you meet $\prod_j j^{j+c}$ or the shape $n^{n^2/2}e^{-n^2/4}$ **DO** recognise the hyperfactorial and estimate $\sum(j+c)\log j$ by comparison with $\int x \log x\, dx$ **BECAUSE** the same monotone comparison that proves Stirling proves this, and the error is one term, which is almost always inside the tolerance of such problems
- **WHEN** a stated asymptotic claim fails at the smallest index **DO** say so explicitly, prove the claim for $n \geq n_0$, and note that finitely many cases are absorbed by enlarging the constant when the bound is nontrivial there **BECAUSE** honesty about the edge case costs one sentence and protects the whole write up; concealing it invites a grader to assume the rest is equally careless

Backtesting

1. First useful cue: the free constant C in the exponent on both sides. That is a tolerance statement, so the first line is to take logarithms and measure it.
2. Method the expression suggested: bracket the lattice point set between a bounding box and an inscribed box, both being products of factorial ratios, then estimate the resulting sums by one integral comparison.
3. Where you would hesitate, and the thought to reinforce: at the lower bound, where it feels as though a careful construction is needed. Reinforce this: the coin of value 1 means completion is free, so the only requirement on the other coordinates is that they cannot overshoot; give yourself a factor of n of headroom per coordinate, since the budget already permits it.
4. The trap the question sets: the temptation to be accurate. Every sharp tool here, Stirling's series, saddle point, Ehrhart theory, costs more than the problem pays. The second trap is $n = 1$, where the statement is false; a solution that quietly ignores it is incomplete, and one sentence fixes it.
5. What makes you recognise it faster next time: the pairing of a factorially spaced weight set with an answer of hyperfactorial shape. That pairing says "product of factorial ratios", which says "box count".
6. The reusable decision rule: WHEN an estimate is requested with a free constant in an exponent, DO convert to logarithms, compute the permitted error per coordinate, and select the crudest method that fits inside it, BECAUSE precision beyond the tolerance is wasted ink and the tolerance is usually calibrated to one specific crude method.

Mirror drills

1. Prove that $\prod_{j=1}^{n} j^j = n^{n^2/2+O(n)}e^{-n^2/4}$. *Answer:* take logarithms and compare $\sum j\log j$ with $\int_1^n x\log x\,dx = \frac{1}{2}n^2\log n - \frac{1}{4}n^2 + \frac{1}{4}$; monotonicity gives an error of at most one term, $n\log n$, which is $O(n)$ in the exponent of n.
2. Let $g(n)$ be the number of ways to write 2^n as an unordered sum of powers 2^k with $0 \leq k \leq n$. Show that $\log_2 g(n) = \frac{1}{2}n^2 + O(n\log n)$. *Answer:* same two boxes. Above, $a_k \leq 2^{n-k}$ gives $\log_2 g \leq \sum_k (n-k) + O(n) = \frac{1}{2}n^2 + O(n)$. Below, restrict to $1 \leq k \leq n-1$ with $a_k < 2^{n-k}/(2n)$, so the total is at most $(n-1)2^n/(2n) < 2^n$ and the unit part absorbs the rest, giving $\log_2 g \geq \frac{1}{2}n^2 - O(n\log n)$.
3. Suppose the same problem had been stated with $e^{\pm Cn}$ in place of $n^{\pm Cn}$. Which steps of the solution would still be legal, and which would break? *Answer:* the replacement of $n!/i! + 1$ by

$2\,n!/i!$ survives, since it costs $e^{n \log 2}$; the shrinking of the ranges from $n!/i!$ to $(n-1)!/i!$ does not, since it costs $n^{\Theta(n)} = e^{\Theta(n \log n)}$. The exercise is the point: read what each wasteful step costs in the currency of the stated tolerance.

Chapter 20

Putnam 2006

20.1 A1. Volume of a Torus from a Quartic Inequality

Find the volume of the region of points (x, y, z) such that

$$(x^2 + y^2 + z^2 + 8)^2 \leq 36(x^2 + y^2).$$

Read it out loud

This is a three variable inequality that only ever touches z through z^2 and only ever touches x, y through the combination $x^2 + y^2$; nowhere does it ask about x or y individually, and nowhere does it mix z with x or y except by both appearing squared inside the same bracket. That is the whole grammar of the problem: whatever region this describes is invariant under any rotation about the z axis, and invariant under nothing else in general. The fields wired into the question are solid geometry, specifically the volume of a region defined implicitly by an inequality, and elementary algebra, specifically isolating a familiar shape by completing a square, dressed up to look harder than it is by the stray constant 8 and the outer square on the left.

The first useful cue, the one you should have circled within the first ten seconds, is that x and y never appear separately, only as $x^2 + y^2$. Any time that happens in three dimensions, write $r = \sqrt{x^2 + y^2}$ before doing anything else; you have converted a genuinely three dimensional inequality into a two dimensional one in the pair (r, z), at the cost of remembering that $r \geq 0$.

The second cue is the shape of the inequality itself: something squared on the left, compared against something else squared on the right. That is not a shape you integrate blindly, it is a shape you unsquare, because $A^2 \leq B^2$ with $B \geq 0$ is nothing more than $-B \leq A \leq B$, and once you write that out you will usually find, as you do here, that one of the two resulting inequalities is vacuous and the other is the entire content of the problem. What survives is a single quadratic inequality in r and z; complete the square in r and you are looking at a disc, at which point the question stops being a calculus question and becomes a question about a named theorem for volumes of revolution.

Approaches, ranked worst to best

1. **Spherical coordinates.** *Cue:* seeing $x^2 + y^2 + z^2$ and reflexively reaching for a radius ρ. *Circuit:* sum of three squares to radius. (why it ranks here: the right hand side depends on $x^2 + y^2$ alone, not $x^2 + y^2 + z^2$, so the boundary surface has no constant ρ, and the angle ϕ never decouples from ρ; you end up with a transcendental relation between ρ and ϕ and no clean region to integrate. This costs ink for no return and should be abandoned the moment you write it down and see it does not simplify.)

2. **Cylindrical coordinates with a washer integral.** *Cue:* the inequality is a function of r and z only, so the solid is a solid of revolution. *Circuit:* invariance under rotation about an axis forces cylindrical coordinates. (why it ranks here: this works completely, and is the honest fallback if you forget the next theorem, but it costs an extra integral. Having reduced the boundary to $(r - 3)^2 + z^2 \leq 1$, you would compute, for each $z \in [-1, 1]$, the annulus $3 - \sqrt{1 - z^2} \leq r \leq 3 + \sqrt{1 - z^2}$ of area $12\pi\sqrt{1 - z^2}$, then integrate that over $z \in [-1, 1]$ using $\int_{-1}^{1} \sqrt{1 - z^2}\, dz = \pi/2$. It gives the right answer but re-derives, by hand, a fact that has a name.)

3. **Cylindrical coordinates plus Pappus's centroid theorem.** *Cue:* after completing the square you are rotating a disc that does not touch the axis of revolution. *Circuit:* solid of revolution of a fixed planar region not meeting the axis. (why it ranks here: this is the theorem the washer integral above is secretly proving from scratch every time; once you recognise the disc and check it stays clear of $r = 0$, the volume is one line, $V = 2\pi(\text{centroid distance})(\text{area})$, with no integral left to compute.)

The ideal solution

Write $r = \sqrt{x^2+y^2} \geq 0$. The given inequality is $(r^2+z^2+8)^2 \leq 36r^2 = (6r)^2$. Since $6r \geq 0$, this is equivalent, by the standard fact that $A^2 \leq B^2$ with $B \geq 0$ unpacks to $-B \leq A \leq B$, to

$$-6r \leq r^2+z^2+8 \leq 6r.$$

The left inequality holds for every $r \geq 0$ regardless of z, because $r^2+z^2+8 \geq 8 > 0 \geq -6r$; it carries no information. The right inequality is the entire content of the problem:

$$r^2+z^2+8 \leq 6r \iff r^2-6r+z^2+8 \leq 0 \iff (r-3)^2+z^2 \leq 1.$$

In the half plane $r \geq 0$, this is the closed disc of radius 1 centred at $(r,z) = (3,0)$. Its centre is at distance 3 from the line $r = 0$, which exceeds the radius 1, so the disc lies entirely in $r > 0$: it never touches, let alone crosses, the axis of revolution.

The three dimensional region is exactly the solid obtained by revolving this disc about the z axis, since a point (x,y,z) with $r = \sqrt{x^2+y^2}$ satisfies the original inequality precisely when (r,z) lies in the disc, and this condition is unchanged by rotating (x,y) about the origin at fixed r,z. This is a solid torus. By Pappus's centroid theorem, the volume of a solid of revolution of a planar region of area A, lying entirely on one side of an axis in its own plane at distance d from its centroid, is $V = 2\pi dA$. Here $A = \pi \cdot 1^2 = \pi$ and, by the disc's symmetry, its centroid is its own centre, at distance $d = 3$ from the axis. Hence

$$V = 2\pi \cdot 3 \cdot \pi = \boxed{6\pi^2}.$$

Layer 1: WHAT

Block 1: notice the inequality depends on x,y only through x^2+y^2, and substitute $r = \sqrt{x^2+y^2}$. Block 2: unpack the squared inequality into its two branches, discard the vacuous one, and complete the square in r. Block 3: recognise $(r-3)^2+z^2 \leq 1$ as a disc not touching the axis $r = 0$, hence a solid torus when revolved. Block 4: apply Pappus's centroid theorem to read off the volume in one line.

The nominal route agrees with the ideal route through Block 3, then splits: instead of naming the theorem, it sets up $V = \int_{-1}^{1} \pi\big((3+\sqrt{1-z^2})^2 - (3-\sqrt{1-z^2})^2\big)\,dz$, expands the difference of squares, and integrates $12\pi\sqrt{1-z^2}$ over $z \in [-1,1]$ by recognising it as the area of a semicircle. It reaches the same answer, $6\pi^2$, but by re-deriving Pappus's theorem for this one case instead of citing it. The bifurcation point is the instant you have written $(r-3)^2+z^2 \leq 1$ and identified it as a disc clear of the axis: from that sentence on, the ideal route costs one line and the nominal route costs an integral.

Layer 2: HOW

Block 1. Cue. The right hand side of the inequality, $36(x^2+y^2)$, and the left hand side inside its bracket, $x^2+y^2+z^2+8$, both treat x,y only through x^2+y^2; neither ever isolates x or y. **Why natural.** A region in $\mathbb{R}^3$ is invariant under rotation about the z axis exactly when its defining relation is a function of z and $r = \sqrt{x^2+y^2}$ alone, by definition of what that rotation does to a point (it fixes z and fixes x^2+y^2, and moves nothing else). The substitution is therefore forced by the structure, not chosen for convenience. **Link to the next block.** With the inequality now written purely in r and z, it becomes a two dimensional problem, and the next natural move is to simplify the algebra: unsquare it.

Block 2. **Cue.** The inequality has the shape (something)^2 $\leq$(something else)^2, with the right hand side manifestly non negative. **Why natural.** $A^2 \leq B^2$ for $B \geq 0$ is, by definition of the square root as a monotonic function on non negative reals, equivalent to $-B \leq A \leq B$; you are entitled to write both branches, and checking which is active (rather than guessing) is the only way to be certain nothing was lost. **Link to the next block.** One branch, $-6r \leq r^2 + z^2 + 8$, is seen immediately to be automatic since the right side is at least 8 and the left side is at most 0 for $r \geq 0$; discarding it correctly leaves a single quadratic inequality, which completing the square converts into a geometric statement.

Block 3. **Cue.** Completing the square produced $(r-3)^2 + z^2 \leq 1$, a sum of two squared linear terms bounded above by a constant. **Why natural.** That is, by definition, the set of points within distance 1 of $(3,0)$ in the (r,z) half plane, i.e. a disc of radius 1 and centre $(3,0)$; and since the centre's distance from the line $r = 0$ is 3, which exceeds the radius 1, the whole disc lies strictly in $r > 0$. **Link to the next block.** A disc revolved about an axis it does not meet is precisely the picture a solid of revolution problem wants you to draw, and precisely the hypothesis under which the volume formula for such a solid needs no case analysis.

Block 4. **Cue.** The task is now: find the volume swept out by revolving a known disc, of known area and known centroid, about an axis at a known distance. **Why natural.** Pappus's centroid theorem states this volume is 2π times the distance from the centroid to the axis, times the area; this is itself just the statement that the volume integral $\iiint 2\pi r \, dr \, dz$ over the disc equals 2π times the first moment of the disc about the axis, divided by its area, times the area, i.e. $2\pi \bar{r} A$, which is exactly the definition of the centroid $\bar{r}$. **Link to the next block.** Nothing remains: substitute the known area and distance and the answer follows in one line.

Layer 3: WHY

The mathematical truth underneath the route is that a triple integral over a solid of revolution factors, under Fubini, into an angular integral (which is always 2π, since nothing in the region depends on the angle) times an integral over the meridian region, weighted by r. That weighted integral is, again by definition, the area of the meridian region times its centroid distance. Pappus's theorem is not a separate fact you must independently verify each time; it is this factorisation, packaged with a name so you never have to write the triple integral out again.

The discernment mechanism for choosing this route over direct integration is a habit of asking, the moment you see a rotationally symmetric region, whether the meridian cross section is a shape whose area and centroid you already know without computation: a disc, a rectangle, a triangle. If it is, Pappus finishes the problem in one line. If the meridian region is irregular, or its centroid is unknown, or it touches the axis (so that revolving it would double count or degenerate), you fall back to the washer or shell integral, which is always valid but always costs more ink. Training this discernment is a matter of pattern matching the algebra: whenever completing the square in two of your three variables produces "(linear)^2+(linear)^2 $\leq$ constant", you should see a circle before you see anything else, and immediately check where its centre sits relative to any axis of revolution in the problem.

The examiners built the numbers to reward exactly this discernment. The constant 8 is placed so that, after moving to cylindrical coordinates and completing the square, the resulting circle has centre distance 3 and radius 1, comfortably clear of the axis, guaranteeing a genuine, non self intersecting torus rather than a degenerate case; and the coefficient $36 = 6^2$ is chosen so the centre and radius come out as small integers, so that a candidate who reaches Pappus's theorem is rewarded with an immediate, clean numerical answer, while a candidate who does not is still able to finish, at greater cost, by direct integration.

Revolving concepts

Cylindrical coordinates: writing a point in $\mathbb{R}^3$ as (r, θ, z) with $r = \sqrt{x^2 + y^2} \geq 0$. You reach for them whenever a defining relation in a three dimensional geometry problem is manifestly unchanged by rotation about a fixed axis, which you recognise from the relation depending on x, y only through $x^2 + y^2$. See Book 1, field: Geometry, topic: coordinate systems adapted to symmetry.

Pappus's centroid theorem, also called the Pappus-Guldinus theorem: for a plane region of area A lying entirely on one side of a line L in its plane, at distance d from its centroid, the solid formed by revolving the region about L has volume $V = 2\pi dA$. There is a companion first theorem for surface area, that the surface generated by revolving a plane curve of length ℓ about an external axis at distance d from the curve's centroid has area $S = 2\pi d\ell$; both are proved by the same Fubini factorisation argument. See Book 1, field: Geometry, topic: solids and surfaces of revolution.

Completing the square as the standard reroute for recognising a circle or sphere disguised inside a polynomial inequality: whenever an inequality is quadratic in two variables with equal leading coefficients and no cross term, complete the square in each variable in turn and read off centre and radius directly, rather than plotting points or reasoning about the inequality numerically. See Book 1, field: Geometry, topic: conic and quadric recognition.

Washer and shell integration for solids of revolution: the general purpose tools, always available, that Pappus's theorem is a packaged special case of; useful in particular when the meridian region touches the axis, self intersects when revolved, or has no known centroid, none of which occur in this problem but all of which should be checked for before reaching for Pappus. See Book 1, field: Calculus and integrals, topic: volumes of revolution.

Circuit Bank

- **WHEN** an inequality in x, y, z depends on x, y only through $x^2 + y^2$ **DO** substitute $r = \sqrt{x^2 + y^2} \geq 0$ and treat the problem as two dimensional in (r, z) **BECAUSE** the region is then invariant under rotation about the z axis, and cylindrical coordinates are exactly the coordinates adapted to that invariance
- **WHEN** you must take a square root across an inequality of the form $A^2 \leq B^2$ **DO** write it out as $-B \leq A \leq B$ for $B \geq 0$ and check both branches explicitly **BECAUSE** silently keeping only the branch that looks useful is a skipped calculation error, and the discarded branch is often itself informative about the shape of the region
- **WHEN** a known planar shape (disc, rectangle, triangle) is being revolved about an axis it provably does not touch **DO** apply Pappus's centroid theorem, $V = 2\pi dA$, instead of setting up a washer or shell integral **BECAUSE** the theorem is the same integral pre packaged; recomputing it by hand costs ink and invites an arithmetic slip for no extra rigour
- **WHEN** completing the square yields "(linear)^2+(linear)$^2 \leq$ constant" **DO** name it immediately as a disc of that centre and radius, and check whether it meets any axis of symmetry in the problem **BECAUSE** this single check determines the topology of the final solid, torus versus ball versus degenerate case, before any volume is computed

Backtesting

1. First useful cue: the right hand side, $36(x^2 + y^2)$, mentions only x and y, telling you the region is a solid of revolution about the z axis before you write a single further line.
2. Method the expression suggested: substitute $r = \sqrt{x^2 + y^2}$, collapsing a three variable inequality into a two variable one.

3. Where you would hesitate, and the thought to reinforce: deciding which branch of $A^2 \le B^2$ survives after the square root; reinforce with the habit of writing $-B \le A \le B$ in full every time, never trusting instinct on which half to keep.
4. The trap the question sets: the added constant 8 tempts you into expanding and grinding through a washer integral rather than pausing, after completing the square, to notice you are looking at a disc that clears the axis, which is exactly when Pappus's theorem applies for free.
5. What makes you recognise it faster next time: seeing "(sum of squares plus a constant) squared, bounded by a constant times a partial sum of squares" as the fingerprint of a circle hiding inside a norm inequality, to be exposed by completing the square immediately after unsquaring.
6. The reusable decision rule: rotational symmetry in the defining inequality routes you to cylindrical coordinates, completing the square routes you to a named planar shape, and a shape clear of the axis routes you to Pappus before any integral is written.

Mirror drills

1. Find the volume of the region $(x^2+y^2+z^2+3)^2 \le 16(x^2+y^2)$. *Answer:* complete the square exactly as above to get $(r-2)^2+z^2 \le 1$, a disc of radius 1 at distance 2 from the axis, so $V = 2\pi(2)(\pi) = 4\pi^2$.
2. A disc $(x-5)^2+y^2 \le 4$ in the xy plane is revolved about the y axis. Find the swept volume without integrating. *Answer:* Pappus directly, area 4π, centroid distance 5, so $V = 2\pi(5)(4\pi) = 40\pi^2$.
3. Decide, without computing a volume, what kind of solid $(x^2+y^2+z^2+2)^2 \le 4(x^2+y^2)$ describes. *Answer:* the same reduction gives $(r-1)^2+z^2 \le -1$, which has no real solutions, so the region is empty.

20.2 A2. Infinitude of Bob's Winning Heap Sizes

Alice and Bob play a game in which they take turns removing stones from a heap that initially has n stones. The number of stones removed at each turn must be one less than a prime number. The winner is the player who takes the last stone. Alice plays first. Prove that there are infinitely many n such that Bob has a winning strategy. (For example, if $n = 17$, then Alice might take 6 leaving 11; then Bob might take 1 leaving 10; then Alice can take the remaining stones to win.)

Read it out loud

This is a combinatorial game dressed as a stone heap, and the command term is the whole plan: you are asked to *prove there are infinitely many n* with a property, not to *find* them or to *determine* the full set of them. Whenever a Putnam problem asks for infinitude of some implicitly defined set rather than its exact description, that is a nudge toward an indirect argument, because it is telling you the set itself may have no usable closed form, only a property you can exploit. The fields wired into this problem are combinatorial game theory, specifically the standard machinery of positions from which the player to move wins or loses, and number theory, specifically the distribution of primes, since the legal move sizes are exactly the numbers one less than a prime.

The second thing to read off is what those legal move sizes look like: $p - 1$ for p prime gives $1, 2, 4, 6, 10, 12, 16, 18, 22, \ldots$, a set with no arithmetic progression structure and no periodicity, because the primes themselves have none. That single observation should kill, before you write a line of proof, any plan that involves computing the winning and losing heap sizes by hand and spotting a repeating pattern; there is nothing here that is forced to repeat.

The first useful cue to write down within the first thirty seconds is the translation of the English statement into game theory language: define a heap size n to be a losing position for the player about to move, and observe that "Bob has a winning strategy at heap n, with Alice moving first" is exactly the statement "n is a losing position for its first mover". Once that translation is made, the whole problem becomes: show the set of losing positions is infinite, and the recursive structure of losing and winning positions is exactly what the rest of the solution manipulates.

Approaches, ranked worst to best

1. **Direct computation and pattern spotting.** *Cue:* a subtraction game with a small explicit move set often has an eventually periodic set of losing positions. *Circuit:* subtraction game, compute small cases, guess and prove a period. (why it ranks here: the move set here is $\{p - 1 : p \text{ prime}\}$, which is infinite and irregular; there is no reason to expect, and in fact no evidence, that the losing positions are periodic, so this plan either never terminates or leads you to assert an unproven regularity, which the hard rules forbid you from presenting as fact.)

2. **Explicit construction of an infinite family of losing positions.** *Cue:* the command term "infinitely many" can be read as an invitation to construct a family directly. *Circuit:* strong induction defining new losing positions from old ones. (why it ranks here: to add a new losing position n to your family you must check that *every* legal move from n, i.e. subtracting $p - 1$ for every prime $p \leq n$, lands on a winning position; since there are infinitely many primes, this requires controlling infinitely many cases at once for every single n you add, which is exactly as hard as knowing the full set of losing positions, so this buys nothing over the direct approach.)

3. **Contradiction on the size of the losing set, using a factorial gadget.** *Cue:* the phrase "infinitely many" together with a move set built from primes signals that the contradiction should exploit a structural fact about primes rather than any specific strategy. *Circuit:* assume finiteness, derive a universal representation requirement, falsify it with a run of consecutive composites.

(why it ranks here: this never needs to know what the losing positions actually are, only that a finite list of them cannot possibly account for every larger heap size, and the falsification is a single classical gadget, $(t+1)!+2, \ldots, (t+1)!+(t+1)$, that defeats an arbitrarily large finite list at once.)

The ideal solution

Call a heap size $n \geq 0$ a *P-position* if the player about to move at a heap of n stones loses under optimal play, with the convention that $n = 0$ is a P-position, since the player to move there has already lost (there are no stones left to take, and by the rules the previous player took the last stone and won). This is a standard recursive notion for impartial games under normal play: for $n \geq 1$, n is a P-position exactly when every legal move from n leads to a position that is *not* a P-position (an N-position), and n is an N-position exactly when *some* legal move from n leads to a P-position. Here the legal moves from a heap of n stones are exactly the sizes $s = p - 1$ for p prime and $1 \leq s \leq n$.

Since Alice moves first, "Bob has a winning strategy at heap n" means precisely that n is a P-position. Let B be the set of P-positions, including 0 by the convention above; the problem asks to show B is infinite.

Suppose, for contradiction, that B is finite, and write $B = \{b_1, \ldots, b_m\}$. By the recursive description above, every non negative integer $n \notin B$ is an N-position, so it admits some legal move to a P-position: there exist $b \in B$ and a prime p with $n - (p-1) = b$, i.e.

$$n = b + p - 1.$$

That is, every non negative integer outside the finite set B can be written in this form for some choice of $b \in B$ and prime p. We now exhibit a specific integer n outside B for which no such choice exists, a contradiction.

Let t be any integer exceeding every element of B (for instance $t = \max(B) + 1$). For each integer k with $2 \leq k \leq t+1$, k divides $(t+1)!$ (since $k \leq t+1$) and k divides k, so k divides $(t+1)!+k$; since $(t+1)!+k > k \geq 2$, this shows $(t+1)!+k$ is composite. Thus

$$(t+1)!+2,\ (t+1)!+3,\ \ldots,\ (t+1)!+(t+1)$$

are t consecutive composite integers.

Set $n = (t+1)!+t$. Since $n > t > \max(B)$, $n \notin B$, so n is an N-position, and by the argument above there must exist $b \in B$ and prime p with $p = n - b + 1$. But for every $b \in B$ we have $0 \leq b \leq \max(B) < t$, so $t - b + 1$ ranges over integers in $[2, t+1]$ as b ranges over B, and

$$n - b + 1 = (t+1)! + t - b + 1 = (t+1)! + (t - b + 1)$$

is exactly one of the composite integers exhibited above. Hence $n - b + 1$ is composite, so it cannot equal p for any prime p, for any $b \in B$. This contradicts the requirement that n admits some representation $n = b + p - 1$.

The contradiction shows B cannot be finite. There are therefore infinitely many n for which Bob has a winning strategy. ■

Layer 1: WHAT

Block 1: translate the game into P-positions and N-positions, and rewrite the target claim as "the set B of P-positions is infinite". Block 2: assume for contradiction that B is finite, and unpack what that forces for every integer outside B: a representation $n = b + p - 1$. Block 3: build, from an arbitrary finite list B, a run of consecutive composite integers long enough to defeat every element of B simultaneously. Block 4: choose n so that $n - b + 1$ lands inside that composite run for every $b \in B$ at once, contradicting the required representation.

The nominal route also starts with the P/N translation, since that is forced by the problem's language, but then tries to compute B directly, or to construct an explicit infinite subfamily of B term by term. Both stall on the same obstruction: neither ever gets to use a property of *all primes at once*, only of individual primes, and the move set $\{p-1 : p \text{ prime}\}$ offers no periodic or arithmetic handle for that. The bifurcation point is the sentence "suppose B is finite": the ideal route treats this supposition itself as the object to be attacked, using only its cardinality, while the nominal route never introduces this supposition and instead tries to fight the infinite, irregular move set directly.

Layer 2: HOW

Block 1. **Cue.** The statement to prove is about the existence of infinitely many heap sizes with a stated player winning, and the win condition is defined recursively (Bob wins now depends on what happens after every possible move). **Why natural.** Any game with a "last player to move wins" rule has, by definition, a recursive split of positions into those where the mover loses (P-positions, previous-player-wins) and those where the mover wins (N-positions, next-player-i.e.-mover-wins); "Bob wins with Alice moving first at n" is definitionally the statement "n is a P-position", so this translation is not a choice, it is what the English sentence means once game theory vocabulary is attached to it. **Link to the next block.** Having reduced the target to "the set of P-positions is infinite", the next natural step is to ask what would have to be true if it were not.

Block 2. **Cue.** The set B of P-positions is being tested for infinitude, and the recursive definition says every N-position has some move to a P-position. **Why natural.** If B were finite, then every one of the infinitely many non negative integers not in B (an N-position) is still required, by the very definition of N-position, to have at least one legal move landing in B; writing that move as subtracting $p-1$ turns "has a move into B" into the concrete algebraic statement $n = b + p - 1$ for some $b \in B$, p prime, which is now a number theory sentence rather than a game theory one. **Link to the next block.** This representation must hold for *every* large n, which is an extremely strong, and therefore fragile, requirement; the natural next move is to try to break it for one specific n.

Block 3. **Cue.** Breaking the representation $n = b + p - 1$ for a fixed finite list $b_1, \ldots, b_m$ means making $n - b_i + 1$ composite for every i at once. **Why natural.** A single classical gadget produces an arbitrarily long run of consecutive composites: for k from 2 to $t+1$, $(t+1)! + k$ is divisible by k and strictly larger than k, hence composite; this is just the definition of divisibility applied to the two terms of the sum, and it works uniformly for every k in that range without needing to know anything else about k. **Link to the next block.** Because the gadget covers every integer in a whole range $[2, t+1]$ at once, and the finitely many values $b \in B$ can be arranged, by choice of t, to make $t - b + 1$ land inside exactly that range, the gadget defeats the entire finite list B in one stroke rather than one b at a time.

Block 4. **Cue.** A specific n is needed that lies outside B and for which $n - b + 1$ falls inside the composite run for every $b \in B$. **Why natural.** Choosing $n = (t+1)! + t$ makes $n - b + 1 = (t+1)! + (t - b + 1)$, and since b ranges over $B \subset [0, t)$, the quantity $t - b + 1$ ranges exactly over $[2, t+1]$, which is precisely the range the gadget covers; this is a matter of matching the shift in the construction to the range of B, arranged by the choice $t > \max(B)$. **Link to the next block.** With $n - b + 1$ composite for every $b \in B$, no prime p can satisfy $n = b + p - 1$ for any $b \in B$, directly contradicting Block 2's requirement, which closes the proof.

Layer 3: WHY

The underlying truth is that finiteness is an extremely strong hypothesis to hand an opponent, and the honest way to use it is to make the contradiction depend only on the *size* of the finite set, never

on its specific contents. That is what happens here: the construction never asks what B actually is, only that it is bounded, via $t > \max(B)$, and the composite run is built purely from t. This is what makes the argument uniform rather than a case analysis that would need to grow with $m = |B|$.

The discernment mechanism for reaching for this kind of argument is recognising when a problem asks you to disprove finiteness of an implicitly, recursively defined set with no closed form: in that situation direct computation is rarely viable, because you would be trying to establish a regularity (periodicity, an arithmetic pattern) that the move set gives you no reason to expect, and a proof of a non existent pattern is not available. Assuming finiteness converts the target from an infinite, open ended search into a single, fixed, checkable obstruction. You train this by noticing, whenever a recursive win or lose condition is stated over an infinite, irregular move set (here, primes), that the productive question is not "what is the answer" but "what would have to be true if there were only finitely many answers, and can I break that".

