DAWN PUTNAM SESSION

Putnam 1971 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1971 — Paper A **5**
1971 A1 5
1971 A2 7
1971 A3 9
1971 A4 11
1971 A5 14
1971 A6 16

Putnam 1971 — Paper B **21**
1971 B1 21
1971 B2 23
1971 B3 26
1971 B4 28
1971 B5 31
1971 B6 33

Putnam 1971 — Paper A

1971 A1

1971 A1. Let there be given nine lattice points (points with integral coordinates) in three dimensional Euclidean space. Show that there is a lattice point on the interior of one of the line segments joining two of these points.

Fields: combinatorics, pigeonhole principle, lattice points **Position:** A1 (opener) **Answer:** Among any nine lattice points in $\mathbb{R}^3$, some pair has a lattice point strictly between them – the midpoint of that pair.

Layer 1 THE WHAT — surface analysis

The natural attempt. Nine points in three-dimensional space, and a claim that some segment between two of them passes through a third lattice point. The word "nine" is unusual enough to register as a count, and the first instinct for a strong student is geometric: try to visualise where nine points could sit so that no segment between any pair contains a lattice point in its interior, and see what goes wrong. You start sketching small cases – four points at the corners of a square, then a cube's eight vertices – and notice that a cube's eight corners are exactly the points where this fails: no edge, face-diagonal, or space-diagonal of the unit cube $\{0,1\}^3$ contains a lattice point strictly between its endpoints, because every coordinate difference between two cube vertices is 0 or 1, never 2. That is a genuinely useful discovery – it tells you eight is achievable without collision – but it is also exactly where many students stall: you have found the extremal configuration for eight points, and the problem is about nine, so you know intuitively that the ninth point must break the pattern, yet "why" is not yet a proof. Trying to argue directly ("the ninth point must go somewhere, and everywhere is bad") is correct in spirit but needs the actual mechanism made explicit, and re-deriving it from the geometric picture of a cube, rather than from the arithmetic that made the cube special in the first place, costs time and risks missing why eight is the exact threshold rather than seven or ten.

The model chain. Do not think about geometric position at all – think about parity. A lattice point in $\mathbb{R}^3$ has three integer coordinates, and each coordinate is independently either odd or even. That gives exactly $2^3 = 8$ possible parity patterns: (odd,odd,odd), (odd,odd,even), and so on through all combinations. Every lattice point, no matter where it sits in space, belongs to exactly one of these eight classes.

With nine lattice points distributed among only eight classes, the Pigeonhole Principle forces two of them, say $P = (p_1, p_2, p_3)$ and $Q = (q_1, q_2, q_3)$, into the same class: $p_i \equiv q_i \pmod 2$ for each $i = 1, 2, 3$. That means each difference $p_i - q_i$ is even, so the midpoint

$$M = \left(\frac{p_1 + q_1}{2}, \frac{p_2 + q_2}{2}, \frac{p_3 + q_3}{2} \right)$$

has every coordinate an integer – M is itself a lattice point. Since $P \neq Q$ (they are two of the nine given points, listed as distinct), M is strictly between P and Q, i.e. in the interior of segment PQ.

This produces the required lattice point directly, with no case-work on geometric configuration at all.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains correctly identify that eight points can be placed lattice-collision-free (the cube's vertices) and that nine cannot. They split on what characterises "collision-free": the natural attempt reaches for the geometric picture of the cube and tries to argue from its shape why a ninth point must fail; the model instead asks what algebraic property made the cube's eight vertices safe in the first place – their coordinates only ever differ by 0 or 1 – and generalises that property (parity) to a classification that applies to every lattice point in space, not just cube corners.

