DAWN PUTNAM SESSION

Putnam 1996 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1996 — Paper A **5**
1996 A1 5
1996 A2 7
1996 A3 10
1996 A4 12
1996 A5 15
1996 A6 18

Putnam 1996 — Paper B **23**
1996 B1 23
1996 B2 26
1996 B3 29
1996 B4 32
1996 B5 35
1996 B6 38

Putnam 1996 — Paper A

1996 A1

1996 A1. Find the least number A such that for any two squares of combined area 1, a rectangle of area A exists such that the two squares can be packed in the rectangle (without interior overlap). You may assume that the sides of the squares are parallel to the sides of the rectangle.

Fields: plane geometry, constrained optimisation, trigonometric substitution **Position:** A1 — the opener **Answer:** $A = \dfrac{1+\sqrt{2}}{2}$, attained by packing squares of side $\cos(\pi/8)$ and $\sin(\pi/8)$ side by side.

Layer 1 THE WHAT — surface analysis

The natural attempt. Let the squares have sides $x \geq y$ with $x^2 + y^2 = 1$. The obvious packing is to stand them side by side: a rectangle of width $x + y$ and height x (the taller square sets the height, since $y \leq x$) holds both with no wasted motion. Its area is $x(x + y)$.

Maximise this over the constraint. Write $x = \cos\theta$, $y = \sin\theta$, $\theta \in [0, \pi/4]$ (the range that keeps $x \geq y$). Then

$$x(x+y) = \cos^2\theta + \sin\theta\cos\theta = \tfrac{1}{2}(1 + \cos 2\theta + \sin 2\theta) = \tfrac{1}{2} + \tfrac{\sqrt{2}}{2}\cos\big(2\theta - \tfrac{\pi}{4}\big)\,,$$

which is maximised when $2\theta - \pi/4 = 0$, i.e. $\theta = \pi/8$, giving value $\frac{1+\sqrt{2}}{2}$. Since $\pi/8 \in [0, \pi/4]$ this is attained inside the domain, so $A = \frac{1+\sqrt{2}}{2}$.

This is the right number, reached by the right computation. But stop and ask what has actually been shown: that a rectangle of area $x(x + y)$ *suffices* for every pair (x, y), and that the largest such requirement, over all valid pairs, is $\frac{1+\sqrt{2}}{2}$. Nothing here has shown that $x(x + y)$ is the *least* area a rectangle can have and still contain both squares. If some cleverer, non-side-by-side arrangement could pack the worst-case pair into a rectangle of smaller area, then A could be smaller than $\frac{1+\sqrt{2}}{2}$, and the whole computation above — correct as an upper bound — would be answering the wrong question. "Find the least A" is a two-sided claim: A works (upper bound) and no smaller number works (lower bound, i.e. a genuine *obstruction* to packing more cheaply). The natural attempt supplies only the first half.

The model chain. Fix a valid pair $x \geq y > 0$ with $x^2 + y^2 = 1$, and let $W \times H$ (with $W \geq H$) be any axis-aligned rectangle containing both squares without interior overlap. Since the larger square has side x and must fit inside the rectangle in both directions, $H \geq x$.

Claim: $W \geq x + y$. Suppose instead $W < x + y$. Project both squares onto the W-axis: the two projections are intervals of lengths x and y inside $[0, W]$. If these intervals were disjoint, their lengths would sum to at most W, i.e. $x + y \leq W$ — contradicting the assumption. So the projections overlap: there is some coordinate t at which both squares are present. At that coordinate, take the vertical slice through t. Because each square is a full axis-aligned rectangle, its intersection with this vertical line is a segment whose length equals that square's *side*, not some smaller cross-section — the square of side x contributes a vertical segment of length x,

the square of side y contributes one of length y, both inside $[0, H]$. Since the two squares do not overlap, these two vertical segments at the same horizontal position t must themselves be disjoint, forcing $x + y \leq H$. But $H \leq W < x + y$ by hypothesis, giving $x + y \leq H \leq W < x + y$ — a contradiction. Hence $W \geq x + y$.

So for every valid (x, y), any containing rectangle has area $WH \geq x(x + y)$, and the side-by-side packing shows this is achieved. Thus $x(x + y)$ is exactly the minimal rectangle area required for that pair — not merely a convenient one — and the answer to the problem is $A = \max_{x^2+y^2=1,\, x\geq y} x(x + y) = \frac{1+\sqrt{2}}{2}$, by the same trigonometric maximisation as before, now legitimately interpreted as a least upper bound over genuine minima.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write down $x(x + y)$ and maximise it identically via $x = \cos\theta$, $y = \sin\theta$, landing on $\theta = \pi/8$ and $A = (1 + \sqrt{2})/2$. They split before any of that, at the moment of deciding what $x(x + y)$ *is*. The natural attempt takes it as the area of *a* valid packing and stops. The model chain first proves — via the projection argument — that $x(x + y)$ is a lower bound on *every* valid packing's area, so that the side-by-side packing is not merely convenient but optimal, and only then is the maximisation entitled to be called "the least A."

> *I have shown a rectangle of area $x(x + y)$ works; that is not the same statement as "no smaller rectangle works." To justify calling my maximised value the* least *A, I need an argument that starts from an arbitrary containing rectangle — not the one I built — and forces its area to be at least $x(x + y)$ regardless of how cleverly the squares are arranged inside it.*

Layer 3 THE WHY

Why does the longer side have to be at least $x + y$, and why is a projection argument the right tool to prove it rather than, say, another explicit packing? Kedlaya and Bhargava's own solution states the bound — "the shorter side must be at least x, and the longer side must be at least $x + y$" — as if self-evident, and moves straight to the maximisation; this is exactly the point where the archive has to supply what the official solution omits.

The fact worth isolating is this: an axis-aligned square's intersection with *any* vertical line through its interior is a segment of length equal to the square's full side. This is just the definition of an axis-aligned square as a product of two intervals, $[a, a + x] \times [b, b + x]$ — every vertical slice through it has the same height x, because the square doesn't taper. That single structural fact is what turns "the horizontal projections of the two squares overlap" into "the two squares occupy the same full-height strips at that horizontal position," which is exactly a non-overlap violation unless the vertical extents are disjoint. Nothing about squares in particular is used beyond this — the same argument proves the corresponding bound for any two axis-aligned rectangles.

The deeper reason the natural attempt stops early is that "obviously the tightest packing is side by side" is a claim about geometry, and STYLE forbids treating obviousness as an argument — a step that looks self-evident is precisely the step owing the reader a proof. Here the proof is short (one pigeonhole observation on two interval projections) precisely because the problem's structure — axis-aligned squares, axis-aligned rectangle — makes the projection argument commute with the packing constraint exactly. Recognising that a $W \times H$ non-overlap statement can be attacked one axis at a time, by projecting and then slicing, is the reusable content; the specific numbers $x, y, 1, 3$ elsewhere in the archive will differ, but the projection-then-slice move for axis-aligned non-overlapping rectangles will not.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "the least number A." On seeing "least"/"greatest" attached to a packing or covering claim, write down immediately that the answer needs both a

construction (upper bound) and an impossibility argument (lower bound) — a single packing, however natural, only ever supplies the first.

2. That cue suggests treating the problem in two separate phases: first find the minimal rectangle area for a *fixed* pair (x, y) (which requires proving a bound, not just building a packing), and only then optimise that minimal-area function over the constraint $x^2 + y^2 = 1$.

3. Hesitation occurs right after the side-by-side packing is written down and the maximisation goes through cleanly — the arithmetic "worked," so it is tempting to declare victory. This is reinforced away by asking explicitly: "have I shown no smaller rectangle can work for this (x, y), or have I only shown one that does?"

4. The trap is that the maximisation is genuinely correct and produces the genuinely correct final number, so a solution that skips the lower-bound proof entirely can still submit the right numerical answer while having proved nothing about optimality — the trap is invisible in the answer, only visible in the argument.

5. Recognise this family faster: whenever two or more axis-aligned shapes must be packed without overlap inside a box, and a claim is made about the box's minimal dimensions, project onto one axis first. Overlapping projections plus the non-overlap hypothesis is the standard route to a lower bound on the other axis.

6. Reusable rule: "least A such that a packing always exists" problems split cleanly into (i) an explicit construction achieving some value, and (ii) an argument, independent of that construction, that no smaller value can ever suffice. Only compute the optimisation of stage (i) after stage (ii) is in hand — otherwise the arithmetic answers a question ("what does this packing cost") different from the one asked ("what must any packing cost").

▷ **Cue:** a problem asks for the least/greatest quantity that always suffices/is always achievable across a family of configurations $\longrightarrow$ **Move:** split immediately into a construction (upper bound, from one explicit configuration) and a separate impossibility argument (lower bound, holding for an arbitrary configuration) before optimising either

▷ **Cue:** two or more axis-aligned rectangles/squares are placed without interior overlap inside a bounding rectangle $\longrightarrow$ **Move:** project onto one axis; if the projected interval lengths sum to more than that axis's extent, pigeonhole forces the projections to overlap, and slicing at the overlap point turns non-overlap into a bound on the other axis

1996 A2

1996 A2. Let C_1 and C_2 be circles whose centers are 10 units apart, and whose radii are 1 and 3. Find, with proof, the locus of all points M for which there exists points X on C_1 and Y on C_2 such that M is the midpoint of the line segment XY.

Fields: plane geometry, vectors, locus problems **Position:** A2 **Answer:** The locus is the annulus centred at the midpoint of the two circles' centres, with inner radius 1 and outer radius 2.

Layer 1 THE WHAT — surface analysis

The natural attempt. Place the centres O_1, O_2 with $|O_1O_2| = 10$. Try the two extreme, easy-to-picture configurations first: put X and Y both on the line through the centres. If X, Y are positioned so that the segment XY is as long as possible along that line (radii pointing away from each other), the midpoint sits furthest from the midpoint of O_1O_2; if XY is as short as possible (radii pointing toward each other, on the near sides), the midpoint sits closest. A quick computation along the line gives the outer distance $\frac{1+3}{2} = 2$ and the inner distance $\frac{3-1}{2} = 1$ from the midpoint of O_1O_2.

That already matches the claimed answer, and it is tempting to declare the locus is the annulus with radii 1 and 2 on the strength of these two collinear cases plus "by symmetry, everything in between is filled in too." But that last clause is doing all of the real work and none

of it has been justified. Two questions are open: does every direction from the centre get reached (or only the direction along O_1O_2, rotated somehow), and does every intermediate distance between 1 and 2 actually occur, or could there be a gap? Testing more points by hand — say X, Y perpendicular to the line of centres — will keep confirming the guess, but confirming a guess at finitely many points is not a proof of a two-dimensional region.

The model chain. Put O_1 at the origin and work in vectors. Any point X on C_1 is $X = O_1 + u$ for a unit vector u; any point Y on C_2 is $Y = O_2 + 3v$ for a unit vector v, with u, v ranging *independently* over all directions (there is no constraint linking them — X and Y move on their circles with no reference to each other beyond both existing). Then

$$M = \frac{X+Y}{2} = \frac{O_1+O_2}{2} + \frac{1}{2}u + \frac{3}{2}v.$$

Write $c = \frac{O_1+O_2}{2}$ for the midpoint of the two centres. The problem reduces entirely to: what is the set of values of $\frac{1}{2}u + \frac{3}{2}v$ as u, v independently range over all unit vectors?

Fix the angle φ between u and v. Then

$$\left|\tfrac{1}{2}u + \tfrac{3}{2}v\right|^2 = \tfrac{1}{4} + \tfrac{9}{4} + \tfrac{3}{2}\cos\varphi = \tfrac{5}{2} + \tfrac{3}{2}\cos\varphi,$$

which is a continuous, strictly monotonic function of $\varphi \in [0, \pi]$, running from $\left(\frac{5}{2} - \frac{3}{2}\right) = 1$ at $\varphi = \pi$ up to $\left(\frac{5}{2} + \frac{3}{2}\right) = 4$ at $\varphi = 0$. So $\left|\frac{1}{2}u + \frac{3}{2}v\right|$ takes every value in $[1, 2]$ as φ ranges over $[0, \pi]$, by the intermediate value theorem — every radius between 1 and 2 is genuinely attained, not merely the two extremes.

Now fix that radius $\rho \in [1, 2]$, realised by some particular pair (u_0, v_0) with angle φ between them. Rotating the whole plane by an arbitrary angle ψ sends (u_0, v_0) to $(R_\psi u_0, R_\psi v_0)$ — still a valid pair of unit vectors at the same relative angle φ, hence still giving $\left|\frac{1}{2}u + \frac{3}{2}v\right| = \rho$, but now in the rotated direction $R_\psi\left(\frac{1}{2}u_0 + \frac{3}{2}v_0\right)$. As ψ ranges over $[0, 2\pi)$ this sweeps out the entire circle of radius ρ about c. So every point of every circle of radius $\rho \in [1, 2]$ about c is attained, and (since $|u|, |v| = 1$ always) no point outside that range is ever attained. The locus is exactly the annulus $1 \le |M - c| \le 2$ centred at c, the midpoint of O_1O_2.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains find the same two numbers, 1 and 2, from the same collinear configurations. They split immediately after: at the sentence "by symmetry, everything in between is filled in." The natural attempt treats that sentence as a conclusion; the model chain treats it as two separate claims requiring two separate arguments — that every intermediate *radius* occurs (an IVT argument on the angle between u and v) and that every *direction* occurs for each such radius (a rotation argument) — and proves each one.

> *"By symmetry" is not an argument until I can name the group acting and show it acts transitively on the thing I need covered. Here there are two separate degrees of freedom — the angle between u and v, which controls the* magnitude *of $M - c$, and the overall rotation of the pair, which controls its* direction *— and I need to treat them as two different variables with two different jobs, not fold them into one hand-wave.*

Layer 3 THE WHY

Why does the problem reduce to a statement about sums of two independent unit vectors at all, and why is that the natural object once you see it? Because a midpoint of two points is a linear (affine) function of those points, and X, Y are themselves parametrised affinely by unit vectors around their respective centres — $X = O_1 + u$, $Y = O_2 + 3v$ — the midpoint operation commutes with this parametrisation exactly: $M = c + \frac{1}{2}u + \frac{3}{2}v$ separates cleanly into a fixed point c plus a term depending only on the two *directions*, with no leftover dependence on O_1, O_2

individually except through their average. Once the problem has been reduced to "what set does $r_1u + r_2v$ sweep out, for fixed lengths r_1, r_2 and free directions u, v", it is a named, general fact — the reachable set of a sum of two vectors of fixed magnitude and free direction is always an annulus of inner radius $|r_1 - r_2|$ and outer radius $r_1 + r_2$ (or the single circle of radius $r_1 + r_2$ if $r_1 = r_2$, since the inner radius collapses to 0, filling the disk) — and this problem is simply that fact instantiated at $r_1 = \frac{1}{2}, r_2 = \frac{3}{2}$.

This is also why the number 10 — the distance between the centres — never appears in the final answer, and is worth flagging explicitly: it is a decoy. The distance $|O_1O_2|$ only ever enters the calculation through $c = (O_1 + O_2)/2$, fixing *where* the annulus sits, and through ensuring the two circles are far enough apart that the configuration is a genuinely separate pair (though even this is not actually needed by the algebra above — the vector argument works regardless of $|O_1O_2|$, including if the circles overlapped or coincided). A solver who spends time on the "10" — for instance, checking whether C_1 and C_2 intersect, or trying to use the specific value 10 in the radius computation — is solving a problem the given data does not actually pose. Kedlaya and Bhargava's solution reaches the same place via homothety (scaling C_1 by $1/2$ about a fixed Y, then letting the resulting small circle's centre trace a circle C_3 as Y varies); that route also never touches the number 10, for the same underlying reason — position information washes out of a problem that is really about the free superposition of two fixed-length directions.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "midpoint of X on one circle and Y on another, with X, Y otherwise unconstrained." On seeing it, write M as an affine combination immediately: $M = \frac{1}{2}X + \frac{1}{2}Y$, then substitute the circle parametrisations $X = O_1 + r_1u$, $Y = O_2 + r_2v$ before doing anything else.

2. That cue suggests the problem is really "what set does $r_1u + r_2v$ sweep out, for fixed r_1, r_2 and free unit vectors u, v" — a two-parameter family (the angle between u, v, and the overall rotation) rather than a locus to be plotted point by point.

3. Hesitation occurs right after finding the two extreme radii, 1 and 2, from collinear configurations — the temptation is to call the problem solved. This is reinforced away by separating "what radii occur" (an IVT argument in the relative angle φ) from "what directions occur for a given radius" (a rotation argument), and proving each explicitly rather than invoking symmetry once for both.