The examiners chose the move set $\{p - 1 : p \text{ prime}\}$ rather than something with visible structure, such as $\{1, 2, 3, \dots\}$ or a fixed arithmetic progression, precisely so that a direct, computational attack has nowhere to gain a foothold, while the existence of arbitrarily long runs of consecutive composite integers (a fact about primes, not about the game) is exactly strong enough to finish the indirect argument in a handful of lines.

Revolving concepts

P-positions and N-positions for impartial combinatorial games under normal play: a position is a P-position if the player about to move loses under optimal play, and an N-position if the player about to move wins; the base case is that a position with no legal moves is a P-position, and the recursive rule is that a position is a P-position exactly when every move leads to an N-position, and an N-position exactly when some move leads to a P-position. This machinery generalises unchanged to any subtraction game with any set S of legal move sizes, finite or infinite: compute P-positions upward from 0 using the same recursive rule, and the winner of a heap of size n with a given first mover is determined purely by whether n is a P-position or N-position for that S. See Book 1, field: Games and invariants, topic: subtraction games and the P/N classification.

Consecutive composite integers via factorials: for any integer $t \geq 1$, the t consecutive integers $(t+1)! + 2, \dots, (t+1)! + (t+1)$ are all composite, since $(t+1)! + k$ is divisible by k for every k in that range, by construction of the factorial. This is the standard tool whenever a problem needs an arbitrarily long run of numbers avoiding a property (here, primality) at once, rather than one number at a time. See Book 1, field: Number theory, topic: gaps between primes.

Proof by contradiction on the cardinality of an implicitly defined set: to show a set is infinite, assume it is finite, extract a universal consequence that every excluded integer must satisfy, then falsify that consequence for one specific integer built to defeat every element of the assumed finite list simultaneously. This pattern recurs whenever a problem defines a set recursively (via a game, a sieve, an iterated construction) and asks only for infinitude, not for an explicit description. See Book 1, field: Logic and set theory, topic: contradiction arguments on set size.

The set $\{p - 1 : p \text{ prime}\}$ itself: an infinite, aperiodic set of positive integers with no arithmetic progression structure, in contrast to move sets such as $\{1, 2, \dots, k\}$ that produce cleanly periodic P-position sets; recognising which kind of move set you are given, before attempting any computation, tells you whether a direct or an indirect proof strategy is appropriate. See Book 1, field: Number theory, topic: irregularity of the primes.

Circuit Bank

- **WHEN** a Putnam problem asks you to prove a set defined by a game or a recursive rule is infinite, rather than to describe the set **DO** treat this as a signal for an indirect argument: assume the set is finite and look for a universal consequence to falsify **BECAUSE** a request for

infinitude rather than description usually means the set has no usable closed form, and direct computation offers nothing to generalise from

- **WHEN** the legal moves in a subtraction game are defined via primes, or any other set with no periodicity **DO** do not attempt to compute the P-positions by hand and guess a period **BECAUSE** primes (and sets built from them, such as $p-1$) have no arithmetic progression structure, so there is no period to find, and time spent searching for one is wasted
- **WHEN** a contradiction argument needs to defeat every element of an assumed finite set B at once **DO** build the falsifying construction so it depends only on a bound on B (such as $t = \max(B) + 1$), never on the individual elements **BECAUSE** a construction tailored to the specific elements of B turns into unbounded casework as $|B|$ grows; a construction tailored to a single bound handles every possible finite B uniformly
- **WHEN** you need an arbitrarily long run of consecutive composite integers **DO** use $(t+1)! + 2, \ldots, (t+1)! + (t+1)$, each divisible by its own offset from $(t+1)!$ **BECAUSE** divisibility of a sum by a common factor of both terms is immediate from the definition of the factorial, and needs no further number theoretic input

Backtesting

1. First useful cue: the command "prove there are infinitely many n" rather than "find all n" or "determine the winner for each n", signalling that the set of interest has no expected closed form.
2. Method the expression suggested: translate the win condition into the standard recursive P-position and N-position machinery for impartial games, since "last player to move wins" is exactly the normal play convention that machinery is built for.
3. Where you would hesitate, and the thought to reinforce: whether it is legitimate to include $n = 0$ in B as a base case when the game is never actually played from an empty heap; reinforce by noting this is purely a bookkeeping convention for the recursion, matching the standard convention that a position with no legal moves is a loss for the player to move.
4. The trap the question sets: the explicit worked example with $n = 17$ invites you to start computing small cases by hand, which can eat significant time before you notice the move set has no period to find; recognise the trap by checking early whether the move set itself (here, $p - 1$ for primes p) has any arithmetic structure to exploit.
5. What makes you recognise it faster next time: seeing "prove infinitely many" attached to a recursively, not explicitly, defined set as the fingerprint of a finite-versus-infinite contradiction argument, before attempting any direct computation.
6. The reusable decision rule: assume the target set is finite, translate that assumption into a universal representation every excluded integer must satisfy, then build one gadget, dependent only on a bound for the assumed finite set, that falsifies the representation for every element at once.

Mirror drills

1. In the same game but with legal move sizes $\{p : p \text{ prime}\}$ instead of $\{p-1\}$, sketch why the same contradiction strategy still works. *Hint:* the representation becomes $n = b + p$ for $b \in B$, p prime; use the same factorial run of composites, shifted by one, to defeat any finite B.
2. Prove there are infinitely many positive integers n that cannot be written as the sum of a prime and a positive perfect square. *Hint:* suppose finitely many exceptions to "cannot be written"; a

run of consecutive composites again forces most large n into the excluded case, or attack the complementary set directly with the same factorial gadget.

3. Explain, without doing the computation, why trying to prove Bob wins for infinitely many n by exhibiting an explicit formula for B (such as a residue class mod some fixed integer) is not a promising plan for this specific move set. *Answer:* B would have to be compatible with subtracting $p-1$ for every prime p at every step, and since primes are not confined to any fixed residue pattern relative to a single modulus in the way needed, no single residue class formula for B can be consistent with the recursive P-position rule.

20.3 A3. 2005 Consecutive Multiples in a Sparse Recursion

Let $1, 2, 3, \ldots, 2005, 2006, 2007, 2009, 2012, 2016, \ldots$ *be a sequence defined by* $x_k = k$ *for* $k = 1, 2, \ldots, 2006$ *and* $x_{k+1} = x_k + x_{k-2005}$ *for* $k \geq 2006$. *Show that the sequence has* 2005 *consecutive terms each divisible by* 2006.

Read it out loud

This is a linear recursion, but not the small, tidy kind you are trained on: x_{k+1} is the sum of the immediately preceding term and a term 2006 steps further back, so as a recursion "of fixed order" it has order 2006, since you need a full block of 2006 consecutive terms to determine the next one. No characteristic polynomial of degree 2006 is going to be solved by hand, and the problem does not ask you to solve one; it asks only for a divisibility statement about a bounded run of terms, which is a much weaker demand than a closed form and should be attacked with correspondingly lighter machinery. The fields wired into the question are sequences and recurrences, specifically the behaviour of a fixed order integer recursion under reduction modulo an integer, and elementary number theory, specifically what it means for a run of terms to be divisible by 2006.

The first useful cue is arithmetic, not structural: rearrange the defining relation $x_{k+1} = x_k + x_{k-2005}$ to isolate the oldest term instead of the newest, giving $x_{k-2005} = x_{k+1} - x_k$. Every coefficient here is ± 1, an integer, so this rearrangement costs nothing and needs no division; the recursion runs backward exactly as cleanly as it runs forward. That single fact, checked in one line, is what decides the whole strategy: a recursion you can run backward with integer coefficients can be extended to a doubly infinite sequence, and a doubly infinite integer sequence reduced modulo a fixed integer, from a recursion of fixed order, is forced by a pigeonhole argument to be periodic, not merely eventually periodic.

The second cue is the initial data itself: $x_k = k$ for $k = 1, \ldots, 2006$ is an arithmetic progression with common difference 1, the simplest possible input to a differencing operation. Since the backward recursion is exactly a difference, $x_{k-2005} = x_{k+1} - x_k$, running it backward from this block is almost free: every difference of consecutive terms in an arithmetic progression with common difference 1 is 1. That observation, not any deep property of 2005 or 2006 individually, is what makes the whole problem collapse to a handful of lines.

Approaches, ranked worst to best

1. **Characteristic polynomial closed form.** *Cue:* a linear recursion suggests solving $t^{2006} = t^{2005} + 1$ for its roots. *Circuit:* linear recurrence to characteristic equation to closed form. (why it ranks here: the characteristic polynomial has degree 2006 and no tractable closed form roots; even granting formal roots, extracting an exact divisibility statement about specific terms from them is not a realistic route, and the problem never requires a closed form at all, so this is solving a harder problem than the one asked.)

2. **Direct forward simulation modulo 2006.** *Cue:* the target is a statement about divisibility by 2006, so work entirely modulo 2006 from the start. *Circuit:* compute residues forward from $k = 1$ and search for the required run. (why it ranks here: the residues modulo 2006 are eventually periodic by a pigeonhole argument, but the argument alone gives no bound on how far forward the first repeat, let alone the first run of 2005 consecutive zero residues, occurs; establishing a workable bound requires essentially the same theory as the ideal route, so this approach either needs an infeasible amount of computation or quietly imports the ideal route's machinery anyway.)

3. **Reverse the recursion, extend to a doubly infinite periodic sequence, search backward from the known block.** *Cue:* the recursion is invertible with unit coefficients, and the known initial

data collapses instantly under differencing. *Circuit:* reversible recursion to full periodicity, then exploit the cheapest available data. (why it ranks here: invertibility upgrades eventual periodicity to full, two sided periodicity, which licenses searching for the target pattern anywhere in the doubly infinite sequence rather than only forward from an unknown point; searching backward from the arithmetic progression $x_k = k$ finds the pattern in a fixed, small number of algebraic steps, and periodicity alone then guarantees the same pattern reappears somewhere in the actual forward sequence.)

The ideal solution

Write the recursion, for $k \geq 2006$, in its given form $x_{k+1} = x_k + x_{k-2005}$, and note it rearranges, using only integer coefficients, to

$$x_{k-2005} = x_{k+1} - x_k.$$

Lemma 1 (eventual periodicity). If integers $y_1, y_2, \ldots$ satisfy $y_j = f(y_{j-1}, \ldots, y_{j-n})$ for $j > n$, where n is fixed and f is a fixed polynomial with integer coefficients, then for any modulus $N \geq 1$ the residue sequence $y_j \bmod N$ is eventually periodic. *Proof.* A window of n consecutive residues mod N takes only N^n possible values, so among the infinitely many windows $(y_{j-n+1}, \ldots, y_j) \bmod N$ for $j > n$, two must coincide, say at $j = j_1 < j_2$. Since f has integer coefficients, congruent inputs give congruent outputs, so the entire windows agree from j_1 onward: $y_{j_1+i} \equiv y_{j_2+i} \pmod N$ for every $i \geq 0$. Hence the residue sequence is periodic with period $j_2 - j_1$ from index $j_1 - n + 1$ onward.

Lemma 2 (full periodicity under reversibility). With the same hypotheses, suppose further there is a fixed integer coefficient polynomial g with $y_{j-n} = g(y_{j-n+1}, \ldots, y_j)$ for $j > n$. Then the sequence extends uniquely to a doubly infinite sequence $\ldots, y_{-1}, y_0, y_1, \ldots$, and this doubly infinite sequence is purely periodic modulo every N. *Proof.* Extend by defining $y_{j-n} := g(y_{j-n+1}, \ldots, y_j)$ downward for $j = n, n-1, n-2, \ldots$. The pigeonhole argument of Lemma 1 still finds $j_1 < j_2$ with matching windows; but now, since g also has integer coefficients, congruent windows propagate backward as well as forward, $y_{j_1-1} = g(\ldots) \equiv g(\ldots) = y_{j_2-1} \pmod N$, and inductively for every earlier index. So the whole doubly infinite sequence is periodic mod N with period $j_2 - j_1$, with no eventually needed.

Both hypotheses hold here, with $n = 2006$: $f(y_n, \ldots, y_1) = y_n + y_1$ recovers x_{k+1} from the window $(x_{k-2005}, \ldots, x_k)$, and $g(y_n, \ldots, y_1) = y_n - y_{n-1}$ recovers x_{k-2005} from $(x_{k-2004}, \ldots, x_{k+1})$; both are integer coefficient polynomials. By Lemma 2, the sequence extends to a doubly infinite sequence, purely periodic modulo 2006. It therefore suffices to exhibit 2005 consecutive terms divisible by 2006 anywhere in this doubly infinite sequence: periodicity will relocate the same run into the region $k \geq 1$, which is the actual sequence given in the problem.

Run the recursion backward from the known block $x_1, x_2, \ldots, x_{2006} = 1, 2, \ldots, 2006$. For $k = 1, \ldots, 2005$, both k and $k+1$ lie in $[1, 2006]$, so $x_{k-2005} = x_{k+1} - x_k = (k+1) - k = 1$. As k ranges over $1, \ldots, 2005$, the index $k - 2005$ ranges over $-2004, \ldots, 0$; combined with $x_1 = 1$, this gives

$$x_{-2004} = x_{-2003} = \cdots = x_0 = x_1 = 1,$$

a block of 2006 consecutive terms all equal to 1. Now take $k = -2004, \ldots, 0$: both k and $k+1$ lie in $[-2004, 1]$, so $x_{k-2005} = x_{k+1} - x_k = 1 - 1 = 0$. As k ranges over $-2004, \ldots, 0$, the index $k - 2005$ ranges over $-4009, \ldots, -2005$, giving

$$x_{-4009} = x_{-4008} = \cdots = x_{-2005} = 0,$$

a run of 2005 consecutive terms, each equal to 0 and hence divisible by 2006.

Let $T \geq 1$ be the period of the doubly infinite sequence modulo 2006 from Lemma 2, and choose a positive integer c with $cT \geq 4010$ (possible since $T \geq 1$). For each integer j in $\{-4009, \ldots, -2005\}$, periodicity gives $x_{j+cT} \equiv x_j \equiv 0 \pmod{2006}$, and the shifted indices $j + cT$ range over 2005 consecutive integers, all at least $-4009 + 4010 = 1$. So the actual sequence $x_1, x_2, \ldots$ contains 2005 consecutive terms, namely $x_{-4009+cT}, \ldots, x_{-2005+cT}$, each divisible by 2006, as required. ■

Layer 1: WHAT

Block 1: rearrange the recursion to isolate the oldest term, $x_{k-2005} = x_{k+1} - x_k$, and check the coefficients are integers, so the recursion is invertible. Block 2: state and prove the general fact that any fixed order integer polynomial recursion is eventually periodic modulo any N, by pigeonhole on the finitely many windows. Block 3: strengthen this, using invertibility, to full two sided periodicity of a doubly infinite extension of the sequence. Block 4: run the (trivial, since the initial data is an arithmetic progression) backward recursion explicitly from the known block $x_1, \ldots, x_{2006}$ to exhibit 2005 consecutive zero terms. Block 5: use periodicity to relocate that run of zeros into the region $k \geq 1$, the actual given sequence, completing the proof.

The nominal route also notices the recursion has order 2006, but then either tries to extract a closed form from the characteristic polynomial, or starts computing $x_{2007}, x_{2008}, \ldots$ modulo 2006 forward, hoping to meet a run of 2005 zero residues by direct search. Both stall for the same reason: neither ever uses the fact that the recursion is reversible, so neither gains access to the arithmetic progression block $x_1, \ldots, x_{2006}$, which is exactly where the backward computation is free. The bifurcation point is the single rearrangement $x_{k-2005} = x_{k+1} - x_k$: written down and recognised as an integer coefficient identity, it opens the whole backward, periodic route; left unwritten, every subsequent route is forced forward, into either an intractable characteristic polynomial or an unbounded search.

Layer 2: HOW

Block 1. **Cue.** The defining relation $x_{k+1} = x_k + x_{k-2005}$ has every coefficient equal to 1, with no division needed to isolate any term. **Why natural.** Testing whether a recursion is reversible is a fixed check, independent of the specific problem, that decides which of two strictly different periodicity results (eventual only, or full two sided) is available; since a unit coefficient recursion can always be solved for its oldest term over the integers, by definition of what it means to divide by 1, the check costs one line and is always worth making immediately. **Link to the next block.** Having confirmed reversibility, the next natural step is to state precisely what reversibility buys you, which is the upgrade from eventual to full periodicity, a standard fact about fixed order integer recursions.

Block 2. **Cue.** The recursion has fixed order 2006 (a window of 2006 consecutive terms determines the next), and the target statement concerns divisibility by a fixed modulus 2006. **Why natural.** A window of n consecutive residues modulo N can take only N^n values, a finite number by definition of "modulo N", so by the pigeonhole principle two windows among the infinitely many produced by the recursion must coincide, and because the recursion is a fixed integer coefficient function of the window, everything after the first coincidence must repeat identically. **Link to the next block.** Eventual periodicity alone only guarantees repetition somewhere far enough forward, with no control over where; the reversibility already checked in Block 1 is what turns this into a statement usable near the boundary of the known data.

Block 3. **Cue.** The oldest term recursion $x_{k-2005} = x_{k+1} - x_k$ is, like the forward recursion, a fixed integer coefficient polynomial in the window. **Why natural.** Since congruence is preserved by integer coefficient polynomials in both directions, the same pigeonhole coincidence that propagates forward (Lemma 1) also propagates backward once the sequence is extended to negative indices, by exactly the same argument applied to g instead of f; this is not a new principle, only the same one applied twice. **Link to the next block.** Full periodicity means the target pattern can be sought anywhere in the doubly infinite sequence, so the natural next move is to search where the algebra is cheapest, not where the problem statement happens to start counting.

Block 4. **Cue.** The known initial block $x_1, \ldots, x_{2006} = 1, \ldots, 2006$ is an arithmetic progression with common difference 1, and the backward recursion is exactly a first difference. **Why natural.** Differencing an arithmetic progression with common difference 1 produces the constant sequence 1, by the definition of arithmetic progression; running the backward recursion once more on a block that is constantly 1 produces the constant sequence 0, by the same definition applied to a degenerate ("difference of equal terms") case. **Link to the next block.** A run of 2005 consecutive zero terms has now been exhibited explicitly, at indices -4009 through -2005; all that remains is to move this run into the positive indices the problem actually asks about.

Block 5. **Cue.** The zero run sits at negative indices, but the problem's sequence is only defined, as stated, for $k \geq 1$. **Why natural.** Full periodicity modulo 2006, established in Block 3, means every residue value recurs at every shift by a multiple of the period T; choosing a sufficiently large multiple moves the entire window of 2005 indices past 1 without altering any of the residues, by the definition of periodicity. **Link to the next block.** Nothing remains: the shifted window lies entirely within $k \geq 1$ and every term in it is divisible by 2006, which is the statement to be proved.

Layer 3: WHY

The mathematical truth underneath the route is that a fixed order recursion, reduced modulo a fixed integer, is a walk on a finite state space (the set of possible windows), and any walk on a finite state space that can be reversed as well as advanced must be purely periodic, not merely eventually so: this is nothing more than the pigeonhole principle applied twice, once forward and once backward, and the reversal is what removes the word "eventually". The discernment mechanism for reaching for this, rather than a closed form or a forward search, is recognising two separate signals at once: an order large enough that a closed form is not viable, and a coefficient pattern (here, all ± 1) that costs nothing to invert. Whenever both are present, the productive question stops being "what is x_k" and becomes "what does the periodic residue sequence look like, and where is it cheapest to compute a window of it".

You train this discernment by treating "can I solve this recursion for its oldest term without dividing by anything but ± 1" as a reflex check performed the instant a recursion of unfamiliar order appears, before any attempt at a closed form. You also train it by noticing when the GIVEN initial data has extra structure (here, a literal arithmetic progression) that trivialises differencing; such structure is usually placed exactly where the backward recursion should be run, not left to be discovered by accident.

The examiners chose the numbers 2005 and 2006, and the initial block $x_k = k$, with real care: making the initial segment the identity function guarantees the first backward difference is the constant 1, and iterating the same backward step once more on a constant sequence guarantees the second backward difference is the constant 0, producing the required run of zeros in exactly two applications of the backward recursion rather than an unpredictable number. The gap of 2005 (so that the recursion has order 2006, matching the modulus 2006) is what makes the closed form route hopeless while leaving the periodicity route completely untouched by the size of the order, since the pigeonhole argument never needs to enumerate the 2006^{2006} possible windows, only to know that this number is finite.

Revolving concepts

Eventual periodicity of integer polynomial recursions modulo N: if $y_j = f(y_{j-1}, \ldots, y_{j-n})$ for a fixed n and a fixed integer coefficient polynomial f, the residues $y_j \bmod N$ are eventually periodic, because there are only N^n possible windows of n consecutive residues, forcing a repeat by pigeonhole, after which the deterministic recursion forces everything to repeat. This generalisation applies verbatim to any fixed order integer recursion, including the Fibonacci sequence modulo N (whose

period is called the Pisano period) and any other "gap" recursion $x_{k+1} = x_k + x_{k-m}$ for fixed m. See Book 1, field: Sequences and recurrences, topic: periodicity modulo an integer.

Full, two sided periodicity for reversible recursions: if in addition $y_{j-n} = g(y_{j-n+1}, \ldots, y_j)$ for a fixed integer coefficient polynomial g, the sequence extends to a doubly infinite sequence that is purely periodic modulo every N, with no eventually needed, because the pigeonhole coincidence propagates backward as well as forward. This generalisation applies whenever a linear (or polynomial) integer recursion has leading coefficients equal to ± 1, so that solving for the oldest term never requires division; the Fibonacci recursion again qualifies, via $F_{j-1} = F_{j+1} - F_j$. See Book 1, field: Sequences and recurrences, topic: reversible recursions and bi-infinite extensions.

Sparse, wide gap linear recursions: a recursion such as $x_{k+1} = x_k + x_{k-m}$ for a large fixed gap m has, as a linear recursion, order $m + 1$ and a characteristic polynomial of that degree, making a closed form solution impractical for large m even though only two terms appear explicitly in each step; recognising this order before attempting a closed form saves the ink of trying one. See Book 1, field: Sequences and recurrences, topic: high order linear recursions.

Divisibility as vanishing residue: an integer term x_k is divisible by N exactly when $x_k \equiv 0 \pmod{N}$, so a run of consecutive terms divisible by N is exactly a run of consecutive zero residues in the sequence reduced mod N; this is the translation that turns the problem's divisibility requirement into a statement purely about the periodic residue sequence tracked throughout the solution. See Book 1, field: Number theory, topic: divisibility and congruence.

Circuit Bank

- **WHEN** a recursion $x_k = f(x_{k-1}, \ldots, x_{k-n})$ has a large order n and you need a mod N divisibility statement about it **DO** reach for eventual periodicity modulo N via pigeonhole on the finitely many windows, before attempting a closed form **BECAUSE** the state space of windows modulo N is always finite, size N^n, regardless of how large n is, while a genuine closed form for order n recursions with n large is rarely available
- **WHEN** a recursion can be algebraically solved for its oldest term using only integer coefficients (in particular, coefficients ± 1, needing no division) **DO** extend the sequence backward to a doubly infinite sequence **BECAUSE** this upgrades eventual periodicity to full periodicity modulo every N, which lets you search for a target pattern anywhere in the sequence, not only forward from an unknown starting point
- **WHEN** you need to exhibit a specific pattern inside a periodic sequence but do not know where it first occurs going forward **DO** run the recursion backward from the known, explicit initial data instead **BECAUSE** initial data is often unusually simple (constant, arithmetic, or otherwise structured), so differencing or otherwise inverting it collapses several recursive steps into one, and periodicity guarantees the pattern found there also occurs somewhere forward
- **WHEN** a linear recursion's gap is large compared to the simplicity of its initial data **DO** exploit the special structure of the initial data (constant differences, ratios, etc.) before treating the recursion generically **BECAUSE** special initial structure usually collapses several steps of the general recursion into a single closed computation, which is exactly what happened here with an arithmetic progression differencing to a constant, then to zero

Backtesting

1. First useful cue: the recursion links a term 2006 steps forward directly to one 2006 steps further back ($x_{k+1} = x_k + x_{k-2005}$), a wide fixed gap that immediately marks the recursion as high order and unsolvable in closed form.

2. Method the expression suggested: rearrange for the oldest term, $x_{k-2005} = x_{k+1} - x_k$, since every coefficient is ± 1, so the recursion is invertible with no division and can be run backward as cleanly as forward.
3. Where you would hesitate, and the thought to reinforce: whether it is legitimate to extend a sequence given only for $k \geq 1$ backward to negative indices; reinforce by treating the extension as a purely algebraic device used to access periodicity, since the final claim concerns only $k \geq 1$, which is recovered afterward via periodicity, so no illegitimate step is taken anywhere in the argument.
4. The trap the question sets: the number 2005, appearing both as the gap in the recursion and as the length of run required, invites hunting for period 2005 structure directly in the forward sequence, when in fact the exact period of the residue sequence is never computed or needed; only one run of length 2005 is required, and it is found painlessly at the far backward end of the extension.
5. What makes you recognise it faster next time: seeing a recursion of the shape $x_{k+1} = x_k + x_{k-m}$, a two term, wide gap recursion, and immediately checking invertibility (unit coefficients) as the very first move, before anything else is attempted.
6. The reusable decision rule: an order n integer recursion, reduced modulo N, that is also reversible extends to a doubly infinite, purely periodic sequence modulo N; search for any required pattern wherever the algebra is cheapest, most often against the known initial data, rather than forward from the first index the problem happens to start at.

Mirror drills

1. The Fibonacci sequence $F_1 = F_2 = 1$, $F_{k+1} = F_k + F_{k-1}$, is reduced modulo an integer N. Show the residue sequence is purely periodic (the Pisano period), not merely eventually periodic. *Hint:* the same two lemmas apply directly, since $F_{k-1} = F_{k+1} - F_k$ is the backward recursion with unit coefficients.
2. A sequence satisfies $y_k = y_{k-1} + y_{k-10}$ for $k > 10$, with $y_1 = \cdots = y_{10} = 1$. Show it has 9 consecutive terms divisible by 10. *Hint:* identical method with gap 10 in place of 2005; the initial block is constant, so the first backward difference is already 0, one step shorter than the original problem.
3. Explain, without computing anything, why attempting a closed form for x_k in the original problem via the characteristic polynomial of $t^{2006} - t^{2005} - 1 = 0$ is not a productive route for this question. *Answer:* the question asks only for a bounded divisibility statement, which the periodicity argument extracts in a handful of algebraic steps, while a degree 2006 characteristic polynomial has no accessible closed form roots and gives no direct route to an integer divisibility conclusion even if its roots were known.

20.4 A4. Average number of peaks of a permutation

Let $S = \{1, 2, \ldots, n\}$ for some integer $n > 1$. Say a permutation π of S has a local maximum *at $k \in S$ if*

(i) $\pi(k) > \pi(k+1)$ *for* $k = 1$;

(ii) $\pi(k-1) < \pi(k)$ *and* $\pi(k) > \pi(k+1)$ *for* $1 < k < n$;

(iii) $\pi(k-1) < \pi(k)$ *for* $k = n$.

(For example, if $n = 5$ and π takes values at $1, 2, 3, 4, 5$ of $2, 1, 4, 5, 3$, then π has a local maximum of 2 at $k = 1$, and a local maximum of 5 at $k = 4$.) What is the average number of local maxima of a permutation of S, averaging over all permutations of S?

Read it out loud

This is an averaging question over a finite symmetric family. You are given $n!$ objects, each object carries a count, and you are asked for the mean of that count. Say that sentence to yourself before you write anything, because it already tells you the method: a mean of a count is never a question about the count's distribution. It is a question about one number, and one number is cheaper to get than a whole distribution.

Now read the thing being counted. A local maximum at k is decided by looking at $\pi(k)$ and its immediate neighbours and nothing else. Positions 2 through $n-1$ look at a window of three consecutive values; positions 1 and n look at a window of two, because the window is truncated by the end of the row, not because a special rule is in force. That is the second structural fact, and it is the one that makes the problem collapse: the count is a sum of n local conditions, each involving at most three entries.

The fields wired into this problem are the combinatorics of permutation statistics (peaks, descents, records) and the elementary theory of expectation, specifically linearity and the notion of an indicator random variable. Nothing else. What is being tested is whether you know that expectation is linear whether or not the summands are independent, and whether you know that a uniformly random permutation looks uniformly random when you restrict attention to any fixed set of positions. The peak indicators here are violently dependent (positions k and $k+1$ can never both be peaks), and that dependence is exactly what makes the distribution of the count hard and the mean of the count easy. The examiner chose the statistic so that the dependence is obvious and irrelevant.

Within thirty seconds you should have written two lines. First, $X(\pi) = \sum_{k=1}^{n} \mathbf{1}[\pi \text{ has a peak at } k]$. Second, $\mathrm{E}[X] = \sum_{k=1}^{n} \mathrm{P}(\text{peak at } k)$. Everything after that is arithmetic on three numbers.

Approaches, ranked worst to best

1. **Compute the whole distribution, then take its mean.** *Cue:* the word "average" misread as "I must know how the count is spread out". *Circuit:* none; this is the zombie route. You would be counting, for each j, the number of permutations of S with exactly j peaks. Those numbers are the coefficients of the peak polynomial, a genuinely intricate object with no one line closed form. You would then have to sum j times that count and simplify. This costs an enormous amount of ink to produce a fact you were never asked for. Rank it last not because it fails but because it answers a harder question than the one on the paper.

2. **Brute force small cases and guess.** *Cue:* a formula in n is wanted and the cases are computable by hand. *Circuit:* pattern extraction. For $n = 2, 3, 4, 5$ you get $1, \frac{4}{3}, \frac{5}{3}, 2$, so the answer is visibly $\frac{n+1}{3}$. Worth thirty seconds as insurance, and worth nothing as a proof. On a "what is the average" problem the Putnam grader wants the derivation, so this is a check, not a solution. Keep it in your pocket and never submit it alone.