> *The cube's eight vertices work not because of their shape, but because their coordinates only take two values each,* 0 *and* 1 *– and "only two values" is really a parity statement in disguise: I should classify every lattice point by whether each coordinate is odd or even, since that's a property every point in space has, not one particular to the cube.*

Layer 3 THE WHY

The parity classification is not an ad hoc trick found by staring at the cube; it is the direct consequence of one fact: a midpoint of two integers is itself an integer if and only if the two integers have the same parity. $\frac{p+q}{2} \in \mathbb{Z} \iff p \equiv q \pmod 2$. This single one-dimensional fact is the entire mechanism, applied independently to each of the three coordinates. "Avoiding a lattice-point midpoint" is therefore, coordinate by coordinate, exactly "avoiding two points of matching parity" – and a set of points avoids all same-parity pairs precisely when it hits each of the $2^3 = 8$ parity classes at most once. The cube's eight vertices are simply one full transversal of the eight classes (each vertex has coordinates in $\{0,1\}$, and the eight sign patterns $\{0,1\}^3$ are exactly the eight parity classes, one representative each) – that is why eight is achievable and why it is the exact extremal number, not a coincidence about cubes. The ninth point, wherever placed, must repeat a class already used, because there are only eight classes to place it in.

This also explains why geometric intuition about the cube, while suggestive, cannot by itself finish the proof: the cube shows one witness configuration for $n = 8$, but the argument needed is about *every* possible configuration of nine points, anywhere in space, and only an argument phrased in terms of the parity classes (which exist independent of any particular geometric arrangement) has that generality built in. The problem's phrasing – "nine lattice points," full stop, no further geometric constraint – is itself a signal that whatever mechanism is used must not depend on how the points are arranged, which rules out a geometric argument and all but requires an invariant (parity) that is arrangement-independent.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the number nine paired with the phrase "lattice points" and a claim about segment interiors. On seeing it, write down: "midpoint of two lattice points is a lattice point exactly when the two points agree in parity, coordinate by coordinate – count the parity classes."

2. That cue suggests the Pigeonhole Principle: with $2^3 = 8$ parity classes available in three dimensions and nine points to place, some class must receive two points.

3. Hesitation occurs right after noticing the cube's eight vertices avoid the problem – it is tempting to keep reasoning geometrically about cube edges, faces, and diagonals. Reinforce against it by asking: what algebraic property, not what shape, made those eight points safe? The answer (coordinates only 0 or 1, hence one point per parity class) redirects to the classification argument.

4. The trap is treating this as a geometry problem about a specific configuration (the cube) rather than a counting problem about an arrangement-independent invariant (parity class). A

geometric argument about the cube proves only that eight points *can* avoid a collision; it says nothing about why nine points *always* collide, which is what pigeonhole on parity classes delivers directly.

5. Recognise this faster next time by treating "how many objects can I place before two must share a property" as an immediate cue for counting the number of possible values of that property and applying pigeonhole – here the property is the 2^3 parity pattern of a point's coordinates.

6. Decision rule: when a problem gives a lattice-point (or other discretely-classifiable) hypothesis together with a threshold count one more than a power of 2 (or more generally one more than the number of classes under some natural finite classification), look for the classification whose class-count matches, and apply pigeonhole – the specific number in the hypothesis is almost always exactly (number of classes) $+1$.

1971 A2

1971 A2. Determine all polynomials $P(x)$ such that $P(x^2+1) = (P(x))^2 + 1$ and $P(0) = 0$.

Fields: functional equations, polynomials, root-counting **Position:** A2 **Answer:** $P(x) = x$ is the unique such polynomial.

Layer 1 THE WHAT — surface analysis

The natural attempt. The two conditions are concrete enough to try a direct ansatz. Since $P(0) = 0$, guess P is linear: $P(x) = ax$. Substituting,

$$P(x^2+1) = a(x^2+1) = ax^2 + a, \qquad (P(x))^2 + 1 = a^2x^2 + 1.$$

Matching coefficients of x^2 gives $a = a^2$, and matching constants gives $a = 1$; both together force $a = 1$, so $P(x) = x$ works. This feels like the problem is solved, but it has only checked the linear case. A careful student then tries to rule out higher degree by comparing leading coefficients in general: if P has degree n with leading coefficient c, then $P(x^2+1)$ has leading term cx^{2n}, while $(P(x))^2 + 1$ has leading term c^2x^{2n}, so $c = c^2$, forcing $c = 1$ (since $c \neq 0$). That pins the leading coefficient to 1, for every degree n – but it does not pin down n itself. At this point the natural attempt is stuck checking degree 2, then degree 3, by direct substitution and coefficient-matching, an approach that could in principle rule out each degree in turn but never actually proves there is no valid degree, because there are infinitely many degrees to check and no finite amount of case-work closes the argument.