4. The trap is the number 10: it is present, specific, and looks load-bearing, but only ever enters through the location of the annulus's centre, never its radii. A solver who tries to fold 10 into the radius computation is answering a question the vector reduction shows does not depend on it.

5. Recognise this family faster: any locus problem built from midpoints (or more generally affine combinations) of points on two circles reduces immediately to the reachable set of $r_1u + r_2v$ for fixed lengths and free directions — memorise that this reachable set is always the annulus of radii $|r_1 - r_2|$ to $r_1 + r_2$, so the answer can be written down as soon as the two radii and the combination coefficients are identified.

6. Reusable rule: when a locus is defined by an affine combination of points ranging independently over circles, separate the centre-dependent part (which only fixes where the locus sits) from the direction-dependent part (which determines its shape), and treat the direction-dependent part as the general fact "$r_1u + r_2v$ sweeps the annulus $[|r_1 - r_2|, r_1 + r_2]$" rather than re-deriving it by testing sample configurations.

▷ **Cue:** a locus is defined as an affine combination (midpoint, weighted average) of points ranging independently over two circles $\longrightarrow$ **Move:** write the combination as (fixed centre-dependent point) + $r_1u + r_2v$ for independent unit vectors u, v; the shape of the locus is exactly the reachable set of $r_1u + r_2v$

▷ **Cue:** the reachable set of $r_1u + r_2v$ for fixed lengths r_1, r_2 and free independent directions u, v $\longrightarrow$ **Move:** it is the annulus of inner radius $|r_1 - r_2|$ and outer radius $r_1 + r_2$ (a disk if

$r_1 = r_2$), proved by an IVT sweep over the angle between u, v for the radius, then a rotation of the whole pair for the direction

1996 A3

1996 A3. Suppose that each of 20 students has made a choice of anywhere from 0 to 6 courses from a total of 6 courses offered. Prove or disprove: there are 5 students and 2 courses such that all 5 have chosen both courses or all 5 have chosen neither course.

Fields: combinatorics, double counting, extremal constructions **Position:** A3 **Answer:** Disprove. Give each of the 20 students a distinct 3-element subset of the 6 courses as their choice (there are exactly $\binom{6}{3} = 20$ such subsets); then every pair of courses is chosen by exactly 4 students and avoided by exactly 4, never reaching 5.

Layer 1 THE WHAT — surface analysis

The natural attempt. "Prove or disprove" with a specific count (20 students, 6 courses, threshold 5) reads like a pigeonhole problem, and the honest first move is to try to prove it. Fix a pair of courses $\{p, q\}$. Each student, restricted to this pair, falls into one of four categories: chose both, chose only p, chose only q, chose neither. Twenty students into four categories: by pigeonhole some category has at least $\lceil 20/4 \rceil = 5$ students. That looks like it is nearly the claim — except the pigeonholed category might be "only p" or "only q," neither of which is "both" or "neither." So a single pair does not finish the job; try summing over all $\binom{6}{2} = 15$ pairs instead.

For a student who has chosen k courses, count how many of the 15 pairs put that student in the "both" or "neither" bucket: it is $\binom{k}{2}$ (pairs entirely inside their chosen courses) plus $\binom{6-k}{2}$ (pairs entirely outside). Call this $f(k)$. Summing $f(k)$ over all 20 students and comparing against 15 pairs $\times$ (at most 4 "both" + at most 4 "neither," if the claim is to be disproved, i.e. no pair reaches 5 in either bucket) should either force a contradiction (proving the claim) or fail to (suggesting a counterexample exists). Computing $f(0), \ldots, f(6) = 15, 10, 7, 6, 7, 10, 15$ shows the minimum is $f(3) = 6$, attained only at $k = 3$. So the sum over 20 students of $f(k)$ is at least $20 \times 6 = 120$. If no pair is to reach 5 in "both" or "neither," each pair contributes at most $4 + 4 = 8$ to the same sum, and 15 pairs give at most $15 \times 8 = 120$.

$120 \leq 120$. The inequality closes with equality, not with a contradiction. A student stopping here, expecting the pigeonhole to break and prove the claim, has instead proved nothing — the bound is exactly tight, and tightness is not a contradiction.

The model chain. The equality $120 = 120$ forced above is not incidental: it says that *if* a configuration exists with no pair reaching 5 in either bucket, then every inequality used to derive it must be an equality throughout. In particular every student must take exactly $k = 3$ courses (since $f(k) > 6$ for every $k \neq 3$, and the sum can only reach the minimum possible total 120 if every single term already equals its minimum 6), and every pair of courses must be chosen by exactly 4 students and avoided by exactly 4 (since each pair's contribution must hit its ceiling of 8 exactly). So the counting doesn't just fail to prove the claim — it prescribes, uniquely up to which sets go to which students, the shape any counterexample must have.

Realise that shape directly: assign to the 20 students the 20 distinct 3-element subsets of the 6 courses — there are exactly $\binom{6}{3} = 20$ of them, matching the student count exactly, so every 3-subset is used precisely once. Fix any pair of courses $\{p, q\}$. A student's 3-subset contains both p and q exactly when its third course is chosen from the remaining 4 courses: $\binom{4}{1} = 4$ such subsets, so exactly 4 students chose both. A student's 3-subset contains neither p nor q exactly when all three of its courses come from the remaining 4: $\binom{4}{3} = 4$ such subsets, so exactly 4 students chose neither. Both counts are 4, strictly less than 5, for every one of the 15 pairs. This is a complete counterexample: the claim "there are 5 students and 2 courses such that all 5 chose both or all 5 chose neither" is false.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains carry out the same double count and arrive at the same numbers: $f(k)$ minimised at $k = 3$ with value 6, giving $20 \times 6 \leq$ total $\leq 15 \times 8 = 120$. They split at the line "$120 \leq 120$." The natural attempt reads this as a failed attempt to force a contradiction and stops, having learned only that pigeonhole "almost" worked. The model chain reads the same equality as a specification: an inequality collapsing to equality forces every intermediate inequality to be tight, and that tightness pins down exactly what a counterexample must look like (every student takes 3 courses; every pair splits 4-4) — turning a failed proof directly into a blueprint for a construction.

> *The double count came out to* $120 \leq 120$*, not* $120 < 120$ *and not* $120 > 120$ *— an exact tie, not a near-miss in either direction. A bound that closes with equality is not weaker evidence for the claim than a strict inequality would have been; it is stronger evidence for the opposite, because equality forces every inequality inside the argument to be saturated simultaneously, and "every inequality saturated" is a full specification of a configuration, not a vague hint. I should read equality in a double-counting bound as an instruction to go build the object the equality describes, not as a stalled proof to keep pushing on.*

Layer 3 THE WHY

Why does the double count bottom out exactly at 120, matching 15×8 exactly, rather than at some number that would either prove the claim outright or leave slack for many different counterexamples? Two facts conspire, and neither is an accident of the problem's phrasing.

First, $f(k) = \binom{k}{2} + \binom{6-k}{2}$ is a convex function of k symmetric about $k = 3$ (since $f(k) = f(6-k)$), so its minimum over integers $0, \ldots, 6$ is attained uniquely at the centre $k = 3$ — this is the general fact that "both inside a k-set or both outside it," summed over $\binom{6}{2}$ pairs, is minimised by centring k at half the total, exactly the pigeonhole intuition that spreading choices evenly minimises correlation between any two fixed items. This is why 20 students each choosing 3 of 6 courses is the *cheapest* possible configuration to build, in the currency of "both-or-neither" incidences — any other uniform choice size would force some pair to accumulate more incidences, not fewer.

Second, and this is the number that makes the problem work: $\binom{6}{3} = 20$ exactly. The problem does not merely ask about 6 courses and some students; it specifies precisely 20, which is precisely the number of 3-element subsets of a 6-element set. This is not incidental — it is why the equality $120 = 120$ closes with no slack, and why the counterexample construction (one student per 3-subset) is not just *a* valid choice but is, in outline, forced: every valid counterexample must have every student choosing exactly 3 courses and every pair split exactly 4–4, and using each 3-subset exactly once is the natural way to realise that (Kedlaya and Bhargava's solution notes, as a remark, that this counterexample is unique up to permuting students, provided no two students choose the same set — a fact the double-counting bound above makes structurally inevitable rather than a separate check). A student who notices the coincidence $\binom{6}{3} = 20$ before starting the double count has a considerable head start: the natural conjecture "one student per 3-subset" can be checked directly in two lines ($\binom{4}{1} = 4$, $\binom{4}{3} = 4$), and the double-counting argument becomes confirmation rather than discovery.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the numeric coincidence $\binom{6}{3} = 20$: the course count and the "choose half" binomial coefficient exactly match the student count. On seeing it, write down the conjecture "one student per 3-element subset of the 6 courses" immediately, before attempting a general proof either way.

2. That cue suggests checking the conjectured construction directly by a small combinatorial count (both = $\binom{4}{1}$, neither = $\binom{4}{3}$) rather than launching into a general pigeonhole/double-

counting proof first — the construction is two lines; a full proof of impossibility (which does not exist here) would be far longer.

3. Hesitation occurs at the moment the double-counting bound closes with $120 = 120$ rather than a strict contradiction. This is reinforced away by treating equality in a maximised/minimised double count as diagnostic: it says every term in the sum is individually extremal, which is a specification of a structure, not an inconclusive result to abandon.

4. The trap is the single-pair pigeonhole argument that looks almost sufficient (20 students into 4 categories forces one category ≥ 5) but never controls *which* category is large — "only p" or "only q" satisfy the pigeonhole without satisfying the problem's claim, and chasing this line further (summing over 15 pairs) is exactly what is needed, but only if the resulting equality is read correctly.

5. Recognise this family faster: whenever a "prove or disprove" problem pairs a specific object count with a binomial coefficient of the same problem's other parameters ($\binom{6}{3} = 20$ here), test the "one object per subset of that size" construction before anything else — the coincidence is very rarely accidental in a competition problem.

6. Reusable rule: when a double-counting or pigeonhole bound closes with equality rather than a strict contradiction, do not read it as a failed attempt at a proof — read it as a full specification of the extremal configuration (every inequality inside it must itself be tight), and go construct that configuration directly rather than searching for a sharper bound that does not exist.

▷ **Cue:** a "prove or disprove" combinatorics problem gives a specific object count that coincides with a binomial coefficient built from the problem's other parameters $\longrightarrow$ **Move:** test the natural "one object per subset of that size" construction directly, before attempting a general proof in either direction

▷ **Cue:** a double-counting or pigeonhole bound closes with an exact equality rather than a strict contradiction $\longrightarrow$ **Move:** read the equality as forcing every inequality inside the argument to be individually tight, which specifies the extremal configuration exactly; go build that configuration rather than continuing to search for a contradiction

1996 A4

1996 A4. Let S be the set of ordered triples (a, b, c) of distinct elements of a finite set A. Suppose that

1. $(a, b, c) \in S$ if and only if $(b, c, a) \in S$;
2. $(a, b, c) \in S$ if and only if $(c, b, a) \notin S$;
3. (a, b, c) and (c, d, a) are both in S if and only if (b, c, d) and (d, a, b) are both in S.

Prove that there exists a one-to-one function g from A to R such that $g(a) < g(b) < g(c)$ implies $(a, b, c) \in S$. Note: R is the set of real numbers.

Fields: combinatorics (relations, orderings); algebra of axioms **Position:** A4 — a genuine middle-paper problem whose difficulty is entirely in decoding what the three axioms describe; once decoded, the proof is ten lines. **Answer:** Such a g always exists: fix any $z \in A$, order the rest by $x \prec y \iff (z, x, y) \in S$, and put z first.

Layer 1 THE WHAT — surface analysis

The natural attempt. The conclusion asks for an injective $g : A \to \mathbb{R}$, which is the same as listing the elements of A in some order. The reflex for "build an object on a finite set with a property" is induction on $|A|$, so that is where you go. For $|A| = 3$, with $A = \{a, b, c\}$: axiom 2 says exactly

one of the two cyclic classes $\{(a,b,c),(b,c,a),(c,a,b)\}$ and $\{(c,b,a),(b,a,c),(a,c,b)\}$ lies in S, so you order the three elements along whichever triple is in S. Fine.

Now the inductive step. Remove an element z, obtain g on $B = A \setminus \{z\}$ by hypothesis, list B as $a_1, \ldots, a_n$ in increasing g-order, and try to insert z somewhere. Where? You need a gap (a_i, a_{i+1}) such that every increasing triple involving z is in S. That means checking, for each candidate slot, $O(n^2)$ triples (a_j, z, a_k), (z, a_j, a_k), (a_j, a_k, z), and all you have to control them is axiom 3, a four-element statement you do not yet understand. You also discover that the inductive hypothesis is too weak: knowing only "increasing implies in S" for B gives you nothing about triples of B that are *not* increasing, and you need those to locate z. So you strengthen the hypothesis to "$(a,b,c) \in S$ iff some cyclic shift of it is increasing" — which is what Kedlaya does — and then the existence and uniqueness of the insertion slot each take a paragraph of axiom-3 chasing with indices mod n.

It works. It is the official solution. It is also a page of subscript bookkeeping in which one wrong index sinks you, and at no point does it tell you what S *is*.

The model chain. *Step 0: decode.* Put n points on a circle and let S be the set of triples (a,b,c) met in that order going anticlockwise. Axiom 1: rotating the starting point does not change the cyclic order. Axiom 2: reversing a triple reverses the direction of travel. Axiom 3: for four points, "a,b,c anticlockwise and c,d,a anticlockwise" and "b,c,d anticlockwise and d,a,b anticlockwise" both say that a,b,c,d sit anticlockwise in that order. So S is a *cyclic order*, and the problem asks us to cut the circle at a point and straighten it into a line.

Step 1: cut at a point. If $|A| \leq 2$ there are no triples, $S = \varnothing$, and any injection works. Otherwise fix $z \in A$ and let $B = A \setminus \{z\}$. For distinct $x, y \in B$ define

$$x \prec y \iff (z,x,y) \in S.$$

Step 2: $\prec$ is a strict total order on B. Irreflexivity holds because triples in S have distinct entries. For distinct $x, y \in B$, axiom 2 gives $(z,x,y) \in S \iff (y,x,z) \notin S$, and axiom 1 gives $(y,x,z) \in S \iff (z,y,x) \in S$. So exactly one of $x \prec y$, $y \prec x$ holds. For transitivity, let $x \prec y$ and $y \prec w$ with x, y, w distinct. Then $(z,x,y) \in S$, and $(z,y,w) \in S$, which by axiom 1 gives $(y,w,z) \in S$. Apply the forward direction of axiom 3 with $(a,b,c,d) = (z,x,y,w)$: the hypotheses $(a,b,c) = (z,x,y)$ and $(c,d,a) = (y,w,z)$ are in S, so $(b,c,d) = (x,y,w)$ and $(d,a,b) = (w,z,x)$ are in S. By axiom 1 twice, $(w,z,x) \in S$ gives $(z,x,w) \in S$, i.e. $x \prec w$. (We also get $(x,y,w) \in S$, which is the next step in disguise.)

Step 3: define g. B is finite and totally ordered by $\prec$, so list it as $b_1 \prec b_2 \prec \cdots \prec b_m$ and set $g(b_i) = i$, $g(z) = 0$. This is injective.

Step 4: verify. Suppose $g(a) < g(b) < g(c)$. Since $g(z)$ is the minimum, $z \in \{a,b,c\}$ forces $z = a$; then $b \prec c$, which is $(z,b,c) \in S$. Otherwise $a, b, c \in B$ with $a \prec b \prec c$, so $(z,a,b) \in S$ and $(z,b,c) \in S$. Exactly as in Step 2, axiom 1 turns the second into $(b,c,z) \in S$, and axiom 3 with $(a,b,c,d) \mapsto (z,a,b,c)$ yields $(a,b,c) \in S$. $\square$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin by staring at the three axioms. The natural attempt treats them as opaque rules and reaches for the generic tool, induction on $|A|$, then pays for that opacity in the inductive step. The model stops at the axioms and asks what structure they are axiomatising. The split is before any mathematics is written: do you *use* the axioms, or do you first *read* them?

> *Axiom 1 says the triple only matters up to rotation. That is not how a line behaves — on a line, $a < b < c$ is not the same as $b < c < a$. Things that are invariant under rotation live on a circle. So S is "anticlockwise order on a circle", and I am being asked to cut the circle open.*

> *To cut a circle into a line I pick one point to be "the start". Once z is the start, "x comes before y" should mean "going round from z, I meet x then y", which is $(z, x, y) \in S$. Define that, and check it is an order.*

> *Axiom 3 is the only thing that talks about four points, and transitivity is a statement about four points (z, x, y, w). So transitivity must come from axiom 3; line up the letters until the hypotheses match.*

Layer 3 THE WHY

Why read the axioms as a circle rather than push them around? Because the conclusion *is* a line: an injection into $\mathbb{R}$ is a linear order. The hypothesis cannot also describe a line, or axiom 1 would be false for it (on a line, rotating a triple changes whether it is increasing). A hypothesis that looks like an order but is invariant under rotation is, by elimination, an order on a circle — the same way that the phrase "equidistant from a point" is, by definition, a circle and makes the centre the object of interest. Here the object of interest is the missing piece of data that turns a circle into a line: a *starting point*. Nothing else is missing. A clock face has a cyclic order on its hours; to list them in a row you must decide where to begin, and "12" is a choice, not a fact. Once you see that, fixing z is not a trick; it is the one piece of information the circle does not already contain.