3. **Split at the maximum and recurse.** *Cue:* the value n is always a peak, and deleting it cuts the row into two independent shorter rows. *Circuit:* divide and conquer on the position of the largest element. Condition on the position j of the value n. The row splits into a left block of length $j-1$ and a right block of length $n-j$, the relative orders inside the blocks are independent and uniform, and you weight by $\binom{n-1}{j-1}$. This does work and it is instructive, but the boundary conventions ruin the elegance: the left block now ends against a value larger than everything in it, so its right hand boundary rule is not the rule in the problem, and you must set up a second statistic for one sided blocks. Two coupled recursions to compute one number.

4. **Linearity of expectation over the values.** *Cue:* ask which numbers $1, \dots, n$ appear as peaks rather than which positions carry peaks. *Circuit:* indicator sum indexed by value, plus a symmetrising device. Write the permutation as the symbols $1, \dots, n$ and one extra symbol $*$ arranged around a circle; the row is recovered by cutting at $*$, and the two truncated windows at the ends become full windows containing $*$, which acts as a value smaller than everything. Then i is a peak precisely when both circular neighbours of i lie in $\{*, 1, \dots, i-1\}$, which happens with probability $\frac{i(i-1)}{n(n-1)}$, and $\sum_{i=1}^{n} \frac{i(i-1)}{n(n-1)} = \frac{2}{n(n-1)}\binom{n+1}{3} = \frac{n+1}{3}$. This is strictly more informative than the winner, since it tells you which values are peaks and not merely how many. It ranks second only because the circle device has to be justified and the winner needs no device at all.

5. **Linearity of expectation over the positions.** *Cue:* the statistic is a sum of conditions each involving at most three adjacent entries. *Circuit:* indicator sum plus local pattern uniformity. Three probabilities, one sum, three lines. This is the ideal route: no device, no case split beyond the two ends, no recursion, and every step is forced by the shape of the definition.

The ideal solution

Let π be chosen uniformly at random from the $n!$ permutations of S, and for $k \in S$ let I_k be the indicator of the event that π has a local maximum at k. The number of local maxima is $X = I_1 + I_2 + \cdots + I_n$, and the average asked for is $\mathrm{E}[X]$. Expectation is linear on any finite collection of random variables, with no independence required, so

$$\mathrm{E}[X] = \sum_{k=1}^{n} \mathrm{E}[I_k] = \sum_{k=1}^{n} \mathrm{P}(\pi \text{ has a local maximum at } k).$$

Everything now rests on one fact about uniform permutations.

Lemma 20.4.1. *Fix positions $i_1 < i_2 < \cdots < i_m$ in S. If π is uniform on the permutations of S, then the relative order of $\pi(i_1), \dots, \pi(i_m)$ is uniformly distributed over the $m!$ possible orderings.*

Proof. For τ a permutation of $\{1, \dots, m\}$, define Φ_τ on permutations of S by $(\Phi_\tau \pi)(i_j) = \pi(i_{\tau(j)})$ for $j = 1, \dots, m$ and $(\Phi_\tau \pi)(x) = \pi(x)$ for every other x. Then Φ_τ permutes the values sitting in the chosen positions and fixes all other values, so it maps permutations of S to permutations of S, and $\Phi_{\tau^{-1}}$ is its inverse. Hence Φ_τ is a bijection. It sends the set of π whose chosen entries have relative order p bijectively onto the set of π whose chosen entries have relative order $\tau \cdot p$. The $m!$ possible relative orders therefore label $m!$ equinumerous classes which partition the $n!$ permutations, so each has probability $1/m!$. □

Apply the lemma three times.

For $k = 1$ the condition is $\pi(1) > \pi(2)$, a statement about the relative order of the entries in the two positions $1, 2$. By the lemma with $m = 2$ the two orders are equally likely, so $\mathrm{P}(I_1 = 1) = \frac{1}{2}$. For $k = n$ the condition is $\pi(n-1) < \pi(n)$, again a statement about two positions, so $\mathrm{P}(I_n = 1) = \frac{1}{2}$. These two events are distinct because $n > 1$.

For $1 < k < n$ the condition is $\pi(k-1) < \pi(k)$ and $\pi(k) > \pi(k+1)$, that is, $\pi(k)$ is the largest of the three entries in positions $k-1, k, k+1$. By the lemma with $m = 3$ the six relative orders are equally likely, and exactly two of them place the middle entry largest (the remaining two entries may be in either order). Hence $\mathrm{P}(I_k = 1) = \frac{2}{6} = \frac{1}{3}$.

There are two boundary positions and $n-2$ interior positions, and $n-2 \geq 0$ since $n > 1$, so

$$\mathrm{E}[X] = 2 \cdot \frac{1}{2} + (n-2) \cdot \frac{1}{3} = 1 + \frac{n-2}{3} = \boxed{\frac{n+1}{3}}.$$

The degenerate case is $n = 2$, where there are no interior positions and the formula reads $\mathrm{E}[X] = 1$; both permutations of $\{1, 2\}$ have exactly one local maximum, so this is correct. As a further check, $n = 3$ gives $\frac{4}{3}$, and the six permutations of $\{1, 2, 3\}$ have $1, 1, 2, 1, 2, 1$ local maxima, total 8, mean $\frac{8}{6} = \frac{4}{3}$.

Layer 1: WHAT

Block 1: read "average number of" as "expected value of a count" and write $X = \sum_k I_k$.

Block 2: observe that I_k depends only on the relative order of at most three entries in fixed positions.

Block 3: state and prove the local pattern uniformity lemma, which converts each $\mathrm{P}(I_k = 1)$ into a count of orderings.

Block 4: evaluate the three probabilities: $\frac{1}{2}$ at each end, $\frac{1}{3}$ in the interior.

Block 5: add up, confirm the boundary case $n = 2$, box $\frac{n+1}{3}$.

The nominal route runs like this: read "average", decide you need the distribution of the number of peaks, start classifying permutations by peak count, notice that adjacent peaks are forbidden so the classification involves compositions of n, and disappear into peak polynomials. The bifurcation point is precisely the sentence after "I want the average". The zombie writes "so I need $\mathrm{P}(X = j)$ for each j". You write "so I need $\mathrm{E}[X]$, and X is a sum, and expectation of a sum is the sum of expectations". That single sentence is the whole problem. Everything downstream is forced.

Layer 2: HOW

Block 1. Cue. The question says "average, over all permutations", and the quantity being averaged is a number of things. A number of things is always a sum of indicators, one per thing that might occur. **Why natural.** The definition of the statistic is already a list of n separate conditions, one per position k. The question has handed you the decomposition; you only have to accept it. Refusing to accept it and instead treating X as a single opaque number is what forces the distribution route. **Link to the next block.** Once X is a sum, linearity of expectation applies, and the problem becomes n separate one position questions.

Block 2. Cue. Each condition in the definition mentions only $\pi(k-1), \pi(k), \pi(k+1)$, and the comparisons are all inequalities between values, never the values themselves. **Why natural.** Because only inequalities appear, the event depends on the permutation solely through the relative order of those two or three entries. The actual sizes are irrelevant, which is the standard signal that a symmetry argument will finish the job. **Link to the next block.** If the relative order is all that matters, you now need to know the distribution of that relative order, which is the next block.

Block 3. Cue. A uniform permutation has a symmetry group acting on it: you may permute the values in any fixed set of positions and stay inside the set of permutations. **Why natural.** This is the definition level reason. A permutation of S is a bijection $S \to S$; rearranging the values assigned to a fixed set of positions produces another bijection, so the rearrangement is a bijection on the sample

space and therefore preserves the uniform measure. There is nothing probabilistic happening here at all: it is a counting statement dressed up. **Link to the next block.** With the lemma, each probability becomes "how many of the $m!$ orderings satisfy the condition", which is immediate.

Block 4. Cue. Two of the conditions involve two positions and $n-2$ involve three. **Why natural.** Counting orderings: one of two puts the first entry higher, so $\frac{1}{2}$; two of six put the middle entry highest, so $\frac{1}{3}$. Notice that the ends are not a special case of the mathematics, only a truncation of the window: the same lemma with a smaller m. **Link to the next block.** The three probabilities are constants, so the sum over k is a count of positions of each type.

Block 5. Cue. The sum has two terms of one kind and $n-2$ of another. **Why natural.** The arithmetic is trivial, but the write up must state that $n > 1$ guarantees positions 1 and n are distinct and that $n-2 \geq 0$, otherwise the count of interior positions is not justified. A grader looks for exactly this sentence. **Link to the next block.** Nothing follows; you have the answer. Spend the remaining ten seconds checking $n = 2$ and $n = 3$ by hand.

Layer 3: WHY

The mathematical truth underneath is that expectation is a linear functional on the vector space of real valued functions on a finite sample space, and linearity has no hypothesis. Independence is a hypothesis you need to multiply expectations, never to add them. Most students carry an unexamined belief that "you can only split $\mathrm{E}[X+Y]$ when X and Y are independent", absorbed from variance calculations where independence really is needed. Kill that belief now. Here the indicators are as dependent as they could be, since $I_k I_{k+1} = 0$ always, and the computation does not notice.

The second truth is that the average of a count is a double count. Write the average as

$$\frac{1}{n!}\sum_{\pi} X(\pi) = \frac{1}{n!}\sum_{\pi}\sum_{k} I_k(\pi) = \frac{1}{n!}\sum_{k}\sum_{\pi} I_k(\pi),$$

and the exchange of the two finite sums is the entire method. Expectation language is a convenience; the content is Fubini for finite sums, that is, counting the pairs (π, k) such that π has a peak at k in two orders. Whenever a problem asks for a total or an average over a symmetric family, ask which of the two summation orders is cheap. Here summing over π first is cheap because the inner condition is local; summing over k first is expensive because it asks you to understand a whole permutation at once.

The third truth is exchangeability. A uniform random permutation restricted to any fixed window is a uniform random pattern. That is why the answer for interior positions is $\frac{1}{3}$ and not something depending on k: no position is distinguished except by how much of its window has been cut off by the ends of the row. Recognising that the only thing distinguishing $k = 1$ from $k = 7$ is a truncated window, and not a different rule, is the discernment being trained here.

How do you train the discernment? Build the habit of asking, of every counting or averaging problem, "is the quantity a sum of local pieces?" If yes, linearity finishes it, whatever the dependence. If no, ask whether it can be made into one. The examiner's nudge in this problem was to define the statistic as a list of three cases rather than as a single global property; that presentation is a guiding light pointing at the indicator decomposition. When a definition arrives as a list of local conditions, the author is telling you to sum indicators.

Revolving concepts

Linearity of expectation. For random variables $X_1, \ldots, X_m$ on a common probability space with finite expectations, $\mathrm{E}[\sum X_i] = \sum \mathrm{E}[X_i]$, and $\mathrm{E}[cX] = c\,\mathrm{E}[X]$. No independence, no identical dis-

tribution, no positivity. The orbiting results are Markov's inequality (from E alone you get tail bounds), the first moment method (if $\mathrm{E}[X] < 1$ and X is a nonnegative integer, then $\mathrm{P}(X = 0) > 0$, which proves existence results), and the fact that variance is not linear without independence, since $\mathrm{Var}(X + Y) = \mathrm{Var}\, X + \mathrm{Var}\, Y + 2\,\mathrm{Cov}(X, Y)$. See Book 1, field: Probability, topic: linearity of expectation and indicators.

Indicator random variables. $\mathbf{1}[A]$ has $\mathrm{E}[\mathbf{1}[A]] = \mathrm{P}(A)$. Every counting statistic is a sum of indicators; the art is choosing what to index by. Here indexing by position beat indexing by value only on ink, not on power.

Permutation statistics. A *descent* of π is a $k < n$ with $\pi(k) > \pi(k+1)$; a *peak* is what this problem calls a local maximum; a *record* or left to right maximum is a k with $\pi(k) > \pi(j)$ for all $j < k$. The same three line method gives: expected descents $= (n-1)/2$, expected records $= H_n = \sum_{i=1}^{n} 1/i$, expected fixed points $= 1$, expected cycles $= H_n$. Note how the record calculation uses the lemma with $m = k$ rather than $m \le 3$: $\mathrm{P}(\text{record at } k) = 1/k$ because the largest of the first k entries is equally likely to sit in any of those k positions. See Book 1, field: Combinatorics, topic: permutation statistics.

The distributions behind the means, and the recursion. The number of permutations of $[n]$ with k descents is the Eulerian number $A(n, k)$, and it satisfies the recursion

$$A(n,k) = (k+1)A(n-1,k) + (n-k)A(n-1,k-1), \qquad A(1,0) = 1,$$

obtained by inserting the value n into a permutation of $[n-1]$ and asking whether the insertion site creates a new descent. The generalised form of this move: if a statistic st on permutations changes in a controlled way when the largest value is inserted or deleted, then the generating function $P_n(t) = \sum_{\pi} t^{\mathrm{st}(\pi)}$ satisfies a recursion in n with polynomial coefficients in t and ∂_t; for descents this is $P_n(t) = ntP_{n-1}(t) + t(1-t)P'_{n-1}(t)$. The peak analogue exists but needs a two variable refinement to handle the truncated windows at the ends, which is exactly the obstruction that sank approach 3 above. The lesson generalises: a recursion built by inserting the largest element is the default tool for the full distribution of a permutation statistic, and it is the wrong tool when only the mean is wanted. See Book 1, field: Combinatorics, topic: Eulerian numbers and insertion recursions.

Splitting at the maximum. The other standard recursion for permutation problems: condition on the position j of the value n, giving $f_n = \sum_{j=1}^{n} \binom{n-1}{j-1}(\text{something}(j-1) \cdot \text{something}(n-j))$, a binomial convolution whose natural transform is the exponential generating function $\sum f_n x^n/n!$. Binomial convolution corresponds to multiplication of exponential generating functions; that correspondence is the single most useful fact about labelled counting.

Circuit Bank

- **WHEN** a problem asks for the average, expected, or total value of a count over a finite symmetric family **DO** write the count as a sum of indicators, one per candidate object, and take expectations term by term **BECAUSE** expectation is linear with no independence hypothesis, so the dependence structure of the indicators is irrelevant to the mean
- **WHEN** an event about a uniform random permutation is stated purely as inequalities among entries in fixed positions **DO** replace it by a count of relative orderings: with m positions involved, each of the $m!$ patterns has probability $1/m!$ **BECAUSE** permuting the values in a fixed set of positions is a bijection of the set of permutations, so it preserves the uniform measure
- **WHEN** a locally defined statistic has different rules at the first and last positions **DO** treat the ends as the same rule with a truncated window, that is, the same lemma with a smaller m **BECAUSE** the boundary is not a new phenomenon, it is a window running off the edge of the row, so the mathematics is identical with fewer entries

- **WHEN** you catch yourself computing a distribution in order to extract its mean **DO** stop and ask whether the quantity is a sum of local pieces **BECAUSE** the distribution is almost always a strictly harder object than the mean, and the extra work earns nothing on the paper
- **WHEN** you have an answer that is a formula in n **DO** test it on the smallest legal n and on one more case by hand **BECAUSE** formulas of this shape fail exactly at the degenerate boundary, here $n = 2$, and the check costs twenty seconds

Backtesting

1. First useful cue: the phrase "average number of" applied to a count over a finite family. That phrase alone forces $X = \sum_k I_k$ and $\mathrm{E}[X] = \sum_k \mathrm{P}(I_k = 1)$.
2. Method the expression suggested: the definition arrived as three local conditions indexed by position, which is an explicit invitation to index the indicators by position rather than by value.
3. Where you would hesitate, and the thought to reinforce: the hesitation is "the peak events are dependent, so surely I cannot just add". Reinforce: independence is required to multiply expectations, never to add them. Say it out loud once per week until it is reflex.
4. The trap the question sets: the three case definition looks like it demands three separate arguments, and the example with $n = 5$ tempts you into enumeration. The three cases are one argument with $m = 2, 3, 2$.
5. What makes you recognise it faster next time: the shape "statistic defined by a sliding window over a random arrangement". Peaks, descents, ascents, records, runs, matching pairs, adjacent equal cards, all of them fall to the same two lines.
6. The reusable decision rule: count of local events $\Rightarrow$ sum of indicators $\Rightarrow$ local pattern uniformity $\Rightarrow$ arithmetic. Never touch the distribution.

Mirror drills

1. For a uniform random permutation of $\{1, \ldots, n\}$, find the expected number of descents, that is, of indices k with $1 \le k < n$ and $\pi(k) > \pi(k+1)$. *Answer:* each of the $n-1$ indices contributes $\frac{1}{2}$ by the lemma with $m = 2$, so the mean is $\frac{n-1}{2}$.
2. For a uniform random permutation of $\{1, \ldots, n\}$, find the expected number of records (indices k with $\pi(k)$ larger than every earlier entry). *Answer:* the largest of the first k entries is equally likely to be in any of those k positions, so $\mathrm{P}(\text{record at } k) = 1/k$ and the mean is $H_n = 1 + \frac{1}{2} + \cdots + \frac{1}{n}$.
3. Arrange n distinct real numbers uniformly at random in the n positions of a circle. Find the expected number of positions that hold a local maximum, meaning a value larger than both of its circular neighbours. *Answer:* every position now has a full window of three and no position is truncated, so each contributes $\frac{1}{3}$ and the mean is $n/3$. Compare with $\frac{n+1}{3}$ for the row and see that cutting the circle open is worth exactly one third of a peak.

20.5 A5. Ratio of symmetric functions of equally spaced tangents

Let n be a positive odd integer and let θ be a real number such that θ/π is irrational. Set $a_k = \tan(\theta + k\pi/n)$, $k = 1, 2, \ldots, n$. Prove that

$$\frac{a_1 + a_2 + \cdots + a_n}{a_1 a_2 \cdots a_n}$$

is an integer, and determine its value.

Read it out loud

A sum divided by a product, over a family of n numbers. Say that aloud and the translation is automatic: this is e_1/e_n, the first elementary symmetric function of $a_1, \ldots, a_n$ divided by the last. Nobody asks for e_1/e_n unless they intend you to use Vieta, and Vieta needs one thing: the polynomial whose roots are exactly $a_1, \ldots, a_n$. So the problem is not "evaluate these tangents". The problem is "find the degree n polynomial that these n tangents satisfy". That reroute is the entire question, and it should be on your page inside thirty seconds.

Next, read the indexing. The angles $\theta + k\pi/n$ for $k = 1, \ldots, n$ form an arithmetic progression with common difference π/n, and tan has period π. Multiply an angle by n and the progression collapses: $n(\theta + k\pi/n) = n\theta + k\pi$, which is $n\theta$ modulo π for every single k. In other words all n of these angles have the same n-fold multiple modulo π, so all n numbers a_k satisfy the same equation $\tan(n \arctan x) = \tan(n\theta)$. That is the polynomial you are looking for, and it exists before you have written a word of trigonometry. Common difference π/n is the guiding light: it is planted so that the n-th multiple collapses, which is the same phenomenon as n-th roots of unity.

The fields wired in are complex numbers via de Moivre, polynomials and Vieta, roots of unity, and one Mobius map. The hypotheses are doing precise work and you must say what: n odd controls the parity signs and guarantees the polynomial really has degree n; θ/π irrational guarantees that no a_k is infinite (no angle hits $\pi/2$ modulo π), that no a_k is zero (so the denominator of the stated expression is nonzero), and that a certain root of unity factor is not degenerate. A hypothesis you cannot account for is a hypothesis you have not understood.

Finally, note the shape of the demand: "prove it is an integer, and determine its value". Being an integer is not a separate task; you will produce the value and it will visibly be $\pm n$. The mild surprise is that the answer does not depend on θ at all, and a full marks write up should explain why that is structural rather than accidental.

Approaches, ranked worst to best

1. **Expand each tangent with the addition formula and try to sum.** *Cue:* $\tan(\theta + \phi)$ is a formula you know. *Circuit:* none; this is grinding. You write $a_k = \frac{\tan\theta + \tan(k\pi/n)}{1 - \tan\theta\tan(k\pi/n)}$ and then stare at a sum of n fractions with different denominators. There is no route from here that does not amount to reconstructing the polynomial anyway, and you will have spent a page. Last place.

2. **Compute small cases and guess.** *Cue:* the answer is a single value, so it is guessable. *Circuit:* pattern extraction. $n = 1$ gives 1, $n = 3$ gives -3, $n = 5$ gives 5, $n = 7$ gives -7. So the answer is $(-1)^{(n-1)/2} n$. Do this, in your head or on scrap, because it tells you the target and it catches sign errors later. It is not a proof and there is no induction step available: passing from n to $n + 2$ changes the entire grid of angles, so nothing is inherited. Insurance only.

3. **Real polynomial route through the tangent multiple angle formula.** *Cue:* $\tan(n\phi)$ is a rational function of $\tan\phi$ with binomial coefficients. *Circuit:* de Moivre, then Vieta. From $(\cos\phi + i\sin\phi)^n$ you get

$$\tan(n\phi) = \frac{\binom{n}{1}x - \binom{n}{3}x^3 + \cdots}{1 - \binom{n}{2}x^2 + \cdots}, \qquad x = \tan\phi,$$

so the a_k are the roots of $\mathrm{Num}(x) - \tan(n\theta)\,\mathrm{Den}(x) = 0$, a polynomial of degree n for n odd. Vieta then gives the answer. This is completely correct and costs perhaps twice the ink of the winner, because you carry two binomial sums and must track the alternating signs in both. Rank it third: right idea, clumsy packaging.

4. **Collapse the products directly.** *Cue:* a product over an arithmetic progression of angles is a classical root of unity product. *Circuit:* the factorisation $\prod_{k=0}^{n-1}(1 - w\zeta^k) = 1 - w^n$. The identities $\prod_{k=0}^{n-1} 2\sin(\phi + k\pi/n) = 2\sin(n\phi)$ and, for n odd, $\prod_{k=0}^{n-1} 2\cos(\phi + k\pi/n) = (-1)^{(n-1)/2} 2\cos(n\phi)$ give the denominator immediately as $(-1)^{(n-1)/2}\tan(n\theta)$, and taking the logarithmic derivative of the cosine identity gives the numerator as $n\tan(n\theta)$. Very short if you own the two identities. It ranks second only because it needs two separate facts plus a differentiation, where the winner needs one object.

5. **Cayley transform: one polynomial, two coefficients.** *Cue:* tangent plus equally spaced angles plus symmetric functions. *Circuit:* $x \mapsto (1+ix)/(1-ix)$ turns angle addition into multiplication, so the a_k become the n-th roots of a fixed number. Build $Q_n(x) = (1+ix)^n - \omega(1-ix)^n$ with $\omega = e^{2in\theta}$, read off two coefficients, and watch the factor $1-\omega$ cancel. Shortest ink, no case analysis, and it delivers every elementary symmetric function of the a_k as a by product rather than just the two you were asked for. This is the ideal route.

The ideal solution

Write $\phi_k = \theta + k\pi/n$ for $k = 1, \ldots, n$, and put $\omega = e^{2in\theta}$.

Step 0: the expression is well defined. If some ϕ_k were an odd multiple of $\pi/2$ then $\theta/\pi = \frac{1}{2} + m - k/n$ for an integer m, which is rational; so every a_k is a finite real number. If some a_k were 0 then $\phi_k \in \pi\mathbb{Z}$ and again θ/π would be rational; so $a_1 a_2 \cdots a_n \neq 0$ and the quotient makes sense. Likewise $\omega \neq 1$, since $\omega = 1$ forces $n\theta \in \pi\mathbb{Z}$, and $\omega \neq -1$, since $\omega = -1$ forces $\theta/\pi = (2m+1)/(2n)$. Both would make θ/π rational. Finally the ϕ_k are pairwise incongruent modulo π, because $\phi_k - \phi_j = (k-j)\pi/n$ is a multiple of π only when $n \mid k-j$, impossible for distinct $j, k \in \{1, \ldots, n\}$; since tan is injective on any interval of length π where it is defined, the numbers $a_1, \ldots, a_n$ are pairwise distinct.

Step 1: turn tangent into a point on the unit circle. For any ϕ with $\cos\phi \neq 0$,

$$1 + i\tan\phi = \frac{\cos\phi + i\sin\phi}{\cos\phi} = \frac{e^{i\phi}}{\cos\phi}, \qquad 1 - i\tan\phi = \frac{e^{-i\phi}}{\cos\phi},$$

which is nothing more than dividing the polar form of $e^{i\phi}$ by its real part. Dividing the two,

$$\frac{1 + i\tan\phi}{1 - i\tan\phi} = e^{2i\phi}.$$

Step 2: the a_k are roots of one explicit polynomial. Apply Step 1 at $\phi = \phi_k$ and raise to the n-th power:

$$\left(\frac{1+ia_k}{1-ia_k}\right)^n = e^{2in\phi_k} = e^{2i(n\theta + k\pi)} = e^{2in\theta} = \omega.$$

Clearing the denominator (legitimate since $1 - ia_k \neq 0$ for real a_k),

$$Q_n(a_k) = 0, \qquad \text{where } Q_n(x) := (1+ix)^n - \omega(1-ix)^n.$$

The coefficient of x^n in Q_n is $i^n - \omega(-i)^n = i^n(1+\omega)$, using $(-i)^n = -i^n$ for n odd; this is nonzero because $\omega \neq -1$. So Q_n has degree exactly n, and since $a_1, \ldots, a_n$ are n distinct roots of it, they are precisely its roots.

Step 3: Vieta. Write $Q_n(x) = c_n x^n + \cdots + c_1 x + c_0$. Then

$$a_1 + \cdots + a_n = -\frac{c_{n-1}}{c_n}, \qquad a_1 a_2 \cdots a_n = (-1)^n \frac{c_0}{c_n} = -\frac{c_0}{c_n},$$

the last equality because n is odd. Dividing, the factor c_n cancels and

$$\frac{a_1 + \cdots + a_n}{a_1 a_2 \cdots a_n} = \frac{c_{n-1}}{c_0}.$$

Now read off the two coefficients from the binomial expansions. Since $n-1$ is even, $(-i)^{n-1} = i^{n-1}$, so

$$c_{n-1} = \binom{n}{n-1} i^{n-1} - \omega \binom{n}{n-1} (-i)^{n-1} = n\, i^{n-1}(1-\omega), \qquad c_0 = 1 - \omega.$$

Step 4: conclude. Since $\omega \neq 1$ the factor $1-\omega$ is nonzero and cancels:

$$\frac{a_1 + \cdots + a_n}{a_1 a_2 \cdots a_n} = n\, i^{n-1} = n\left(i^2\right)^{(n-1)/2} = \boxed{(-1)^{(n-1)/2}\, n}$$

which is n when $n \equiv 1 \pmod 4$ and $-n$ when $n \equiv 3 \pmod 4$. In particular it is an integer, and it does not depend on θ.

By-products worth banking. Since $c_n = i^n(1+\omega)$, the same two lines give the individual values

$$a_1 + \cdots + a_n = n \tan(n\theta), \qquad a_1 \cdots a_n = (-1)^{(n-1)/2} \tan(n\theta),$$

using $\frac{1-\omega}{1+\omega} = \frac{e^{-in\theta} - e^{in\theta}}{e^{-in\theta} + e^{in\theta}} = -i\tan(n\theta)$. More generally $c_j = \binom{n}{j} i^j \left(1 - \omega(-1)^j\right)$, so the elementary symmetric functions of the a_k are

$$e_{2r} = (-1)^r \binom{n}{2r}, \qquad e_{2r+1} = (-1)^r \binom{n}{2r+1} \tan(n\theta).$$

Every even one is a fixed integer independent of θ; every odd one is an integer multiple of $\tan(n\theta)$. That is the real content of the problem, and the quotient asked for is the cheapest visible consequence.

Layer 1: WHAT

Block 1: read "sum over product" as e_1/e_n, hence resolve to find the polynomial with roots $a_1, \ldots, a_n$.

Block 2: notice that the common difference π/n makes $n\phi_k \equiv n\theta \pmod \pi$ for every k, so all n numbers satisfy one common equation.

Block 3: convert tan into a Mobius function of $e^{2i\phi}$ by dividing the polar form of $e^{i\phi}$ by $\cos\phi$.

Block 4: raise to the n-th power, clear denominators, obtain $Q_n(x) = (1+ix)^n - \omega(1-ix)^n$; check degree n and distinctness of roots to know these are all of them.

Block 5: Vieta, note c_n cancels in the quotient, extract c_{n-1} and c_0.

Block 6: see the common factor $1-\omega$ cancel, hence the θ independence; finish with $i^{n-1} = (-1)^{(n-1)/2}$.

The nominal route: expand $\tan(\theta + k\pi/n)$ with the addition formula and attempt to sum and multiply n fractions. The bifurcation point is the very first decision after reading the expression. The zombie asks "what is a_k?". You ask "what equation do all the a_k satisfy at once?". Symmetric functions never want the members of a family, only the family's defining equation, and the question told you it wants symmetric functions by presenting exactly e_1 over e_n.

Layer 2: HOW

Block 1. Cue. A quotient whose numerator is the sum of all the objects and whose denominator is the product of all the objects. **Why natural.** e_1 and e_n are, by definition, coefficients of the monic polynomial with those roots, up to sign. So the quotient is a ratio of two coefficients of one object you have not yet built, and building that object is strictly easier than evaluating n tangents. **Link to the next block.** Building the polynomial requires an equation satisfied by every a_k, so look at what the index k does.

Block 2. **Cue.** The angles form an arithmetic progression with common difference exactly π/n, and the function has period π. **Why natural.** Multiplying by n turns the increment $k\pi/n$ into $k\pi$, which is invisible to tan. This is the same collapse that makes n-th roots of unity useful: an equally spaced family becomes a single value after taking the n-th power. The designer put π/n there for exactly this reason, and n odd is there so that the collapse works with the right parity. **Link to the next block.** Now you know the common equation in the form $\tan(n \arctan x) = \tan(n\theta)$; you need to write it as a polynomial without dragging in the multiple angle expansion by hand.