The model chain. Use the functional equation itself as a machine for generating known values of P, rather than trying to determine P's shape by its coefficients. Define the sequence

$$x_0 = 0, \qquad x_{n+1} = x_n^2 + 1 \quad (n \geq 0),$$

so $x_0 = 0,\ x_1 = 1,\ x_2 = 2,\ x_3 = 5,\ x_4 = 26, \ldots$, a strictly increasing sequence of distinct nonnegative integers. Claim: $P(x_n) = x_n$ for every n. Proof by induction. Base case: $P(x_0) = P(0) = 0 = x_0$, given directly. Inductive step: assume $P(x_n) = x_n$; then

$$P(x_{n+1}) = P(x_n^2 + 1) = (P(x_n))^2 + 1 = x_n^2 + 1 = x_{n+1},$$

using the functional equation in the middle step and the inductive hypothesis in the third. This establishes $P(x_n) = x_n$ for all $n \geq 0$.

Now consider the polynomial $Q(x) = P(x) - x$. It vanishes at every x_n, and since the x_n are strictly increasing they are pairwise distinct, so Q has infinitely many roots. A nonzero polynomial has only finitely many roots (at most its degree), so Q must be the zero polynomial:

$P(x) = x$ identically. This holds regardless of what degree P was assumed to have – the argument never needed to know n in advance, because it produces infinitely many fixed points directly from the functional equation and the single seed value $P(0) = 0$.

Finally check $P(x) = x$ genuinely satisfies both conditions: $P(0) = 0$ holds, and $P(x^2+1) = x^2+1 = (P(x))^2+1$ holds since $P(x) = x$. So $P(x) = x$ is the unique solution.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same two facts – $P(0) = 0$, and that the leading coefficient of any solution must equal 1 (from matching leading terms in $P(x^2+1) = (P(x))^2 + 1$). They split on what to do next: the natural attempt tries to pin down the *degree* n of P by testing degrees one at a time, an open-ended case search. The model instead abandons the question of degree entirely and iterates the substitution $x \mapsto x^2+1$ starting from the known value $P(0) = 0$, generating an infinite sequence of points where P is forced to agree with the identity function.

> *I don't need to guess or bound the degree of P and check each one – if I can show $P(x) = x$ at infinitely many points, the polynomial $P(x) - x$ must be the zero polynomial regardless of what degree P was, because a nonzero polynomial can only have finitely many roots.*

Layer 3 THE WHY

The move to generate an infinite orbit is not an external trick imported from elsewhere; it is handed directly by the shape of the functional equation itself. $P(x^2+1) = (P(x))^2+1$ says precisely: "if you know the value of P at some point t, you immediately know the value of P at t^2+1, namely $(P(t))^2+1$." That is a recursive rule for producing new known values from old ones, and $P(0) = 0$ supplies the one seed value needed to set it running. So the "infinite orbit" idea is not a clever import – it is what the functional equation is, read as an instruction rather than as an equation to solve for coefficients.

The reason this orbit approach settles the problem completely, while the degree-matching approach never can, is the root-counting theorem for polynomials: a polynomial of degree d has at most d roots unless it is identically zero. This is itself a direct consequence of the factor theorem (each root r contributes a factor $(x-r)$, and these factors are coprime for distinct roots, so d of them cannot divide a polynomial of degree less than d unless that polynomial is the zero polynomial). Because the orbit $x_0, x_1, x_2, \ldots$ is infinite, no finite degree d can survive having $P(x) - x$ vanish at more than d of these points – the argument closes off every possible degree *simultaneously*, which is exactly what the degree-by-degree case search was trying and failing to do one degree at a time.