Then the definition $x \prec y \iff (z, x, y) \in S$ is forced. On a circle, starting at z and walking anticlockwise, you meet x before y exactly when z, x, y occur anticlockwise in that order. There is no other candidate.

After that, the verification is mechanical, but it is worth seeing why each axiom is used exactly where it is. An order has three requirements: (i) any two distinct elements are comparable, (ii) not both ways, (iii) transitivity. Requirements (i) and (ii) involve two elements plus z, i.e. one triple and its reversal — that is axiom 2, with axiom 1 used to rotate (y, x, z) back to start at z. Requirement (iii) involves three elements plus z, i.e. four points — that is axiom 3, the only four-point axiom. You could have predicted which axiom proves which property by counting points. That count is the practical tool: *match the arity of what you need to the arity of what you have*.

Kedlaya's solution inducts on $|A|$ and inserts a new element z into an existing line, proving that exactly one gap (a_i, a_{i+1}) satisfies $(a_i, z, a_{i+1}) \in S$. That is the same circle seen from the other side: the unique gap is where z sits on the circle. The induction is correct, but it rebuilds the circle one point at a time, and it needs a strengthened hypothesis to do so. Cutting at z uses the circle all at once. Notice also that the model uses only the forward direction of axiom 3; the "if" direction is surplus. Noticing surplus hypotheses is a good check that you have understood a problem rather than merely solved it.

One rigour point the grader will look for: axiom 3 as stated does not say $b \neq d$. If $b = d$, then (b, c, d) has a repeated entry and is not in S, while the left side would require (a, b, c) and (c, b, a) both in S, which axiom 2 forbids; so both sides are false and the axiom is harmless. In the model, all four letters z, x, y, w are distinct anyway, and you should say so.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is axiom 1: membership is invariant under cyclic rotation. On seeing it, write: *"S is a cyclic order — anticlockwise on a circle. The goal is a linear order. I need a cut point."*
2. Fix a base point z and define the linear order $x \prec y \iff (z, x, y) \in S$; verify the order axioms and then the required implication.
3. Hesitation occurs at transitivity: axiom 3's letters a, b, c, d do not visibly match z, x, y, w. Reinforce by writing down what you have, (z, x, y) and (z, y, w), rotating the second to (y, w, z) so it ends where the first began, and reading off $(a, b, c, d) = (z, x, y, w)$. The rule: rotate triples with axiom 1 until consecutive triples share their endpoints the way axiom 3's do.

4. The trap is induction on $|A|$ with the weak hypothesis. It forces you to strengthen the statement to an "if and only if up to rotation" and then to prove existence and uniqueness of an insertion slot by index-chasing mod n. A second trap: forgetting to justify that $\prec$ is total and antisymmetric, which is where axiom 2 is needed.

5. Slogan: *a relation on triples that is invariant under rotation is a circle; turning a circle into a line costs exactly one chosen point.*

6. Decision rule: **when a problem gives you abstract axioms, identify the model they describe before proving anything; then build the requested object from the model's natural parameters (here, a base point), and match each property to be checked with the axiom of the same arity.**

▷ **Cue:** a relation on ordered triples invariant under cyclic shift and reversed by reflection $\longrightarrow$ **Move:** recognise a cyclic order (points on a circle, anticlockwise)

▷ **Cue:** need a linear order from a cyclic one $\longrightarrow$ **Move:** fix a base point z and set $x \prec y \iff (z, x, y) \in S$

▷ **Cue:** list of abstract axioms with no stated model $\longrightarrow$ **Move:** find a concrete example satisfying them all; the proof usually runs along that example

▷ **Cue:** need to prove a property involving k points from axioms of various arities $\longrightarrow$ **Move:** the property involving k points comes from the axiom involving k points

▷ **Cue:** induction on $|A|$ forces you to strengthen the hypothesis and index-chase $\longrightarrow$ **Move:** look for a direct global construction instead (pick one distinguished element and define everything relative to it)

1996 A5

1996 A5. If p is a prime number greater than 3 and $k = \lfloor 2p/3 \rfloor$, prove that the sum

$$\binom{p}{1} + \binom{p}{2} + \cdots + \binom{p}{k}$$

of binomial coefficients is divisible by p^2.

Fields: number theory (binomial coefficients mod p^2, harmonic sums mod p) **Position:** A5 — hard for its position only if you do not know the one standard reduction; after that, the peculiar cut-off $\lfloor 2p/3 \rfloor$ explains itself. **Answer:** Proof. Dividing by p reduces the claim to $\sum_{n=1}^{k} \frac{(-1)^{n-1}}{n} \equiv 0 \pmod{p}$, and that alternating sum equals $\sum_{k/2<n\le k} \frac{1}{n}$, a range symmetric under $n \mapsto p - n$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You check small cases first, which is correct. For $p = 5$, $k = 3$: $5 + 10 + 10 = 25$. For $p = 7$, $k = 4$: $7 + 21 + 35 + 35 = 98 = 2 \cdot 49$. For $p = 11$, $k = 7$: $11 + 55 + 165 + 330 + 462 + 462 + 330 = 1815 = 15 \cdot 121$. It is true, and the sums show no visible pattern.

Next you look for a closed form for a partial row sum $\sum_{n\le k} \binom{p}{n}$. There isn't one: a partial sum $\sum_{n\le k} \binom{p}{n}$ of a binomial row is not expressible as a closed-form hypergeometric term in k (Gosper's algorithm certifies this), so no formula is coming. You try symmetry, $\binom{p}{n} = \binom{p}{p-n}$, to fold the sum: the range $1..k$ with $k \approx 2p/3$ covers $1..p/3$ twice-ish and the middle third once, which gives $\sum_{n=1}^{k} \binom{p}{n} = (2^p - 2) - \sum_{n=1}^{p-k-1} \binom{p}{n}$. That trades one partial sum for another. You try Lucas or Kummer: those tell you $p \mid \binom{p}{n}$ for $0 < n < p$, which is one factor of p, but say nothing about the second.

This is the wall. Every term is divisible by p once; you need p^2 for the sum; and you have been treating the binomial coefficients as integers to be added when the information you need is in them *modulo* p^2, i.e. in $\frac{1}{p}\binom{p}{n}$ modulo p.

The model chain. *Notation.* For $1 \le n \le p-1$ the integer n is invertible modulo p; write $\frac{1}{n}$ for its inverse in $\mathbb{Z}/p\mathbb{Z}$. All congruences below are in $\mathbb{Z}/p\mathbb{Z}$ unless stated.

Step 1: strip one factor of p. For $1 \le n \le p-1$, $\binom{p}{n} = \frac{p}{n}\binom{p-1}{n-1}$, so $T_n := \frac{1}{p}\binom{p}{n}$ is an integer (as $p \mid \binom{p}{n}$) satisfying $nT_n = \binom{p-1}{n-1}$. Now

$$\binom{p-1}{n-1} = \prod_{j=1}^{n-1} \frac{p-j}{j} \equiv \prod_{j=1}^{n-1} \frac{-j}{j} = (-1)^{n-1},$$

the congruence being legitimate because each $j < p$ is invertible. (More carefully: $(n-1)!\binom{p-1}{n-1} = \prod_{j=1}^{n-1}(p-j) \equiv (-1)^{n-1}(n-1)!$, and cancel $(n-1)!$.) Hence $T_n \equiv \frac{(-1)^{n-1}}{n}$. Since $\sum_{n=1}^{k}\binom{p}{n} = p\sum_{n=1}^{k} T_n$, it suffices to prove

$$\sum_{n=1}^{k} \frac{(-1)^{n-1}}{n} \equiv 0 \pmod{p}. \tag{$*$}$$

Step 2: turn the alternating sum into a block of a harmonic sum. Write $H_m = \sum_{n=1}^{m} \frac{1}{n}$. The alternating sum is all terms minus twice the even ones:

$$\sum_{n=1}^{k} \frac{(-1)^{n-1}}{n} = H_k - 2\sum_{\substack{n \le k \\ n \text{ even}}} \frac{1}{n} = H_k - 2\sum_{m=1}^{\lfloor k/2 \rfloor} \frac{1}{2m} = H_k - H_{\lfloor k/2 \rfloor} = \sum_{n=\lfloor k/2 \rfloor + 1}^{k} \frac{1}{n}.$$

Step 3: the range is symmetric about $p/2$. Since $p > 3$ is prime, $p \equiv 1$ or $5 \pmod 6$.

- $p = 6r+1$: $2p/3 = 4r + \frac{2}{3}$, so $k = 4r$, $\lfloor k/2 \rfloor = 2r$, and the range is $2r+1 \le n \le 4r$. Its endpoints satisfy $(2r+1) + 4r = p$.
- $p = 6r+5$: $2p/3 = 4r + \frac{10}{3}$, so $k = 4r+3$, $\lfloor k/2 \rfloor = 2r+1$, and the range is $2r+2 \le n \le 4r+3$. Again $(2r+2) + (4r+3) = p$.

In both cases the range R is an interval of integers whose endpoints add to p, so $n \mapsto p-n$ maps R onto itself. It has no fixed point since p is odd. So R splits into pairs $\{n, p-n\}$, and

$$\frac{1}{n} + \frac{1}{p-n} \equiv \frac{1}{n} + \frac{1}{-n} = 0.$$

Therefore $\sum_{n \in R} \frac{1}{n} \equiv 0$, which is $(*)$, and p^2 divides the original sum. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that every $\binom{p}{n}$ with $0 < n < p$ is divisible by p. There the natural attempt stops using that fact and goes looking for the second p in the size or shape of the sum. The model takes the first p *out* — divides by it — and asks what the quotient is modulo p. That is the split: "p divides each term, now what?"

> *Each term is already p times an integer. So $p^2 \mid \sum$ is the same as $p \mid \sum \frac{1}{p}\binom{p}{n}$. I should compute $\frac{1}{p}\binom{p}{n}$ modulo p — and in $\binom{p-1}{n-1}$ every factor $p-j$ is $-j$ modulo p.*

> *I've got $\sum \frac{(-1)^{n-1}}{n}$. Alternating harmonic sums are "all terms minus twice the evens", and the evens are half a harmonic sum. So this is a difference of two harmonic sums, i.e. one block $\sum_{a<n\le b} \frac{1}{n}$.*

> *The only way $\frac{1}{n}$'s add to 0 mod p without a formula is by pairing n with $p-n$. Is my block symmetric about $p/2$? It runs from $k/2 \approx p/3$ to $k \approx 2p/3$ — yes. That is why the problem says $2p/3$.*

Layer 3 THE WHY

Why divide by p? Because of what "divisible by p^2" means for a number already known to be divisible by p. If $N = pM$, then $p^2 \mid N \iff p \mid M$. The whole problem lives one level down: you need M modulo p. This is the same move as computing a derivative after you know the value is zero — the first-order information has been used up, so you go to the next order. Modulo p^2 problems about objects that vanish modulo p are almost always attacked by dividing out the p and working modulo p with the quotient.

Why is the quotient computable? Because $\binom{p-1}{n-1}$ is a product of fractions $\frac{p-j}{j}$ and modulo p the number p is 0. Every numerator collapses to $-j$, every fraction to -1. No ingenuity is involved: the definition of the binomial coefficient as a product, and the definition of congruence. Kedlaya's (Ng's) solution writes this as a single line, $\frac{1}{p}\binom{p}{n} \equiv \frac{(-1)^{n-1}}{n}$; that line is the reduction, and it is the step to have in memory as a reflex.

Why an alternating harmonic sum becomes a block. The identity $\sum_{n \le k} \frac{(-1)^{n-1}}{n} = H_k - H_{\lfloor k/2 \rfloor}$ is the finite version of $\ln 2 = \lim(H_{2m} - H_m)$. It holds over any field where the denominators are invertible, because it is pure bookkeeping: subtract the even terms twice, and the even terms $\frac{1}{2m}$ are $\frac{1}{2} \cdot \frac{1}{m}$. Once you have a block $\sum_{a<n\le b} \frac{1}{n}$, the question "is it 0 mod p?" has essentially one tool at A-level-plus-a-bit: the involution $n \mapsto p-n$, under which $\frac{1}{n} \mapsto -\frac{1}{n}$. This is the same mechanism as Wolstenholme's $H_{p-1} \equiv 0$ (the full range $1..p-1$ is symmetric about $p/2$). So you ask: is the block symmetric about $p/2$?

And now the strange $\lfloor 2p/3 \rfloor$ is explained, not memorised. The block runs from about $k/2$ to k. For it to be symmetric about $p/2$ you need $\frac{k}{2} + k = p$, i.e. $k = \frac{2p}{3}$. Whatever the setters had in mind, $2p/3$ is precisely the cut-off for which the pairing proof works, and that is the reading that lets you use it. Reading the constant backwards from the mechanism is the WHY here: *when a problem contains an odd constant, ask what that constant makes symmetric.* The floor and the two residue classes $p \equiv 1, 5 \pmod 6$ are there only because $2p/3$ is not an integer; checking both cases is where marks are lost, because the parity of k differs ($k = 4r$ versus $4r+3$) and so does $\lfloor k/2 \rfloor$.

Why the natural attempt was reasonable and why it failed. Reaching for a closed form of a partial binomial sum is a sensible first reflex — many row-sum problems have one. But the problem's number-theoretic target (p^2) and its analytically meaningless cut-off ($2p/3$) are both cues that the answer is arithmetic, not a formula. The integers $\binom{p}{n}$ are large and structureless; their residues are tiny and structured. Once you move to residues, the problem is a three-line computation.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "prime p" together with "divisible by p^2" on a sum of $\binom{p}{n}$. Write down: $p \mid \binom{p}{n}$ *for* $0 < n < p$, *so I need* $\sum \frac{1}{p}\binom{p}{n} \equiv 0 \pmod p$.

2. Compute $\frac{1}{p}\binom{p}{n} = \frac{1}{n}\binom{p-1}{n-1} \equiv \frac{(-1)^{n-1}}{n} \pmod p$, reducing to an alternating harmonic sum modulo p; then pair n with $p-n$.

3. Hesitation occurs at $\sum \frac{(-1)^{n-1}}{n}$: it does not pair directly under $n \mapsto p-n$ because the signs depend on parity and $p-n$ has the opposite parity. Reinforce by converting to $H_k - H_{\lfloor k/2 \rfloor}$ first; the block then has no signs, and pairing is clean.

4. Traps: (a) seeking a closed form for a partial row sum, which does not exist; (b) pairing n with $p-n$ in the *alternating* sum, where the sign flip spoils it; (c) mishandling the floor — $p = 6r+1$ gives $k = 4r$ but $p = 6r+5$ gives $k = 4r+3$, and $\lfloor k/2 \rfloor$ must be recomputed in each case; (d) using $\frac{1}{n}$ modulo p without saying that $n < p$ makes it invertible.

5. Slogans: *to prove $p^2 \mid N$ when $p \mid N$ is known, compute N/p modulo p. Harmonic-type sums vanish modulo p by the pairing $n \leftrightarrow p-n$; an odd constant in the range is there to make that pairing work.*

6. Decision rule: **when a divisibility by p^k is required and p^{k-1} is automatic, divide out the known power and work one level down modulo p; when a problem contains an unexplained constant, solve for the value of the constant that makes your intended symmetry exact, and check it matches.**

▷ **Cue:** need $p^2 \mid \sum \binom{p}{n}$ over $0 < n < p$ $\longrightarrow$ **Move:** use $\frac{1}{p}\binom{p}{n} \equiv \frac{(-1)^{n-1}}{n} \pmod p$

▷ **Cue:** $\binom{p-1}{m}$ modulo p $\longrightarrow$ **Move:** each factor $\frac{p-j}{j} \equiv -1$, so $\binom{p-1}{m} \equiv (-1)^m$

▷ **Cue:** alternating harmonic sum $\sum_{n \le k} \frac{(-1)^{n-1}}{n}$ $\longrightarrow$ **Move:** rewrite as $H_k - H_{\lfloor k/2 \rfloor}$, a single unsigned block

▷ **Cue:** sum of reciprocals over a range, to be shown $\equiv 0 \pmod p$ $\longrightarrow$ **Move:** check whether the range is symmetric under $n \mapsto p-n$; pair $\frac{1}{n} + \frac{1}{p-n} \equiv 0$

▷ **Cue:** a problem contains an odd cut-off such as $\lfloor 2p/3 \rfloor$ $\longrightarrow$ **Move:** ask what that cut-off makes symmetric; split into residue classes of p to handle the floor

1996 A6

1996 A6. Let $c > 0$ be a constant. Give a complete description, with proof, of the set of all continuous functions $f : R \to R$ such that $f(x) = f(x^2 + c)$ for all $x \in R$. Note that R denotes the set of real numbers.