Block 3. **Cue.** You want $e^{2i\phi}$, and $\tan\phi$ is a ratio of the imaginary and real parts of $e^{i\phi}$. **Why natural.** Divide $e^{i\phi} = \cos\phi + i\sin\phi$ through by $\cos\phi$: that single division produces $1 + i\tan\phi = e^{i\phi}/\cos\phi$. There is nothing to invent. The conjugate gives $1 - i\tan\phi$, the unwanted $\sec\phi$ cancels in the ratio, and you land on the Cayley transform $x \mapsto (1+ix)/(1-ix)$. **Link to the next block.** The Cayley transform sends the real line to the unit circle and sends a_k to $e^{2i\phi_k}$, so the collapse of Block 2 becomes the statement that these images are the n-th roots of ω.

Block 4. **Cue.** The images satisfy $z^n = \omega$; you want an equation in x, not z. **Why natural.** Substitute $z = (1+ix)/(1-ix)$ into $z^n - \omega = 0$ and multiply by $(1-ix)^n$. That is the standard way to turn a Mobius substitution into a polynomial identity, and it is safe because $1 - ix$ never vanishes for real x. The degree check and the distinctness check are not decoration: without them you know only that Q_n vanishes at the a_k, not that it has no other roots, and Vieta needs all the roots. **Link to the next block.** With all roots known, the coefficients are the symmetric functions.

Block 5. **Cue.** You need only two coefficients, and both sit at the ends of the binomial expansion. **Why natural.** c_{n-1} and c_0 are the cheapest coefficients in any binomial expansion, and the quotient you want kills c_n automatically since n is odd. That is why the write up never needs to evaluate ω or $\tan(n\theta)$ at all. **Link to the next block.** Extract them and compare.

Block 6. **Cue.** Both coefficients carry the factor $1 - \omega$. **Why natural.** This is the structural reason the answer is independent of θ: all of the θ dependence entered through the single scalar ω, and ω enters c_{n-1} and c_0 through the same factor. When a quotient is claimed to be parameter free, always locate the factor that cancels and name it. Then i^{n-1} is a square of i raised to an integer power because n is odd, giving the sign $(-1)^{(n-1)/2}$. **Link to the next block.** Done; state the two congruence cases explicitly since the problem said "determine its value".

Layer 3: WHY

The first truth: elementary symmetric functions are coefficients. If you are asked for any symmetric expression in a family of numbers, you should never compute the members. You should find the polynomial they satisfy. The cost of finding one polynomial is $O(1)$; the cost of evaluating n members is $O(n)$ and usually impossible in closed form. This is the same principle that makes Newton's identities useful, and it is the reason Vieta is not a piece of school trivia but a primary weapon.

The second truth: the Cayley transform is not a device pulled from the air, it is the map that diagonalises angle addition. Adding an angle to ϕ acts on $x = \tan\phi$ by the Mobius map $T(x) = \frac{x+t}{1-tx}$ where $t = \tan$ of the added angle. Solve $T(x) = x$ and you get $x^2 = -1$, so the fixed points are $\pm i$ for every t. A Mobius map is understood by conjugating it so that its fixed points go to 0 and ∞, whereupon it becomes plain multiplication $z \mapsto \lambda z$. The conjugating map that sends $i \mapsto 0$ and $-i \mapsto \infty$ is precisely $x \mapsto \frac{1+ix}{1-ix}$ up to a constant. So the substitution in Step 1 is forced by the structure of the problem: you are handed a family generated by repeated angle addition, and you

move to the coordinate in which angle addition is multiplication. Nothing here is a rabbit out of a hat, and if you ever feel it is, re-derive the fixed points and the feeling goes.

The third truth: the arithmetic progression of angles with step π/n is the tangent-world image of the n-th roots of unity. Under the Cayley transform the family $\{\phi + k\pi/n\}$ maps to $\{e^{2i\theta}\zeta^k\}$ with $\zeta = e^{2\pi i/n}$, a coset of the n-th roots of unity in the unit circle. Every classical identity about equally spaced angles, from $\prod_{k=1}^{n-1} 2\sin(k\pi/n) = n$ to the multiple angle formulas, is the single factorisation $\prod_{k=0}^{n-1}(X - \zeta^k) = X^n - 1$ wearing different clothes.

The discernment mechanism, then, is a two question filter. Question one: is the quantity requested a symmetric function of a family? If yes, hunt for the family's polynomial. Question two: is the family generated by a group action (translation of an angle, multiplication by a root of unity, iteration of a map)? If yes, move to coordinates where the action is multiplication, and the polynomial writes itself. Train it by re-deriving, from scratch and without notes, the tangent multiple angle formula, the product formulas for $\prod \sin(\phi + k\pi/n)$, and the fixed points of the tangent addition map. Three derivations, once, and the circuit fires forever.

Revolving concepts

Elementary symmetric functions and Vieta. For monic $p(x) = \prod_{k=1}^{n}(x - a_k) = x^n - e_1x^{n-1} + e_2x^{n-2} - \cdots + (-1)^n e_n$, the coefficient of x^{n-m} is $(-1)^m e_m$. For a non-monic $\sum c_j x^j$, $e_m = (-1)^m c_{n-m}/c_n$. Orbiting results: Newton's identities linking power sums $p_m = \sum a_k^m$ to the e_m; the fundamental theorem of symmetric polynomials (every symmetric polynomial in the roots is a polynomial in the e_m with integer coefficients), which is what makes "it is an integer" predictable in advance. See Book 1, field: Polynomials, topic: symmetric functions and Vieta.

Roots of unity and the master factorisation. $\prod_{k=0}^{n-1}(X - \zeta^k) = X^n - 1$ with $\zeta = e^{2\pi i/n}$, equivalently $\prod_{k=0}^{n-1}(1 - w\zeta^k) = 1 - w^n$. From this single identity: $\prod_{k=0}^{n-1} 2\sin(\phi + k\pi/n) = 2\sin(n\phi)$ for all n, and $\prod_{k=0}^{n-1} 2\cos(\phi + k\pi/n) = (-1)^{(n-1)/2} 2\cos(n\phi)$ for n odd, hence $\prod_{k=0}^{n-1} \tan(\phi + k\pi/n) = (-1)^{(n-1)/2}\tan(n\phi)$, which is the denominator of this problem. Taking the logarithmic derivative of the cosine product converts it into the sum identity $\sum_{k=0}^{n-1} \tan(\phi + k\pi/n) = n\tan(n\phi)$, which is the numerator. Logarithmic differentiation is the standard bridge from a product identity to a sum identity, and you should reach for it every time you own a product formula and want a sum. See Book 1, field: Complex numbers, topic: roots of unity products.

The Cayley transform and Mobius maps. $C(x) = \frac{1+ix}{1-ix}$ is a bijection from $\mathbb{R} \cup \{\infty\}$ to the unit circle, with $C(\tan\phi) = e^{2i\phi}$, $C(i) = 0$, $C(-i) = \infty$. A Mobius map with two distinct fixed points is conjugate to $z \mapsto \lambda z$; a Mobius map with one fixed point is conjugate to $z \mapsto z + 1$. Conjugating to normal form is the universal way to iterate a Mobius map, and it subsumes the tangent addition formula.

de Moivre and the multiple angle recursion, generalised. Define $C_n(x), S_n(x)$ by $(1 + ix)^n = C_n(x) + iS_n(x)$. Then

$$\begin{pmatrix} C_{n+1} \\ S_{n+1} \end{pmatrix} = \begin{pmatrix} 1 & -x \\ x & 1 \end{pmatrix} \begin{pmatrix} C_n \\ S_n \end{pmatrix}, \qquad \begin{pmatrix} C_0 \\ S_0 \end{pmatrix} = \begin{pmatrix} 1 \\ 0 \end{pmatrix},$$

so $\tan(n\phi) = S_n(x)/C_n(x)$ with $x = \tan\phi$. This is the general shape of a two term linear recursion with constant matrix M: solve it by diagonalising M, whose eigenvalues here are $1 \pm ix$ with eigenvectors $(1, \mp i)$. The complex substitution $(1 \pm ix)^n$ used throughout this problem is nothing but that diagonalisation written out, and the same move solves every constant coefficient linear recursion: Fibonacci through φ and $\bar{\varphi}$, Chebyshev through $e^{\pm i\phi}$, and $\tan(n\phi)$ through $1 \pm ix$. Whenever a sequence is defined by applying the same linear map repeatedly, the general solution is a combination of λ_j^n over the eigenvalues λ_j, with polynomial-in-n corrections when eigenvalues repeat. See Book 1, field: Sequences and recurrences, topic: linear recurrences by diagonalisation.

Circuit Bank

- **WHEN** a problem asks for a symmetric expression in an indexed family (a sum, a product, a ratio of the two, or any e_m) **DO** stop evaluating the members and instead construct one polynomial having exactly that family as its roots, then read Vieta **BECAUSE** elementary symmetric functions are literally the coefficients, so one object answers every symmetric question at once
- **WHEN** the family is indexed by an arithmetic progression of angles with common difference π/n or $2\pi/n$ **DO** expect the n-th power or n-th multiple to collapse the family to a single value, and expect roots of unity **BECAUSE** $\prod_{k=0}^{n-1}(X-\zeta^k) = X^n - 1$ is the one identity behind every equally spaced angle formula
- **WHEN** you have tan and want algebraic structure **DO** write $1 + i\tan\phi = e^{i\phi}/\cos\phi$ and pass to $x \mapsto (1+ix)/(1-ix)$ **BECAUSE** that Mobius map sends the fixed points $\pm i$ of the tangent addition map to 0 and ∞, converting angle addition into plain multiplication
- **WHEN** an answer is claimed to be independent of a parameter **DO** identify the single algebraic factor carrying the parameter and show it cancels **BECAUSE** a cancellation you can name is a structural reason, whereas a cancellation you merely observe is an accident waiting to be a sign error
- **WHEN** the expression to be evaluated has a denominator built from the data, or a hypothesis you have not used **DO** spend two lines proving the denominator is nonzero and every object is finite, naming the hypothesis that guarantees it **BECAUSE** here θ/π irrational is exactly what forces $\tan(n\theta) \neq 0$, all a_k finite and nonzero, and $\omega \neq \pm 1$; a grader looks for those lines

Backtesting

1. First useful cue: the quotient is e_1/e_n for the family $a_1, \ldots, a_n$. Symmetric functions demand a polynomial, not an evaluation.
2. Method the expression suggested: common difference π/n in the angles plus period π in tan means all n numbers satisfy the single equation $\tan(n \arctan x) = \tan(n\theta)$; package it as $(1+ix)^n = \omega(1-ix)^n$.
3. Where you would hesitate, and the thought to reinforce: the hesitation is at $1 + i\tan\phi = e^{i\phi}\sec\phi$, which looks like an invented identity. Reinforce that it is one division: take $e^{i\phi} = \cos\phi + i\sin\phi$ and divide by $\cos\phi$. Derive it, never memorise it.
4. The trap the question sets: forgetting to check that Q_n has degree exactly n (so that $a_1, \ldots, a_n$ are all of its roots), and forgetting to use the irrationality hypothesis. Both are places to drop marks while having the right final number.
5. What makes you recognise it faster next time: the pairing "equally spaced angles" with "symmetric function". That pair always means one polynomial, and almost always means the Cayley transform if the function is tan, or $z + z^{-1}$ if it is cos.
6. The reusable decision rule: symmetric function of a family $\Rightarrow$ find the family's polynomial; family generated by a translation $\Rightarrow$ change to the coordinate where translation becomes multiplication.

Mirror drills

1. Let n be odd and $a_k = \tan(\theta + k\pi/n)$ as above. Show that $\sum_{i<j} a_i a_j = -\binom{n}{2}$, independent of θ. *Answer:* from $c_j = \binom{n}{j} i^j (1 - \omega(-1)^j)$, the coefficient c_{n-2} carries the factor $1+\omega$, the same factor as c_n, so $e_2 = c_{n-2}/c_n = \binom{n}{2} i^{-2} = -\binom{n}{2}$.

2. Prove $\prod_{k=1}^{n-1} 2\sin(k\pi/n) = n$ for every integer $n \geq 2$. *Answer:* put $X = 1$ in $\prod_{k=1}^{n-1}(X - \zeta^k) = (X^n - 1)/(X - 1)$ to get $\prod(1 - \zeta^k) = n$, then use $1 - e^{2i\alpha} = -2i\sin(\alpha)e^{i\alpha}$ and check that the phases multiply to 1.

3. Let n be odd. Prove $\sum_{k=0}^{n-1} \tan(\phi + k\pi/n) = n\tan(n\phi)$ directly. *Answer:* take log and differentiate the product identity $\prod_{k=0}^{n-1} 2\cos(\phi + k\pi/n) = (-1)^{(n-1)/2} 2\cos(n\phi)$, giving $-\sum \tan(\phi + k\pi/n) = -n\tan(n\phi)$.

20.6 A6. Sylvester four point problem in a disc

Four points are chosen uniformly and independently at random in the interior of a given circle. Find the probability that they are the vertices of a convex quadrilateral.

Read it out loud

Four independent uniform points in a disc, and a question about their configuration. Say the negation out loud, because the negation is the workable object: four points in general position in the plane either form a convex quadrilateral, or exactly one of them sits inside the triangle formed by the other three. There is no third possibility. So the question is not about quadrilaterals at all; it is about the probability that a random point lands inside a random triangle, and by symmetry that is four copies of one event.

Now push one step further. Conditional on the first three points, the probability that the fourth lands inside their triangle is exactly the triangle's area divided by the disc's area, because a uniform point hits a region with probability proportional to its area. Averaging over the first three converts a probability into an expectation of an area. That is the whole reduction, and it is worth saying in one sentence: the answer is $1 - 4\,\mathrm{E}[T]/|D|$ where T is the area of the triangle on three uniform points and $|D|$ is the area of the disc. Everything after that is a definite integral.

The fields wired in are convexity in the plane (the four point classification, which is the planar case of Radon's theorem), geometric probability (uniform measure, hitting probabilities, Fubini), and multivariable calculus (polar coordinates, a Jacobian, Wallis integrals). What is being tested is whether you can convert a configuration question into an area expectation, and then whether you can organise the resulting integral so that it is mechanical rather than a swamp of absolute values.

Be honest about where this problem sits. This is Sylvester's four point problem, posed in 1864, and the disc case is a named classical result. The reduction is genuinely available to you in the arena within two minutes. The evaluation of $\mathrm{E}[T]$ is a real computation that will take twenty to thirty minutes of careful work, and there is no route that avoids at least one nontrivial integral. A6 is where you find out whether your integration hygiene survives pressure. Within thirty seconds you should have written: no three collinear almost surely; exactly one of four configurations fails; $\mathrm{P}(\text{convex}) = 1 - 4\,\mathrm{E}[T]/\pi$ for the unit disc.

Approaches, ranked worst to best

1. **Integrate the convexity indicator over all eight coordinates.** *Cue:* the problem is a probability, so write the probability as an integral. *Circuit:* none; this is the zombie route. An eight dimensional integral whose integrand is the indicator of a configuration property, with no separable structure. There is no way to start. Last place, and it is last by a mile.

2. **Reduce to $\mathrm{E}[T]$, then integrate the area directly over three points.** *Cue:* after the reduction you need the mean of $\frac{1}{2}|(Q-P)\times(R-P)|$. *Circuit:* brute polar integration. In polar coordinates the area is $\frac{1}{2}\left|r_1 r_2 \sin(\alpha_1-\alpha_2) + r_2 r_3 \sin(\alpha_2-\alpha_3) + r_3 r_1 \sin(\alpha_3-\alpha_1)\right|$. That absolute value around a three term sum is the killer: you cannot integrate it without splitting into the cases where the bracket changes sign, and the sign regions are exactly the configurations you were trying to avoid thinking about. Correct in principle, unmanageable in practice.

3. **Signed area decomposition from the centre.** *Cue:* a triangle's area is a signed sum of the three triangles it makes with any chosen point, here the centre O. *Circuit:* $[PQR] = \pm[OPQ] \pm [OQR] \pm [ORP]$. Each of $[OPQ] = \frac{1}{2} r_1 r_2 \sin\angle POQ$ is easy to average, and the correction terms involve whether the three radial directions lie in a semicircle. The integrals are small; the sign bookkeeping is not. You must classify by semicircle position and by whether the middle point falls inside the triangle OPR, and get every sign right. This is a strong route with a treacherous middle, and treacherous middles are where you lose four points out of ten.

4. **Region decomposition plus the chord distance distribution.** *Cue:* the three lines PQ, QR, RP cut the disc into seven regions, and the union of the triangle with the three far regions is exactly the bad region for the fourth point. *Circuit:* comparing two expectations of the same quantity. Naming the expected areas A (triangle), B (three regions adjacent to a side), C (three far regions) gives $A + 3B + 3C = \pi$, and comparing expectations of the bad area gives $3C + A = 4A$, hence $A = C$. That identity is the same statement as Efron's identity and it is a real insight. But you still have to compute $B + 2A$, which requires the distribution of the distance from the centre to the line through two random points, namely $\frac{16}{3\pi}(1 - h^2)^{3/2}$, and deriving that is a chunk of work in its own right. Second place: elegant structure, expensive execution.
5. **Scale by the farthest point, then one double integral.** *Cue:* the disc is scale invariant, areas scale as the square of the radius, and the radial distribution of a uniform point is explicit. *Circuit:* condition on the largest of the three radii and rescale to a unit disc with one vertex on the boundary. After conditioning, the problem is the expected area of a triangle with one vertex fixed on the boundary and two uniform inside, and in polar coordinates centred at that boundary vertex the disc has the equation $r = 2\sin\theta$, which makes both radial integrals into single powers. No signs, no case analysis, one double integral in θ and θ' that reduces entirely to Wallis integrals. This is the ideal route: every step is mechanical after one structural idea.

The ideal solution

By scale invariance take the circle to have radius 1; write D for the open unit disc, O for its centre, $|D| = \pi$, and $[XYZ]$ for the area of triangle XYZ. Let P, Q, R, S be independent and uniform on D.

Step 1: reduce to the expected area of a triangle.

Lemma 20.6.1. *Let P, Q, R, S be four points in the plane, no three collinear. Then either they are the vertices of a convex quadrilateral, or exactly one of them lies in the interior of the triangle formed by the other three, and not both.*

Proof. Consider the convex hull H of the four points. Since no three are collinear, H is either a quadrilateral or a triangle. If H is a quadrilateral, all four points are vertices of it and they form a convex quadrilateral. If H is a triangle, its three vertices are three of the points and the fourth lies in H; it is not on the boundary, since being on an edge would make three of the points collinear, so it lies in the interior. Exactly one point is omitted from the vertex set in that case, so the two alternatives are exclusive and the interior point is unique. □

Three of four independent uniform points are collinear with probability zero (given two of them, the third must land on a line, a null set), so we may work on the complement of a null event. The four events "S is inside triangle PQR", "R is inside PQS", and so on, are mutually exclusive by the lemma and have equal probability by symmetry of the four points. Hence

$$\mathrm{P}(\text{convex}) = 1 - 4\ \mathrm{P}\big(S \in \triangle PQR\big).$$

Conditioning on P, Q, R and using that S is uniform on D independently of them,

$$\mathrm{P}\big(S \in \triangle PQR\big) = \mathrm{E}\left[\frac{[PQR]}{\pi}\right] = \frac{\mathrm{E}[T]}{\pi}, \qquad T := [PQR].$$

So the whole problem is the value of $\mathrm{E}[T]$, and

$$\mathrm{P}(\text{convex}) = 1 - \frac{4\ \mathrm{E}[T]}{\pi}.$$

Step 2: condition on the farthest of the three points and rescale.

Let X denote the expected area of a triangle with one vertex fixed at a point of the unit circle and the other two independent and uniform on the unit disc. (By rotational symmetry this does not depend on which boundary point is fixed.)

For a uniform point of D, $\mathrm{P}(|OP| \le r) = r^2$, so $|OP|$ has density $2r$ on $[0,1]$. Ties among $|OP|, |OQ|, |OR|$ have probability zero, so by symmetry

$$\mathrm{E}[T] = 3\,\mathrm{E}\big[T \cdot \mathbf{1}[|OR| > |OP|,\ |OR| > |OQ|]\big].$$

Condition on $|OR| = c$. Given this, the direction of R is uniform on the circle and independent of c, by rotational invariance of the uniform measure on D. The event $\{|OP| \le c,\ |OQ| \le c\}$ has probability $c^2 \cdot c^2 = c^4$, and conditionally on it P and Q are independent and uniform on the disc of radius c, because conditioning a uniform distribution on a subregion returns the uniform distribution on that subregion. Rescaling that whole picture by the factor $1/c$ maps it to the unit disc with one vertex on the unit circle and two uniform interior points, and areas are multiplied by c^{-2}. Therefore the conditional expectation of T on this event is $c^2 X$, and

$$\mathrm{E}[T] = 3\int_0^1 2c \cdot c^4 \cdot c^2 X\, dc = 6X\int_0^1 c^7\, dc = \frac{3X}{4}.$$

Step 3: compute X in polar coordinates centred at the boundary vertex.

Put the fixed vertex V at the origin of a polar system and the centre of the disc at the point $(0,1)$, so that the unit circle has polar equation

$$r^2 - 2r\sin\theta = 0, \qquad \text{that is} \qquad r = 2\sin\theta, \quad \theta \in [0,\pi],$$

and the disc is $\{(r,\theta) : 0 \le \theta \le \pi,\ 0 \le r \le 2\sin\theta\}$. A uniform point of the disc has density $\frac{1}{\pi} r\, dr\, d\theta$ in these coordinates. If $P = (\rho,\theta)$ and $Q = (\rho',\theta')$ then the angle at V between VP and VQ is $|\theta - \theta'| \in [0,\pi]$, so

$$[VPQ] = \tfrac{1}{2}\rho\rho' \sin|\theta - \theta'|.$$

Hence

$$X = \frac{1}{2\pi^2}\int_0^\pi\int_0^\pi \sin|\theta-\theta'| \left(\int_0^{2\sin\theta} \rho^2 d\rho\right)\left(\int_0^{2\sin\theta'} \rho'^2 d\rho'\right) d\theta\, d\theta'.$$

Each radial integral is $\frac{8}{3}\sin^3$, so

$$X = \frac{1}{2\pi^2}\cdot\frac{64}{9}\int_0^\pi\int_0^\pi \sin^3\theta \sin^3\theta' \sin|\theta-\theta'|\, d\theta'\, d\theta = \frac{64}{9\pi^2}\, I, \quad I := \int_0^\pi\int_0^\theta \sin^3\theta \sin^3\theta' \sin(\theta-\theta')\, d\theta'\, d\theta,$$

the last step using the symmetry of the integrand in $\theta \leftrightarrow \theta'$ to fold the square onto the triangle $\theta' < \theta$, where $|\theta - \theta'| = \theta - \theta'$.

Expand $\sin(\theta-\theta') = \sin\theta\cos\theta' - \cos\theta\sin\theta'$ and integrate in θ' first, using $\int_0^\theta \sin^3\theta' \cos\theta'\, d\theta' = \frac{1}{4}\sin^4\theta$ and $\int_0^\theta \sin^4\theta'\, d\theta' = \frac{3\theta}{8} - \frac{\sin 2\theta}{4} + \frac{\sin 4\theta}{32}$:

$$I = \frac{1}{4}\int_0^\pi \sin^8\theta\, d\theta \;-\; \frac{3}{8}\int_0^\pi \theta\sin^3\theta\cos\theta\, d\theta \;+\; \frac{1}{4}\int_0^\pi \sin^3\theta\cos\theta\sin 2\theta\, d\theta \;-\; \frac{1}{32}\int_0^\pi \sin^3\theta\cos\theta\sin 4\theta\, d\theta.$$

Evaluate the four pieces. The Wallis values are $\int_0^\pi \sin^{2m}\theta\, d\theta = \pi\binom{2m}{m}4^{-m}$, giving $\int_0^\pi \sin^4 = \frac{3\pi}{8}$, $\int_0^\pi \sin^6 = \frac{5\pi}{16}$, $\int_0^\pi \sin^8 = \frac{35\pi}{128}$, hence also $\int_0^\pi \sin^4\cos^2 = \frac{3\pi}{8} - \frac{5\pi}{16} = \frac{\pi}{16}$ and $\int_0^\pi \sin^6\cos^2 = \frac{5\pi}{16} - \frac{35\pi}{128} = \frac{5\pi}{128}$. Then:

$$\int_0^\pi \theta\sin^3\theta\cos\theta\, d\theta = \Big[\theta\,\tfrac{\sin^4\theta}{4}\Big]_0^\pi - \frac{1}{4}\int_0^\pi \sin^4\theta\, d\theta = -\frac{3\pi}{32},$$

$$\int_0^\pi \sin^3\theta\cos\theta\sin 2\theta\, d\theta = 2\int_0^\pi \sin^4\theta\cos^2\theta\, d\theta = \frac{\pi}{8},$$

$$\int_0^\pi \sin^3\theta\cos\theta\sin 4\theta\,d\theta = 4\int_0^\pi \sin^4\theta\cos^2\theta\,d\theta - 8\int_0^\pi \sin^6\theta\cos^2\theta\,d\theta = \frac{\pi}{4} - \frac{5\pi}{16} = -\frac{\pi}{16},$$

where the last line uses $\sin 4\theta = 4\sin\theta\cos\theta(1-2\sin^2\theta)$. Substituting,

$$I = \frac{35\pi}{512} + \frac{9\pi}{256} + \frac{\pi}{32} + \frac{\pi}{512} = \frac{35+18+16+1}{512}\pi = \frac{70\pi}{512} = \frac{35\pi}{256}.$$

Step 4: assemble. Therefore

$$X = \frac{64}{9\pi^2}\cdot\frac{35\pi}{256} = \frac{35}{36\pi}, \qquad \mathrm{E}[T] = \frac{3}{4}X = \frac{35}{48\pi}, \qquad \mathrm{P}(\text{convex}) = 1 - \frac{4}{\pi}\cdot\frac{35}{48\pi} = \boxed{1 - \frac{35}{12\pi^2}}.$$

Numerically this is 0.70448, which passes the reality check: a little over 70 per cent, comfortably above the value $2/3$ that the same formula gives for a triangular region and below 1.

Layer 1: WHAT

Block 1: classify four points in general position into exactly two cases, convex position or one inside the triangle of the others.

Block 2: use symmetry and exclusivity to write $\mathrm{P}(\text{not convex}) = 4\,\mathrm{P}(S \in \triangle PQR)$.

Block 3: condition on P, Q, R and turn that probability into $\mathrm{E}[T]/\pi$, converting the configuration question into an area expectation.

Block 4: exploit scale invariance: condition on the farthest of the three points, rescale so that it sits on the unit circle, and reduce to the single number X with $\mathrm{E}[T] = \frac{3}{4}X$.

Block 5: compute X by polar coordinates centred at the boundary vertex, where the disc is $r \le 2\sin\theta$.

Block 6: fold the double integral onto $\theta' < \theta$, expand $\sin(\theta - \theta')$, and finish with Wallis integrals and one integration by parts.

Block 7: assemble and sanity check numerically.

The nominal route: set up the probability as an integral over the eight coordinates, or reduce to $\mathrm{E}[T]$ and then attack the area formula with its absolute value head on. The bifurcation point is at Block 4. Everyone who gets Blocks 1 to 3 has the same problem in front of them, namely $\mathrm{E}[T]$. The zombie starts integrating $\left|\sum r_i r_j \sin(\alpha_i - \alpha_j)\right|$ over six variables. You ask instead: what symmetry of the disc have I not yet spent? The answer is scaling, and scaling is spent by conditioning on the largest radius, because that is the one statistic that a rescaling normalises to 1.

Layer 2: HOW

Block 1. Cue. The problem asks about a configuration of four points, and the property is qualitative. **Why natural.** Qualitative configuration properties of finitely many points in the plane always have short classifications, and the classification is by the convex hull. With four points and no three collinear the hull has three or four vertices, so there are exactly two cases. Note that "no three collinear" is not a convenience assumption, it is an almost sure event, and you must say so. **Link to the next block.** The classification produces a complement that is a union of four events, which invites symmetry.

Block 2. Cue. Four events, exchangeable under relabelling the points, and mutually exclusive by the lemma. **Why natural.** Exclusivity is what lets you add probabilities without inclusion and exclusion, and exchangeability is what lets you compute only one of them. Both facts come from the lemma, which is why the lemma is proved and not assumed. **Link to the next block.** You are left with one event about a point landing inside a triangle built from the other points.

Block 3. **Cue.** A uniform point landing in a region. **Why natural.** By definition of the uniform distribution, the probability is the region's area over the total area. The region here is random, but it is independent of S, so condition on it first and average afterwards. This is the standard move: a probability involving a random region becomes an expectation of an area. **Link to the next block.** The problem is now a purely geometric expectation, $\mathrm{E}[T]$, with no probability language left in it.

Block 4. **Cue.** The disc is invariant under rotations about O and under scaling about O, and you have used rotations but not scalings. **Why natural.** An unused symmetry is an unused reduction. To use scaling you need a statistic that scaling acts on transitively, and the natural one is the largest of the three distances from O: rescale by that and it becomes 1, putting one vertex on the boundary. The distribution $2c\,dc$ of a single radius and the factor c^4 for the other two both come from $\mathrm{P}(|OP| \le r) = r^2$, which is itself just an area ratio. **Link to the next block.** After the reduction there is one unknown constant X, defined by a configuration with a fixed boundary vertex, which is a strictly easier object because one of the three points is no longer random.