The phrasing of the problem is itself the tell: it hands you a value at one point ($P(0) = 0$) together with a rule that propagates known values forward ($x \mapsto x^2+1$ inside P). Whenever a functional equation pairs an explicit seed value with a substitution rule of the form $f(g(x)) = h(f(x))$, the seed and the substitution together generate an orbit on which f is pinned down – and root-counting turns "pinned down on an infinite set" into "pinned down everywhere," for a polynomial.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is having a concrete seed value, $P(0) = 0$, combined with a functional equation of the shape $P(g(x)) = h(P(x))$. On seeing it, write down: "plug in the seed value and see what new value the equation produces for free."
2. That cue suggests iterating the substitution: define $x_{n+1} = x_n^2 + 1$ from $x_0 = 0$, and prove by induction that $P(x_n) = x_n$ for all n.
3. Hesitation occurs right after verifying $P(x) = x$ satisfies the equation for the linear ansatz – it is tempting to stop there. Reinforce against it by asking explicitly: "have I shown no other

polynomial works, or only that this one does?" – matching a candidate is not the same as proving uniqueness.

4. The trap is trying to bound the degree of P by comparing leading coefficients and then checking degrees one at a time; this only pins the leading coefficient to 1 and never terminates as a proof of uniqueness, because infinitely many degrees remain unchecked.

5. Recognise this faster next time by treating any functional equation of the form $f(g(x)) = h(f(x))$ paired with one known value $f(x_0)$ as an immediate invitation to generate the orbit $x_0, g(x_0), g(g(x_0)), \ldots$ and test whether it is infinite and distinct.

6. Decision rule: to prove a functional equation pins down a polynomial uniquely, do not analyse coefficients degree by degree; instead find a seed value and a self-feeding substitution that together generate infinitely many points of agreement with the candidate solution, then invoke that a nonzero polynomial has only finitely many roots.

1971 A3

1971 A3. The three vertices of a triangle of sides a, b, and c are lattice points and lie on a circle of radius R. Show that $abc \geq 2R$. (Lattice points are points in the Euclidean plane with integral coordinates.)

Fields: plane geometry, lattice points, circumradius **Position:** A3 **Answer:** $abc \geq 2R$ always, with equality possible; the bound comes from $2 \times$ (area) being a positive integer.

Layer 1 THE WHAT — surface analysis

The natural attempt. The circumradius R immediately suggests the Law of Sines: $a = 2R\sin A$, $b = 2R\sin B$, $c = 2R\sin C$, where A, B, C are the angles opposite a, b, c. Multiplying,

$$abc = 8R^3 \sin A \sin B \sin C.$$

The target inequality $abc \geq 2R$ then becomes $8R^3 \sin A \sin B \sin C \geq 2R$, i.e.

$$R^2 \sin A \sin B \sin C \geq \tfrac{1}{4}.$$

This is a clean-looking trigonometric inequality, and it is tempting to attack it directly – perhaps via AM-GM on the angles, or by recalling that $\sin A \sin B \sin C$ is maximised for the equilateral triangle. But working the inequality this way runs into a wall almost immediately: nothing in it refers to lattice points at all. R, A, B, C can be any circumradius and angles of any triangle; the inequality $R^2 \sin A \sin B \sin C \geq \frac{1}{4}$ is simply false in general (take R small and any fixed angles, or a very obtuse triangle with a large circumradius and small angles – the left side can be made far smaller than $\frac{1}{4}$). Something essential – the fact that the vertices are lattice points – has been thrown away by passing to angles and never used. The natural attempt can chase this dead end for some time before recognising that trigonometric substitution alone cannot recover a hypothesis it discarded.

The model chain. Keep the connection between abc and R that goes through the triangle's *area*, since area is the one quantity with a direct, computable link to lattice-point coordinates. Starting from the same Law-of-Sines relation $a = 2R\sin A$ (angle A opposite side a), and writing the area as $\text{Area} = \frac{1}{2}bc\sin A$, eliminate $\sin A$ between the two:

$$\sin A = \frac{a}{2R}, \qquad \text{Area} = \tfrac{1}{2}bc \cdot \frac{a}{2R} = \frac{abc}{4R}.$$

So $abc = 4R \cdot \text{Area}$ – an identity relating abc to R through area alone, with no sines left in it.