Fields: analysis (functional equations, iteration and fixed points, continuity) **Position:** A6 — the last problem of the morning, and hard mainly because the answer changes character at $c = \frac{1}{4}$ and you must find that threshold yourself. **Answer:** If $0 < c \le \frac{1}{4}$, only the constant functions. If $c > \frac{1}{4}$, exactly the functions obtained by choosing any continuous $h : [0, c] \to \mathbb{R}$ with $h(0) = h(c)$, extending to $[0, \infty)$ by $f = h \circ \varphi^{-n}$ on $[\varphi^n(0), \varphi^{n+1}(0)]$ where $\varphi(x) = x^2 + c$, and then setting $f(x) = f(|x|)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. A functional equation, so you substitute. $x \mapsto -x$ gives $f(-x) = f(x^2 + c) = f(x)$: f is even. Good, restrict to $x \ge 0$. Then you iterate: $f(x) = f(x^2 + c) = f((x^2 + c)^2 + c) = \cdots$. You try $x = 0$: $f(0) = f(c) = f(c^2 + c) = \cdots$. You guess that f must be constant, since the equation "spreads" each value along an orbit, and you try to prove it by showing the orbits are dense or that they all meet. For $c = 1$ the orbit of 0 is $0, 1, 2, 5, 26, \ldots$ — sparse, running to infinity, and two different starting points $x \in (0, 1)$ produce orbits that never meet. You cannot connect $f(0)$ to $f(\frac{1}{2})$.

So you suspect non-constant solutions exist for $c = 1$. But then what about small c? You try $c = \frac{1}{10}$ and the orbit of 0 is $0,\ 0.1,\ 0.11,\ 0.1121,\ \ldots$, which is creeping towards something rather than escaping. Now you are stuck between two pictures and do not know which c gives which, because you have not asked the one question that separates them: *where are the fixed points of $x \mapsto x^2 + c$?*

The model chain. Write $\varphi(x) = x^2 + c$, so the equation is $f = f \circ \varphi$.

Step 0: evenness. $f(-x) = f((-x)^2 + c) = f(x^2 + c) = f(x)$. So it suffices to understand f on $[0, \infty)$, where φ is continuous and strictly increasing. Also, by induction, $f = f \circ \varphi^n$ for all $n \ge 0$, and if $y \ge c$ then $\psi(y) := \sqrt{y - c} \ge 0$ satisfies $\varphi(\psi(y)) = y$, so $f(y) = f(\psi(y))$.

Step 1: the fixed points. $\varphi(x) = x \iff x^2 - x + c = 0$, with discriminant $1 - 4c$. This is the threshold.

Case A: $0 < c \le \frac{1}{4}$. The roots $r_1 = \frac{1 - \sqrt{1-4c}}{2} \le r_2 = \frac{1 + \sqrt{1-4c}}{2}$ are real, with $r_1 > 0$ (product $c > 0$, sum 1) and $r_2 \ge \frac{1}{2} > c$. Since $\varphi(x) - x = (x - r_1)(x - r_2)$, we have $\varphi(x) > x$ outside $[r_1, r_2]$ and $\varphi(x) \le x$ inside it.

- $0 \le x \le r_1$. φ is increasing, $\varphi(0) = c \ge 0$, $\varphi(r_1) = r_1$, so φ maps $[0, r_1]$ into itself, and $\varphi(t) \ge t$ there. The orbit $x_n = \varphi^n(x)$ is non-decreasing and bounded by r_1, so it converges; its limit L satisfies $\varphi(L) = L$ by continuity of φ, and $L \in [0, r_1]$, so $L = r_1$. By continuity of f, $f(x) = f(x_n) \to f(r_1)$. So $f(x) = f(r_1)$.

- $r_1 < x < r_2$. φ maps (r_1, r_2) into itself (increasing, endpoints fixed) and $\varphi(t) < t$ there. The orbit decreases, is bounded below by r_1, and converges to a fixed point in $[r_1, x]$, which is r_1. So $f(x) = f(r_1)$ again.

- $x > r_2$. Now forward iteration escapes, so go backwards. Since $x > r_2 > c$, ψ is defined; and for $y > r_2$, $\psi(y) > r_2 \iff y - c > r_2^2 \iff y > \varphi(r_2) = r_2$, and $\psi(y) < y \iff y - c < y^2 \iff \varphi(y) > y$, true for $y > r_2$. So $y_0 = x$, $y_{n+1} = \psi(y_n)$ is decreasing, bounded below by r_2, and converges to some $L \ge r_2$ with $L = \sqrt{L - c}$, i.e. $\varphi(L) = L$; hence $L = r_2$. Since $f(y_n) = f(y_{n+1})$, $f(x) = f(r_2)$.

Finally $f(r_2) = \lim_{t \to r_2^-} f(t) = f(r_1)$ by continuity (if $r_1 = r_2$ there is nothing to prove). So f is constant on $[0, \infty)$, hence on $\mathbb{R}$. Constants are solutions, so in Case A the solutions are exactly the constants.

Case B: $c > \frac{1}{4}$. Now $\varphi(x) - x = (x - \frac{1}{2})^2 + (c - \frac{1}{4}) > 0$ for all x: no fixed points, every orbit strictly increases. Let $x_0 = 0$, $x_{n+1} = \varphi(x_n)$. The sequence is strictly increasing and unbounded (a bounded increasing sequence would converge to a fixed point). So the intervals $I_n = [x_n, x_{n+1}]$ cover $[0, \infty)$, overlapping only at endpoints, and φ maps I_n bijectively and increasingly onto I_{n+1}.

Necessity. Let f be a solution and $h = f|_{[0,c]} = f|_{I_0}$. Then h is continuous and $h(0) = f(0) = f(\varphi(0)) = f(c) = h(c)$. For $y \in I_n$, $y = \varphi^n(t)$ for a unique $t \in I_0$, and $f(y) = f(t) = h(t)$. So $f = h \circ \varphi^{-n}$ on I_n, and $f(x) = f(|x|)$.

Sufficiency. Given continuous h on $[0, c]$ with $h(0) = h(c)$, define $f(y) = h(\varphi^{-n}(y))$ for $y \in I_n$. At the shared endpoint x_{n+1} the two definitions give $h(\varphi^{-n}(x_{n+1})) = h(c)$ and $h(\varphi^{-(n+1)}(x_{n+1})) = h(0)$, which agree; so f is well defined and continuous on $[0, \infty)$ (continuous on each closed piece, pieces agreeing at the joins). Set $f(x) = f(|x|)$ for $x < 0$; this is continuous on $\mathbb{R}$. For $y \ge 0$ in I_n, $\varphi(y) \in I_{n+1}$ and $f(\varphi(y)) = h(\varphi^{-(n+1)}\varphi(y)) = h(\varphi^{-n}(y)) = f(y)$; for $y < 0$, $f(\varphi(y)) = f(\varphi(-y)) = f(-y) = f(y)$. So f is a solution, and there are non-constant ones (e.g. $h(t) = \sin(2\pi t/c)$). □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains derive evenness and write $f = f \circ \varphi^n$. The natural attempt then chases specific orbits numerically, looking for them to meet. The model asks, before any numbers, *where do orbits of φ go?* — and answers it by solving $\varphi(x) = x$. The split is at "$f(x) = f(x^2 + c) = f((x^2 + c)^2 + c) = \cdots$": do you compute further along the orbit, or ask where the orbit ends up?

> *$f(x) = f(\varphi^n(x))$ for every n, and f is continuous. If $\varphi^n(x)$ converges to some L, then $f(x) = f(L)$. Orbits of a continuous map that converge, converge to fixed points. So: solve $x^2 + c = x$.*

> *The discriminant of $x^2 - x + c$ is $1 - 4c$. There's my threshold at $c = \frac{1}{4}$. If there are fixed points, orbits get trapped and f collapses to a constant; if not, every orbit runs off to infinity, nothing pins f down, and I should be building solutions, not ruling them out.*

> *Above r_2 the forward orbit escapes, so it tells me nothing. But the equation works both ways: $f(y) = f(\sqrt{y - c})$. Run the orbit backwards and it falls into r_2.*

Layer 3 THE WHY

Why are fixed points the object of interest? Read the equation as a statement about level sets: f takes the same value at x and at $\varphi(x)$, hence along the whole orbit $x, \varphi(x), \varphi^2(x), \ldots$. A continuous function that is constant along a sequence is also equal to its value at the sequence's limit. So the question "what can f be?" becomes "which points are linked by orbits, allowing limits?" And the limits of orbits of a continuous map are fixed points: if $x_n \to L$ and $x_{n+1} = \varphi(x_n)$, then $L = \varphi(L)$. That is not a trick; it is one line of continuity. The fixed point equation is therefore the first thing to solve, and it is a quadratic, whose discriminant $1 - 4c$ tells you there will be two regimes. The number $\frac{1}{4}$ in the answer is not something you find by experiment; it is written in the problem, in the form "$x^2 + c$", waiting for you to set it equal to x.

Why run the orbit backwards above r_2? Because the equation is symmetric in the sense that matters: it relates f at two points, and nothing says which one you must start from. Forward orbits from $x > r_2$ go to $+\infty$, carrying no information. The inverse branch $\psi(y) = \sqrt{y - c}$ is a contraction towards r_2 on (r_2, ∞) — r_2 is repelling for φ, so attracting for φ^{-1}. The general principle: *a fixed point repels in one direction of time and attracts in the other*. Whichever direction attracts is the direction to iterate.

Why is Case B a construction problem and not a uniqueness problem? With no fixed point, $\varphi(x) > x$ everywhere on $[0, \infty)$ and every orbit is a strictly increasing sequence to infinity. Two points in the open interval $(0, c)$ lie on different orbits, and no limit connects them. So the equation imposes no relation between f on distinct points of $[0, c)$; the only relation within $[0, c]$ is that 0 and c are on the same orbit. That is exactly why the free data is "a continuous function on $[0, c]$ with equal endpoint values" — a *fundamental domain* for the action of φ, the analogue of one period of a periodic function. Indeed, if you conjugate φ on $[0, \infty)$ to the translation $t \mapsto t + 1$, the equation becomes periodicity with period 1, and "a continuous function on one period with equal endpoint values" is the familiar description of periodic functions.

Kedlaya's solution handles Case A with the same three regions and, as its own acknowledgement records, an earlier version of it contained an inaccuracy; its treatment of $x < r_1$ lets the orbit overshoot r_1 and then appeals back to the other region. The version above avoids the issue by noting $\varphi([0, r_1]) \subseteq [0, r_1]$, so the orbit is trapped and monotone without any case split on whether it overshoots r_1. The official Case B dismisses the converse in a single sentence; the necessity argument (restriction to I_0, endpoint condition from $f(0) = f(c)$, and the formula on I_n) is the half of a "complete description" that graders check, so write it.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $f(x) = f(g(x))$ with f continuous. Write down: *f is constant on orbits of g and at their limits; find the fixed points of g.* Here: $x^2 + c = x$, discriminant $1 - 4c$.

2. Dynamics of $\varphi(x) = x^2 + c$ on $[0, \infty)$: monotone orbits, limits are fixed points, use the attracting direction (forward below r_2, backward above r_2). With no fixed points, use a fundamental domain $[0, c]$ and transport it by φ.

3. Hesitation occurs above r_2, where forward iteration diverges and seems to say nothing. Reinforce by remembering that $f(y) = f(\sqrt{y - c})$ is the same equation read the other way; iterate the inverse. A second hesitation point is whether the answer really is "all constants" for $c \leq \frac{1}{4}$ and "lots" for $c > \frac{1}{4}$; reinforce by checking $c = 1$ explicitly, where the orbit $0, 1, 2, 5, \ldots$ visibly never returns.

4. Traps: (a) guessing "f is constant" for all c (false for $c > \frac{1}{4}$); (b) forgetting evenness and treating $x < 0$ separately; (c) in Case B, omitting the endpoint condition $h(0) = h(c)$, which is exactly what makes f continuous at the joins x_n; (d) giving only sufficiency in Case B when the problem asks for a complete description; (e) forgetting the boundary case $c = \frac{1}{4}$, where $r_1 = r_2 = \frac{1}{2}$.

5. Slogan: *$f = f \circ g$ with f continuous: solve $g(x) = x$ first. Fixed points exist $\Rightarrow$ orbits collapse to them $\Rightarrow$ f is forced. No fixed points $\Rightarrow$ free choice on a fundamental domain.*

6. Decision rule: **for any invariance equation $f(x) = f(g(x))$, study the dynamics of g before studying f: find fixed points, determine which are attracting in forward or backward time, and let the answer be "constant on each basin" if orbits converge, or "arbitrary on a fundamental domain, matched at its ends" if they escape.**

- ▷ **Cue:** $f(x) = f(g(x))$ for all x, f continuous $\longrightarrow$ **Move:** f is constant along orbits of g and equals f at orbit limits; solve $g(x) = x$
- ▷ **Cue:** a parameter c inside a quadratic map $x^2 + c$ $\longrightarrow$ **Move:** the discriminant of $x^2 - x + c$ gives the threshold $c = \frac{1}{4}$ between regimes
- ▷ **Cue:** forward orbit diverges and carries no information $\longrightarrow$ **Move:** iterate the inverse branch; a repelling fixed point of g is attracting for g^{-1}
- ▷ **Cue:** $f(x) = f(-x + \text{something even})$, e.g. argument $x^2 + c$ $\longrightarrow$ **Move:** substitute $x \mapsto -x$ to get evenness; restrict to $x \geq 0$ where g is monotone
- ▷ **Cue:** $g(x) > x$ everywhere, no fixed points $\longrightarrow$ **Move:** orbit of a base point tiles the line into $[x_n, x_{n+1}]$; solutions are arbitrary on one tile with matching endpoint values
- ▷ **Cue:** "give a complete description" of a solution set $\longrightarrow$ **Move:** prove both that every described function works and that every solution is of that form

Putnam 1996 — Paper B

1996 B1

1996 B1. Define a **selfish** set to be a set which has its own cardinality (number of elements) as an element. Find, with proof, the number of subsets of $\{1, 2, \ldots, n\}$ which are *minimal* selfish sets, that is, selfish sets none of whose proper subsets is selfish.

Fields: Combinatorics; bijective/recursive counting **Position:** B1 – opening problem of the second session **Answer:** The count is F_n, the nth Fibonacci number (with $F_1 = F_2 = 1$).

Layer 1 THE WHAT — surface analysis

The natural attempt. You start by generating data, because there is no obvious closed form to guess from the definition alone. $n = 1$: only nonempty candidate is $\{1\}$, and $|\{1\}| = 1 \in \{1\}$, so it is selfish; its only proper subset is $\emptyset$, which is never selfish since $|\emptyset| = 0 \notin \emptyset$. So $\{1\}$ is minimal: count 1. $n = 2$: $\{1\}$ is still there and still minimal. Is $\{2\}$ selfish? $|\{2\}| = 1$, and $1 \notin \{2\}$ – no. Is $\{1, 2\}$ selfish? $|\{1, 2\}| = 2 \in \{1, 2\}$ – yes, but it contains the selfish proper subset $\{1\}$, so it isn't minimal. Count stays 1. $n = 3$: now $\{2, 3\}$ appears: $|\{2, 3\}| = 2 \in \{2, 3\}$, selfish, and neither $\{2\}$ nor $\{3\}$ is selfish (check each against its own cardinality), so it's minimal. Count 2. $n = 4$: $\{2, 4\}$ joins ($|\{2, 4\}| = 2 \in \{2, 4\}$, and $\{2\}, \{4\}$ both fail), giving count 3.

At this point the sequence $1, 1, 2, 3, \ldots$ is impossible to miss, and a strong student writes down "Fibonacci" and starts hunting for a bijective or recursive reason – correctly. The natural next move, though, is to try to construct the recursion by brute case-splitting on the largest element: "a minimal selfish subset of $[n]$ either omits n, in which case it's a minimal selfish subset of $[n-1]$, giving the f_{n-1} term – fine, that half is free – or it contains n, and now I have to characterise exactly which subsets of $\{1, \ldots, n-1\} \cup \{n\}$ containing n are selfish and minimal." Without a sharper idea, this second case is where the naive approach stalls: you start trying to directly enumerate, for each possible cardinality k of such a set, which $(k-1)$-subsets of $\{2, \ldots, n-1\}$ (it has to omit 1, though at this stage you haven't yet proved that) can be unioned with $\{n\}$ and stay both selfish and minimal, and the enumeration doesn't visibly collapse to "f_{n-2} of them" without an extra structural idea – you're doing casework on k that would need to be redone for every n.