Block 5. **Cue.** Two random points, one fixed vertex, and an area formula that wants distances and an included angle from that fixed vertex. **Why natural.** Choose the origin of coordinates where the geometry is simplest, which is the fixed vertex, not the centre. The reward is immediate: a circle through the origin has polar equation $r = 2R\sin\theta$, so the region is described by a single inequality with no case split, and the area formula $\frac{1}{2}\rho\rho'\sin|\theta - \theta'|$ separates into a radial part and an angular part. Choosing the centre as origin instead would give a messy radial limit and no separation. **Link to the next block.** The radial integrals evaluate to powers of sin, leaving a pure double integral in the angles.

Block 6. **Cue.** An absolute value $\sin|\theta - \theta'|$ in a symmetric double integral. **Why natural.** The integrand is symmetric under swapping θ and θ', so restrict to $\theta' < \theta$ and double. This is the correct way to kill an absolute value: not by case analysis inside the integrand, but by using a symmetry to halve the domain. Then $\sin(\theta - \theta')$ expands by the addition formula into products that integrate in θ' in closed form. **Link to the next block.** What remains is four one dimensional integrals, all standard.

Block 7. **Cue.** A closed form answer to a classical question. **Why natural.** Check it numerically. $35/(12\pi^2) \approx 0.2955$, so the answer is about 0.7045, and comparison with the known triangle value $2/3$ and square value $25/36$ places it correctly. A numerical check costs thirty seconds and catches every dropped factor of two. **Link to the next block.** Done.

Layer 3: WHY

The first truth is that a probability about a configuration of independent points should be converted, as early as possible, into an expectation of a geometric quantity. The mechanism is always the same: single out one point, condition on the others, and use that a uniform point hits a region with probability equal to the normalised area. Every problem of this family, from Buffon's needle to Sylvester's, runs on this exchange. The corollary is that the hard part is never the probability; it is the geometric mean value that pops out.

The second truth is that you must spend every symmetry the region has. The disc has rotations, reflections, and scalings about its centre. Rotations are spent by declaring the angular coordinate of one point to be uniform and independent. Scalings are spent by conditioning on a scale statistic, and the largest radius is the natural choice because rescaling by it produces a normalised configuration with one vertex on the boundary. That is what turns a three point problem into a one point problem plus a one dimensional integral in c, and the integral in c is trivial because the scaling acts on

the area with weight c^2 and on the probability with weight $c^4 \cdot 2c$. The general shape is: if a quantity is homogeneous of degree d under the symmetry and the symmetry parameter has a known distribution, the conditional expectation factors as a moment times a normalised constant.

The third truth is about absolute values. Every geometric probability integrand carries one, because areas and distances are unsigned. There are only two honest ways to remove it: use a symmetry of the integrand to restrict to the region where the sign is known, or decompose the quantity into pieces with fixed sign. Route 3 in the ranking takes the second option and pays with case analysis; the ideal route takes the first and pays nothing. Train the reflex: when you see $|f(u) - f(v)|$ inside a symmetric double integral, halve the domain immediately, before doing anything else.

What were the examiners nudging toward? The word "convex" with exactly four points is the nudge to the hull classification. The word "circle" rather than "square" or "triangle" is the nudge to polar coordinates and to the scaling symmetry. And the request for a probability rather than a proof is the signal that the answer is a specific constant that will not be pretty, which tells you in advance not to expect the integral to collapse.

Knowledge gaps, stated plainly. To have a real chance at this problem in the arena you must already own: the hull classification of four points, that is, the planar case of Radon's theorem; the fact that a uniform point in a disc has radial density $2r$, which is the polar Jacobian; the polar equation $r = 2R\sin\theta$ of a circle through the origin; the Wallis values of $\int_0^\pi \sin^{2m}$; and the discipline of checking a final constant numerically. You do not need Crofton's formula, which is the classical tool for this problem, but you should know it exists and what it does, because it is the standard machinery for turning geometric probability questions about lines and chords into one dimensional integrals. Naming that gap is better than pretending the ideal solution above is the only way in.

Revolving concepts

Convex position and Radon's theorem. Radon's theorem: any $d + 2$ points in $\mathbb{R}^d$ can be partitioned into two sets whose convex hulls intersect. In the plane, $d + 2 = 4$, and in general position the partition is either $2 + 2$ (the two diagonals cross, so the four points are in convex position) or $1 + 3$ (one point inside the triangle of the others). That dichotomy is precisely the lemma used above. See Book 1, field: Geometry, topic: convex hulls and Radon.

Uniform measure, hitting probabilities and Fubini. For U uniform on a region K and any measurable $A \subseteq K$, $\mathrm{P}(U \in A) = |A|/|K|$; if A is random and independent of U, then $\mathrm{P}(U \in A) = \mathrm{E}\,|A|/|K|$. The interchange is Fubini for a bounded nonnegative integrand, so no justification beyond measurability is needed.

Radial distribution and scaling in the disc. For U uniform on the disc of radius R, $\mathrm{P}(|OU| \le r) = r^2/R^2$, so the density is $2r/R^2$ on $[0, R]$, and the polar angle is independent and uniform. Consequences worth banking: $\mathrm{E}\,|OU| = \frac{2}{3}R$, $\mathrm{E}\,|OU|^2 = \frac{1}{2}R^2$, and for n independent points the largest radius has distribution function $(r/R)^{2n}$ and density $2nr^{2n-1}/R^{2n}$. The maximum of n independent copies always has distribution function F^n; that is the whole of order statistics that this problem needs.

Polar coordinates at a boundary point. A circle of radius R through the origin, with centre in the direction $\theta = \pi/2$, is $r = 2R\sin\theta$ for $\theta \in [0, \pi]$. This is the single change of viewpoint that makes Step 3 mechanical, and it is worth installing permanently: for any problem about chords or triangles anchored at a point of a circle, put the origin at that point.

Wallis integrals. $\int_0^\pi \sin^{2m}\theta\,d\theta = \pi\binom{2m}{m}4^{-m}$ and $\int_0^{\pi/2} \sin^a\theta\cos^b\theta\,d\theta = \frac{1}{2}B\left(\frac{a+1}{2}, \frac{b+1}{2}\right)$ with $B(x, y) = \Gamma(x)\Gamma(y)/\Gamma(x + y)$. Every integral in Step 3 is one of these after a two line reduction.

Sylvester's four point problem and its family. For a bounded convex region K, the same reduction gives $\mathrm{P}(\text{convex}) = 1 - 4\,\mathrm{E}[T]/|K|$. Known values: triangle $2/3$ (since $\mathrm{E}[T] = |K|/12$), square $25/36$ (since $\mathrm{E}[T] = 11|K|/144$), disc $1 - 35/(12\pi^2) \approx 0.7045$. Blaschke's theorem says these are extreme: over all bounded convex K the probability is minimised by triangles at $2/3$

and maximised by ellipses at $1 - 35/(12\pi^2)$. Crofton's formula, which expresses integrals over families of lines in terms of a measure invariant under rigid motions, is the classical machine for such computations. See Book 1, field: Probability, topic: geometric probability and Sylvester's problem.

Efron's identity, and the recursion that generalises this problem. Let V_n be the number of vertices of the convex hull of n independent uniform points in K and A_n the area of that hull. Then

$$\mathrm{E}[V_n] = n\left(1 - \frac{\mathrm{E}[A_{n-1}]}{|K|}\right),$$

because a point is a hull vertex exactly when it lies outside the hull of the other $n-1$, and linearity of expectation over the n points does the rest. This is a genuine recursion linking hull vertex counts at size n to hull areas at size $n-1$, and it contains this problem as the case $n = 4$: since $V_4 \in \{3, 4\}$ we get $\mathrm{E}[V_4] = 4 - \mathrm{P}(\text{not convex})$, and Efron gives $\mathrm{E}[V_4] = 4\big(1 - \mathrm{E}[A_3]/\pi\big) = 4 - 35/(12\pi^2)$, in agreement. Read in the other direction it also recovers the observation $A = C$ from ranked approach 4. The general lesson: for random convex hulls, vertex counts and areas are two faces of one recursion, and whichever of the two is easier to compute determines the other. See Book 1, field: Probability, topic: random convex hulls.

Circuit Bank

- **WHEN** a probability question asks whether k random points form a particular configuration **DO** classify the configurations by the convex hull, check the alternatives are mutually exclusive, and use exchangeability to reduce to one representative event **BECAUSE** configuration properties of finitely many points have short hull based classifications, and exclusivity plus symmetry turns a union of k events into k times one probability
- **WHEN** a uniform point must land inside a region built from the other random points **DO** condition on those other points and replace the probability by the expected area over the total area **BECAUSE** the defining property of the uniform distribution is that hitting probability equals normalised area, so the probability becomes a purely geometric expectation
- **WHEN** the region is a disc and you have already used rotational symmetry **DO** spend the scaling symmetry by conditioning on the largest of the radii and rescaling it to 1 **BECAUSE** areas are homogeneous of degree two and the radius has an explicit power law distribution, so the conditional expectation factors into a single moment integral times a normalised constant
- **WHEN** you must integrate over a disc with one distinguished point on its boundary **DO** place the polar origin at that boundary point, where the circle reads $r = 2R\sin\theta$ **BECAUSE** the region then needs one inequality instead of a case split, and the radial integrals collapse to powers of $\sin\theta$, leaving Wallis integrals
- **WHEN** an absolute value such as $|\theta - \theta'|$ sits inside a symmetric multiple integral **DO** halve the domain using the symmetry rather than splitting the integrand into sign cases **BECAUSE** a symmetry that fixes the integrand and reverses the sign of the argument is exactly what an absolute value is asking for, and case analysis in geometric probability is the main source of dropped marks

Backtesting

1. First useful cue: "four points" plus "convex". That pairing forces the hull dichotomy, and the dichotomy immediately converts the question into one about a point inside a triangle.
2. Method the expression suggested: a uniform point inside a random triangle means expected area over total area, so the target became $\mathrm{E}[T]$ within two minutes of reading.

3. Where you would hesitate, and the thought to reinforce: at $\mathrm{E}[T]$, staring at an absolute value over six variables. Reinforce: list the symmetries the region has and ask which ones you have not yet spent. Rotation was spent; scaling was not; conditioning on the largest radius spends it.
4. The trap the question sets: attacking $\mathrm{E}[T]$ by direct integration, or attempting the signed decomposition from the centre and losing control of the sign cases. Both are routes that look like progress for twenty minutes and then collapse.
5. What makes you recognise it faster next time: the phrase "probability that random points are in convex position" is a named classical family (Sylvester). Recognise the family, write $1 - 4\,\mathrm{E}[T]/|K|$ at once, and spend all your remaining time on the area.
6. The reusable decision rule: configuration probability $\Rightarrow$ hull classification $\Rightarrow$ expected area; then, in a disc, condition on the largest radius and put the polar origin at the boundary vertex.

Mirror drills

1. Three points are chosen independently and uniformly in a triangle of area 1; it is known that the expected area of the triangle they form is $1/12$. Deduce Sylvester's answer for a triangular region. *Answer:* $\mathrm{P}(\text{convex}) = 1 - 4 \cdot \frac{1}{12} = \frac{2}{3}$, using the same reduction as in Step 1, which never used the shape of the region.
2. For a point U uniform in the unit disc, compute $\mathrm{E}\,|OU|$ and $\mathrm{E}\,|OU|^2$, and the density of max of three independent copies. *Answer:* density $2r$ gives $\mathrm{E}\,|OU| = \int_0^1 2r^2 dr = \frac{2}{3}$ and $\mathrm{E}\,|OU|^2 = \frac{1}{2}$; the maximum has distribution function r^6 and density $6r^5$.
3. Using Efron's identity, find the expected number of vertices of the convex hull of four independent uniform points in a disc, and check it against the answer above. *Answer:* $\mathrm{E}[V_4] = 4(1 - \mathrm{E}[A_3]/\pi) = 4 - 4 \cdot \frac{35}{48\pi^2} = 4 - \frac{35}{12\pi^2}$, which equals $3\,\mathrm{P}(\text{not convex}) + 4\,\mathrm{P}(\text{convex})$, as it must.

20.7 B1. equilateral triangle hiding inside a cubic curve

Show that the curve $x^3 + 3xy + y^3 = 1$ contains only one set of three distinct points, A, B, and C, which are vertices of an equilateral triangle, and find its area.

Read it out loud

The claim is uniqueness: exactly one equilateral triangle sits on this curve. A genuine smooth irreducible cubic (an elliptic curve, for instance) has continuous symmetries you can rotate an inscribed triangle along, so it would carry either zero or infinitely many equilateral triangles, never exactly one. The word "only one" is therefore not a fact about triangles, it is a fact about the curve: it is telling you, before you have done a single calculation, that this cubic must be reducible into pieces so rigid that a triangle can barely fit. That is the first cue, and it tells you to stop thinking about the curve as one object and start hunting for a factorisation.

Look at the shape of the left side: $x^3 + 3xy + y^3$. Two cubes and a cross term with coefficient exactly 3. That coefficient is not decorative. It is the signature of the symmetric identity $x^3 + y^3 + z^3 - 3xyz$, which factors completely over the reals. You are one substitution away from turning an algebraic geometry question into a line-and-point question.

The fields embedded here are polynomial factorisation and Euclidean geometry, and what is really being tested is whether you reroute immediately, before any parametrisation attempt, into checking a curve against a small stock of standard symmetric identities. Within the first thirty seconds you should have written $z = -1$ next to the expression and started checking whether $x^3 + 3xy + y^3 - 1$ matches $x^3 + y^3 + z^3 - 3xyz$ at that value of z.

Approaches, ranked worst to best

1. **brute rotation.** *Cue:* "three points, equilateral". *Circuit:* rotate by 120° about an unknown centroid using complex numbers, substitute into the cubic, solve. (this treats the curve as one irreducible object and sets up a nonlinear system in the centroid and the rotation, with no guarantee it simplifies; it costs a great deal of ink and gives you no independent way to confirm uniqueness, since you would be searching blind rather than seeing the rigid structure first)

2. **singularity hunting.** *Cue:* a cubic that behaves this rigidly. *Circuit:* compute ∂_x, ∂_y of the defining polynomial and look for common zeros, to test whether the curve is singular. (this correctly flags that something is wrong with the curve, but a singular point alone does not hand you the factorisation; it tells you WHERE to look, not the shape of the answer, so it is at best a warm up to the real approach)

3. **symmetric identity match.** *Cue:* cube terms plus a $3xy$ cross term. *Circuit:* recognise $x^3 + 3xy + y^3 - 1$ as $x^3 + y^3 + z^3 - 3xyz$ with $z = -1$, factor immediately. (this converts the whole problem into elementary algebra in one line, then pure Euclidean geometry finishes it; it is the ideal route because it costs the least ink and leaves nothing to search for)

The ideal solution

Set $z = -1$. Then $z^3 = -1$ and $-3xyz = 3xy$, so

$$x^3 + 3xy + y^3 - 1 = x^3 + y^3 + z^3 - 3xyz\Big|_{z=-1}.$$

Recall the symmetric factorisation

$$x^3 + y^3 + z^3 - 3xyz = (x + y + z)(x^2 + y^2 + z^2 - xy - yz - zx).$$

(This is verified either by direct expansion of the right side, which collects to $x^3+y^3+z^3-3xyz$ after every cross term cancels in pairs, or by noting it is the determinant of the circulant matrix with rows the cyclic shifts of (x,y,z), whose eigenvalues at the eigenvectors $(1,\omega^i,\omega^{2i})$, $i=0,1,2$, for ω a primitive cube root of unity, are $x+\omega^i y+\omega^{2i}z$; the determinant is the product of these three eigenvalues, and multiplying the two complex conjugate factors for $i=1,2$ gives the real quadratic above.)

Substituting $z=-1$,

$$x^3+3xy+y^3-1=(x+y-1)\left(x^2-xy+y^2+x+y+1\right).$$

Doubling the second factor and regrouping,

$$2\left(x^2-xy+y^2+x+y+1\right)=(x-y)^2+(x+1)^2+(y+1)^2,$$

a sum of three squares, so the second factor is nonnegative and vanishes only when $x-y=0$, $x+1=0$, and $y+1=0$ hold simultaneously, that is, only at $(x,y)=(-1,-1)$.

Hence the curve is exactly

$$\{(x,y):x+y=1\}\ \cup\ \{(-1,-1)\}.$$

These two pieces are disjoint, since $(-1)+(-1)=-2\neq 1$.

Three distinct points cannot all lie on the line $x+y=1$, since a triangle needs non collinear vertices. Since the only point off the line is the single point $(-1,-1)$, and the three vertices are distinct, at most one of them can equal $(-1,-1)$. Combining these two restrictions, exactly one vertex, say A, equals $(-1,-1)$, and the other two, B and C, lie on the line $x+y=1$.

Write M for the foot of the perpendicular from $A=(-1,-1)$ to the line $x+y=1$. For ABC to be equilateral it must in particular be isosceles with apex A, and the altitude from the apex of an isosceles triangle bisects the base, so M must be the midpoint of BC. This pins down M as a single fixed point on the line, independent of any choice, and the altitude length

$$h=\frac{|(-1)+(-1)-1|}{\sqrt{1^2+1^2}}=\frac{3}{\sqrt{2}}=\frac{3\sqrt{2}}{2}$$

is likewise fixed. In an equilateral triangle of side s, the height satisfies $h=\frac{\sqrt{3}}{2}s$ (drop the altitude, apply Pythagoras to a half triangle of base $s/2$ and hypotenuse s), so s is forced to be $s=\dfrac{2h}{\sqrt{3}}$; conversely with M and h fixed and s chosen this way, the two points on the line at distance $s/2$ from M are automatically at distance s from A by the same Pythagorean identity run in reverse, so the resulting triangle genuinely is equilateral. This leaves no freedom: B and C are the unique pair of points on the line at distance $s/2$ either side of M, so there is exactly one such triangle, as a set of three points.

Its area is

$$\text{Area}=\frac{\sqrt{3}}{4}s^2=\frac{\sqrt{3}}{4}\cdot\frac{4h^2}{3}=\frac{h^2}{\sqrt{3}}.$$

With $h^2=\dfrac{9}{2}$,

$$\text{Area}=\frac{9/2}{\sqrt{3}}=\frac{9}{2\sqrt{3}}=\boxed{\frac{3\sqrt{3}}{2}}.$$

Layer 1: WHAT

Block 1: read "only one triangle" as a claim about the curve's rigidity, not about triangles. Block 2: match the left side against $x^3+y^3+z^3-3xyz$ with $z=-1$ and factor. Block 3: show the quadratic

factor is a sum of three squares, locate its unique real zero. Block 4: the curve is a line union one external point, so exactly one vertex is forced to that point. Block 5: turn the geometry into one scalar, the distance from that point to the line, and finish with the height-to-area formula for an equilateral triangle.

Nominal route: treat the cubic as one irreducible curve and search for three points related by a 120° rotation directly. Ideal route: factor first. The bifurcation point is the instant you decide whether to parametrise the curve or to test it against the sum-of-cubes-minus-$3xyz$ identity; everything downstream of that decision is either a nonlinear system in disguise or three lines of algebra followed by pure geometry.

Layer 2: HOW

Block 1. **Cue.** "only one set of three points" attached to a cubic that looks otherwise generic. **Why natural.** Uniqueness of an inscribed equilateral triangle is not a property a smooth irreducible cubic can have, by the continuous symmetry argument, so the curve must be degenerate. **Link to the next block.** Degeneracy of a cubic in two variables is almost always exposed by a known symmetric factorisation, so look for one.

Block 2. **Cue.** two cube terms and a cross term with coefficient exactly 3. **Why natural.** This is precisely the shape of $x^3 + y^3 + z^3 - 3xyz$ with the third variable fixed at a constant; the constant term -1 on the curve fixes $z = -1$ since $z^3 = -1$ matches and $-3xyz = 3xy$ matches simultaneously. **Link to the next block.** Once matched, the standard factorisation of that identity applies immediately.

Block 3. **Cue.** a quadratic factor that is claimed, or suspected, to vanish at only one real point. **Why natural.** A quadratic form that is a sum of squares is visibly nonnegative, and its zero set is exactly where every square vanishes at once; this replaces an inspection claim with an equation you can solve directly. **Link to the next block.** Locating the single zero of the quadratic factor tells you exactly which point is not on the line, which is the missing piece needed to pin down the triangle's apex.

Block 4. **Cue.** the curve splits into a line and one isolated point. **Why natural.** Three collinear points cannot form a triangle, and only one copy of the isolated point is available since the vertices are distinct, so the case analysis collapses to exactly one configuration: one vertex at the point, two on the line. **Link to the next block.** With the apex fixed and the opposite side confined to a fixed line, the triangle has no geometric freedom left except its size, which the equilateral condition also fixes.

Block 5. **Cue.** a fixed apex and a fixed supporting line for the opposite side. **Why natural.** The altitude from the apex of an isosceles or equilateral triangle bisects the base, so the foot of the altitude is forced to be the perpendicular projection of the apex onto the line; the altitude length is then just a distance computation, and the equilateral relation $h = \frac{\sqrt{3}}{2}s$ converts that one number into the side length and hence the area. **Link to the next block.** This is the last step: plug the computed height into the area formula.

Layer 3: WHY

The route is ideal because it exploits a general truth about degree bounds: a genuine irreducible cubic curve, viewed as a real algebraic curve, does not admit a distinguished finite set of inscribed equilateral triangles unless it is forced to by a rigid, low-dimensional configuration, and the only

way a cubic in two variables becomes that rigid is by factoring into lower degree pieces. The moment you see a uniqueness claim about a geometric configuration ON a curve of degree three or higher, the discernment mechanism is to ask whether the curve is secretly reducible before attempting any direct geometric attack, because reducibility is what converts a continuum of possibilities into a finite, checkable set.

You train this discernment by keeping a short mental stock of two and three variable identities that commonly hide inside contest cubics: the sum of cubes minus $3xyz$ used here, the difference of cubes, and sum-of-squares completions of quadratic forms. Seeing a coefficient of exactly 3 next to a cross product, sitting alongside matching cube terms, should fire the identity match before any other method is attempted, precisely because testing it costs one line and either succeeds spectacularly or fails harmlessly.

Revolving concepts

Symmetric factorisation $x^3+y^3+z^3-3xyz=(x+y+z)(x^2+y^2+z^2-xy-yz-zx)$**.** Provable by direct expansion, or as the determinant of the circulant matrix with rows the cyclic shifts of (x,y,z), whose eigenvalues at the eigenvectors $(1,\omega^i,\omega^{2i})$ for $i=0,1,2$ and ω a primitive cube root of unity are $x+\omega^i y+\omega^{2i}z$; the two complex conjugate eigenvalue factors for $i=1,2$ multiply to the real quadratic factor. See Book 1, field: Polynomials, topic: symmetric functions and circulant matrices. No recursion is present in this problem, so there is nothing to generalise along that axis here; the identity itself is the reusable object.

Sum-of-squares (SOS) certificates. Any positive semidefinite quadratic form in finitely many real variables can be written, after completing the square, as an explicit sum of squares, and its zero set is exactly the common zero set of every square in that decomposition. See Book 1, field: Inequalities, topic: sum of squares.

Equilateral triangle rigidity given one apex and the supporting line of the opposite side. If a triangle's apex A is fixed and its opposite side is required to lie on a fixed line not through A, requiring the triangle to be equilateral determines the other two vertices uniquely as an unordered pair: the altitude foot is the perpendicular projection of A, the altitude length is the fixed distance from A to the line, and the side length follows from $h=\frac{\sqrt{3}}{2}s$. See Book 1, field: Geometry, topic: equilateral and isosceles triangle relations.

Circuit Bank

- **WHEN** a two or three variable cubic has cube terms together with a $3xy$ or $3xyz$ style cross term **DO** test it against $x^3+y^3+z^3-3xyz=(x+y+z)(x^2+y^2+z^2-xy-yz-zx)$, fixing one variable at a constant if needed **BECAUSE** this is the single most common source of unexpected reducibility in a symmetric-looking cubic, and reducibility turns a curve-theoretic question into a line-and-point question

- **WHEN** a quadratic form is claimed or suspected to be positive semidefinite with an isolated zero **DO** complete the square into an explicit sum of squares **BECAUSE** the zero set of a sum of squares is visibly the simultaneous vanishing locus of every square, replacing an unjustified inspection claim with a solvable system

- **WHEN** a geometric configuration fixes one vertex and confines the opposite side to a fixed line **DO** reduce immediately to the single scalar distance from the vertex to the line **BECAUSE** for a prescribed triangle shape, that one scalar determines the entire triangle up to reflection, collapsing a two dimensional search to one number

Backtesting

1. First useful cue: "only one set of three points" attached to a cubic that otherwise looks like a generic curve.
2. Method the expression suggested: test the sum of cubes minus $3xyz$ identity with the constant term supplying the fixed third variable.
3. Where you would hesitate, and the thought to reinforce: confirming the quadratic factor really has a unique real zero rather than a whole conic of them; reinforce by writing it explicitly as a sum of squares rather than asserting positivity by inspection.
4. The trap the question sets: treating the cubic as one smooth irreducible curve and hunting for the triangle by a 120° rotation, which is possible in principle but costs far more ink and invites algebraic error.
5. What makes you recognise it faster next time: the exact coefficient pattern, cubes plus a coefficient 3 cross term, is now a stored shape rather than something to rediscover from scratch.
6. The reusable decision rule: symmetric cubic with cube terms and a 3-times cross product term, test the sum-of-cubes-minus-$3xyz$ identity before anything else.

Mirror drills

1. Factor $x^3 + y^3 - 3xy + 1$ and describe its real zero locus. *Answer:* set $z = 1$, matching $x^3 + y^3 + z^3 - 3xyz$ since $-3xyz = -3xy$; it factors as $(x + y + 1)(x^2 + y^2 + 1 - xy - x - y)$, and the second factor is $\frac{1}{2}[(x - y)^2 + (x - 1)^2 + (y - 1)^2]$, zero only at $(1, 1)$; the curve is the line $x + y = -1$ together with the point $(1, 1)$.
2. An equilateral triangle has one vertex fixed at the origin, with its opposite side required to lie on the line $x = 4$. Find its area. *Answer:* $h = 4$, so area $= h^2/\sqrt{3} = 16/\sqrt{3} = \dfrac{16\sqrt{3}}{3}$.
3. Show $2x^2 + 2y^2 + 2z^2 - 2xy - 2yz - 2zx$ is a sum of three squares, and find every real triple where it vanishes. *Answer:* it equals $(x - y)^2 + (y - z)^2 + (z - x)^2$, vanishing exactly when $x = y = z$.

20.8 B2. pigeonhole on fractional parts of partial sums

Prove that, for every set $X = \{x_1, x_2, \ldots, x_n\}$ of n real numbers, there exists a non-empty subset S of X and an integer m such that

$$\left| m + \sum_{s \in S} s \right| \leq \frac{1}{n+1}.$$

Read it out loud

An "there exists a subset and an integer" statement with a bound of exactly $\frac{1}{n+1}$ sitting against n given numbers is a fingerprint, not a coincidence. Whenever you see a bound of the shape $1/(k+1)$ attached to k objects, the honest first thought is pigeonhole: count out $k+1$ auxiliary quantities and place them on a line or a circle, because $k+1$ items forced into a space of total size 1 guarantees two of them within $1/(k+1)$ of each other. The examiners are not hiding this, they are handing it to you in the denominator.

The fields embedded are elementary number theory adjacent reasoning about fractional parts, and combinatorics in the form of the pigeonhole principle, applied on the continuous circle $\mathbb{R}/\mathbb{Z}$ rather than on a finite set of discrete boxes. What is really being tested is whether you can see that "subset S" secretly wants to be "a contiguous block of the list $x_1, \ldots, x_n$ in a fixed order," because only contiguous blocks come with the free telescoping identity that turns a difference of two auxiliary quantities into exactly the block sum you need.

Within the first thirty seconds you should have written down the $n+1$ partial sums $s_0 = 0, s_1 = x_1, s_2 = x_1 + x_2, \ldots, s_n = x_1 + \cdots + x_n$, because the count $n+1$ matches the denominator exactly, and the inclusion of the empty sum $s_0 = 0$ should already look deliberate: it is there to anchor the pigeonhole argument at a known point.

Approaches, ranked worst to best

1. **ad hoc existence.** *Cue:* "there exists". *Circuit:* try small numerical examples and guess a rounding argument, or hope for a slick continuity claim without any pigeonhole structure. (this gives you no constructive handle on which S and m to use, does not generalise past a single worked example, and is very likely to stall on a clean general proof)

2. **brute subset pigeonhole.** *Cue:* "subset S". *Circuit:* enumerate all 2^n subset sums, take their fractional parts, and pigeonhole among those. (this is not wrong in spirit, but the 2^n subset sums have no natural order and no automatic telescoping structure between any two of them, so recovering a specific S and a clean bound of exactly $1/(n+1)$ from an unstructured family of 2^n points is far more ink than the bound requires)

3. **prefix sum pigeonhole on the circle.** *Cue:* denominator $n+1$ against n numbers. *Circuit:* build the $n+1$ partial sums including the empty one, take fractional parts, sort them around $[0,1)$, and pigeonhole the $n+1$ gaps between consecutive values, which sum to exactly 1. (this is ideal because partial sums come with a built in telescoping identity, so the gap you find translates immediately into a genuine contiguous block S and an explicit integer m)

The ideal solution

For $i = 0, 1, \ldots, n$ define the partial sums $s_i = x_1 + x_2 + \cdots + x_i$, so $s_0 = 0$ is the empty sum. Write $\{t\} = t - \lfloor t \rfloor \in [0,1)$ for the fractional part of a real number t. Sort the $n+1$ values $\{s_0\}, \{s_1\}, \ldots, \{s_n\}$ into ascending order as $0 \leq t_0 \leq t_1 \leq \cdots \leq t_n < 1$; since $\{s_0\} = \{0\} = 0$ is among the values being sorted and no fractional part is negative, in fact $t_0 = 0$ exactly.