Now bring in the lattice-point hypothesis. For a triangle with integer-coordinate vertices $(x_1, y_1), (x_2, y_2), (x_3, y_3)$, the shoelace formula gives

$$\text{Area} = \frac{1}{2} \left| x_1(y_2 - y_3) + x_2(y_3 - y_1) + x_3(y_1 - y_2) \right|.$$

Every term inside the absolute value is a product and sum of integers, hence an integer; so $2 \times$ Area is a nonnegative integer. Since the triangle is non-degenerate (it has a well-defined circumradius R, so its three vertices are not collinear), Area > 0, and a positive integer is at least 1: $2 \times \text{Area} \geq 1$, i.e. Area $\geq \frac{1}{2}$.

Combining: $abc = 4R \cdot \text{Area} \geq 4R \cdot \frac{1}{2} = 2R$, exactly the required inequality.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin from the same Law-of-Sines fact, $a = 2R \sin A$. They split immediately after: the natural attempt multiplies all three side relations together to get $abc = 8R^3 \sin A \sin B \sin C$, eliminating every trace of area and lattice structure in the process. The model instead combines the single relation $a = 2R \sin A$ with the area formula $\text{Area} = \frac{1}{2}bc \sin A$ to eliminate $\sin A$ and land on $abc = 4R \cdot \text{Area}$ – a form that still has area sitting in it, ready to receive the lattice-point hypothesis.

> *The bound is an exact, specific constant – $2R$, not "roughly R" or "at least cR for some constant c" – and that kind of exactness in a competition inequality almost always comes from an integer somewhere being at least 1; the only place an integer can come from here is the lattice-point hypothesis, so I need a formula that turns "vertices are lattice points" into "some quantity is an integer," and area is the natural candidate since it's computable directly from coordinates.*

Layer 3 THE WHY

Why does "lattice points" force $2 \times$ Area to be an integer? Because the shoelace formula expresses area as one-half of a sum of products of integer coordinates – and any expression built entirely from integers using only addition, subtraction, and multiplication is itself an integer. This is not a fact about triangles specifically; it is the closure of $\mathbb{Z}$ under these operations, applied to the determinant that computes area. The formula $2 \cdot \text{Area} = \det \begin{pmatrix} x_1 - x_3 & y_1 - y_3 \\ x_2 - x_3 & y_2 - y_3 \end{pmatrix}$ (equivalent to the shoelace expression above) makes this transparent: a determinant of an integer matrix is an integer, full stop, regardless of what the matrix represents geometrically.

This also explains precisely why the trigonometric path cannot work, structurally rather than just by counterexample: $\sin A, \sin B, \sin C$ are continuous functions of the triangle's shape with no memory of whether the vertices happen to be lattice points. Passing from coordinates to angles is a lossy transformation for this problem – it forgets the one piece of information (integrality) that the proof needs. The area, by contrast, is computed *directly* from the coordinates via a formula that is manifestly integer-valued when the inputs are integers; it is the quantity that transports the lattice-point hypothesis forward without loss. The identity $abc = 4R \cdot \text{Area}$ is the bridge that lets that integrality reach all the way to abc and R, because it is an *equality*, not an inequality – nothing is lost or approximated in passing from area to abc, so the integer content of the area survives intact into the final bound.

The phrasing of the problem is itself a signal: an exact threshold like $2R$ (not an asymptotic or approximate one) attached to a lattice-point hypothesis should make you look for the identity that converts "lattice points" into "some specific quantity is a positive integer," rather than reach for inequalities among continuous quantities like angles.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the exact constant 2 in "$abc \geq 2R$" sitting next to the hypothesis "vertices are lattice points." On seeing it, write down: "an exact integer-flavoured bound plus a

lattice-point hypothesis means look for a quantity forced to be a positive integer, not for a continuous trigonometric inequality."

2. That cue suggests finding the identity linking abc, R, and area ($abc = 4R \cdot \text{Area}$), then computing area directly from the lattice coordinates via the shoelace/determinant formula to show $2 \times$ Area is a positive integer.

3. Hesitation sets in right after writing $a = 2R\sin A$, $b = 2R\sin B$, $c = 2R\sin C$ and multiplying them together – it feels like progress because it directly involves abc and R. Reinforce against it by checking whether the lattice-point hypothesis has been used yet; if the current expression could be evaluated without knowing the vertices are lattice points, the hypothesis has been discarded and the route needs to change.