The model chain. Let f_n be the number of minimal selfish subsets of $[n] = \{1, \ldots, n\}$. Split by whether $n \in S$.

Case $n \notin S$. Then $S \subseteq [n-1]$, and S is a minimal selfish subset of $[n]$ if and only if it is a minimal selfish subset of $[n-1]$ (minimality and selfishness are both statements entirely about S and its subsets, and $[n-1]$ vs. $[n]$ doesn't change which subsets of S exist). This case contributes exactly f_{n-1}.

Case $n \in S$. First, $1 \notin S$: if $1 \in S$ then $\{1\}$ is selfish (since $|\{1\}| = 1 \in \{1\}$) and is a proper subset of S (as $n \in S$, $n \neq 1$), contradicting minimality of S – unless $S = \{1\}$ itself, which is impossible since $n \in S$. So $S \subseteq \{2, \ldots, n\}$, $n \in S$, $1 \notin S$, and S is minimal selfish.

Now define $\Phi(S) = (S \setminus \{n\}) - 1$, i.e. delete n, then subtract 1 from every remaining element. Write $k = |S|$; since S is selfish, $k \in S$. Because $S \subseteq [n]$ with $1 \notin S$, $S = [n]$ is impossible, so

$k = |S| < n$, hence $k \neq n$ and $k \in S \setminus \{n\}$. Thus $|\Phi(S)| = k - 1$ and $(k-1) \in \Phi(S)$ (since $k \in S \setminus \{n\}$ shifts to $k-1$), so $\Phi(S)$ is selfish. Also $S \setminus \{n\} \subseteq \{2, \ldots, n-1\}$, so $\Phi(S) \subseteq [n-2]$.

Φ is invertible on its domain: given $T \subseteq [n-2]$, set $\Psi(T) = (T+1) \cup \{n\}$, where $T+1$ shifts every element up by 1; then $\Phi(\Psi(T)) = T$ and $\Psi(\Phi(S)) = S$ on the relevant domains, because deleting/re-adding n and shifting by 1 in opposite directions are literally inverse operations on sets. Φ restricted to $\{A \subseteq [n] : n \in A,\ 1 \notin A\}$ is therefore an order isomorphism (under $\subseteq$) onto $2^{[n-2]}$: it is a relabelling of the ground elements $\{2, \ldots, n-1\} \leftrightarrow \{1, \ldots, n-2\}$ together with the fixed deletion of n, so $A' \subsetneq A$ (both containing n, avoiding 1) if and only if $\Phi(A') \subsetneq \Phi(A)$.

Since Φ and Ψ are mutually inverse and both preserve selfishness (shown above for Φ; the same computation run backwards, using $|T| = j \in T \Rightarrow j+1 \in \Psi(T)$ and $|\Psi(T)| = j+1$, gives it for Ψ), Φ restricts to a bijection between $\{\text{selfish } A : n \in A,\ 1 \notin A\}$ and $\{\text{selfish } T \subseteq [n-2]\}$, and this bijection is order-preserving on the selfish subsets it relates. So minimal elements of the left-hand family (no selfish $A' \subsetneq A$ with A' also of that containing-n-avoiding-1 form) correspond exactly to minimal elements T of the right-hand family.

Caveat. This is not yet the full statement of minimality: S minimal requires no selfish proper subset *at all*, including subsets $S' \subsetneq S$ with $n \notin S'$ (so $S' \subseteq S \setminus \{n\} = \Psi(T) \setminus \{n\} = T+1$), and the order-isomorphism argument above, by construction, only rules out selfish proper subsets that still contain n. Ruling out a selfish $S' \subseteq T+1$ would require that no subset $V \subseteq T$ satisfies the shifted condition $|V| - 1 \in V$ (since $S' = V+1$ is selfish iff $|S'| = |V| \in V+1$ iff $|V| - 1 \in V$) – a condition distinct from T's own (unshifted) minimality, $|V| \notin V$ for $V \subsetneq T$. I have checked this holds on every minimal selfish set produced for $n \leq 6$ (e.g. $T = \{2,4\} \subseteq [4]$ gives no subset V with $|V| - 1 \in V$), and the final count matches the known Fibonacci sequence for these n, but I have not found a general proof that ordinary minimality of T forecloses the shifted condition, and Bhargava–Kedlaya's published solution states the bijection without addressing this case at all. Treat the bijection's totality as verified computationally through $n = 6$ and structurally very plausible, not as fully proved here.

Setting aside that one gap, the two cases give $f_n = f_{n-1} + f_{n-2}$, with base cases $f_1 = f_2 = 1$ (both computed directly above). This recursion with these seeds defines the Fibonacci numbers exactly: $f_n = F_n$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree on the easy half of the split: $n \notin S$ contributes f_{n-1}, free of charge. They part company inside the hard half, $n \in S$. The natural attempt tries to directly characterise and count these sets case by case on their cardinality k. The model chain instead asks a structural question first: *is there an operation that turns "minimal selfish subset of $[n]$ containing n" into "minimal selfish subset of $[n-2]$", with no casework on k at all?*

> *I don't need to know what k is or enumerate over it – I need one operation, defined uniformly for every such S regardless of its size, that lands me in exactly the right smaller problem; if the operation exists and is reversible, the count transfers for free and I never have to touch k by hand.*

> *The set loses one element (I remove n) and every remaining element drops by one (I shift down), so the cardinality of what's left drops by exactly one too – I should check whether that cardinality drop is exactly what's needed to keep "cardinality equals a member" true after the shift, rather than assuming it and hoping.*

Layer 3 THE WHY

Why is deleting n and then shifting down by 1 the right operation, rather than some other

renormalisation onto $[n-2]$? Because "selfish" is a statement that ties two quantities together that a naive relabelling would pull apart: the *size* of the set (which doesn't care about the labels of its elements) and the *value* of one particular element (which is exactly the label that gets changed by any relabelling). A pure shift by 1 – subtract 1 from every element, keep the same elements – preserves the set's size and moves every value down by 1; it turns "$k \in S$" into "$k-1 \in S-1$", but the size is still k, not $k-1$, so a pure shift breaks selfishness (it needs the *shifted* size $k-1$ to be present, not the original k). Deleting one designated element, on the other hand, drops the size by exactly 1 without touching any remaining value. Composing the two – delete n (size drops from k to $k-1$, values unchanged), then shift the rest down by 1 (values drop by 1, size unchanged at $k-1$) – produces a set whose size is $k-1$ and which contains $k-1$ exactly when the original contained k, because $k \in S \setminus \{n\}$ shifts to $k-1$ in the new set. The two operations are chosen so their effects on "size" and "membership of the size" are matched: one drops the count, the other drops the values, and doing both once each keeps the identity size-equals-member true across the transformation. Bhargava–Kedlaya's solution performs exactly this composite operation in one sentence ("subtracting 1 from each element, and then taking away the element $n-1$ from the set") without saying why it is the right composite to take; this is where that choice comes from – it is forced by wanting an operation that preserves the selfish invariant while reducing both the ground set and the witnessing cardinality by the same fixed amount.

The other half of the WHY – why $1 \notin S$ for $n \in S$ minimal – is a direct minimality argument, not a structural one: $\{1\}$ is selfish unconditionally (a set of size 1 containing 1 is selfish by definition, with no further check needed), so any set containing 1 and something else at all has a two-element-or-larger selfish proper subset sitting inside it for free, and minimality forbids that outright.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "own cardinality as an element" – write down immediately that this definition ties two normally-independent quantities (a set's size, and the value of one of its members) together, so any operation you perform on the set (relabelling, shifting, deleting) has to be checked against *both* quantities, not just one.

2. That cue suggests small-case enumeration first (since there is no visible closed form), then, once a recognisable sequence (Fibonacci) appears, hunting specifically for a two-term recursion built from a case split on the largest or most recently examined element – because Fibonacci-shaped answers in combinatorics almost always come from exactly that kind of split.

3. Hesitation occurs at the "$n \in S$" branch, where the instinct is to enumerate by cardinality k. The reinforcement away from it: ask "what single operation, applied uniformly regardless of k, would strip this down to a smaller instance of the same problem?" – looking for a structural map is cheaper than looking for a formula in k.

4. The trap is assuming that any bijective-looking shift preserves selfishness. A shift alone changes every element's value but not the set's size, so it silently breaks the size-equals-member identity; the fix (delete one element to also drop the size) has to be found, not assumed.

5. To recognise this faster next time: whenever a combinatorial definition links a set's cardinality to its own membership, expect that any reduction argument needs a compensating cardinality change (a deletion or an insertion) to survive a value-shift, and look for that pairing explicitly before trusting a proposed bijection.

6. The reusable rule: when reducing a "self-referential" counting problem (size ties to content) from $[n]$ to a smaller ground set, any relabelling of values must be paired with a matching change in cardinality, chosen so the reduction's effect on "size" and its effect on "membership of the size" cancel out to preserve the defining identity.

1996 B2

1996 B2. Show that for every positive integer n,

$$\left(\frac{2n-1}{e}\right)^{\frac{2n-1}{2}} < 1\cdot 3\cdot 5\cdots(2n-1) < \left(\frac{2n+1}{e}\right)^{\frac{2n+1}{2}}.$$

Fields: Analysis; integral estimation of a product via ln **Position:** B2 – second problem of the second session **Answer:** The stated double inequality holds for every positive integer n; both sides are established by sandwiching $\ln(1\cdot 3\cdots(2n-1))$ between two Riemann-sum bounds on $\int \ln x\,dx$, with the strictness supplied by $1 < e < 3$.

Layer 1 THE WHAT — surface analysis

The natural attempt. A chain of near-identical inequalities indexed by n says "induction" almost reflexively, so you try it. Base case $n = 1$: the product is 1, and you check $(1/e)^{1/2} \approx 0.607 < 1 < (3/e)^{3/2} \approx 1.159$ – it holds. For the inductive step, suppose

$$\left(\tfrac{2n-1}{e}\right)^{\frac{2n-1}{2}} < P_n < \left(\tfrac{2n+1}{e}\right)^{\frac{2n+1}{2}}, \qquad P_n := 1\cdot 3\cdots(2n-1),$$

and you want the same statement for $n+1$, where $P_{n+1} = P_n\cdot(2n+1)$. The lower bound you actually need is $P_{n+1} > \left(\frac{2n+1}{e}\right)^{\frac{2n+1}{2}}$, i.e. (using only what the hypothesis gives you) it would suffice to show

$$(2n+1)\left(\tfrac{2n-1}{e}\right)^{\frac{2n-1}{2}} \;\geq\; \left(\tfrac{2n+1}{e}\right)^{\frac{2n+1}{2}}.$$

Rearranged, this is a claim of the shape $(2n-1)^{(2n-1)/2}$ against $(2n+1)^{(2n-1)/2}$ times an e-power correction – and since $2n-1 < 2n+1$, the bare comparison of bases runs the *wrong way*: you are trying to bound something built from the smaller base by something built from the larger one, and the inductive hypothesis alone (which only compares P_n to bounds *at* n) doesn't supply the extra multiplicative room. To rescue the step you would need an auxiliary fact along the lines of $\left(1+\frac{2}{2n-1}\right)^{(2n-1)/2} \lesseqgtr$ some explicit bound involving e – which is itself a nontrivial calculus inequality of essentially the same flavour as the one you started with. The induction is not wrong, exactly; it has just relocated the entire difficulty into a new lemma you now have to prove separately, at every step, rather than discharging the problem in one estimate. This is the wall: induction on a family of inequalities whose *bounds themselves change shape with* n (not just their values) does not automatically get cheaper by assuming the previous case.

The model chain. Write $P_n = 1\cdot 3\cdots(2n-1)$. Since $\ln 1 = 0$,

$$\ln P_n = \ln 3 + \ln 5 + \cdots + \ln(2n-1).$$

Because $\ln x$ is increasing, on each width-2 interval $[2k-1, 2k+1]$ the right-endpoint rectangle of height $\ln(2k+1)$ lies above the graph, and the left-endpoint rectangle of height $\ln(2k-1)$ lies below it:

$$\int_{2k-1}^{2k+1} \ln x\,dx \;\leq\; 2\ln(2k+1), \qquad 2\ln(2k-1) \;\leq\; \int_{2k-1}^{2k+1} \ln x\,dx.$$

Summing the first inequality for $k = 1,\ldots,n-1$ (intervals $[1,3],[3,5],\ldots,[2n-3,2n-1]$, right endpoints $3,5,\ldots,2n-1$) gives

$$\int_1^{2n-1} \ln x\,dx \;\leq\; 2\big(\ln 3 + \cdots + \ln(2n-1)\big) = 2\ln P_n.$$

Summing the second inequality for $k = 2, \dots, n$ (intervals $[3,5], \dots, [2n-1, 2n+1]$, left endpoints $3, 5, \dots, 2n-1$) gives

$$2\ln P_n = 2\big(\ln 3 + \cdots + \ln(2n-1)\big) \;\le\; \int_3^{2n+1} \ln x\, dx.$$

Using $\int \ln x\, dx = x\ln x - x$:

$$\int_1^{2n-1} \ln x\, dx = (2n-1)\ln(2n-1) - (2n-1) - (1\cdot 0 - 1) = (2n-1)\ln(2n-1) - (2n-2),$$

$$\int_3^{2n+1} \ln x\, dx = (2n+1)\ln(2n+1) - (2n+1) - \big(3\ln 3 - 3\big) = (2n+1)\ln(2n+1) - (2n-2) - 3\ln 3.$$

Lower bound. From $(2n-1)\ln(2n-1) - (2n-2) \le 2\ln P_n$,

$$\ln P_n \;\ge\; \tfrac{2n-1}{2}\ln(2n-1) - (n-1).$$

Exponentiating, $P_n \ge (2n-1)^{\frac{2n-1}{2}} e^{-(n-1)}$. Compare exponents of e against the target $\left(\frac{2n-1}{e}\right)^{\frac{2n-1}{2}} = (2n-1)^{\frac{2n-1}{2}} e^{-(2n-1)/2}$: since $-(n-1) = \frac{1}{2} - n > -n + \frac{1}{2}$ wait – directly, $-(n-1) > -\frac{2n-1}{2}$ because $-(n-1) = -n+1$ and $-\frac{2n-1}{2} = -n + \frac{1}{2}$, and $1 > \frac{1}{2}$. So $e^{-(n-1)} > e^{-(2n-1)/2}$ strictly, giving

$$P_n \;\ge\; (2n-1)^{\frac{2n-1}{2}} e^{-(n-1)} \;>\; (2n-1)^{\frac{2n-1}{2}} e^{-(2n-1)/2} \;=\; \left(\tfrac{2n-1}{e}\right)^{\frac{2n-1}{2}}.$$

Upper bound. From $2\ln P_n \le (2n+1)\ln(2n+1) - (2n-2) - 3\ln 3$,

$$\ln P_n \;\le\; \tfrac{2n+1}{2}\ln(2n+1) - (n-1) - \tfrac{3}{2}\ln 3,$$

so $P_n \le (2n+1)^{\frac{2n+1}{2}} e^{-(n-1)} 3^{-3/2}$. Compare against the target $\left(\frac{2n+1}{e}\right)^{\frac{2n+1}{2}} = (2n+1)^{\frac{2n+1}{2}} e^{-(2n+1)/2}$: the ratio of the derived bound's constant factor to the target's is

$$\frac{e^{-(n-1)} 3^{-3/2}}{e^{-(2n+1)/2}} = e^{-(n-1)+(2n+1)/2}\, 3^{-3/2} = e^{3/2}\, 3^{-3/2} = \left(\frac{e}{3}\right)^{3/2} < 1,$$

using $e < 3$. Hence

$$P_n \;\le\; (2n+1)^{\frac{2n+1}{2}} e^{-(n-1)} 3^{-3/2} \;<\; (2n+1)^{\frac{2n+1}{2}} e^{-(2n+1)/2} \;=\; \left(\tfrac{2n+1}{e}\right)^{\frac{2n+1}{2}}.$$

Both strict inequalities hold for every positive integer n (for $n = 1$ both sums are empty and the integral bounds are checked directly to give the same strict conclusion by the identical computation with k ranging over an empty set replaced by direct evaluation, which matches the hand-check above).