Consider the $n+1$ nonnegative real numbers

$$g_1 = t_1 - t_0, \quad g_2 = t_2 - t_1, \quad \ldots, \quad g_n = t_n - t_{n-1}, \quad g_{n+1} = 1 - t_n.$$

These telescope:

$$g_1 + g_2 + \cdots + g_{n+1} = (t_n - t_0) + (1 - t_n) = 1 - t_0 = 1,$$

using $t_0 = 0$. Since $n+1$ nonnegative reals sum to exactly 1, their average is $\frac{1}{n+1}$, so at least one of them satisfies $g_k \le \frac{1}{n+1}$; fix such an index $k \in \{1, \ldots, n+1\}$.

Case A: $1 \le k \le n$**.** Here $g_k = t_k - t_{k-1}$ is an internal gap. Since the $n+1$ sorted values $t_0, \ldots, t_n$ arise from a permutation of the $n+1$ indices $0, \ldots, n$, the positions $k-1$ and k correspond to two distinct indices $p \ne q$ in $\{0, \ldots, n\}$, with $t_{k-1} = \{s_p\}$ and $t_k = \{s_q\}$. Let $i = \min(p, q)$ and $j = \max(p, q)$, so $0 \le i < j \le n$. Put

$$S = \{x_{i+1}, x_{i+2}, \ldots, x_j\} \quad \text{(nonempty, since } i < j\text{)}, \qquad m = \lfloor s_i \rfloor - \lfloor s_j \rfloor.$$

Then

$$m + \sum_{s \in S} s = m + (s_j - s_i) = \big(s_j - \lfloor s_j \rfloor\big) - \big(s_i - \lfloor s_i \rfloor\big) = \{s_j\} - \{s_i\},$$

which equals $\pm(t_k - t_{k-1})$ depending on whether $j = q$ or $j = p$. Either way,

$$\left| m + \sum_{s \in S} s \right| = g_k \le \frac{1}{n+1}.$$

Case B: $k = n+1$**.** Here the minimal gap is the wraparound gap $g_{n+1} = 1 - t_n \le \frac{1}{n+1}$. Since $n \ge 1$, $\frac{1}{n+1} < 1$, so $t_n = 1 - g_{n+1} > 0$, meaning s_n is not an integer and $\lceil s_n \rceil = \lfloor s_n \rfloor + 1$. Take $S = \{x_1, \ldots, x_n\} = X$ (nonempty) and $m = -\lceil s_n \rceil$. Then

$$m + s_n = s_n - \lceil s_n \rceil = \big(s_n - \lfloor s_n \rfloor\big) - 1 = t_n - 1,$$

so $|m + s_n| = 1 - t_n = g_{n+1} \le \frac{1}{n+1}$.

In either case a nonempty $S \subseteq X$ and an integer m have been produced with $|m + \sum_{s \in S} s| \le \frac{1}{n+1}$. ■

Layer 1: WHAT

Block 1: build the $n+1$ partial sums, deliberately including the empty sum $s_0 = 0$. Block 2: pass to fractional parts, sort them on $[0, 1)$, and note $t_0 = 0$ makes the $n+1$ gaps between consecutive sorted values, including the wraparound gap, telescope to exactly 1. Block 3: average the $n+1$ nonnegative gaps to get one gap of length at most $\frac{1}{n+1}$. Block 4: if that gap is internal, translate it into a contiguous block S of the original list and an integer m built from the two relevant floors. Block 5: if that gap is the wraparound one, take S to be the whole list and m the negative ceiling of the full sum.

Nominal route: enumerate all 2^n subset sums and pigeonhole among those, or attempt an unstructured existence argument. Ideal route: restrict to the $n+1$ ordered partial sums. The bifurcation point is the choice of WHICH auxiliary family to pigeonhole: an unstructured family of 2^n subset sums with no natural differences to exploit, against a structured family of $n+1$ partial sums whose pairwise differences are automatically contiguous block sums.

Layer 2: HOW

Block 1. Cue. a bound of $\frac{1}{n+1}$ next to a set of n numbers. **Why natural.** the denominator $n+1$ is the exact count of auxiliary objects a pigeonhole argument on n items typically needs, and including the empty sum is what anchors the sorted list at a known value. **Link to the next block.** once you have $n+1$ real numbers, the natural next move is to place them on some bounded space and measure gaps.

Block 2. **Cue.** the auxiliary numbers are fractional parts, which live in $[0,1)$, a space of total length 1. **Why natural.** sorting them and taking consecutive differences, plus the wraparound difference back to 1, is the only way to make $n+1$ numbers in $[0,1)$ account for the entire length 1 without leftover. **Link to the next block.** this telescoping identity is exactly what lets you invoke averaging.

Block 3. **Cue.** $n+1$ nonnegative numbers summing to a fixed total. **Why natural.** the minimum of any collection of nonnegative reals cannot exceed their average, so at least one gap is at most $\frac{1}{n+1}$. **Link to the next block.** the found gap corresponds to a pair of the original partial sums, and the two cases (internal gap or wraparound gap) must be handled separately because only internal gaps come with two genuine indices already present in the list.

Block 4. **Cue.** an internal gap between two sorted fractional parts. **Why natural.** the identity $\{s_j\} - \{s_i\} = (s_j - s_i) - (\lfloor s_j \rfloor - \lfloor s_i \rfloor)$ converts a difference of fractional parts directly into a contiguous block sum plus an integer correction. **Link to the next block.** choosing $m = \lfloor s_i \rfloor - \lfloor s_j \rfloor$ makes that correction explicit and finishes the case.

Block 5. **Cue.** the wraparound gap, where no second partial sum index is available past s_n. **Why natural.** the identity here is between s_n and the nearest integer above it, since $g_{n+1} \le \frac{1}{n+1} < 1$ guarantees s_n is not itself an integer, so the ceiling is well defined. **Link to the next block.** this closes the only remaining case, completing the proof.

Layer 3: WHY

The underlying truth is a form of Dirichlet's pigeonhole principle on the circle group $\mathbb{R}/\mathbb{Z}$: any $n+1$ points placed on a circle of circumference 1 leave $n+1$ arcs between them that sum to exactly the circumference, so some arc is at most $\frac{1}{n+1}$ long. This is the same mechanism behind Dirichlet's classical approximation theorem, where the $n+1$ points are the fractional parts of $0, \alpha, 2\alpha, \ldots, n\alpha$ for a single real α. Here the points are instead the fractional parts of the partial sums of a whole list of numbers, and the reason partial sums, rather than arbitrary subset sums, are the right auxiliary family is that the DIFFERENCE of any two partial sums is automatically a contiguous block sum, which is exactly the kind of subset sum the problem allows.

The discernment mechanism to train is this: whenever a bound of the shape $\frac{1}{n+1}$ (or more generally $\frac{1}{k+1}$) sits next to a set of n (or k) objects, count out $k+1$ auxiliary quantities to pigeonhole, and ask "$k+1$ of WHAT." The "+1" is almost never accidental; it is the fingerprint of needing one extra, often trivial or empty, term to close a telescoping identity exactly at the boundary. Drilling this pattern across problems, rather than rediscovering it from scratch each time, is what makes this recognition instantaneous in the arena.

Revolving concepts

Dirichlet's approximation theorem. For any real α and positive integer N, there exist integers p, q with $1 \le q \le N$ and $|q\alpha - p| < \frac{1}{N}$, obtained by pigeonholing the $N+1$ fractional parts $\{0\}, \{\alpha\}, \ldots, \{N\alpha\}$ either into N bins of width $\frac{1}{N}$ or, as here, by sorting and taking the smallest gap. See Book 1, field: Number theory, topic: Dirichlet approximation.

Prefix sum trick for contiguous block existence. For a sequence $x_1, \ldots, x_n$, define $s_0 = 0$, $s_i = s_{i-1} + x_i$; any contiguous block sum $x_{i+1} + \cdots + x_j$ equals $s_j - s_i$, so a question about the existence of a contiguous block with a numeric property reduces to a question about a pair of partial sums, usually attacked via pigeonhole on residues in the discrete case or fractional parts in the continuous case. *Recursion and its generalisation:* the defining recursion is $s_i = s_{i-1} + x_i$, $s_0 = 0$; its general solution method is telescoping, $s_j - s_i = x_{i+1} + \cdots + x_j$, and this generalises to any associative combination, sums, exclusive-ors, products, wherever a problem asks for the

existence of a contiguous block satisfying a property, by pigeonholing the $n+1$ partial values under whichever equivalence (equal residue, close fractional part) the property demands. See Book 1, field: Sequences and recurrences, topic: prefix sum pigeonhole.

Averaging as a weak pigeonhole. If nonnegative reals $g_1, \ldots, g_k$ sum to a fixed total T, then $\min_i g_i \le T/k$; this is the continuous cousin of the discrete pigeonhole statement that $k+1$ items into k boxes force a repeat, and it is the exact tool that extracts "some gap is small" from "the gaps sum to 1." See Book 1, field: Inequalities, topic: averaging and extremal principle.

Fractional part and floor or ceiling identities. $\{t\} = t - \lfloor t \rfloor \in [0,1)$; for non integer t, $\lceil t \rceil = \lfloor t \rfloor + 1$; and $\{a\} - \{b\} = (a-b) - (\lfloor a \rfloor - \lfloor b \rfloor)$, the identity that converts a difference of fractional parts into an integer adjusted difference of the original reals. See Book 1, field: Real analysis, topic: floor and fractional part manipulation.

Circuit Bank

- **WHEN** a claimed bound has denominator $n+1$ attached to a set of n objects **DO** build $n+1$ auxiliary quantities, such as partial sums including a trivial empty one, and pigeonhole them **BECAUSE** the "+1" is almost always the exact count of auxiliary objects the pigeonhole argument needs, and the trivial term is what closes the telescoping identity at the boundary
- **WHEN** you need "some subset has property P" where P is really a numeric proximity claim **DO** restrict to contiguous blocks via prefix sums rather than all 2^n subsets **BECAUSE** differences of prefix sums are exactly contiguous block sums, giving a structured, telescoping family instead of an unstructured one
- **WHEN** you have k nonnegative reals known to sum to a fixed total **DO** immediately conclude the minimum one is at most the total divided by k **BECAUSE** this averaging step is the cheapest possible pigeonhole and is very often the entire existence content of the problem
- **WHEN** sorting real numbers into a cyclic order via fractional parts on $[0,1)$ **DO** deliberately include the value 0 itself among the numbers being sorted **BECAUSE** it anchors the smallest sorted value at 0 and makes the wraparound gap close the telescoping sum to exactly the full circumference rather than to an unknown leftover

Backtesting

1. First useful cue: a denominator of $n+1$ sitting next to a set of n numbers.
2. Method the expression suggested: pigeonhole the $n+1$ partial sums, including the empty one, by fractional part around the unit circle.
3. Where you would hesitate, and the thought to reinforce: justifying that the wraparound case still yields a genuine nonempty S and a well defined integer m; reinforce by checking $g_{n+1} \le \frac{1}{n+1} < 1$ forces s_n to be non integer, so the ceiling step is legitimate.
4. The trap the question sets: enumerating all 2^n subset sums instead of the $n+1$ structured partial sums, which loses the automatic contiguous block structure and makes recovering a specific S and m far harder than necessary.
5. What makes you recognise it faster next time: the phrase "there exists a nonempty subset and an integer" next to a $\frac{1}{n+1}$ bound is now a stored trigger for prefix sum pigeonhole on a circle.
6. The reusable decision rule: n numbers, bound $\frac{1}{n+1}$, build $n+1$ prefix sums, pigeonhole fractional parts on the circle, read the contiguous block off the smallest gap.

Mirror drills

1. Show that among any 5 integers, some nonempty subset has sum divisible by 5. *Answer:* pigeonhole the 6 partial sums $s_0, \dots, s_5$ modulo 5; two share a residue, and their difference is a contiguous run divisible by 5.
2. For any 4 reals, exhibit the guarantee from the ideal solution explicitly. *Answer:* it is the case $n = 4$ of the theorem just proved, giving a nonempty subset S and integer m with $|m + \sum_{s \in S} s| \le \frac{1}{5}$.
3. Prove Dirichlet's approximation theorem: for irrational α and positive integer N, some $q \in \{1, \dots, N\}$ has $\|q\alpha\| \le \frac{1}{N+1}$, where $\|\cdot\|$ denotes distance to the nearest integer. *Answer:* apply the ideal solution with $x_i = \alpha$ for every i and $n = N$, so the partial sums are $s_i = i\alpha$, and the produced $S = \{x_1, \dots, x_q\}$ for some q gives $|m + q\alpha| \le \frac{1}{N+1}$.

20.9 B3. counting linear partitions of a planar point set

Let S be a finite set of points in the plane. A linear partition of S is an unordered pair $\{A, B\}$ of subsets of S such that $A \cup B = S$, $A \cap B = \emptyset$, and A and B lie on opposite sides of some straight line disjoint from S (A or B may be empty). Let L_S be the number of linear partitions of S. For each positive integer n, find the maximum of L_S over all sets S of n points.

Read it out loud

"Find the maximum of L_S over all sets S of n points" is an extremal combinatorics question wearing a geometry costume. Extremal questions of this exact shape, maximise a count over all configurations of size n, split cleanly into two independent jobs: an upper bound valid for every configuration, and a matching construction that actually attains it. You should read the problem already planning to write two separate arguments, not one.

The fields embedded are combinatorial geometry, in particular the behaviour of a separating line as it varies continuously, and induction. What is really being tested is whether you can set up a controlled removal, take away one point and ask how much the count can possibly grow back when it is reinstated, since a direct global count over an arbitrary configuration of n points has no natural handle to grab.

The first cue worth writing down within thirty seconds is the phrase "A or B may be empty," which tells you the trivial partition, everything on one side, always exists, so $L_S \geq 1$ for every S; and the phrase "opposite sides of some straight line" tells you every other partition is indexed by a direction or position of a line, a continuously varying object, which is exactly the kind of structure a rotating sweep argument is built for.

Approaches, ranked worst to best

1. **direct arrangement count.** *Cue:* the target $\binom{n}{2} + 1$ resembles the classical regions-cut-by-lines formula. *Circuit:* draw all $\binom{n}{2}$ lines through pairs of points and count the regions they cut the plane into. (the target number is suggestively similar, but L_S counts partitions by an ARBITRARY separating line, not regions cut specifically by those $\binom{n}{2}$ lines through pairs of points; turning one count into the other needs a bijection this approach does not supply for free, and chasing it directly is a harder dual-arrangement computation than the problem requires)

2. **single-point rotation only.** *Cue:* a line can be rotated about one fixed point. *Circuit:* fix one point P, rotate a line through it by 180°, and count how the partition of the rest changes. (this correctly captures part of the mechanism and will be reused, but on its own it only counts partitions realised via lines through that one fixed point, undercounting the true maximum; it needs to be wired into an induction as a lemma, not used as the whole argument)

3. **removal induction with a rotating-line lemma, plus a convex position construction.** *Cue:* "maximum over all sets of n points" splits into a bound and a construction. *Circuit:* remove one point P, relate linear partitions of S to those of $S \setminus \{P\}$ via a restriction map, bound how many outputs of that map have two preimages using the single-point rotation from the previous approach, close with Pascal's identity, then separately exhibit a convex n-gon attaining equality via a bijection with pairs of edges. (this is ideal because the bound and the construction are proved by two independently clean arguments, each using the tool suited to it)

The ideal solution

The claim is $\max_S L_S = \binom{n}{2} + 1$.

Upper bound, by induction on n. For $n = 1$, a single point has exactly one linear partition, the trivial one, since any line disjoint from S leaves the point on one side; $L_S = 1 = \binom{1}{2} + 1$.

Suppose the bound $L_{S'} \le \binom{n-1}{2} + 1$ holds for every set S' of $n-1$ points, where $n \ge 2$, and let S be a set of n points. Fix any $P \in S$ and put $S' = S \setminus \{P\}$. Every line disjoint from S is also disjoint from S', so it induces a linear partition of S' as well as of S; this defines a restriction map f from linear partitions of S to linear partitions of S', by deleting P from whichever side it fell on.

Fix a partition $\pi' = \{A', B'\}$ of S' and let $\mathcal{F}(\pi')$ be the (nonempty) family of all lines disjoint from S' that induce π', keeping A' on a fixed side and B' on the other. As a line varies continuously within $\mathcal{F}(\pi')$, by translating or rotating while staying disjoint from S' and preserving which side A' and B' are on, the signed distance from P to the line varies continuously.

If every line in $\mathcal{F}(\pi')$ places P on the same side, then every S-partition restricting to π' places P on that same side too, so $f^{-1}(\pi')$ has at most one element. If instead two lines of $\mathcal{F}(\pi')$ place P on opposite sides, the signed distance from P to the line changes sign somewhere along a continuous path within $\mathcal{F}(\pi')$ joining them, so by the intermediate value theorem some line in this family passes exactly through P; call such a π' special. Only for special π' can $f^{-1}(\pi')$ have two elements, the two extensions obtained by attaching P to either side, so in every case $|f^{-1}(\pi')| \le 2$.

Now bound the number of special partitions. Parametrise lines through P by direction $\theta \in [0, \pi)$, a line and its reverse direction being the same line. For each of the $n-1$ points $Q \in S'$, there is exactly one direction, namely that of the segment PQ, at which the line through P passes through Q; at every other direction the line through P is disjoint from S' and induces some partition of S', and this partition can only change as θ crosses one of these $n-1$ critical directions. Hence at most $n-1$ distinct partitions of S' are induced by lines through P, so there are at most $n-1$ special partitions.

Combining these bounds,

$$L_S = \sum_{\pi' \in \mathrm{image}(f)} |f^{-1}(\pi')| \le L_{S'} + (\text{number of special } \pi' \text{ in the image}) \le \binom{n-1}{2} + 1 + (n-1),$$

since every non special π' contributes exactly its own count of 1 toward $L_{S'}$, and every special one contributes at most one unit more than that. By Pascal's identity, $\binom{n-1}{2} + (n-1) = \binom{n}{2}$, so $L_S \le \binom{n}{2} + 1$. This closes the induction, so $L_S \le \binom{n}{2} + 1$ for every set S of n points.

A construction attaining equality. For $n \ge 3$, let S be the vertices $v_1, \dots, v_n$ of a convex n-gon in cyclic order, with edges $e_1 = v_1v_2, \dots, e_n = v_nv_1$. Any line disjoint from S either misses the convex hull entirely, giving the trivial partition, or crosses the boundary at two points interior to two distinct edges e_i, e_j (it cannot pass through a vertex, since it is disjoint from S), splitting the vertices into the two arcs of the polygon boundary strictly between e_i and e_j on each side. Conversely, for any pair of distinct edges e_i, e_j, a line can be drawn crossing exactly those two edges, close to a segment joining interior points of each, inducing the partition into the two arcs they bound; distinct pairs of edges give distinct arc splits, since the pair of edges crossed is recoverable from the split as the two edges each arc exits through. So the nontrivial linear partitions of S are in exact bijection with the $\binom{n}{2}$ unordered pairs of distinct edges, giving $L_S = \binom{n}{2} + 1$, the extra 1 for the trivial partition. For $n = 1, 2$ this value is checked directly: $L_S = 1 = \binom{1}{2} + 1$ for a single point, and $L_S = 2 = \binom{2}{2} + 1$ for two points, one trivial partition and one separating them.

Since the upper bound is attained, $\max_S L_S = \boxed{\binom{n}{2} + 1}$.

Layer 1: WHAT

Block 1: base case $n = 1$. Block 2: build the restriction map f from linear partitions of S to linear partitions of $S \setminus \{P\}$. Block 3: show every fibre of f has size at most 2, and size 2 is only possible for

partitions achieved by some line through P, using continuity and the intermediate value theorem. Block 4: bound the number of such special partitions by $n-1$, by rotating a line through P and counting critical directions. Block 5: combine via Pascal's identity to close the induction on the upper bound. Block 6: separately construct a convex n-gon and biject its nontrivial partitions with pairs of edges to show the bound is attained.

Nominal route: attempt to count L_S directly via the arrangement of all $\binom{n}{2}$ pairwise-point lines and the standard regions-cut-by-lines formula. Ideal route: compare S to the $(n-1)$-point problem through a restriction map with a provably small correction term, then separately verify tightness. The bifurcation point is the decision of whether to attack the n-point configuration directly, which has no natural recursive handle since a general point set is not in special position, or to compare it to the $(n-1)$-point configuration and control exactly how much one extra point can add.

Layer 2: HOW

Block 1. **Cue.** "maximum of L_S over all sets S of n points." **Why natural.** an extremal count over configurations rarely has a clean closed form from a direct global count; it almost always wants an inductive comparison to the $(n-1)$ case. **Link to the next block.** the natural way to compare to the $(n-1)$ case is to remove one point and ask how partitions restrict.

Block 2. **Cue.** removing a point P from S. **Why natural.** every separating line for S is automatically a separating line for the smaller set $S \setminus \{P\}$, so restriction is a well defined map with no extra work. **Link to the next block.** the whole upper bound now rests on understanding how many partitions of S can restrict to the same partition of $S \setminus \{P\}$.

Block 3. **Cue.** a partition of S' might be achieved by several different separating lines. **Why natural.** the family of lines realising a fixed partition of S' is a connected set, since any two such lines can be joined by a continuous deformation staying within that family, so the signed distance from P to the line varies continuously along it; a continuous quantity that takes both signs must vanish somewhere, by the intermediate value theorem, which is exactly a line through P. **Link to the next block.** this shows the only partitions with more than one preimage are those achieved by some line through P, converting an open ended geometric worry into a bounded count.

Block 4. **Cue.** lines through the single fixed point P. **Why natural.** rotating a line through P by a full 180° only changes which side each other point of S' falls on at one critical direction per point, since a point's side flips exactly when the line's direction matches the direction from P to that point. **Link to the next block.** counting critical directions, rather than counting partitions directly, bounds the number of special partitions by $n-1$, the number of other points.

Block 5. **Cue.** an inductive bound of the shape $\binom{n-1}{2}+1+(n-1)$. **Why natural.** Pascal's identity $\binom{n-1}{2}+(n-1)=\binom{n}{2}$ collapses this immediately to $\binom{n}{2}+1$, confirming the induction is exactly tight rather than merely an upper estimate. **Link to the next block.** having proved the universal bound, a separate matching construction is needed to know the bound is not wasteful.

Block 6. **Cue.** convex position, and the phrase "A or B may be empty" recalling the trivial partition. **Why natural.** placing the points in convex position makes every pair of boundary edges correspond to a genuinely distinct achievable separation, with no coincidences from interior points, so counting edge pairs directly counts nontrivial partitions. **Link to the next block.** this closes the problem, since the bound and the construction now agree exactly.

Layer 3: WHY

The underlying truth is that a linear partition is indexed, many-to-one, by the direction and position of a separating line, and the space of lines disjoint from a fixed finite point set, modulo "induces the same partition," is a disjoint union of open connected cells. Connectivity within a cell is precisely what licenses the intermediate value theorem step: if a fixed point P sits on different sides for two lines in the same cell, some line in that cell must pass through P. Whenever an extremal count is asked for objects induced by a continuously varying separator, a line, a hyperplane, a threshold, the discernment mechanism is to look first for a removal induction with a bounded correction term, because a direct global count over a continuum of separators rarely closes without first discretising the problem this way.

You train this by explicitly writing the restriction map for a candidate induction and asking "what is the largest number of preimages any single output can have," since that number is exactly the correction term the induction needs; and by keeping "convex or general position is usually the extremiser" as a standing first guess whenever the construction half of an extremal geometry problem is needed, since it is confirmed correct far more often than not for this style of problem.

Revolving concepts

Extremal induction via a restriction map with bounded fibre size. Remove one generator from an n-object configuration, relate the count to the $(n-1)$-object case through a restriction map, bound the fibre size of that map, here at most 2, and bound how many outputs actually achieve the larger fibre size, here at most $n-1$ via a geometric argument, then add the correction term, often closing via a combinatorial identity such as Pascal's rule. See Book 1, field: Combinatorics, topic: incremental extremal counting.

Sweeping or rotating line through a fixed point. Rotating a line about a fixed point through a finite point set changes the induced combinatorial partition only at finitely many critical directions, one per other point, giving an immediate upper bound on the number of distinct partitions realised. See Book 1, field: Discrete structures and graphs, topic: line arrangements and sweep arguments.

Maximum regions cut by n lines. n lines in general position, no two parallel, no three concurrent, divide the plane into $\binom{n}{2}+n+1$ regions. *Recursion and its generalisation:* adding the k-th line to $k-1$ previously placed lines crosses them in $k-1$ distinct points, splitting the new line into k pieces and so adding k new regions, giving the recursion $R(k) = R(k-1)+k$ with $R(0)=1$, solved by summation to $R(n) = 1+\sum_{k=1}^{n} k = \binom{n}{2}+n+1$; the same additive recursion, plus summation, pattern applies whenever a new generic object crosses all previous ones in distinct points. This closely related shape is why counts like $\binom{n}{2}+1$ recur across discrete geometry extremal problems. See Book 1, field: Discrete structures and graphs, topic: arrangements of lines.

Convex position as the extremiser for linear-separation counting. For problems asking for the maximum number of distinct ways a straight line can separate a finite point set, convex, general position, configurations typically attain the extremal count, since every pair of boundary edges or points then yields a genuinely distinct achievable separation, free of coincidences from interior points. See Book 1, field: Geometry, topic: convex position extremal configurations.

Circuit Bank

- **WHEN** an extremal problem counts combinatorial objects induced by a continuously varying separator, a line, a hyperplane, a threshold **DO** remove one generator and build the restriction map to the $(n-1)$ problem, bounding how many outputs have two or more preimages **BECAUSE** this converts an unbounded looking global count into an induction with an explicitly boundable correction term
- **WHEN** you need to bound how many distinct combinatorial states are visible from lines or

rays through one fixed point **DO** rotate through the full range of directions and count only the critical directions where the state can change **BECAUSE** a sweeping parameter changes the combinatorics only at finitely many critical values, one per other object, and counting those values is always easier than counting states directly

- **WHEN** an inductive bound produces an expression of the shape $\binom{n-1}{2} + (n-1)$ **DO** simplify immediately via Pascal's identity to $\binom{n}{2}$ **BECAUSE** recognising this pattern instantly confirms the induction is tight and closes the algebra in one line
- **WHEN** you need to prove an upper bound is attained, not merely valid **DO** reach for the convex or general position construction first **BECAUSE** convex configurations are almost always the extremiser for "number of distinct linear separations" style problems, and the bijection with pairs of boundary features is usually immediate

Backtesting

1. First useful cue: "find the maximum of L_S over all sets S of n points," an extremal-over-configurations phrasing.
2. Method the expression suggested: a removal induction with a rotating-line lemma to bound the correction term.
3. Where you would hesitate, and the thought to reinforce: justifying that a partition can have at most 2 preimages under restriction, and that only special, through-P, partitions achieve 2; reinforce with the explicit intermediate value theorem argument on the signed distance from P to a continuously varying line.
4. The trap the question sets: counting L_S directly via the arrangement of all $\binom{n}{2}$ point-pairwise lines using the standard regions formula, which resembles the target closely but is not counting the same object without an extra justification step.
5. What makes you recognise it faster next time: the shape "$\binom{n}{2} + 1$" is now stored as "one trivial case plus one case per pair of boundary features," so it should immediately suggest an edge-pair or point-pair bijection for the construction half, and a remove-one-point induction for the bound half.
6. The reusable decision rule: extremal count of line-induced separations, split into a removal induction with a rotating-line correction bound, and a convex position construction with an edge-pair bijection.

Mirror drills

1. n lines in general position, no two parallel, no three concurrent, divide the plane into how many regions? *Answer:* $\binom{n}{2} + n + 1$, by the same style of induction, the k-th line adds k new regions.
2. How many nontrivial linear partitions of a convex n-gon isolate exactly one vertex? *Answer:* n, one for each vertex, corresponding to the n pairs of adjacent edges.
3. If S has n points, no three collinear, and $L_S = \binom{n}{2} + 1$, must S be in convex position? *Hint:* yes; if a point lay strictly inside the hull of the others, some pair of edges could no longer be simultaneously crossed by a line without ambiguity about that interior point's side, collapsing two of the $\binom{n}{2}$ candidate partitions into one, so equality forces convex position.

20.10 B4. Counting cube vertices inside a subspace

Let Z denote the set of points in $\mathbb{R}^n$ whose coordinates are 0 or 1. (Thus Z has 2^n elements, which are the vertices of a unit hypercube in $\mathbb{R}^n$.) Given a vector subspace V of $\mathbb{R}^n$, let $Z(V)$ denote the number of members of Z that lie in V. Let k be given, $0 \leq k \leq n$. Find the maximum, over all vector subspaces $V \subseteq \mathbb{R}^n$ of dimension k, of the number of points in $V \cap Z$.

Read it out loud

This is a counting question wearing a linear algebra coat. A finite set, the 2^n vertices of the unit cube, and a continuous object, a k dimensional subspace: how much of the finite set can the continuous object swallow? The fields wired in are linear algebra (dimension, rank, spanning sets) and finite combinatorics (counting points over a two letter alphabet). There is no metric geometry here. Any line you write involving a length or an angle is wasted ink.

Now read the parameters, because they are the loudest guiding light on this paper. You are given n and k with $0 \leq k \leq n$. Does the answer depend on both? Test $k = n$: the count is 2^n. Test $k = 0$: the count is 1. Test $k = 1$ in $\mathbb{R}^2$: the line $y = x$ takes $(0,0)$ and $(1,1)$, and no line through the origin takes three cube vertices, since the origin plus two more would already force dimension 2. So the data reads $1, 2, 4, \ldots$, that is 2^k, with no trace of n. That absence is the examiners telling you where the bound comes from. A bound that cannot see n cannot depend on how the cube sits inside $\mathbb{R}^n$, so it must come from dimension alone, and your proof should never mention n.