4. The trap is multiplying all three Law-of-Sines relations together, which produces an expression in abc and R that looks relevant but has erased every trace of the lattice-point hypothesis, leading to a false general trigonometric inequality.

5. Recognise this faster next time by checking, after any substitution in a geometry problem, whether the special hypothesis given (lattice points, integer coordinates, rational coordinates, etc.) is still present in the resulting expression; if it has vanished, the substitution has thrown away necessary information.

6. Decision rule: when a problem combines a lattice-point (or other integrality) hypothesis with an exact numerical threshold, look first for the identity relating the target quantities through *area* (via $abc = 4R \cdot \text{Area}$ or the shoelace/determinant formula), since area is the quantity that transmits integrality directly from coordinates; avoid routes through angles or trigonometric functions, which discard integrality.

1971 A4

1971 A4. Show that for $0 < \epsilon < 1$ the expression $(x+y)^n(x^2-(2-\epsilon)xy+y^2)$ is a polynomial with positive coefficients for n sufficiently large and integral. For $\epsilon = .002$ find the smallest admissible value of n.

Caveat. The MAA OCR prints the parameter as "e" and flattens the exponent; the statement is restored as $(x+y)^n(x^2-(2-\epsilon)xy+y^2)$ with $0<\epsilon<1$ and $\epsilon = .002$. Identity check: the book's solution under 1971 A-4 expands exactly this product, minimises the bracket $\phi(k)$ at $k = n/2$, and concludes $n = 1999$, so label and statement match. The answer 1999 was also confirmed by exact rational arithmetic: every coefficient is positive for $n = 1999, 2000, 2001$ and some coefficient is ≤ 0 for $n = 1996, 1997, 1998$.

Fields: algebra (binomial coefficients), inequalities **Position:** A4 — mid-paper; routine once the coefficient is written as a binomial times a bracket, messy if you try to compare three binomials directly
Answer: All coefficients are positive exactly when $\epsilon > 4/(n+2)$ for even n and $\epsilon > 4/(n+3)$ for odd n; for $\epsilon = .002$ the smallest admissible n is 1999 (and every $n \geq 1999$ works; $n = 1998$ gives a zero coefficient).

Layer 1 THE WHAT — surface analysis

The natural attempt. You expand. Write the product as $\sum_j c_j x^j y^{n+2-j}$. The x^j term collects three contributions — x^2 times $\binom{n}{j-2}x^{j-2}$, the middle term times $\binom{n}{j-1}x^{j-1}$, and y^2 times $\binom{n}{j}x^j$:

$$c_j = \binom{n}{j-2} - (2-\epsilon)\binom{n}{j-1} + \binom{n}{j}.$$

So far so good, and correct. Now you have to decide when three binomial coefficients with a negative middle weight sum to something positive. The instinct is to compare them pairwise:

$\binom{n}{j-2} + \binom{n}{j}$ against $2\binom{n}{j-1}$. Near the edges of Pascal's row the outer ones are small and the middle one dominates... no, at the very edge ($j = 1$) you get $0 - (2-\epsilon) + n$, which is positive. In the middle of the row the three binomials are nearly equal and the sum is nearly $\epsilon\binom{n}{n/2}$ — but "nearly" is exactly the issue. You try $n = 4$, $n = 6$ by hand, get some coefficients negative, and have no way to extrapolate to $\epsilon = .002$ without a formula. The student who stops here has the right expression and no handle on it.

The model chain. Put $k = j - 1$, so $c_{k+1} = \binom{n}{k-1} - (2-\epsilon)\binom{n}{k} + \binom{n}{k+1}$ for $0 \le k \le n$, and the two extreme coefficients $c_0 = c_{n+2} = 1$ are positive. Factor out the middle binomial using the ratios

$$\binom{n}{k-1} = \binom{n}{k}\frac{k}{n-k+1}, \qquad \binom{n}{k+1} = \binom{n}{k}\frac{n-k}{k+1},$$

(valid also at $k = 0$ and $k = n$, where the vanishing terms give 0). Hence

$$c_{k+1} = \binom{n}{k}\big(\phi(k) - (2-\epsilon)\big), \qquad \phi(k) = \frac{k}{n-k+1} + \frac{n-k}{k+1}.$$

Since $\binom{n}{k} > 0$, all coefficients are positive iff $\phi(k) > 2 - \epsilon$ for every integer $0 \le k \le n$.