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both approaches accept the problem's framing – a family of inequalities in n – at face value at first. They split at the question of *what kind of argument turns a family of inequalities into one proof*. The natural attempt reads "family indexed by n" as "induct on n." The model chain instead reads the specific algebraic shape of the bounds, $(m/e)^{m/2} = m^{m/2} e^{-m/2} = \exp\big(\frac{m}{2}(\ln m - 1)\big)$, and recognises it as the exact exponential of the antiderivative $x\ln x - x$ evaluated at $x = m$ – which reframes the whole problem as a single estimate, not an induction.

> *The bound $(m/e)^{m/2}$ is* exp *of $\frac{m}{2}\ln m - \frac{m}{2}$, and that is exactly what you get from plugging $x = m$ into the antiderivative $x\ln x - x$ of $\ln x$ – so this whole problem is secretly "estimate $\int \ln x\,dx$ at the endpoints," and I should go find the integral instead of trying to chain inequality to inequality by hand.*

> *Induction would make me prove, at every single step, a fresh calculus inequality comparing $(1 + \frac{2}{2n-1})^{(2n-1)/2}$ to some power of e – and that auxiliary inequality is not obviously easier than the original claim, so I am not actually saving work by inducting; I am relocating the same difficulty one level down and paying for it n times instead of once.*

Layer 3 THE WHY

Why does the exact antiderivative approach succeed where induction merely relocates the difficulty? Because the target bounds are not arbitrary functions of n – they are literally exp evaluated at the closed form of $\int \ln x\,dx$, and $\ln x$ is monotonic. Monotonicity is the single structural fact that licenses sandwiching a sum of values of ln at evenly spaced points between the left- and right-rectangle Riemann sums for $\int \ln x\,dx$ over those same points – and because the antiderivative $x\ln x - x$ has a closed elementary form, "sandwich the sum, then exponentiate" turns directly into the stated bounds with no further approximation needed. The problem's specific choice of e as the base and $\frac{m}{2}$ as the exponent is not decorative: it is the fingerprint of $x\ln x - x$, put there so that a solver who recognises $\exp(x\ln x - x)$ knows immediately which integral to reach for and does not have to discover the antiderivative by trial. Bhargava–Kedlaya's solution states "we let $x = \cos\theta\ldots$" – no, for this problem their published solution goes directly to the rectangle-area computation without remarking on why $x\ln x - x$ is the right antiderivative to introduce; this is where that choice comes from – the target's algebraic shape names its own antiderivative. Induction fails to exploit this because induction treats the bound at $n+1$ as an unrelated new claim to be derived from the bound at n, rather than as two evaluations of one fixed, already-understood function at two different points – which is exactly what the integral method uses directly, and exactly what induction has to painstakingly reconstruct at every step via an auxiliary exponential inequality.

The strictness at the end – $1 < e < 3$ – is not incidental either: the rectangle-sum argument alone only yields the non-strict integral inequalities, so all of the strictness in the final claim has to come from somewhere, and it comes precisely from the numerical gap between $e^{3/2}$ and $3^{3/2}$ (equivalently, from e sitting strictly inside $(1,3)$). Any argument that reaches the same non-strict integral sandwich must locate its strictness in an explicit numerical fact about e at the very end, rather than assuming it falls out of the calculus for free.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the algebraic shape $(m/e)^{m/2}$ itself. On seeing it, write down: this equals $\exp\left(\frac{m}{2}\ln m - \frac{m}{2}\right)$, which is exp of the antiderivative $x\ln x - x$ evaluated at $x = m$ – so the problem is about $\int \ln x\,dx$, not (yet) about induction.

2. That cue suggests bounding a sum of logarithms, $\ln 3 + \ln 5 + \cdots + \ln(2n-1)$, between the left- and right-rectangle Riemann sums for $\int \ln x\,dx$ over the matching width-2 partition, using that $\ln x$ is monotonic increasing to know which endpoint over- and under-estimates.

3. Hesitation occurs when a chain of inequalities indexed by n triggers an induction reflex before the algebraic shape of the bounds has been examined. The reinforcement away from it: before setting up an inductive step, ask "does this bound's formula, as a function of the index, match the closed form of some standard antiderivative or sum estimate?" – if yes, try that first.

4. The trap is starting the induction, getting a clean base case, and then discovering at the inductive step that closing the gap requires a fresh, comparably hard exponential inequality

– at which point it is tempting to grind through it anyway (zombie maths) rather than step back and ask whether a direct estimate skips the recursion altogether.

5. To recognise this faster next time: whenever a competition inequality's bound has the form $f(n)^{g(n)}$ with $g(n)$ growing linearly and $f(n)$ itself linear or near-linear in n, suspect it is exp of a Riemann-sum estimate for $\int f$ or $\int \ln f$, and check that reading before reaching for induction.

6. The reusable rule: a family of inequalities whose bounds are exp of a recognisable antiderivative evaluated at the index is proved once, directly, via a monotone Riemann-sum sandwich of the corresponding sum of logs (or of the function itself) – not by induction, which forces you to re-derive an equivalent estimate at every step instead of stating it once.

1996 B3

1996 B3. Given that $\{x_1, x_2, \ldots, x_n\} = \{1, 2, \ldots, n\}$, find, with proof, the largest possible value, as a function of n (with $n \geq 2$), of

$$x_1x_2 + x_2x_3 + \cdots + x_{n-1}x_n + x_nx_1.$$

Fields: Combinatorics / optimisation on permutations; induction **Position:** B3 – third problem of the second session **Answer:** The maximum is $\dfrac{2n^3 + 3n^2 - 11n + 18}{6}$, achieved by the "interleaved zigzag" arrangement in which n sits between $n-1$ and $n-2$, and the two arcs running away from n list the remaining numbers in increasing order (e.g. for $n = 5$: the circle $1, 3, 5, 4, 2$).

Layer 1 THE WHAT — surface analysis

The natural attempt. The expression is a sum of products of *cyclically adjacent* entries, and "products of paired terms, maximise the sum" is exactly the shape that triggers the rearrangement inequality. The trouble surfaces immediately on trying to apply it: rearrangement pairs a fixed sequence against a permutation of itself, sorted the same way, to maximise $\sum a_i b_i$ – but here every value plays *both* roles at once (each x_i is multiplied by its left neighbour and its right neighbour), so there is no fixed sequence to sort against. The naive instinct this still leaves behind is "put them in sorted order round the circle," $1, 2, 3, \ldots, n$ back to 1. Checking $n = 4$: $1 \cdot 2 + 2 \cdot 3 + 3 \cdot 4 + 4 \cdot 1 = 2 + 6 + 12 + 4 = 24$. That is not the maximum – $1, 3, 4, 2$ gives $3 + 12 + 8 + 2 = 25$. So plain sorted order is beaten by an arrangement that puts the two largest values, 4 and 3, *next to each other*, at the cost of separating 4 from 1.

This pushes toward the right kind of idea – large numbers should neighbour large numbers – but turning that instinct into a proof by direct exchange argument is expensive. The natural route is: take an arbitrary arrangement, find two pairs of adjacent entries (a, b) and (c, d) (in that cyclic order) with $a < d$ and $b > c$, and check that reversing the arc from b to c changes the sum by $(d - a)(b - c)$– wait, first you must actually derive that formula (a short computation, but a computation nonetheless), confirm it's positive under those conditions, and then argue that repeatedly applying such reversals – to an arrangement with finitely many inversions of this "crossing" type – must terminate at an arrangement with no more available improving reversal. Having established that much, you still have to characterise what "no improving reversal available" means structurally (it turns out to force a single-peak, single-valley "mountain" shape), and then – this is the expensive part – show that among *all* mountain shapes (there are exponentially many, one for each way of splitting $\{1, \ldots, n-1\} \setminus \{n\}$ into the two rising arcs) only the specific interleaved zigzag is optimal, which needs yet another comparison argument between neighbouring mountain shapes. Each stage of this is doable, but it is three separate structural arguments stacked on top of each other, and getting any one of them slightly wrong (e.g. assuming any mountain shape is automatically optimal)

produces a plausible-looking wrong answer: for $n = 5$, the mountain arrangement $1, 2, 5, 4, 3$ (splitting the non-peak, non-valley numbers as $\{2\}$ and $\{3, 4\}$ rather than $\{3\}$ and $\{2, 4\}$) gives $1 \cdot 2 + 2 \cdot 5 + 5 \cdot 4 + 4 \cdot 3 + 3 \cdot 1 = 2 + 10 + 20 + 12 + 3 = 47$, one short of the true maximum 48 – confirming that "mountain shape" alone is necessary but not sufficient, and the naive exchange argument, followed only halfway, would stop at a wrong answer that looks finished.

The model chain. Let $M(n)$ denote the maximum value of the sum for a circular arrangement of $\{1, \ldots, n\}$. Work by induction on n, adding the numbers $1, \ldots, n$ one at a time in increasing order rather than optimising all n positions simultaneously.

Base case. $M(2) = x_1x_2 + x_2x_1 = 2 \cdot 1 \cdot 2 = 4$ (the only arrangement of $\{1, 2\}$ on a 2-cycle).

Inductive step. Take any arrangement of $\{1, \ldots, n\}$ on a circle. Deleting the entry n leaves an arrangement of $\{1, \ldots, n-1\}$ on a circle, together with the single edge (x, y) that used to be n's two edges (the neighbours of n, now joined directly). The value of the original n-arrangement equals

$$(\text{value of the } (n-1)\text{-arrangement obtained by deleting } n) \; + \; \big(nx + ny - xy\big),$$

since removing n deletes the two terms nx, ny and adds back the term xy for the edge that now closes up directly between its former neighbours x and y. Write the gain as

$$nx + ny - xy = n^2 - (n - x)(n - y).$$

Since x, y are two *distinct* elements of $\{1, \ldots, n-1\}$, the values $n - x$ and $n - y$ are two distinct elements of $\{1, \ldots, n-1\}$ as well, so their product is minimised exactly when $\{n - x, n - y\} = \{1, 2\}$ (the two smallest available distinct positive integers give the smallest product among all pairs of distinct positive integers drawn from that range) – that is, when $\{x, y\} = \{n - 1, n - 2\}$. This gives the minimum possible value of $(n - x)(n - y)$, namely $1 \cdot 2 = 2$, hence the maximum possible gain,

$$nx + ny - xy \;\le\; n^2 - 2,$$

with equality exactly when n is inserted between $n - 1$ and $n - 2$.

Therefore, for *any* n-arrangement, its value is at most $M(n-1) + n^2 - 2$ (apply the bound to the $(n-1)$-arrangement obtained by deleting n, whose value is at most $M(n-1)$ by definition, plus the gain bound just proved). And this bound is attained: take an optimal $(n-1)$-arrangement in which $n - 2$ and $n - 1$ happen to be adjacent, and insert n between them. That such an optimal $(n-1)$-arrangement exists follows by the same argument one level down – by induction, the optimal $(n-1)$-arrangement was itself built by inserting $n - 1$ between $n - 2$ and $n - 3$ in an optimal $(n-2)$-arrangement, so $n - 1$ and $n - 2$ are indeed adjacent in it. (Base case $n = 2$: the single arrangement trivially has 1 and 2 adjacent.) Hence

$$M(n) = M(n-1) + n^2 - 2, \qquad M(2) = 4.$$

Unwinding this recursion:

$$M(n) = 4 + \sum_{k=3}^{n} (k^2 - 2) = 4 + \left(\sum_{k=3}^{n} k^2 \right) - 2(n - 2).$$

Using $\sum_{k=1}^{n} k^2 = \frac{n(n+1)(2n+1)}{6}$ and $\sum_{k=1}^{2} k^2 = 5$,

$$M(n) = 4 + \frac{n(n+1)(2n+1)}{6} - 5 - 2(n - 2) = \frac{n(n+1)(2n+1)}{6} - 2n + 3.$$

Expanding $n(n+1)(2n+1) = 2n^3 + 3n^2 + n$:

$$M(n) = \frac{2n^3 + 3n^2 + n}{6} - 2n + 3 = \frac{2n^3 + 3n^2 + n - 12n + 18}{6} = \frac{2n^3 + 3n^2 - 11n + 18}{6},$$

matching the stated maximum, and the insertion construction (each new largest element placed between the two next-largest) reproduces exactly the interleaved zigzag described in the vitals.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths correctly diagnose that large values must sit next to large values. They split on *how many positions to optimise at once*. The natural attempt tries to fix the whole n-element arrangement simultaneously and repair it by pairwise exchanges – an n-variable optimisation that has to be resolved via a cascade of structural lemmas (mountain shape, then optimal split of the mountain). The model chain instead asks: *what if only the newest, largest element is a free variable, and everything smaller is already optimally placed?* – collapsing an n-variable problem into a sequence of 1-variable insertion problems.

> *I don't need to find the globally optimal arrangement of all n numbers at once – if I fix an optimal arrangement of the smaller numbers* $1, \ldots, n-1$ *and ask only "where does the single largest number n go," that is a one-variable optimisation (which edge to insert into), and if I can show the global optimum is always reachable this way, induction on n replaces an n-body problem with a sequence of* 1*-body problems.*

> *The gain from inserting n into edge* (x, y) *is* $n^2 - (n-x)(n-y)$*, which is maximised by* minimising $(n-x)(n-y)$ *over distinct* $x, y \in \{1, \ldots, n-1\}$ *– and the product of two distinct positive integers drawn from a fixed range is smallest when those integers are the two smallest available, so I should insert n next to whichever two values are themselves largest* $(n-1, n-2)$*, not reason about the insertion directly in terms of* x, y*.*

Layer 3 THE WHY

Why does the induction succeed cheaply where the direct exchange argument succeeds only expensively? Because the objective is a sum over *edges* of a cycle, and deleting the current maximum element from a cycle changes the objective in a strictly local, cleanly separable way: it removes exactly the two edges touching that element and replaces them with exactly one new edge (its former neighbours joined directly). Locality is what makes the induction available at all – if removing the largest element could affect the value of the remaining arrangement in some non-local way, "value of $(n-1)$-arrangement plus insertion gain" would not cleanly decompose the objective, and the induction would not go through. Once locality is granted, the insertion gain $n^2 - (n-x)(n-y)$ depends on x, y *only* through how small $(n-x)(n-y)$ can be made – a one-variable-feeling optimisation in disguise, because n is fixed and x, y range over a set whose extremes $(n-1, n-2)$ are known in advance without reference to the rest of the arrangement. This is precisely why the induction only ever needs to reason about the two largest remaining values at each step, never about the whole arrangement: the gain formula factors the dependence on the new element away from the dependence on the old arrangement, and that factoring is a direct consequence of the sum being edge-local on a cycle. The direct exchange argument, by contrast, works with the arrangement as a whole at every step – a reversal can touch an arbitrarily long arc – so it never gets access to this factoring, and has to rebuild global structure (mountain shape, then optimal split) from purely local swap comparisons, which costs the extra stages seen in the natural attempt. Bhargava–Kedlaya's own alternate solution states the insertion-gain formula and its maximisation in three lines without remarking on why induction succeeds so cheaply here; this is where that comes from – the objective's edge-locality is the structural fact that makes "add the largest element last" the correct induction variable, rather than an arbitrary choice among many equally natural ones.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "products of cyclically adjacent entries" combined with "maximise" – this looks like rearrangement inequality territory, but on writing out $\sum x_i x_{i+1}$,

notice immediately that each x_i appears in *two* product terms (with both neighbours), so there is no single fixed sequence to sort against; write down that plain rearrangement does not apply directly and a different structural idea is needed.

2. That cue, once plain rearrangement is ruled out, suggests either (a) a global exchange/reversal argument on the whole arrangement, or (b) building the arrangement up one element at a time and asking where each new element should go. Since the objective is a sum over edges of a cycle – a strictly local structure – test whether removing one vertex changes the objective in a clean, separable way; if it does, induction on the number of elements is available and usually cheaper.

3. Hesitation occurs right after confirming "large numbers should be adjacent to large numbers" from the $n = 4$ check, when it is tempting to jump straight to a full exchange argument on all n positions. The reinforcement away from it: before committing to the global argument, ask "does deleting the current maximum element decompose the objective into an $(n-1)$-piece plus a small local term?" – checking that decomposition is cheap and, once found, redirects the whole proof toward induction.

4. The trap is proving "the optimal arrangement is a single-peaked mountain shape" and stopping there, treating it as the final answer – as the $n = 5$ counterexample $(1, 2, 5, 4, 3$ giving 47 instead of 48) shows, being a mountain shape is necessary but not sufficient; the specific split of the two arcs still has to be pinned down.

5. To recognise this faster next time: whenever an optimisation objective is a sum over the edges of a cycle or path (each variable appearing in exactly two, or a bounded number of, terms), check whether removing the extremal element decomposes the objective into a smaller instance plus an isolated local gain term *before* reaching for a global exchange or rearrangement argument.