Within thirty seconds write two things. The construction: $V_0 = \{x : x_1 = \cdots = x_{n-k} = 0\}$ is k dimensional and meets Z in the 2^k vertices supported on the last k coordinates. And the remaining task: every k dimensional subspace holds at most 2^k cube vertices. On a "find the maximum" the construction is worth almost nothing and the bound is worth almost everything. Then say the sentence that contains the entire solution: a cube vertex is a point each of whose coordinates takes one of two values, and dimension k says only k coordinates are really free. Two values, k free coordinates, 2^k points. Your only job is to turn "only k coordinates are really free" into a theorem, and that theorem already exists: row rank equals column rank.

Approaches, ranked worst to best

1. **Tabulate small cases and assert the pattern.** *Cue:* two free parameters and the words "find the maximum". *Circuit:* compute, guess, hope. You reach 2^k in two minutes and prove nothing. A construction with no upper bound is 1 out of 10. Reconnaissance only.

2. **Chase dependences among the cube vertices.** *Cue:* "dimension k" reads as "any $k+1$ of them are dependent". *Circuit:* Steinitz exchange, maximal independent subsets. Pick independent cube vertices $v_1, \ldots, v_k$ spanning V and try to force every other cube vertex in V to be one of 2^k combinations. It stalls immediately, and you should be able to say why: the coefficients are not restricted. In $\mathbb{R}^3$, $(1,1,1) = (1,1,0) + (0,0,1)$ but also $(1,0,1) + (0,1,1) - (0,0,1)$. This route never uses the alphabet, and the alphabet is where the problem lives.

3. **Slice the cube by a coordinate hyperplane and induct.** *Cue:* $Z = \{0,1\} \times \{0,1\}^{n-1}$ is a product. *Circuit:* split into $x_n = 0$ and $x_n = 1$, induct on $n + k$. Generically each slice drops the dimension by one and contributes 2^{k-1}. Correct, and it proves the stronger affine statement, but it costs three pieces of bookkeeping: you must strengthen the claim to affine planes before the induction closes, since the slice $x_n = 1$ is affine; you must handle the degenerate branch where V lies inside one slice; and you must induct on $n + k$ rather than on either alone.

4. **Reduce modulo 2 and count inside $\mathbb{F}_2^n$.** *Cue:* the coordinates are 0 or 1, exactly the elements of $\mathbb{F}_2$. *Circuit:* integer dependences survive reduction, and an odd minor lifts. Reduction is injective on Z. If the image held $k+1$ vectors independent over $\mathbb{F}_2$, the integer matrix they lift

to would have an odd, hence nonzero, maximal minor, so those $k+1$ cube vertices would be independent over $\mathbb{R}$ and $\dim V > k$. So the image lies in an $\mathbb{F}_2$ subspace of dimension at most k, which has $2^{\dim}$ elements. Short and structurally honest: it explains why the answer is a power of two. It ranks below the next route only because it spends machinery you can avoid.

5. **Take the rank, then choose k coordinates that determine the point.** *Cue:* "dimension k" sitting next to "each coordinate takes two values". *Circuit:* row rank equals column rank, then count coordinatewise. Stack the points of $V \cap Z$ as rows. The row space lies in V, so the rank is at most k, so some k columns span all the columns, so any two rows agreeing on those k coordinates agree everywhere. Each of those coordinates takes two values, so there are at most 2^k rows. Four sentences, no induction, no finite fields, and no mention of n.

The ideal solution

The maximum is $\boxed{2^k}$.
Construction. Let $V_0 = \{x \in \mathbb{R}^n : x_1 = \cdots = x_{n-k} = 0\}$, the span of the last k standard basis vectors, so $\dim V_0 = k$. A point of Z lies in V_0 exactly when its first $n-k$ coordinates vanish and its last k are arbitrary elements of $\{0,1\}$, so $|V_0 \cap Z| = 2^k$.
Bound. Everything rests on one fact about matrices.

Lemma 20.10.1. *Let M be a real matrix of rank r. Then there is a set J of r column indices such that any two rows of M agreeing in every coordinate belonging to J are equal.*

Proof. The rank is the dimension of the column space, so choose J with $|J| = r$ such that the columns C_j, $j \in J$, form a basis of the column space. Every column satisfies $C_i = \sum_{j\in J} \lambda_{ij} C_j$ for real scalars λ_{ij}. If rows s and t agree on J, that is $(C_j)_s = (C_j)_t$ for all $j \in J$, then for every i,

$$(C_i)_s = \sum_{j\in J} \lambda_{ij}(C_j)_s = \sum_{j\in J} \lambda_{ij}(C_j)_t = (C_i)_t,$$

so the two rows agree in every coordinate and are equal. □

Let $\dim V = k$. Since Z is finite, so is $V \cap Z$; let M be the matrix whose rows are its distinct elements in any order. Every row lies in V, so the row space of M is a subspace of V and

$$r := \operatorname{rank} M = \dim(\text{row space of } M) \le \dim V = k.$$

Take J from the lemma. Then the restriction map

$$\pi_J : V \cap Z \to \{0,1\}^J, \qquad x \mapsto (x_j)_{j\in J}$$

is injective: two points with the same image are two rows agreeing on J, hence the same row. The codomain has exactly 2^r elements, since each coordinate of a point of Z is 0 or 1. Therefore $|V \cap Z| \le 2^r \le 2^k$.

The two halves meet, so the maximum is exactly 2^k. Edge cases behave: for $k = 0$ we get $V = \{0\}$, M has the single row 0, $r = 0$, $J = \emptyset$, and the bound reads 1; for $k = n$ it reads 2^n. The matrix always has at least one row, since $0 \in V \cap Z$ for every subspace, so no empty case arises.

Remark 20.10.2. The same proof bounds $|A \cap Z|$ for an *affine* k plane A. If the intersection is empty there is nothing to prove; otherwise fix $p \in A \cap Z$ and let M have rows $q - p$ for $q \in A \cap Z$. These lie in the k dimensional direction space of A, so $\operatorname{rank} M \le k$, and injectivity of $(q-p) \mapsto (q-p)|_J$ with p fixed gives injectivity of $q \mapsto q|_J \in \{0,1\}^J$.

Remark 20.10.3. Coordinate subspaces are not the only maximisers. Partition part of $\{1,\dots,n\}$ into nonempty blocks $B_1,\dots,B_k$ and let V be the span of their indicator vectors; then $\sum c_i \mathbf{1}_{B_i} \in Z$ exactly when every $c_i \in \{0,1\}$, so $|V \cap Z| = 2^k$. In $\mathbb{R}^3$ the span of $(1,1,0)$ and $(0,0,1)$ holds four cube vertices, while the span of $(1,1,0)$ and $(0,1,1)$ holds only three, namely $(0,0,0), (1,1,0), (0,1,1)$.

Layer 1: WHAT

Block 1. Read the parameters: the answer is 2^k with no n in it, so look for a proof that never uses n.

Block 2. Write the coordinate subspace, count 2^k, and declare the construction finished.

Block 3. Name what makes a cube vertex special: every coordinate takes one of two values. So the count will be a power of 2 whose exponent counts coordinates.

Block 4. Convert "dimension k" into "at most k coordinates matter" by stacking the points into a matrix and using row rank equals column rank.

Block 5. Choose k spanning columns, prove the restriction is injective, count 2^k.

Nominal against ideal. The nominal route keeps the points of $V \cap Z$ as vectors and tries to control them through a basis of V, bounding the count by bounding the coefficients. The coefficients are not bounded, so the route does not merely run slowly, it controls the wrong data.

The bifurcation point is Block 4 and it is one word long: transpose. Up to Block 3 the two routes agree. At Block 4 the nominal route studies rows and their coefficients; the ideal route studies columns and their spans. Rows are points, columns are coordinates, and the alphabet restriction is a statement about coordinates. Decide to look at columns and the problem is over.

Layer 2: HOW

Block 1. **Cue.** Two parameters, n and k, and a request for a maximum. **Why natural.** Test the extremes before computing anything: $k = 0$ gives 1, $k = n$ gives 2^n. The pattern 2^k contains no n, and a bound blind to n cannot depend on the embedding of the cube. That rules out in advance every route that begins by exploiting the geometry of $\mathbb{R}^n$. **Link to the next block.** Knowing the answer, split the work the way every maximum question splits: construction, then bound.

Block 2. **Cue.** "Vector subspace of dimension k" with a cube attached. **Why natural.** The cheapest k dimensional subspace is a coordinate subspace, and a coordinate subspace meets the cube in a k dimensional face. That is the definition of a face, so it needs no thought. **Link to the next block.** One line spent, so all remaining attention can go to finding the shortest possible bound.

Block 3. **Cue.** "whose coordinates are 0 or 1". **Why natural.** This clause is not decoration; it says the alphabet has size 2. When a count comes out as a power of 2, the exponent counts independent binary choices. So the shape of the answer already tells you the proof will exhibit k binary choices determining the point. You are no longer hunting a proof, you are hunting the k choices. **Link to the next block.** The only structure producing the number k is $\dim V$, so the choices must come from dimension. What converts dimension into a count of coordinates?

Block 4. **Cue.** The word "dimension" next to the word "coordinates". **Why natural.** Rank is the one theorem that translates between them, being at once the dimension of the row span and of the column span. List the points as rows: the row rank is at most $\dim V = k$ by definition of span, so the column rank is at most k, so at most k coordinate functions on your point set are independent. Note what this move required: treating all of $V \cap Z$ at once as a single matrix. The count is a property of the whole set, so the whole set must be the object. **Link to the next block.** A spanning set of columns yields an injective restriction map, and injectivity plus a two letter alphabet is a count.

Block 5. **Cue.** A set injecting into $\{0,1\}^J$ with $|J| \leq k$. **Why natural.** Injectivity turns a hard count into a trivial one: the domain is no larger than the codomain, and the codomain is a set of binary strings of length at most k. The only algebra in the whole solution is the two line proof that agreeing on the basis columns forces agreeing on all columns. **Link to the next block.** The bounds meet at 2^k, so nothing remains.

Layer 3: WHY

The truth underneath is rank duality, stated in the form the problem needs: if $S \subseteq A^n$ with A a finite alphabet of size q, and the span of S in $\mathbb{R}^n$ has dimension at most k, then $|S| \leq q^k$. The proof is the one above with 2 replaced by q. Read the statement carefully: the ambient dimension has vanished and so has the geometry, and what is left is an exchange rate between dimension and information. One dimension buys one coordinate, one coordinate buys $\log_2 q$ bits, so k dimensions buy q^k points. This problem is a statement about how much information a low dimensional linear space can carry when its points are written in a finite alphabet.

The discernment mechanism you must train is knowing when to transpose. A matrix has two readings, and you should always ask which reading the hypothesis lives in and which the conclusion lives in. Here the hypothesis is about rows, since it says the points lie in a k dimensional space; the constraint on entries is about columns, since it caps the values a coordinate can take. Hypothesis in rows, constraint in columns: transpose. Whenever a problem pairs a rank condition on a family of vectors with an entrywise condition on those vectors, rank duality is the bridge, and the reading in which the constraint is simplest is the one to work in. The examiners nudged three times: they told you parenthetically that Z has 2^n elements, which plants "power of two" before you start; they wrote $0 \leq k \leq n$ with no relation between the two, signalling uniformity in n; and they added an editorial note removing an ambiguity in the phrasing, which is their way of saying the ambiguity was never the difficulty. The same rank duality also explains why the modulo 2 route lands on the same number: both routes compute a rank, one over $\mathbb{R}$ and one over $\mathbb{F}_2$, and the minor lifting argument says exactly that reduction cannot increase rank.

Revolving concepts

Rank. For a real $m \times n$ matrix, the row rank equals the column rank, and this common value also equals the largest size of a nonvanishing minor. Orbiting facts: invariance under elementary operations, $\operatorname{rank} M \leq \min(m, n)$, and rank plus nullity equals n. The minor characterisation is what powers the finite field route. See Book 1, field: Linear algebra, topic: rank.

Span and dimension. If every row of M lies in a subspace V then the row space is contained in V, so $\operatorname{rank} M \leq \dim V$. This is the only property of V the ideal solution uses. See Book 1, field: Linear algebra, topic: subspaces and dimension.

The hypercube. Vertices of $\{0,1\}^n$ correspond to subsets of $\{1, \ldots, n\}$ by indicator vectors, so every cube statement has a set system translation. Fixing some coordinates and freeing the rest gives a face; a k dimensional face has 2^k vertices. See Book 1, field: Discrete structures and graphs, topic: hypercube.

The field $\mathbb{F}_2$. The two element field with $1 + 1 = 0$. A d dimensional subspace of $\mathbb{F}_2^n$ has exactly 2^d elements, since a basis gives d independent binary choices. Lifting lemma: an integer matrix with an odd maximal minor has that minor nonzero over $\mathbb{R}$, so reduction modulo 2 cannot increase rank. See Book 1, field: Abstract algebra, topic: finite fields.

Affine planes. A k dimensional affine plane is $p + W$ with $\dim W = k$. Subtracting one point of a set from all of them converts affine to linear without changing cardinality, and is the standard first move when a problem says "plane" rather than "subspace". See Book 1, field: Linear algebra, topic: affine subspaces.

The recursion, generalised. The slicing route is a recursion and deserves to be written as one. Let $f(n,k)$ be the largest number of points of $\{0,1\}^n$ on an affine k plane. Slicing by $x_n = 0$ and $x_n = 1$ gives

$$f(n,k) \leq \max\{\, 2f(n-1,k-1),\ f(n-1,k) \,\}, \qquad f(n,0) = 1,$$

the first branch generic and the second the degenerate case where the plane lies inside a slice; induction on $n + k$ gives $f(n,k) \leq 2^k$. Install the general shape, not the instance: a quantity defined on $\{0,1\}^n$ is bounded by splitting the cube into its two facets, applying the hypothesis to each, and

adding. The additive sibling $g(n,d) \le g(n-1,d) + g(n-1,d-1)$, $g(n,0) = 1$, solves to $\sum_{i=0}^{d} \binom{n}{i}$ and is the Sauer and Shelah lemma. The present problem is its linear algebraic sibling: "dimension at most k" plays the role of "shatters nothing of size $k+1$". See Book 1, field: Combinatorics, topic: VC dimension and shattering.

Circuit Bank

- **WHEN** you see vectors with entries from a finite alphabet of size q lying in a space of dimension k **DO** stack them as rows, take the rank, choose k spanning columns, and bound the count by q^k **BECAUSE** row rank equals column rank turns "dimension k" into "k coordinates determine the point", and a finite alphabet then caps the number of points
- **WHEN** a hypothesis is about the rows of a natural matrix and the constraint is about its entries columnwise **DO** transpose, that is, study the column space instead **BECAUSE** rank duality is the only bridge between the two readings, and you should work in the reading where the constraint is simplest
- **WHEN** the answer to a "find the maximum" omits one of the given parameters **DO** look for a proof that never mentions that parameter **BECAUSE** the absence tells you which structure carries the bound, so any argument using the absent parameter is doing unnecessary work
- **WHEN** the objects are 0 or 1 vectors **DO** consider reducing modulo 2 into $\mathbb{F}_2^n$ **BECAUSE** 0 and 1 are exactly the elements of $\mathbb{F}_2$, subspaces there have size $2^{\dim}$, and an odd maximal minor lifts, so reduction cannot increase rank
- **WHEN** an induction over a cube stalls because a slice is affine rather than linear **DO** strengthen the statement to affine planes and induct on the sum of the parameters **BECAUSE** a stronger hypothesis is usually the only way to close a slicing recursion, and $n+k$ decreases in both branches

Backtesting

1. First useful cue: the answer 2^k contains no n, visible after two seconds of testing $k = 0$ and $k = n$, which forces a dimension only proof.
2. Method the expression suggested: "coordinates are 0 or 1" beside "dimension k" is an instruction to stack the points into a matrix and use row rank equals column rank.
3. Where you would hesitate, and the thought to reinforce: forming a matrix out of the entire intersection feels like discarding structure. Reinforce the opposite thought, that the count is a property of the whole set, so the whole set must be the algebraic object.
4. The trap the question sets: controlling the coefficients of cube vertices in terms of a chosen basis of cube vertices. The coefficients are unrestricted, so the trap costs you the paper.
5. What makes you recognise it faster next time: a rank hypothesis paired with an entrywise constraint. That pairing means transpose, every time.
6. The reusable decision rule: a subset of A^n spanning a space of dimension k has at most $|A|^k$ elements, and the proof is column selection.

Mirror drills

1. Let $S \subseteq \{0,1,2\}^n$ lie in a k dimensional subspace of $\mathbb{R}^n$. Show $|S| \le 3^k$. *Answer:* identical argument; stack S as rows, rank at most k, choose k spanning columns, and the injection lands in $\{0,1,2\}^k$.

2. What is the maximum number of points of $\{0,1\}^n$ on an affine k plane? *Answer:* 2^k again. Subtract one point of the intersection from all of them and apply the rank argument to the differences.

3. In $\mathbb{R}^3$ exhibit a plane through the origin holding exactly three cube vertices and one holding four, and name the feature that separates them. *Answer:* the span of $(1,1,0)$ and $(0,1,1)$ holds three; the span of $(1,1,0)$ and $(0,0,1)$ holds four. The second is spanned by indicators of *disjoint* coordinate blocks, so every subfamily sum is again a 0 or 1 vector.

20.11 B5. Maximising a weighted quadratic integral functional

For each continuous function $f : [0,1] \to \mathbb{R}$, let $I(f) = \int_0^1 x^2 f(x)\,dx$ and $J(x) = \int_0^1 x\,(f(x))^2\,dx$. Find the maximum value of $I(f) - J(f)$ over all such functions f.

Read it out loud

Say the structure aloud before you touch a pen. You are asked to maximise, over an infinite dimensional space of functions, a quantity that is linear in f minus quadratic in f, with both terms carrying a weight that vanishes at the same endpoint. That is the entire description. The fields wired in are real analysis (integration and continuity), inequalities, and the calculus of variations, which appears here only so that you can recognise it and decline it. Note the harmless slip in the statement: J is defined with the letter x on the left and used as $J(f)$; it is a function of f, and nothing in the problem depends on the typo.

The first structural move is forced. Two integrals over the same interval with the same variable should be one integral, so write

$$I(f) - J(f) = \int_0^1 \left(x^2 f(x) - x f(x)^2\right) dx.$$

Now look at the integrand with x frozen. It is a quadratic in the single real number $f(x)$, with leading coefficient $-x$, and on $[0,1]$ that leading coefficient is at most zero. A quadratic with nonpositive leading coefficient has a maximum, not a minimum, and you find it by completing the square. That is the whole solution, and it should be on the page within thirty seconds.

Before proving anything, run the reconnaissance that tells you the answer. The two weights are x^2 against f and x against f^2, which are balanced when f scales like x, so try $f(x) = cx$. Then $I = c/4$ and $J = c^2/4$, so $I - J = (c - c^2)/4$, maximised at $c = 1/2$ with value $1/16$. You now know the target and, more usefully, you know the centre of the square you are about to complete. The remaining question is not what the answer is but why no other function beats it.

The one thing worth flagging in advance, because it is where the examiners did their design work, is the point $x = 0$. There the quadratic degenerates: the coefficient of $f(x)^2$ vanishes, so at that single point $f(0)$ is unconstrained by the objective. Watch whether that costs you anything. It will not, and the reason it does not is the real content of Layer 3.

Approaches, ranked worst to best

1. **Discretise into a Riemann sum and maximise a finite quadratic form.** *Cue:* "quadratic in f" plus "maximum". *Circuit:* diagonalise a quadratic form, then pass to the limit. You replace f by its values on a grid, maximise a diagonal quadratic form in those values, and let the mesh go to zero. It gives the right number and it is a nightmare to justify: you must show the maxima of the discretisations converge to the supremum of the functional and that the limit is attained by a continuous function. Every one of those steps is harder than the problem.

2. **Euler and Lagrange, the first variation.** *Cue:* "maximum over all such functions" is the phrase that names the calculus of variations. *Circuit:* perturb $f \to f + \varepsilon g$, differentiate at $\varepsilon = 0$, set the derivative to zero. You get $\int_0^1 g(x)\left(x^2 - 2xf(x)\right) dx = 0$ for every continuous g, and the fundamental lemma of the calculus of variations forces $x^2 - 2xf(x) = 0$, so $f(x) = x/2$ on $(0,1]$ and $f(0) = 0$ by continuity. It is correct, but it is overbuilt in two ways: a stationary point is not a maximum until you add the second variation $-\int_0^1 x g(x)^2 dx \le 0$, and you have spent the fundamental lemma on a problem with no derivative of f in it at all. Machinery you did not need.

3. **Guess the one parameter family $f = cx$ and optimise c.** *Cue:* the weights x^2 and x balance exactly when f scales like x. *Circuit:* reduce an infinite dimensional problem to a one dimensional one. Twenty seconds to the number $1/16$, and it proves only that the maximum is at least $1/16$. Indispensable as reconnaissance, worthless as a write up. Never confuse the two.

4. **Cauchy and Schwarz with the weight x.** *Cue:* J contains xf^2, so $\sqrt{x}f$ is the natural object, and $x^2 f = x^{3/2} \cdot \sqrt{x} f$. *Circuit:* split a product so that one factor squares to a given integral. You get

$$I(f) = \int_0^1 x^{3/2} \cdot \sqrt{x} f(x)\, dx \le \Big(\int_0^1 x^3 dx \Big)^{1/2} J(f)^{1/2} = \tfrac{1}{2}\sqrt{J(f)},$$

so with $t = \sqrt{J(f)} \ge 0$ we have $I - J \le t/2 - t^2 \le 1/16$. Complete, short, and honest. It ranks below the next route only because it needs two equality analyses stitched together: equality in Cauchy and Schwarz forces f proportional to x, and $t = 1/4$ then pins the constant to $1/2$.

5. **Complete the square pointwise, inside the integral.** *Cue:* the integrand is a quadratic in $f(x)$ with nonpositive leading coefficient, at each fixed x. *Circuit:* maximise under the integral sign when no derivative of the unknown appears. Six lines, one identity, one sign observation, one exhibited maximiser. This is the ideal solution.

The ideal solution

The maximum is $\boxed{\frac{1}{16}}$, attained only by $f(x) = x/2$.

For every real number y and every x, expanding the right hand side verifies the identity

$$x^2 y - x y^2 = \frac{x^3}{4} - x\Big(y - \frac{x}{2}\Big)^2.$$

Apply it with $y = f(x)$ for each $x \in [0,1]$. Since $x \ge 0$ on $[0,1]$, the subtracted term $x\big(f(x) - x/2\big)^2$ is nonnegative, so

$$x^2 f(x) - x f(x)^2 \le \frac{x^3}{4} \qquad \text{for every } x \in [0,1].$$

Both sides are continuous on $[0,1]$, so integrating preserves the inequality:

$$I(f) - J(f) = \int_0^1 \big(x^2 f(x) - x f(x)^2\big) dx \le \int_0^1 \frac{x^3}{4}\, dx = \frac{1}{16}.$$

The bound is attained. The function $f(x) = x/2$ is continuous on $[0,1]$ and makes the subtracted term identically zero, so for it $I(f) - J(f) = 1/16$. Hence the supremum is achieved and the maximum value is $1/16$.

Remark 20.11.1 (Uniqueness of the maximiser). Suppose $I(f) - J(f) = 1/16$. Subtracting the displayed inequality from the equality gives

$$\int_0^1 x\Big(f(x) - \frac{x}{2}\Big)^2 dx = 0.$$

The integrand is continuous and nonnegative on $[0,1]$, and a continuous nonnegative function with zero integral is identically zero. Hence $x\big(f(x) - x/2\big)^2 = 0$ for all x, so $f(x) = x/2$ for every $x \in (0,1]$, and $f(0) = 0$ follows from continuity at 0. So $f(x) = x/2$ is the unique maximiser, even though the objective never constrains the single value $f(0)$ directly.

Layer 1: WHAT

Block 1. Merge the two integrals into one integral of $x^2 f - x f^2$.

Block 2. Freeze x and read the integrand as a quadratic in the real number $f(x)$, with leading coefficient $-x \leq 0$.

Block 3. Complete the square: the pointwise maximum is $x^3/4$, attained at $f(x) = x/2$.

Block 4. Observe that the pointwise maximiser, the function $x \mapsto x/2$, is itself continuous, hence admissible. Therefore the pointwise optimum and the global optimum coincide and no variational machinery is needed.

Block 5. Integrate the pointwise bound to get $1/16$, exhibit the maximiser, and prove uniqueness from the vanishing of a continuous nonnegative integral.

Nominal against ideal. The nominal route reads "maximum over all continuous functions" and reaches for the calculus of variations, because that is the machine whose label matches the phrase. It then spends four lines on a perturbation argument, one lemma on the fundamental lemma, and one more line on the second variation, to reach a conclusion the ideal route reaches by expanding a bracket.

The bifurcation point is Block 4, and it is a single test: does the integrand contain f'? Here it does not. An integral of $\Phi(x, f(x))$ with no derivative of f has no coupling between different values of x, so the values $f(x)$ can be chosen independently, one x at a time. Maximising inside the integral is then legitimate, provided the resulting pointwise maximiser is an admissible function. If f' had appeared, the values would be coupled and the calculus of variations would be the correct tool. That single test is the reusable content of this problem.

Layer 2: HOW

Block 1. **Cue.** Two integrals, same interval, same variable, subtracted. **Why natural.** Linearity of the integral says a difference of integrals is the integral of the difference, and you always want one integrand rather than two because you can only inspect structure in an expression you can see whole. Nothing is lost and everything becomes visible. **Link to the next block.** With one integrand in front of you, the next question is what kind of object it is.

Block 2. **Cue.** $x^2 f(x) - x f(x)^2$: degree two in $f(x)$, degree at most two in x. **Why natural.** Read it as a polynomial in $f(x)$ with coefficients depending on x, because $f(x)$ is the unknown and x is the parameter. Once you do, it is $-xy^2 + x^2 y$ with $y = f(x)$, and the sign of the leading coefficient decides everything. On $[0,1]$ that coefficient is $-x \leq 0$, so the quadratic is concave and has a maximum. **Link to the next block.** A concave quadratic is maximised by completing the square, which is the one manipulation A level installs and which never stops paying.

Block 3. **Cue.** A concave quadratic in y with a parameter. **Why natural.** Completing the square converts "maximise" into "read off", and it does so with an identity rather than an inequality, which means no hypotheses and no error terms. The vertex sits at $y = x/2$ and the maximum value is $x^3/4$; both come out of the same line of algebra. **Link to the next block.** You now have a pointwise bound. The question is whether the pointwise bound is achievable simultaneously at every x by one admissible function.

Block 4. **Cue.** The pointwise argmax is $x \mapsto x/2$, and the admissible class is the continuous functions. **Why natural.** This is the step that decides which method the problem is. The pointwise maximisers form a function, and that function is continuous, so it lies in the class over which you are maximising. Hence the supremum of the integral equals the integral of the pointwise supremum. This is legitimate precisely because the integrand involves $f(x)$ and not $f'(x)$: with no derivative present, different values of x do not constrain each other, and the constraint of continuity is satisfied

for free by the argmax. **Link to the next block.** Nothing prevents the bound from being attained, so integrate it and evaluate.

Block 5. Cue. $\int_0^1 x^3/4\,dx$ and the phrase "find the maximum". **Why natural.** The word maximum, as opposed to supremum, obliges you to exhibit a function attaining the value. Producing $f(x) = x/2$ costs one sentence and converts a bound into an answer. The uniqueness statement costs one more sentence and uses the standard fact that a continuous nonnegative function with zero integral vanishes identically, which is worth including because it also disposes of the awkward point $x = 0$. **Link to the next block.** The value $1/16$ is both an upper bound and an attained value, so it is the maximum.

Layer 3: WHY

There are two truths beneath this route, and both are worth more than the problem.

The first is the pointwise principle. For a functional $F(f) = \int_a^b \Phi(x, f(x))\,dx$ in which no derivative of f appears, the values of f at different points are not coupled by anything except the admissibility class. Therefore

$$\sup_f F(f) \le \int_a^b \sup_y \Phi(x, y)\,dx,$$

always, with equality whenever the map $x \mapsto \operatorname{argmax}_y \Phi(x, y)$ is single valued and admissible. The calculus of variations exists to handle the opposite case, where f' appears and the values are coupled along the interval, so that you cannot move one value without moving its neighbours. Learn to run the test first: look for f' in the integrand. Absent means pointwise; present means Euler and Lagrange.

The second is the quadratic structure, which is best seen in the language of inner products. On $[0, 1]$ define the weighted inner product $\langle u, v\rangle = \int_0^1 x\,u(x)v(x)\,dx$, with norm $\|u\|^2 = \int_0^1 xu(x)^2dx$, and let $g(x) = x$. Then $J(f) = \|f\|^2$ and $I(f) = \int_0^1 x \cdot f(x) \cdot x\,dx = \langle f, g\rangle$, so the problem reads: maximise $\langle f, g\rangle - \|f\|^2$. In any inner product space,

$$\langle f, g\rangle - \|f\|^2 = \frac{\|g\|^2}{4} - \left\|f - \frac{g}{2}\right\|^2 \le \frac{\|g\|^2}{4},$$

with equality exactly at $f = g/2$. That is one identity covering every problem of this shape, and here it gives $\|g\|^2 = \int_0^1 x^3dx = 1/4$ and the answer $1/16$ immediately. The pointwise completion of the square and this Hilbert space identity are the same computation seen at two zoom levels, which is why the ideal route feels like a triviality once you have seen it: it is a triviality, in the right coordinates.