Treat k as real on $[0, n]$:

$$\phi'(k) = \frac{n+1}{(n-k+1)^2} - \frac{n+1}{(k+1)^2},$$

which is negative for $k < n/2$ and positive for $k > n/2$. So ϕ decreases then increases, and over the integers its minimum is at $k = \lfloor n/2 \rfloor$ or $\lceil n/2 \rceil$ (these give equal values, since $\phi(n-k) = \phi(k)$).

n even, $k = n/2$: $\phi = 2 \cdot \frac{n/2}{n/2+1} = 2 - \frac{4}{n+2}$. Condition: $\epsilon > \frac{4}{n+2}$.

n odd, $k = \frac{n+1}{2}$: $\frac{k}{n-k+1} = 1$ and $\frac{n-k}{k+1} = \frac{n-1}{n+3}$, so $\phi = 2 - \frac{4}{n+3}$. Condition: $\epsilon > \frac{4}{n+3}$.

Both thresholds tend to 0, so for each $\epsilon \in (0, 1)$ all coefficients are positive for all sufficiently large n (any $n > 4/\epsilon - 2$ works). For $\epsilon = .002$: even n needs $n + 2 > 2000$, i.e. $n \ge 2000$; odd n needs $n + 3 > 2000$, i.e. $n \ge 1999$. At $n = 1998$ the middle coefficient is $\binom{1998}{999}(2 - \frac{4}{2000} - 2 + .002) = 0$, not positive. Both thresholds decrease in n, so every $n \ge 1999$ is admissible and the smallest is $n = 1999$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write $c_j = \binom{n}{j-2} - (2-\epsilon)\binom{n}{j-1} + \binom{n}{j}$. At that line the natural attempt tries to compare three binomials; the model divides through by the middle one, turning three factorial-sized numbers into one positive factor times a bracket of two rational functions of k.

> *Three neighbouring binomial coefficients differ by simple ratios. If I divide by the middle one, the sign of the whole thing is the sign of a bracket I can actually minimise.*

> *The question is "positive for all k". That is "the minimum over k is positive". So I need the k where the bracket is smallest, and by symmetry of Pascal's row that has to be the middle.*

> *"Smallest admissible n" means parity may matter: the middle of the row is one term for even n and two for odd n. Compute both before answering.*

Layer 3 THE WHY

Why is dividing by $\binom{n}{k}$ the natural move rather than a trick? Because the question is about a *sign*, and a sign is unchanged by multiplying by a positive number. Whenever you are asked "is

this sum of terms positive?" and the terms share a large common factor, the first act is to remove it; what is left is the part of the expression that actually decides the sign. For neighbours in Pascal's row the common factor is always available, because the definition $\binom{n}{k} = \frac{n!}{k!(n-k)!}$ makes adjacent entries differ by one factor each in numerator and denominator: $\binom{n}{k+1}/\binom{n}{k} = \frac{n-k}{k+1}$. That ratio is the same fact that proves the row increases to the middle and then decreases. So the bracket $\phi(k)$ is nothing exotic — it is the sum of the two "step ratios" out of position k, one to the left and one to the right.

That reading also explains, without calculus, where the minimum sits. Near the edges one step ratio is huge (the row is climbing steeply), so ϕ is large; in the middle the row is flat and both ratios are close to 1, so ϕ is close to 2. The constant $2-\epsilon$ is testing precisely how flat the top of Pascal's row is. The official solution differentiates ϕ; that is fine and I have kept it, but the flatness picture is what tells you in advance where to look.