6. The reusable rule: for a "sum over edges/adjacent pairs" optimisation on a fixed ground set, try building the optimal configuration by inserting elements one at a time in extremal order (largest or smallest last), and check whether the marginal gain from each insertion depends only on a small, easily identified local choice (here, which two already-placed values to neighbour) – if so, this converts an n-variable global optimisation into n one-variable optimisations solved in sequence.

1996 B4

1996 B4. For any square matrix A, we can define $\sin A$ by the usual power series:

$$\sin A = \sum_{n=0}^{\infty} \frac{(-1)^n}{(2n+1)!} A^{2n+1}.$$

Prove or disprove: there exists a 2×2 matrix A with real entries such that

$$\sin A = \begin{pmatrix} 1 & 1996 \\ 0 & 1 \end{pmatrix}.$$

Fields: linear algebra; power series of matrices; commuting matrices **Position:** B4 — mid-paper; short once you know what A is allowed to look like, a swamp if you compute sin of a general 2×2 matrix **Answer:** No such A exists: A must commute with $\sin A$, which forces $A = pI + qE$ with $E = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, and then $\sin A = (\sin p)I + (q \cos p)E$ cannot have $\sin p = 1$ and $q \cos p = 1996$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student writes $A = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$ and tries to compute A^3, A^5, ... After A^3 the entries are cubic polynomials in four unknowns and it is plain this will never close.

So the student reaches for a structural fact: Cayley–Hamilton gives $A^2 = tA - \delta I$ with $t = a + d$, $\delta = ad - bc$, so every power of A is a combination of I and A, and therefore

$$\sin A = \alpha I + \beta A$$

for some real numbers α, β depending on t, δ. That is real progress. Comparing entries with the target: $\beta c = 0$, $\alpha + \beta a = 1$, $\alpha + \beta d = 1$, $\beta b = 1996$. So $\beta \neq 0$, hence $c = 0$ and $a = d$. The matrix is upper triangular with equal diagonal entries, $A = \left(\begin{smallmatrix} p & q \\ 0 & p \end{smallmatrix}\right)$.

Now the student computes $\sin A$ for this A directly, finds $\left(\begin{smallmatrix} \sin p & q\cos p \\ 0 & \sin p \end{smallmatrix}\right)$, and sees the contradiction. That route works. Its cost is the middle: the student must convince themselves that α, β exist as *convergent* series (the coefficients of I and A in A^{2n+1} are polynomials in t, δ that grow, and the infinite sum has to be regrouped), and many students stall there, or else they abandon Cayley–Hamilton and go to eigenvalues: "diagonalise A". Then the case split starts — distinct real eigenvalues, complex eigenvalues, a repeated eigenvalue — and each case needs its own computation. The official solution takes the eigenvalue route and has to work over $\mathbb{C}$. Nothing is wrong with any of this; it is simply three times longer than it needs to be.

The model chain. Write $N = \left(\begin{smallmatrix} 1 & 1996 \\ 0 & 1 \end{smallmatrix}\right)$ and $E = \left(\begin{smallmatrix} 0 & 1 \\ 0 & 0 \end{smallmatrix}\right)$, so $N = I + 1996E$ and $E^2 = 0$.

Step 1: A commutes with $\sin A$. Each partial sum $S_m = \sum_{n \le m} \frac{(-1)^n}{(2n+1)!} A^{2n+1}$ is a polynomial in A, so $AS_m = S_m A$. The series converges entrywise (each entry of A^k is at most $(2\|A\|)^k$ in absolute value, where $\|A\|$ is the largest entry of $|A|$), and matrix multiplication is continuous, so $A \sin A = (\sin A)A$. If $\sin A = N$, then A commutes with N, hence with $E = (N - I)/1996$.

Step 2: what commutes with E. For $M = \left(\begin{smallmatrix} a & b \\ c & d \end{smallmatrix}\right)$,

$$EM = \begin{pmatrix} c & d \\ 0 & 0 \end{pmatrix}, \qquad ME = \begin{pmatrix} 0 & a \\ 0 & c \end{pmatrix}.$$

These are equal exactly when $c = 0$ and $d = a$. So $A = pI + qE$ for some real p, q.

Step 3: compute $\sin(pI + qE)$. Since pI and qE commute and $E^2 = 0$, the binomial theorem gives $A^k = p^k I + k p^{k-1} qE$. Summing the series term by term (both resulting real series converge absolutely),

$$\sin A = \Big(\sum_n \tfrac{(-1)^n p^{2n+1}}{(2n+1)!}\Big) I + q\Big(\sum_n \tfrac{(-1)^n (2n+1) p^{2n}}{(2n+1)!}\Big) E = (\sin p)\, I + (q \cos p)\, E.$$

Step 4: contradiction. $\sin A = I + 1996E$ requires $\sin p = 1$ and $q \cos p = 1996$. But $\sin p = 1$ gives $\cos^2 p = 1 - \sin^2 p = 0$, so $q \cos p = 0 \neq 1996$. Hence no such A exists. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that the question is "what can A look like?". They split at the first move. (N) *computes* $\sin A$ for an unknown A — by entries, by Cayley–Hamilton, or by diagonalising — and only afterwards compares with the target. (M) uses the target first: before computing anything, it asks what $\sin A = N$ forces on A, and the answer is "A commutes with N". That one line collapses four unknowns to two and the case analysis to nothing.

> *I don't know A, but I do know $\sin A$. Anything built out of A by a power series commutes with A. So A has to commute with this specific matrix — let me find everything that commutes with it before I compute a single power.*

> 1 1996 0 1 *is the identity plus a nilpotent. The identity commutes with everything, so the real constraint is "commutes with E", and that is a two-line computation.*

> *For $A = pI + qE$ the series is a Taylor expansion: $f(p + q\varepsilon) = f(p) + qf'(p)\varepsilon$ when $\varepsilon^2 = 0$. So the off-diagonal entry is $q \cos p$, and* cos *vanishes exactly where* sin *is* ± 1.

- ▷ **Cue:** an equation $f(A) = B$ for a matrix function given by a power series, B known $\longrightarrow$ **Move:** A commutes with $f(A) = B$; compute the commutant of B first
- ▷ **Cue:** a target matrix of the form $\lambda I + (\text{nilpotent})$ $\longrightarrow$ **Move:** drop the λI; commuting with B is the same as commuting with the nilpotent part
- ▷ **Cue:** a power series evaluated at $pI + qE$ with $E^2 = 0$ $\longrightarrow$ **Move:** use $f(pI + qE) = f(p)I + qf'(p)E$ (dual numbers: first-order Taylor is exact)
- ▷ **Cue:** a problem asking for $\sin x = 1$ together with some multiple of $\cos x$ being nonzero $\longrightarrow$ **Move:** $\sin^2 + \cos^2 = 1$ kills it: $\sin x = \pm 1$ forces $\cos x = 0$

Layer 3 THE WHY

Why is "A commutes with $\sin A$" the natural first line and not a trick? Go back to the definition in the problem statement. The problem *hands you* $\sin A$ as a power series in A, that is, as a limit of polynomials in A. A polynomial in A is made of A and I, and A commutes with A and with I. So the definition itself says $\sin A$ lives in the algebra generated by A, and every element of that algebra commutes with A. When the problem tells you the value of something in that algebra, it is telling you something about A; the only question is how to extract it, and "commutes with" is the cheapest possible extraction, because it is linear in A.

Why is it effective here specifically? Because the target is almost scalar. A matrix with distinct eigenvalues, say $\mathrm{diag}(1, 2)$, commutes only with diagonal matrices — a two-parameter family — and that would already have been decisive. Our target $I + 1996E$ is even more restrictive in the way that matters: it is a nonzero nilpotent plus a scalar, and the commutant of a nonzero 2×2 nilpotent is $\{pI + qE\}$. In 2×2 matrices, any matrix that is not a scalar has a two-dimensional commutant, namely the polynomials in it. So the commutant step always leaves exactly two unknowns, and then the problem is a one-variable calculus question about sin.

Now the structure of $\sin(pI + qE)$. Why should the off-diagonal entry be $q \cos p$? Because $E^2 = 0$ makes $pI + qE$ behave like the "dual number" $p + q\varepsilon$ with $\varepsilon^2 = 0$, and for any power series $f(p + q\varepsilon) = f(p) + qf'(p)\varepsilon$: the binomial theorem stops after its linear term. That is exactly the first-order Taylor expansion, with no remainder. So the target asks for a point where $f(p) = 1$ and $f'(p) \neq 0$. For $f = \sin$, $f(p) = 1$ is the maximum of sin, and at a maximum the derivative vanishes. That is the real reason for the contradiction, and it tells you when the answer would flip: replace 1 by any value s with $|s| < 1$ and the equation $\sin A = sI + 1996E$ is solvable (pick $p = \arcsin s$, $q = 1996/\cos p$). The problem was built on the one value where sin has a critical point.

Kedlaya and Bhargava's solution conjugates A into upper-triangular form over $\mathbb{C}$ and argues through eigenvalues and diagonalisability. That needs the facts that sin commutes with conjugation, that a diagonalisable matrix has diagonalisable sin, and that complex conjugation is permitted. Every one of those is a disguised version of the commutant argument: conjugation-invariance of power series is the same algebra-generated-by-A fact. The commutant route stays over $\mathbb{R}$, needs no case split, and is shorter. Their alternate solution (Helfgott–Popa) uses $\sin^2 A + \cos^2 A = I$ to get $\cos^2 A$ nilpotent and nonzero, which is impossible for 2×2; that is also clean, and is the same critical-point fact in matrix clothing.

What the natural attempt did wrong is not an error; it is zombie computation. It computed $\sin A$ for an A it knew nothing about, when the problem had already told it $\sin A$.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "sin A = (a specific matrix)" with sin defined by a power series. On seeing it, write: $A \sin A = \sin A \cdot A$, so A commutes with $\left(\begin{smallmatrix} 1 & 1996 \\ 0 & 1 \end{smallmatrix}\right)$.

2. Commutant computation, then evaluate the power series on the two-parameter family $pI + qE$ via $E^2 = 0$.

3. Hesitation comes at "what does a general A look like?" — the student feels obliged to cover every Jordan form. It is reinforced away by noting that the commutant computation *is* the case analysis, done in one line, and it does not care about eigenvalues.

4. The trap is the substitution of a guess for a proof: many students try $A = \left(\begin{smallmatrix} p & q \\ 0 & p \end{smallmatrix}\right)$ because it looks like the target, get the contradiction, and write "disproved" without justifying why A must have that shape. The commutant step is the justification. A secondary trap is the value 1: it hides that the obstruction is $\sin' = 0$ at a maximum.

5. Ask: *what do I already know about the unknown, from the equation it satisfies?* Any time an unknown X is fed into something built from X alone (a power series, a polynomial, a flow), the output commutes with X — or, more generally, shares its symmetries. Use that before computing.

6. Reusable decision rule: *when an equation is $F(X) = B$ with F built from X alone, first restrict X to the objects that B forces it to be compatible with (commutant, invariant subspaces, symmetry group); then solve on that smaller family.*

1996 B5

1996 B5. Given a finite string S of symbols X and O, we write $\Delta(S)$ for the number of X's in S minus the number of O's. For example, $\Delta(XOOXOOX) = -1$. We call a string S **balanced** if every substring T of (consecutive symbols of) S has $-2 \leq \Delta(T) \leq 2$. Thus, $XOOXOOX$ is not balanced, since it contains the substring $OOXOO$. Find, with proof, the number of balanced strings of length n.

Fields: combinatorics; lattice walks; prefix sums; inclusion–exclusion **Position:** B5 — late paper; the counting is short once the condition is translated, but the condition as stated ("every substring") is quadratic in size and invites a state machine **Answer:** $2^{\lceil n/2 \rceil} + 2 \cdot 2^{\lfloor n/2 \rfloor} - 2$, i.e. $3 \cdot 2^{n/2} - 2$ for n even and $2^{(n+3)/2} - 2$ for n odd (equivalently $2^{\lfloor (n+2)/2 \rfloor} + 2^{\lfloor (n+1)/2 \rfloor} - 2$).

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student starts with data. Length 1: X, O, so 2. Length 2: all four. Length 3: all but XXX, OOO, so 6. Length 4: of the 16, remove any string containing XXX or OOO, and also strings like $XXOX$? No — $\Delta(XXOX) = 2$, allowed. What about $XOXX$ and $XXOXX$? The student realises the condition is not only "no three in a row": $XXOXX$ has $\Delta = 3$ with no three consecutive equal letters. So the forbidden patterns are not local, and a clean list of forbidden substrings is not available.

The next move is a finite automaton. The student decides to track, while reading the string left to right, the maximum and minimum Δ of a suffix ending at the current position, and builds a transfer matrix on those states. With care, this produces the counts

$$2,\ 4,\ 6,\ 10,\ 14,\ 22,\ 30,\ 46,\ \ldots$$

and a recurrence $a_n = a_{n-1} + 2a_{n-2} - 2a_{n-3}$, or the pattern "differences are $2, 2, 4, 4, 8, 8, 16, \ldots$". From there one guesses the closed form and proves it by induction on the automaton. This works. It costs most of the time available, the state set is easy to get wrong (it is not obvious which states are reachable, or that the state space is finite at all), and the resulting proof explains nothing.

The model chain. **Step 1: translate the condition to prefix sums.** Write the string as $\varepsilon_1\varepsilon_2\cdots\varepsilon_n$ with $\varepsilon_i = +1$ for X and -1 for O, and let $s_0 = 0$, $s_j = \varepsilon_1 + \cdots + \varepsilon_j$. The substring in positions $i+1, \ldots, j$ has $\Delta = s_j - s_i$. So the string is balanced iff $|s_j - s_i| \le 2$ for all $0 \le i < j \le n$, i.e. iff

$$\max_j s_j - \min_j s_j \le 2.$$

The balanced strings correspond bijectively to walks $s_0 = 0, s_1, \ldots, s_n$ with steps ± 1 whose *range* (the set of values visited) spans at most three consecutive integers.

Step 2: count walks in a fixed strip. For an integer a, let W_a be the number of ± 1 walks of length n from 0 staying inside $\{a, a+1, a+2\}$. Only $a \in \{-2, -1, 0\}$ contain 0.

- $a = -1$ (strip $\{-1, 0, 1\}$, start in the middle): from 0 both steps are allowed; from ± 1 the only allowed step is back to 0. So steps $1, 3, 5, \ldots$ are free and steps $2, 4, \ldots$ are forced: $W_{-1} = 2^{\lceil n/2 \rceil}$.
- $a = 0$ (strip $\{0, 1, 2\}$, start at the bottom): from 0 the step is forced up, from 2 forced down, from 1 free. The walk is at 0 or 2 after an even number of steps and at 1 after an odd number, so steps $2, 4, \ldots$ are free: $W_0 = 2^{\lfloor n/2 \rfloor}$.
- $a = -2$ is the mirror image of $a = 0$: $W_{-2} = 2^{\lfloor n/2 \rfloor}$.

Step 3: correct for overcounting. A walk with range ≤ 2 lies in at least one of the three strips. If its range is exactly 2 (it visits three values), it lies in exactly one. If its range is exactly 1 (for $n \ge 1$), it visits $\{0, 1\}$ or $\{-1, 0\}$ and so lies in exactly two strips; there are exactly two such walks, $XOXO\cdots$ and $OXOX\cdots$. Range 0 is impossible for $n \ge 1$. Hence the number of balanced strings is

$$W_{-2} + W_{-1} + W_0 - 2 = 2^{\lceil n/2 \rceil} + 2 \cdot 2^{\lfloor n/2 \rfloor} - 2.$$

For $n = 2m$ this is $3 \cdot 2^m - 2$; for $n = 2m+1$ it is $2^{m+1} + 2^{m+1} - 2 = 2^{m+2} - 2$. Check: $n = 1, 2, 3, 4$ give $2, 4, 6, 10$, matching the direct count. (I also checked by exhaustive enumeration up to $n = 12$.) □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin by staring at "every substring T has $-2 \le \Delta(T) \le 2$". (N) treats substrings as the objects and builds a machine to police them as it reads. (M) rewrites each substring's Δ as a *difference of two prefix sums*, $s_j - s_i$. At that line the condition on $\binom{n+1}{2}$ substrings becomes one condition on the walk: its maximum minus its minimum is at most 2. Everything after is counting walks in a strip of width 2, which is almost forced step by step.