Now the design of the question, which is where the examiners hid the real difficulty. Write the general integrand as $\Phi(x, y) = b(x)y - a(x)y^2$ with $a \ge 0$. The pointwise maximum is $b(x)^2/(4a(x))$ at $y = b(x)/(2a(x))$, so the answer is $\int_0^1 \frac{b(x)^2}{4a(x)}\,dx$ whenever that makes sense. Here $a(x) = x$ vanishes at the endpoint, but $b(x) = x^2$ vanishes faster, so $b^2/(4a) = x^3/4$ is continuous and integrable and the pointwise maximiser $x/2$ is continuous. Break that balance and the problem changes character entirely: for $\int_0^1 \big(f(x) - xf(x)^2\big)dx$ we have $b^2/(4a) = 1/(4x)$, which is not integrable, and indeed taking the continuous functions $f_M(x) = \min\{1/(2x), M\}$ gives values at least $\frac{1}{4}\log(2M)$, so the supremum is infinite and there is no maximum. So the choice of the weights x^2 and x was not decoration; it was the condition that makes the answer finite and attained. Reading a question's design is how you learn what the setter was protecting against.

Revolving concepts

Completing the square. For $a \neq 0$, $ay^2 + by + c = a(y + b/2a)^2 + c - b^2/4a$, so a concave quadratic ($a < 0$) has maximum $c - b^2/4a$ at $y = -b/2a$. Here it is applied with x as a silent parameter, which is the only new habit required. See Book 1, field: Polynomials, topic: quadratics and the vertex.

Weighted inner products and the L^2 norm. For a weight $w \geq 0$ on $[a, b]$, $\langle u, v\rangle_w = \int wuv$ is an inner product on continuous functions when $w > 0$ on the interior, with the Cauchy and Schwarz inequality $|\langle u, v\rangle_w| \leq \|u\|_w \|v\|_w$ and equality exactly for proportional u, v. The identity $\langle f, g\rangle - \|f\|^2 = \|g\|^2/4 - \|f - g/2\|^2$ solves every "linear minus quadratic" maximisation in one line. See Book 1, field: Real analysis, topic: inner product spaces.

The calculus of variations. For $F(f) = \int \Phi(x, f, f')dx$ the stationarity condition is the Euler and Lagrange equation $\frac{d}{dx}(\partial\Phi/\partial f') = \partial\Phi/\partial f$. When Φ does not involve f' this degenerates to $\partial\Phi/\partial f = 0$, which is exactly the pointwise maximisation. The fundamental lemma of the calculus of variations, that a continuous h with $\int hg = 0$ for all continuous g is identically zero, is what converts the first variation into a pointwise equation. See Book 1, field: Calculus and integrals, topic: variational problems.

A continuous nonnegative function with zero integral vanishes. If $h \geq 0$ is continuous on $[a, b]$ and $\int_a^b h = 0$ then $h \equiv 0$; otherwise $h(x_0) = \varepsilon > 0$ gives an interval of positive length on which $h > \varepsilon/2$, forcing a positive integral. This is the standard tool for turning an equality case into pointwise information. See Book 1, field: Real analysis, topic: positivity and integrals.

Monotonicity of the integral. If $u \leq v$ pointwise on $[a, b]$ and both are integrable then $\int u \leq \int v$. It is the only step in the ideal solution that moves from local to global, and it is worth naming as such: the entire proof is one pointwise inequality plus one application of monotonicity.

No recursion here, and that is informative. This problem contains no sequence, no induction, no telescoping and no self similar structure to generalise. That absence is a signal: a functional with no recursive structure and no derivative of the unknown is a one shot pointwise problem, and any method you reach for that builds an iteration or an approximating sequence is fighting the problem rather than reading it.

Circuit Bank

- **WHEN** an integral functional contains f but no f' **DO** maximise the integrand pointwise with x frozen, then check the pointwise maximiser is an admissible function **BECAUSE** with no derivative present nothing couples different values of x, so the supremum of the integral is the integral of the supremum whenever the argmax is admissible

- **WHEN** the objective is linear in the unknown minus quadratic in it with a nonnegative weight **DO** write it as $\|g/2\|^2 - \|f - g/2\|^2$ in the weighted inner product **BECAUSE** in any inner product space $\langle f, g\rangle - \|f\|^2$ is maximised at $f = g/2$ with value $\|g\|^2/4$, which hands you both the answer and the unique maximiser

- **WHEN** a weight in the objective vanishes at an endpoint **DO** check that the pointwise maximum $b^2/(4a)$ is still integrable there **BECAUSE** that integrability is exactly the line between a finite attained maximum and an infinite supremum, and it is where the setter did the design work

- **WHEN** the question says maximum rather than supremum **DO** exhibit the maximiser and verify it lies in the stated class **BECAUSE** a bound alone answers a different question, and the attainment sentence costs one line

Backtesting

1. First useful cue: two integrals over the same interval, subtracted, giving one integrand that is quadratic in $f(x)$ with leading coefficient $-x \leq 0$.
2. Method the expression suggested: complete the square in $f(x)$ with x as a parameter, then integrate the resulting pointwise bound.
3. Where you would hesitate, and the thought to reinforce: at the claim that maximising inside the integral is allowed. Reinforce the test rather than the anxiety: no f' in the integrand means no coupling between different x, so pointwise optimisation is exact provided the argmax is admissible, and here it is the continuous function $x/2$.
4. The trap the question sets: the phrase "over all such functions", which baits you into the calculus of variations, and the degenerate point $x = 0$, where the coefficient of $f(x)^2$ vanishes and an unwary equality argument leaves $f(0)$ undetermined. Continuity closes that gap.
5. What makes you recognise it faster next time: an integrand that is a polynomial in $f(x)$ alone. That is a pointwise problem wearing variational clothes.
6. The reusable decision rule: for $\int \big(b(x)f - a(x)f^2\big)dx$ with $a \geq 0$, the maximum is $\int b^2/(4a)$ at $f = b/(2a)$, valid whenever that function is admissible and the integral converges.

Mirror drills

1. Find the maximum of $\int_0^1 \big(xf(x) - f(x)^2\big)dx$ over continuous f. *Answer:* $b^2/(4a) = x^2/4$, so the maximum is $\int_0^1 x^2/4\,dx = 1/12$, attained at $f(x) = x/2$.
2. Find the maximum of $\int_0^1 \big(x^3 f(x) - xf(x)^2\big)dx$ over continuous f. *Answer:* $b^2/(4a) = x^5/4$, so the maximum is $1/24$, attained at $f(x) = x^2/2$.
3. Show that $\int_0^1 \big(f(x) - xf(x)^2\big)dx$ has no maximum over continuous f. *Answer:* here $b^2/(4a) = 1/(4x)$ is not integrable at 0; take $f_M(x) = \min\{1/(2x), M\}$, which is continuous, and the integral is at least $\int_{1/(2M)}^1 \frac{dx}{4x} = \frac{1}{4}\log(2M) \to \infty$.

20.12 B6. Asymptotics of a self slowing recursion

Let k be an integer greater than 1. Suppose $a_0 > 0$, and define

$$a_{n+1} = a_n + \frac{1}{\sqrt[k]{a_n}}$$

for $n > 0$. Evaluate

$$\lim_{n\to\infty} \frac{a_n^{k+1}}{n^k}.$$

Read it out loud

This is a differential equation that has been discretised and handed to you as a sequence. The increment $a_{n+1} - a_n$ depends only on the current value, so the process is autonomous, and the increment $a_n^{-1/k}$ decreases as a_n grows, so the sequence is increasing and slowing. Every question about such a sequence is a question about the balance between those two facts. The fields wired in are real analysis (limits of sequences, Cesaro means) and the asymptotics of recursions, with a separable first order ODE as the guiding light.

Read the target next, because the examiners have already told you the answer's shape. They ask for $\lim a_n^{k+1}/n^k$, which is a finite nonzero limit only if a_n grows like $n^{k/(k+1)}$. So the growth exponent is a gift and you never have to discover it. More than that: rewriting $a_n^{k+1}/n^k = \left(a_n^{(k+1)/k}/n\right)^k$ shows that the quantity actually being controlled is $a_n^{(k+1)/k}$ divided by n. The exponent $(k+1)/k$ is being pointed at, and it is the exponent of the substitution.

Now do the twenty second reconnaissance that makes everything else routine. Replace the recursion by the differential equation $y' = y^{-1/k}$, which is what "increment equals a function of the current value" means in continuous time. Separate: $y^{1/k}\,dy = dx$, so $\frac{k}{k+1}y^{(k+1)/k} = x + C$, so $y^{(k+1)/k} \sim \frac{k+1}{k}x$ and $y^{k+1}/x^k \to \left(\frac{k+1}{k}\right)^k$. You now have the answer, and you also have the substitution: the ODE separated into the variable $y^{(k+1)/k}$, so that is the variable in which the recursion will behave. This is the central point of the whole problem and it is not a guess; the substitution is read off the separation.

Two small blemishes in the statement deserve a sentence and no more. The recursion is written "for $n > 0$", which leaves a_1 undefined; it is plainly intended for all $n \geq 0$. And the hypothesis that k is an integer greater than 1 is never used: everything below works for any real $k > 0$.

Approaches, ranked worst to best

1. **Compute numerically for small k and guess.** *Cue:* a limit with a free parameter k. *Circuit:* tabulate and pattern match. For $k = 2$ you get 2.25, and recognising that as $(3/2)^2$ requires you to already suspect the form. Zero marks on a B6, and it does not even suggest a method.

2. **Bootstrap two sided bounds on a_n directly.** *Cue:* the sequence is increasing with shrinking increments. *Circuit:* prove $c_1 n^{k/(k+1)} \leq a_n \leq c_2 n^{k/(k+1)}$, then improve the constants. This works, and it is essentially the official second solution: convexity gives a lower bound for $a_n^{(k+1)/k}$, that lower bound feeds back into the increment to give an upper bound, and the two constants meet. But you must run the bootstrap at least twice and carry explicit constants through every inequality, and you must recognise that the loop converges. Correct, expensive, and error prone under time pressure.

3. **Guess the answer and control the error sequence.** *Cue:* the limit is a clean constant, so posit $a_n \approx Ln^{k/(k+1)}$. *Circuit:* set $b_n = a_n - Ln^{k/(k+1)}$ and show b_n stays bounded. This is the official third solution. It proves more than you were asked, namely a second order statement, and it

costs three separate Taylor estimates plus an $O(n^{-2})$ summation. On a paper where B6 is the last thing you reach, this is the wrong trade.

4. **Substitute $b_n = a_n^{(k+1)/k}$ and bound the increment by Taylor with remainder.** *Cue:* the exponent visible in the target. *Circuit:* make the increments asymptotically constant, then estimate the error. The official first solution. Correct and short, but it invokes Taylor's theorem with an explicit remainder bound in order to prove a statement that is only a limit. You are paying for an error term you will never use.

5. **Substitute $b_n = a_n^{(k+1)/k}$, factor the recursion into $b_{n+1} = b_n(1 + 1/b_n)^{(k+1)/k}$, read the increment limit as a derivative, then apply Cesaro.** *Cue:* the ODE $y' = y^{-1/k}$ separates exactly into the variable $y^{(k+1)/k}$. *Circuit:* change to the variable in which the recursion becomes "add a constant", then convert increments into growth. No remainder terms, no bootstrap, no constants to chase. One substitution, one difference quotient, one standard lemma. This is the ideal solution.

The ideal solution

Write $m = \dfrac{k+1}{k} > 1$. The answer is $\boxed{\left(\dfrac{k+1}{k}\right)^k}$.

Step 1: the sequence increases to infinity. Since $a_0 > 0$ and $a_{n+1} = a_n + a_n^{-1/k}$, induction gives $a_n > 0$ for all n, so the increments $a_n^{-1/k}$ are positive and (a_n) is strictly increasing. Suppose it were bounded, say $a_n \to L$ with $0 < a_0 \le L < \infty$. A convergent sequence has $a_{n+1} - a_n \to 0$, but $a_{n+1} - a_n = a_n^{-1/k} \to L^{-1/k} > 0$, a contradiction. Hence (a_n) is increasing and unbounded, so $a_n \to \infty$.

Step 2: the substitution makes the increments constant. Put $b_n = a_n^m$. Since $a_n^{-1/k} = a_n \cdot a_n^{-(k+1)/k} = a_n \cdot a_n^{-m}$, the recursion factors as

$$a_{n+1} = a_n\big(1 + a_n^{-m}\big),$$

and raising to the power m gives

$$b_{n+1} = a_n^m\big(1 + a_n^{-m}\big)^m = b_n\Big(1 + \frac{1}{b_n}\Big)^m.$$

Therefore, writing $t_n = 1/b_n$,

$$b_{n+1} - b_n = b_n\Big[\Big(1 + \frac{1}{b_n}\Big)^m - 1\Big] = \frac{(1+t_n)^m - 1}{t_n}.$$

By Step 1, $b_n = a_n^m \to \infty$, so $t_n \to 0^+$. The function $t \mapsto (1+t)^m$ is differentiable at $t = 0$ with derivative m, and the displayed quotient is exactly its difference quotient at 0. Hence

$$\lim_{n\to\infty}(b_{n+1} - b_n) = m = \frac{k+1}{k}.$$

No remainder estimate is needed, because a derivative is precisely the limit of this quotient.

Step 3: from increments to growth.

Lemma 20.12.1 (Cesaro mean)**.** *If $x_i \to c$ then $\frac{1}{n}\sum_{i=1}^n x_i \to c$.*

Proof. Let $\varepsilon > 0$. Choose N with $|x_i - c| < \varepsilon/2$ for all $i > N$, and put $C = \sum_{i=1}^N |x_i - c|$, a fixed number. For $n > N$,

$$\Big|\frac{1}{n}\sum_{i=1}^n x_i - c\Big| \le \frac{1}{n}\sum_{i=1}^n |x_i - c| \le \frac{C}{n} + \frac{n-N}{n}\cdot\frac{\varepsilon}{2} < \frac{C}{n} + \frac{\varepsilon}{2},$$

which is less than ε as soon as $n > 2C/\varepsilon$. □

Apply the lemma to $x_i = b_i - b_{i-1}$, which converges to m by Step 2. The sum telescopes, so

$$\frac{b_n - b_0}{n} = \frac{1}{n}\sum_{i=1}^{n}(b_i - b_{i-1}) \longrightarrow m,$$

and since $b_0/n \to 0$ we get $b_n/n \to m$.

Step 4: conclude. Because $mk = k+1$,

$$\frac{a_n^{k+1}}{n^k} = \frac{\left(a_n^m\right)^k}{n^k} = \left(\frac{b_n}{n}\right)^k \longrightarrow m^k = \left(\frac{k+1}{k}\right)^k,$$

using continuity of $t \mapsto t^k$. This is the required limit.

Remark 20.12.2. Nothing above used that k is an integer, or that $k > 1$; the argument holds for every real $k > 0$, and for $k = 1$ it gives $\lim a_n^2/n = 2$ for the recursion $a_{n+1} = a_n + 1/a_n$. The answer written as $\left(1 + 1/k\right)^k$ increases with k and tends to e, and the growth exponent $k/(k+1)$ tends to 1, both reflecting the fact that a larger k makes the increment $a_n^{-1/k}$ decay more slowly and the sequence closer to linear.

Layer 1: WHAT

Block 1. Read the recursion as the discretisation of $y' = y^{-1/k}$, separate it, and obtain both the answer $\left(\frac{k+1}{k}\right)^k$ and the substitution exponent $(k+1)/k$.

Block 2. Prove $a_n \to \infty$, which costs three lines and is the only hypothesis every later step needs.

Block 3. Set $b_n = a_n^{(k+1)/k}$ and factor a_n out of the recursion, so that the recursion for b is autonomous and self contained: $b_{n+1} = b_n(1 + 1/b_n)^m$.

Block 4. Recognise $b_{n+1} - b_n = \left((1+t)^m - 1\right)/t$ with $t = 1/b_n \to 0$ as a difference quotient, so the increment limit is a derivative, not an estimate.

Block 5. Convert "the increments tend to m" into "b_n/n tends to m" by telescoping and the Cesaro mean lemma.

Block 6. Raise to the power k and read off the answer.

Nominal against ideal. The nominal route keeps the sequence in the variable it was handed, guesses $a_n \approx Cn^{k/(k+1)}$, and tries to prove that by induction with explicit inequalities. It can be made to work, but every line carries constants, every constant has to be improved by a second pass, and any slip in a direction of an inequality is invisible until the constants fail to meet.

The bifurcation point is Block 3, and it is the decision to change variable rather than to estimate. The zombie route estimates a_n; the ideal route changes coordinates so that there is nothing left to estimate. Name the principle in the form you will reuse: for $a_{n+1} = a_n + f(a_n)$, the right variable is $F(a) = \int \frac{da}{f(a)}$, because in that variable the recursion becomes "add approximately 1", and "add exactly 1" is the only recursion whose asymptotics are free.

Layer 2: HOW

Block 1. Cue. $a_{n+1} - a_n$ is a function of a_n alone, and the question asks for a limit of a_n^{k+1}/n^k. **Why natural.** An autonomous increment is a discrete derivative, so the continuous shadow $y' = y^{-1/k}$ is not an analogy but the leading order behaviour, and it is separable, which is the easiest class of ODE there is. The separation integrates $y^{1/k}$, producing $y^{(k+1)/k}$, which is exactly the exponent hidden in the requested ratio. Two independent signals point at the same power, so the substitution is forced rather than inspired. **Link to the next block.** You know the answer and the substitution; what remains is to make the argument discrete and rigorous.

Block 2. **Cue.** A positive increasing sequence with no a priori bound. **Why natural.** Every later step is an asymptotic statement, and asymptotics are meaningless unless $a_n \to \infty$. The cheapest proof is the standard one: an increasing sequence either diverges or converges, and if it converged its increments would tend to zero, whereas here they tend to a positive number. Notice this argument uses only that $t \mapsto t^{-1/k}$ is continuous and positive at L, so it is reusable verbatim for any positive increment function. **Link to the next block.** With $a_n \to \infty$ in hand, any expression tending to zero in a_n may be treated as a small parameter.

Block 3. **Cue.** The exponent $(k+1)/k$, delivered by the separation and by the target. **Why natural.** Do not raise the recursion to the power m blindly; first factor a_n out of the right hand side, because $a_n^{-1/k} = a_n \cdot a_n^{-m}$ turns the additive recursion into a multiplicative one, $a_{n+1} = a_n(1 + a_n^{-m})$. Multiplicative recursions survive being raised to a power, additive ones do not. After raising, $a_n^{-m} = 1/b_n$, so the new recursion mentions b_n and nothing else. That closure is the point: the recursion for b is autonomous, and its increment involves only $1/b_n$. **Link to the next block.** The increment is now an explicit function of one small quantity, so take its limit.

Block 4. **Cue.** The shape $\big((1+t)^m - 1\big)/t$ with $t \to 0$. **Why natural.** This is the difference quotient of $t \mapsto (1+t)^m$ at $t = 0$, so its limit is the derivative m by definition. Recognising it as such removes the need for Taylor's theorem with remainder entirely; a remainder bound would be needed only if you wanted a rate, and you do not. Train the reflex: whenever a limit has the form $(\varphi(t) - \varphi(0))/t$ with $t \to 0$, stop estimating and differentiate. **Link to the next block.** Increments converge, so the sequence itself grows linearly, provided you have a theorem that says so.

Block 5. **Cue.** A sequence whose consecutive differences converge. **Why natural.** A sequence is the telescoping sum of its increments, so b_n/n is the average of the first n increments up to the vanishing term b_0/n. Averages preserve limits, which is the Cesaro mean lemma, and it is worth proving in three lines rather than citing, both because a grader may want it and because the proof is the same ε split you use everywhere. This is the special case $b_n = n$ of the Stolz and Cesaro theorem. **Link to the next block.** You have $b_n/n \to m$, and the requested expression is a fixed power of b_n/n.

Block 6. **Cue.** $a_n^{k+1} = \big(a_n^m\big)^k$ because $mk = k+1$. **Why natural.** The problem asked for a_n^{k+1}/n^k rather than $a_n^{(k+1)/k}/n$ purely to disguise the substitution; undoing the disguise is one line, and continuity of $t \mapsto t^k$ transports the limit. **Link to the next block.** Done, with the answer in the form $\big(1 + 1/k\big)^k$.

Layer 3: WHY

The mathematical truth here is that the ODE heuristic stops being a heuristic the moment you name the change of variable, and it is worth stating the general theorem because it swallows the entire problem class.

Proposition 20.12.3. *Let $f : (0,\infty) \to (0,\infty)$ be decreasing, let $a_0 > 0$, and let $a_{n+1} = a_n + f(a_n)$. Suppose $a_n \to \infty$ and $f(a + f(a))/f(a) \to 1$ as $a \to \infty$. Define $F(a) = \int_{a_0}^{a} \frac{dt}{f(t)}$. Then $F(a_n)/n \to 1$.*

Proof. $F(a_{n+1}) - F(a_n) = \int_{a_n}^{a_n + f(a_n)} \frac{dt}{f(t)}$. On that interval of length $f(a_n)$, monotonicity of f gives $f(a_n + f(a_n)) \le f(t) \le f(a_n)$, so

$$1 = \frac{f(a_n)}{f(a_n)} \le F(a_{n+1}) - F(a_n) \le \frac{f(a_n)}{f(a_n + f(a_n))}.$$

The right hand side tends to 1 by hypothesis, so the increments of $F(a_n)$ tend to 1, and the Cesaro mean lemma gives $F(a_n)/n \to 1$. □

Apply this with $f(a) = a^{-1/k}$: then $F(a) = \frac{k}{k+1}\left(a^{(k+1)/k} - a_0^{(k+1)/k}\right)$ and $f(a + f(a))/f(a) = (1 + a^{-(k+1)/k})^{-1/k} \to 1$, so $\frac{k}{k+1}a_n^{(k+1)/k}/n \to 1$, which is Step 3 again. The proposition is worth more than the problem: it tells you that for $a_{n+1} = a_n + e^{-a_n}$ you get $e^{a_n} \sim n$, hence $a_n \sim \log n$, and for $a_{n+1} = a_n + a_n^{-p}$ you get $a_n^{p+1} \sim (p+1)n$, with no new work in either case.

The discernment mechanism, the thing that decides between estimating and substituting, has three signals and you should be able to name all three. First, the increment is a function of the current value alone, which is what makes an ODE comparison legitimate rather than decorative. Second, the increment function is monotone, which is exactly the hypothesis the integral squeeze in the proposition needs; if f oscillated, the squeeze would fail and you would be back to estimating. Third, the question supplied the normalisation a_n^{k+1}/n^k, which tells you which power of a_n is the natural variable, so you are being handed the change of coordinates. When all three are present, never estimate. Change variable and let the recursion become trivial.

There is a second, smaller truth worth extracting: the difference between needing a rate and needing a limit. The official first solution invokes Taylor's theorem with remainder to bound $\left|a_{n+1}^m - a_n^m - m\right|$ by a multiple of a_n^{-m}. That bound is true and it is more than the problem asks, since you only need the difference to converge, and convergence of $\left((1+t)^m - 1\right)/t$ is the definition of a derivative. Every time you reach for a quantitative estimate, ask whether the qualitative statement suffices; if it does, the qualitative route is shorter and has fewer places to make an arithmetic slip. That instinct is worth several marks a paper.

Knowledge gaps this problem exposes

Say plainly what must be installed before this problem is even approachable, because pretending otherwise wastes your training time. You need the Cesaro mean lemma and preferably the Stolz and Cesaro theorem it generalises to; neither is in the A level or Further Maths syllabus, and without one of them you cannot get from "increments converge" to "the sequence is asymptotically linear". You need the derivative as a limit of difference quotients used in reverse, that is, recognising an unfamiliar quotient as a derivative you already know. You need separable first order ODEs, which you have, but used as an asymptotic oracle for a discrete process, which you do not. And you need comfort with the idea that a change of variable can be chosen so that a recursion becomes trivial, rather than being chosen to simplify an expression. That last one is a habit, not a theorem, and it is the one that takes longest to install.

Revolving concepts

Monotone sequences. An increasing sequence either converges or tends to infinity, and a convergent sequence has increments tending to zero. Together these give the standard divergence test used in Step 1. See Book 1, field: Sequences and recurrences, topic: monotone convergence.

The derivative as a limit of difference quotients. $\varphi'(0) = \lim_{t\to 0}\left(\varphi(t) - \varphi(0)\right)/t$. Used forwards this defines the derivative; used backwards it evaluates limits such as $\left((1+t)^m - 1\right)/t \to m$ with no expansion required. See Book 1, field: Calculus and integrals, topic: differentiation from first principles.

Cesaro means and the Stolz and Cesaro theorem. If $x_i \to c$ then the averages converge to c. The general theorem: if (B_n) is positive, strictly increasing and unbounded and $(A_{n+1} - A_n)/(B_{n+1} - B_n) \to L$, then $A_n/B_n \to L$; the case $B_n = n$ is the mean lemma. This is the discrete analogue of l'Hopital's rule, and that analogy is the right way to remember when to reach for it. See Book 1, field: Real analysis, topic: Cesaro summation.

Separable first order ODEs. $y' = g(y)$ separates as $\int dy/g(y) = x + C$. The antiderivative $F(y) = \int dy/g(y)$ is the coordinate in which the flow has unit speed, which is exactly why it is the right substitution for the discrete problem too. See Book 1, field: Calculus and integrals, topic: separable equations.

The general recursion, generalised. The recursion $a_{n+1} = a_n + f(a_n)$ with f positive and decreasing is the whole family this problem lives in, and the proposition above is its general solution: with $F(a) = \int^a dt/f(t)$ you get $F(a_n) \sim n$, that is, $a_n \sim F^{-1}(n)$. Instances to keep on hand: $f(a) = a^{-p}$ gives $a_n \sim \big((p+1)n\big)^{1/(p+1)}$; $f(a) = e^{-a}$ gives $a_n \sim \log n$; $f(a) = 1/(a \log a)$ gives $\frac{1}{2}a_n^2 \log a_n \sim n$. The same idea in reverse handles decreasing sequences such as $a_{n+1} = a_n - a_n^2$, where $F(a) = -1/a$ gives $1/a_n \sim n$, hence $a_n \sim 1/n$; that variant appears on the Putnam and on olympiads at least as often as the increasing one, so install both directions. See Book 1, field: Sequences and recurrences, topic: asymptotics of first order recursions.

Circuit Bank

- **WHEN** a recursion has the form $a_{n+1} = a_n + f(a_n)$ with f monotone **DO** write the ODE $y' = f(y)$, separate it, and use $F(a) = \int da/f(a)$ as the new variable **BECAUSE** in that variable the recursion becomes "add approximately 1", and adding a constant is the only recursion whose asymptotics come for free
- **WHEN** the consecutive differences of a sequence converge **DO** conclude the sequence divided by n converges to the same limit **BECAUSE** the sequence is the telescoping sum of its increments and Cesaro averaging preserves limits, which is the case $B_n = n$ of Stolz and Cesaro
- **WHEN** the question hands you a normalisation such as a_n^{k+1}/n^k **DO** read the growth exponent and the substitution straight off it before doing any work **BECAUSE** a finite nonzero limit forces the growth rate, so the examiner has already told you which power of a_n is the natural variable
- **WHEN** a limit has the shape $(\varphi(t) - \varphi(0))/t$ with $t \to 0$ **DO** call it a derivative rather than expanding **BECAUSE** recognising a difference quotient removes every remainder estimate, and you only need a rate if the problem asks for one
- **WHEN** an additive recursion must be raised to a power **DO** factor the leading term out first to make it multiplicative **BECAUSE** powers pass through products but not through sums, so $a_{n+1} = a_n(1 + a_n^{-m})$ survives being raised to the power m while $a_n + a_n^{-1/k}$ does not

Backtesting

1. First useful cue: the increment depends only on a_n, so the sequence is a discretised separable ODE, and the requested ratio names the exponent $(k+1)/k$.
2. Method the expression suggested: substitute $b_n = a_n^{(k+1)/k}$, the variable into which $y' = y^{-1/k}$ separates, and show its increments converge.
3. Where you would hesitate, and the thought to reinforce: at raising the recursion to a fractional power, which feels illegitimate on a sum. Reinforce the fix rather than the fear: factor a_n out first so the recursion is multiplicative, and then the power passes through cleanly.
4. The trap the question sets: the exponents. Writing the target as a_n^{k+1}/n^k rather than as a power of $a_n^{(k+1)/k}/n$ hides the substitution, and mishandling $mk = k+1$ at the last step throws away the whole solution.
5. What makes you recognise it faster next time: any recursion of the form "current value plus a decreasing function of the current value". That shape means change variable by $\int da/f(a)$, always.
6. The reusable decision rule: for $a_{n+1} = a_n + f(a_n)$ with f positive, decreasing, and satisfying $f(a + f(a))/f(a) \to 1$, you have $F(a_n) \sim n$ where $F(a) = \int^a dt/f(t)$.

Mirror drills

1. Let $a_0 > 0$ and $a_{n+1} = a_n + 1/a_n$. Evaluate $\lim a_n^2/n$. *Answer:* $F(a) = a^2/2$, so $a_n^2/(2n) \to 1$ and the limit is 2, which is the case $k = 1$ of the problem.

2. Let $a_{n+1} = a_n + e^{-a_n}$ with a_0 real. Find the asymptotic behaviour of a_n. *Answer:* increments are positive and $a_n \to \infty$ by the same divergence test, so with $F(a) = e^a$ you get $e^{a_n}/n \to 1$, that is $a_n - \log n \to 0$.

3. Let $x_0 \in (0, 1)$ and $x_{n+1} = x_n - x_n^2$. Find $\lim n x_n$. *Answer:* x_n decreases to 0, and $1/x_{n+1} - 1/x_n = 1/(1 - x_n) \to 1$, so by the Cesaro mean lemma $1/(n x_n) \to 1$ and the limit is 1. Same circuit, decreasing direction.