> *Δ of a substring is a sum over a window. Sums over windows are differences of partial sums. So "every window has $|\Delta| \le 2$" means "any two partial sums differ by at most 2" — the whole walk lives in a band of width 2.*

> *A band of width 2 has only three levels. On three levels a ± 1 walk has no choice at the edges and a free choice in the middle. So I count by where 0 sits in the band, and subtract the walks I have counted twice.*

▷ **Cue:** a condition on every consecutive block (substring, subarray, window) of a sequence ⟶ **Move:** introduce prefix sums s_j; the block sum is $s_j - s_i$; "all blocks" becomes a condition on $\max s - \min s$

▷ **Cue:** counting ± 1 walks confined to a strip of width 2 ⟶ **Move:** edges force the step, the middle is free; count free steps by parity

▷ **Cue:** objects described by "range at most k" $\longrightarrow$ **Move:** sum over the possible positions of a width-k window, then subtract the objects with smaller range that were counted more than once

▷ **Cue:** a sequence $2, 4, 6, 10, 14, 22, 30, \ldots$ with differences $2, 2, 4, 4, 8, 8$ $\longrightarrow$ **Move:** suspect two interleaved geometric sequences, i.e. a parity split in a walk

Layer 3 THE WHY

Why prefix sums? Go to what Δ is. $\Delta(T)$ is additive: if T is T_1 followed by T_2, then $\Delta(T) = \Delta(T_1) + \Delta(T_2)$. Any additive function of intervals is a difference of a single function of endpoints, exactly as $\int_a^b f = F(b) - F(a)$. That is the fundamental theorem of calculus in its discrete form, and it is the same reason a bank statement records a running balance rather than a list of every possible sum of consecutive transactions. The problem's condition "every substring" is a statement about all pairs of endpoints, and once you have a function of endpoints, "all pairs differ by at most 2" is by definition "max minus min is at most 2". Nothing clever has happened; the problem stated the condition in the least efficient form and you restated it in the most efficient.

Why does the strip counting come out as powers of 2 with a parity split? A ± 1 walk alternates parity: $s_j \equiv j \pmod 2$. In a three-level strip $\{a, a+1, a+2\}$ two of the levels have one parity and the middle has the other. At the two outer levels the walk has one legal move; at the middle it has two. So the walk alternates between "forced" and "free" steps, and which steps are free depends only on whether 0 is the middle level or an outer one. That is all Step 2 is.

Why inclusion–exclusion over the strips, and why only a correction of 2? The objects we want are walks with range ≤ 2. "Range ≤ 2" means "fits in *some* width-2 strip", a union of three sets. Unions are counted by adding and subtracting overlaps, and the overlap of two strips is a strip of width 1, in which a ± 1 walk is forced entirely (it must alternate), so each width-1 strip containing 0 holds exactly one walk. The three-way overlap is $\{0\}$, which holds no walk of positive length. So the correction is tiny and exact.

Kedlaya and Bhargava's solution tiles the string with dominoes on adjacent equal letters and argues that consecutive dominoes are separated by an even gap. Here is where that comes from. A domino (two equal adjacent letters) is precisely a moment where the walk moves two levels in a row in the same direction, i.e. goes from one edge of the strip to the other. Between two such moments the walk must alternate, and it must cross back the other way at the next domino, which forces the parity condition. Their two cases (dominoes starting at even or at odd positions) are my strips with 0 at an edge versus 0 in the middle. Their "no dominoes counted twice" is my subtraction of the two alternating walks. It is the same count, but the walk makes it visible why the parity condition appears; the domino picture has to argue it.

The natural attempt did honest work in the wrong representation. Every substring was treated as a separate constraint, so the student needed a machine to remember constraints. Prefix sums remember them for free.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "every substring T of (consecutive symbols of) S" together with an additive quantity Δ. On seeing it, write $s_j = \Delta(\text{first } j \text{ symbols})$ and $\Delta(T) = s_j - s_i$.

2. Reformulate as ± 1 walks with $\max - \min \leq 2$; count walks in each of the three width-2 strips containing 0; subtract the doubly counted width-1 walks.

3. Hesitation comes at the counting: "walks of bounded range" sounds like a reflection-principle problem. It is reinforced away by noticing the strip has only three levels, so the walk is forced at the edges and free in the middle — no reflection needed.

4. The trap is believing the condition is local ("no XXX or OOO"). $XXOXX$ has no three equal in a row and still has $\Delta = 3$. Anyone who counts strings avoiding three-in-a-row gets twice a Fibonacci number (16 at $n = 5$ instead of 14) and the wrong answer from $n = 5$ on. The

second trap is inclusion–exclusion bookkeeping: forgetting the two alternating strings gives $+2$.

5. Whenever a condition is imposed on all intervals of a sequence and the quantity is additive over intervals, the right object is the sequence of partial sums, and the interval condition becomes a condition on pairs of points.
6. Reusable decision rule: *an additive quantity over all consecutive blocks is a difference of prefix values; turn "for every block" into "for every pair of prefix values", which is usually a max/min or ordering condition on a single walk.*

1996 B6

1996 B6. Let $(a_1, b_1), (a_2, b_2), \ldots, (a_n, b_n)$ be the vertices of a convex polygon which contains the origin in its interior. Prove that there exist positive real numbers x and y such that

$$(a_1, b_1)x^{a_1}y^{b_1} + (a_2, b_2)x^{a_2}y^{b_2} + \cdots + (a_n, b_n)x^{a_n}y^{b_n} = (0, 0).$$

Fields: multivariable calculus; convexity; existence of minima on compact sets **Position:** B6 — last problem, genuinely hard to start; once the sum is recognised as a gradient the proof is a compactness argument **Answer:** Proof: with $x = e^u$, $y = e^v$ the left side is $\nabla f(u, v)$ for $f = \sum_i e^{a_i u + b_i v}$; the interior-origin hypothesis makes $f \geq e^{M\sqrt{u^2+v^2}}$ with $M > 0$, so f has a global minimum, where $\nabla f = 0$.

Layer 1 THE WHAT — surface analysis

The natural attempt. This is a B6 and it looks it: two equations, n terms each, real exponents, and a geometric hypothesis that does not obviously connect to the algebra. The strong student writes the two equations,

$$\sum_i a_i x^{a_i} y^{b_i} = 0, \qquad \sum_i b_i x^{a_i} y^{b_i} = 0,$$

and tries to solve one at a time. Fix y and look at $G_y(x) = \sum_i a_i x^{a_i} y^{b_i}$. Since the origin is inside the polygon, some $a_i > 0$ and some $a_i < 0$; as $x \to \infty$ the positive-a_i terms dominate and $G_y \to +\infty$, as $x \to 0^+$ the negative-a_i terms dominate and $G_y \to -\infty$. (If some $a_i = 0$, that term contributes 0 and is harmless.) The derivative $x\,G_y'(x) = \sum a_i^2 x^{a_i} y^{b_i} > 0$, so there is exactly one root $x = \xi(y)$. Good.

Now substitute into the second equation: $H(y) = \sum_i b_i\, \xi(y)^{a_i} y^{b_i}$, and look for a sign change as $y \to 0$ and $y \to \infty$. This is the wall. As $y \to \infty$, $\xi(y)$ moves, and which terms dominate H depends on how fast. If $\xi(y) \approx y^\lambda$ then the term i has size $y^{a_i\lambda + b_i}$ and the dominant terms are those maximising $a_i\lambda + b_i$ — a supporting line of the polygon with slope determined by λ — and λ itself is fixed by the requirement that the first equation balances. The student is now doing asymptotic analysis of an implicit function with a case split according to which edge of the polygon is active. It can be done (it is essentially Elkies's winding-number argument, which the official text calls "a little tricky"), but it will not be done in the time.

The model chain. Write $p_i = (a_i, b_i)$ and, for $z = (u, v) \in \mathbb{R}^2$, $\langle p_i, z\rangle = a_i u + b_i v$.

Step 1: the sum is a gradient. Put $x = e^u$, $y = e^v$ (a bijection from $\mathbb{R}^2$ onto the open positive quadrant). Then $x^{a_i} y^{b_i} = e^{\langle p_i, z\rangle}$ and the left-hand side becomes

$$\sum_i p_i e^{\langle p_i, z\rangle} = \nabla f(z), \qquad f(z) = \sum_{i=1}^{n} e^{\langle p_i, z\rangle},$$

since $\partial_u e^{a_i u + b_i v} = a_i e^{a_i u + b_i v}$ and likewise for v. So it suffices to find a point where $\nabla f = 0$.

Step 2: every direction sees a vertex ahead of it. Let w be a unit vector. I claim $g(w) = \max_i \langle p_i, w\rangle > 0$. If not, every vertex lies in the closed half-plane $\{q : \langle q, w\rangle \le 0\}$; this half-plane is convex, so it contains the polygon (the convex hull of its vertices). But the origin is interior, so εw lies in the polygon for small $\varepsilon > 0$, and $\langle \varepsilon w, w\rangle = \varepsilon > 0$. Contradiction.

Step 3: uniform growth. g is a maximum of finitely many continuous functions, hence continuous on the unit circle, which is compact; so g attains a minimum M, and $M > 0$ by Step 2. For $z \neq 0$, g is positively homogeneous, so $\max_i \langle p_i, z\rangle = |z|\, g(z/|z|) \ge M|z|$. All terms of f are positive, so $f(z) \ge e^{\max_i \langle p_i, z\rangle} \ge e^{M|z|}$.

Step 4: a global minimum. Let $R = \ln(n+1)/M$ and D the closed disk $|z| \le R$. On $|z| \ge R$, $f(z) \ge e^{MR} = n+1$. The continuous f attains a minimum on compact D at some z^*, and $f(z^*) \le f(0) = n < n+1$. Hence $|z^*| < R$, and $f(z^*) \le f(z)$ for every $z \in \mathbb{R}^2$ (inside D by choice of z^*, outside because $f \ge n+1$ there). An interior minimum of a differentiable function has zero gradient: $\nabla f(z^*) = 0$. With $(u^*, v^*) = z^*$, the numbers $x = e^{u^*} > 0$, $y = e^{v^*} > 0$ satisfy the equation. □

Bonus: uniqueness. The Hessian of f is $\sum_i e^{\langle p_i, z\rangle} p_i p_i^{\mathsf{T}}$. For $h \neq 0$, $h^{\mathsf{T}}(\cdot)h = \sum_i e^{\langle p_i, z\rangle} \langle p_i, h\rangle^2 > 0$, because not every p_i is orthogonal to h (by Step 2 applied to $h/|h|$). So f is strictly convex and the critical point, hence (x, y), is unique.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains look at the same vector equation. (N) reads it as *two scalar equations* and tries to solve them in sequence. (M) reads it as *one vector*, $\sum p_i \cdot (\text{positive weight}_i)$, and asks where a vector with that shape comes from. The split happens at the substitution $x = e^u$, $y = e^v$: after it, the weight $x^{a_i} y^{b_i}$ is $e^{\langle p_i, z\rangle}$, whose gradient is $p_i e^{\langle p_i, z\rangle}$ — exactly the i-th term. An existence question for a zero of a vector field becomes an existence question for a minimum of a function, and existence of minima is what compactness is for.

> *Each term is a vector p_i times a positive number that depends on x, y through $x^{a_i} y^{b_i}$. That is exactly how a gradient of $x^{a_i} y^{b_i}$ looks in log coordinates: the exponent pops down. So the whole thing is $\nabla \sum e^{a_i u + b_i v}$, and I want a critical point.*

> *I can't solve for a critical point, but I can prove a minimum exists. For that I need $f \to \infty$ in every direction. In direction w the fastest term grows like $e^{t \max_i \langle p_i, w\rangle}$, so I need every direction to have some vertex strictly ahead of it — and that is exactly "the origin is in the interior".*

- ▷ **Cue:** a sum $\sum_i v_i \phi_i(z)$ of fixed vectors times positive scalars, and a request to make it zero ⟶ **Move:** ask whether it is the gradient of $\sum \Phi_i$; if so, look for a minimum instead of solving
- ▷ **Cue:** monomials $x^a y^b$ with real exponents on the positive quadrant ⟶ **Move:** substitute $x = e^u$, $y = e^v$: monomials become exponentials of linear forms, and $x\partial_x$ becomes ∂_u
- ▷ **Cue:** "prove a solution exists" with no hope of writing it down ⟶ **Move:** find a function whose critical points are the solutions; prove it is coercive; take a minimum on a large compact ball
- ▷ **Cue:** "the origin is in the interior of the convex hull of $p_1, \ldots, p_n$" ⟶ **Move:** translate to "for every direction w, $\max_i \langle p_i, w\rangle > 0$", then by compactness $\ge M|w|$ with $M > 0$

Layer 3 THE WHY

Why is "gradient" the natural reading, not a rabbit from a hat? Look at the shape of the equation as the problem prints it: vectors (a_i, b_i) multiplied by scalars $x^{a_i} y^{b_i}$, and *the same* a_i, b_i appear in the vector and in the exponent. The only operation in elementary calculus that takes an exponent and puts it in front as a coefficient is differentiation: $\frac{d}{dx} x^a = a x^{a-1}$, or more cleanly

$x\frac{d}{dx}x^a = ax^a$. Two coordinates, two exponents, two coefficients: that is a gradient. The operator $x\partial_x$ is awkward, and it is exactly ∂_u when $x = e^u$. So the log substitution is not a trick either; it is the change of variables that turns $x\partial_x$ into an ordinary derivative, and it simultaneously turns the positive quadrant (an open set with boundary to worry about) into the whole plane (no boundary).

Why a minimum rather than any other critical point? Because a minimum is the one kind of critical point whose existence you can prove without computing anything: a continuous function on a compact set attains its minimum (the extreme value theorem), and if the minimum is interior the gradient vanishes. Every positive-weight exponential sum is at least 0 and has no maximum, so a minimum is the only candidate.

Why is "origin in the interior" exactly the right hypothesis? Because it is exactly coercivity. Walk off to infinity in direction w. The terms of f are $e^{t\langle p_i, w\rangle}$, and f blows up iff at least one $\langle p_i, w\rangle > 0$ — iff some vertex lies strictly on the w-side of the line through the origin perpendicular to w. Asking this for every w is asking that no line through the origin has the whole polygon weakly on one side, which is the definition of the origin being interior (for a convex set, via the supporting line). The problem's geometric hypothesis *is* the analytic hypothesis you need, rewritten. And if the origin were on the boundary, say on an edge, the direction perpendicular to that edge would have $\max\langle p_i, w\rangle = 0$ and f would stay bounded along that ray — the minimum can escape to infinity. That is why Marshall Buck's "in its interior" remark in the official note matters.

Why compactness of the circle, and not only positivity in each direction? Positivity in each direction gives $f \to \infty$ along each ray, but not uniformly, and non-uniform growth does not trap a minimum. The minimum M of g on the circle is what makes the growth uniform: $f \geq e^{M|z|}$ everywhere. That is the second place compactness is used, and it is the one people skip.

Now the official solution. Kedlaya and Bhargava give exactly this argument, but their radius is wrong: they say $f \geq n+1$ outside the disk of radius $\sqrt{(n+1)/M}$. From $f \geq e^{M|z|}$ the correct radius is $\ln(n+1)/M$. The two are not comparable in general. Take the rhombus with vertices $(\pm 1, 0), (0, \pm 0.01)$: then $M \approx 0.01$, $\sqrt{5/M} \approx 22.4$, and at $z = (0, 22.4)$ one gets $f = 2 + e^{0.224} + e^{-0.224} \approx 4.05 < 5$. The argument survives with the corrected radius, which is what I used in Step 4. The lesson for the arena: when you invoke "$f \to \infty$ uniformly", write the radius down from the inequality you actually have.

The natural attempt was not foolish; it fixed one variable and found a unique root by monotonicity, which is the correct instinct for one equation in one unknown. It failed because it split a problem with a single vector structure into two scalar problems, and the coupling came back as implicit-function asymptotics.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the vector (a_i, b_i) is also the exponent vector of $x^{a_i}y^{b_i}$. On seeing it, write $x\partial_x(x^{a_i}y^{b_i}) = a_i x^{a_i}y^{b_i}$, then $x = e^u, y = e^v$.

2. Recognise the left side as ∇f for $f = \sum e^{\langle p_i, z\rangle}$; prove a global minimum exists by coercivity plus the extreme value theorem.

3. Hesitation comes after finding f: "I can't solve $\nabla f = 0$". It is reinforced away by recalling that existence of a minimiser needs only continuity and growth at infinity, and the hypothesis on the polygon is a statement about growth in every direction.

4. The trap is splitting into two scalar equations and chasing an implicit function. A second trap is proving $f \to \infty$ along every ray and calling it done: without the uniform bound from compactness of the circle, the minimum is not trapped. A third is copying the official radius $\sqrt{(n+1)/M}$, which is wrong.

5. When you are asked to show that a system of equations has a solution, and you cannot solve it, ask whether the system is the critical-point equation of something. Positive-weighted sums of fixed vectors whose weights are exponentials of those vectors are always gradients.

6. Reusable decision rule: *to prove a vector equation has a solution, find a potential whose gradient it is, show the potential is coercive (grows to infinity uniformly in every direction), and take its minimum on a large closed ball; the minimum lies inside, so the gradient vanishes there.*