Now the deeper structure. Rewrite the quadratic factor:

$$x^2 - (2-\epsilon)xy + y^2 = (x-y)^2 + \epsilon xy.$$

So the product is $(x+y)^n(x-y)^2 + \epsilon\, xy(x+y)^n$. The second piece has all coefficients positive, of size $\epsilon\binom{n}{k}$. The first piece has coefficients equal to the *second difference* of the binomial row, $\binom{n}{k-1} - 2\binom{n}{k} + \binom{n}{k+1}$, which is negative near the middle, where the row is concave. The problem is a race: the concavity at the top of Pascal's row, relative to its height, is about $4/n$; the perturbation is ϵ. Positivity holds exactly when ϵ beats the relative curvature. That is why the answer is $n \approx 4/\epsilon$, and why "sufficiently large n" exists at all: the relative curvature at the peak of the binomial row goes to 0 as the row gets wider. (This is the discrete shadow of the Gaussian shape of the binomial distribution: its width grows like $\sqrt{n}$, and relative curvature at the peak is width$^{-2} \sim 1/n$.)

A last point about the parity split. The book remarks that if the middle coefficients are non-positive for some odd n, they remain so at the next even n — equivalently, $4/(n+3)$ at odd n equals $4/((n+1)+2)$ at the next even n. Odd n gets a slightly easier test because its two central entries are equal, so the row is "flatter" at the top. You do not need that remark if you compute both cases, which is cheaper than being clever about it.

Layer $+\alpha$ BACKTESTING

1. The cue is "positive coefficients". Write down the general coefficient c_j immediately, as a combination of three binomials; everything follows from that one line.

2. A sign question on a combination of neighbouring binomials suggests dividing by the middle binomial, reducing to a rational function of k, and minimising it.

3. Hesitation comes at $k = 0$ and $k = n$, where $\binom{n}{-1}$ or $\binom{n}{n+1}$ appears. Reinforce by checking directly: the edge coefficients are $n - (2-\epsilon)$ and 1, both positive for $n \geq 2$, and the ratio formulas give the right value 0 for the missing terms. The second hesitation is whether the real minimiser $n/2$ is an integer; handle even and odd n separately.

4. The trap is the non-strict inequality. At $n = 1998$ the central coefficient is exactly 0, so "$\epsilon \geq 4/(n+2)$" gives the wrong answer 1998. The second trap is answering 2000 by only doing the even case.

5. Recognise $x^2 - (2-\epsilon)xy + y^2$ as $(x-y)^2 + \epsilon xy$: a perturbation of a square that kills positivity. Whenever a problem perturbs a degenerate object by ϵ and asks for n, expect n of order a power of $1/\epsilon$ and look for the quantity that decays in n.

6. To decide the sign of a sum of large terms sharing a common factor, divide the factor out first; the sign lives in the ratio. Then "positive for all k" becomes "minimum over k positive", and symmetry of the setup tells you where the minimum is before you compute it.

▷ **Cue:** "polynomial with positive coefficients" after multiplying by $(x+y)^n$ $\longrightarrow$ **Move:** write the general coefficient as a short combination of neighbouring binomials $\binom{n}{k-1}, \binom{n}{k}, \binom{n}{k+1}$

▷ **Cue:** a sign question about $\binom{n}{k-1} - \lambda\binom{n}{k} + \binom{n}{k+1}$ ⟶ **Move:** divide by $\binom{n}{k}$ using $\binom{n}{k\pm1}/\binom{n}{k}$ and minimise the bracket over k

▷ **Cue:** a quadratic factor $x^2 - (2-\epsilon)xy + y^2$ ⟶ **Move:** rewrite as $(x-y)^2 + \epsilon xy$: second difference of the row plus a positive perturbation

▷ **Cue:** "for n sufficiently large" with a small parameter ϵ ⟶ **Move:** find the quantity that decays in n (here $4/(n+2)$) and set it below ϵ

▷ **Cue:** "smallest admissible n" with a symmetric row ⟶ **Move:** split into even and odd n and check equality cases at the boundary

1971 A5

1971 A5. A game of solitaire is played as follows. After each play, according to the outcome, the player receives either a or b points (a and b are positive integers with a greater than b), and his score accumulates from play to play. It has been noticed that there are thirty-five non-attainable scores and that one of these is 58. Find a and b.

Caveat. The OCR of the statement is clean apart from "orb" for "or b". Identity check: the book's solution under 1971 A-5 counts non-attainable scores as $\frac{1}{2}(a-1)(b-1)$ and concludes $a = 11$, $b = 8$; the statement is also quoted verbatim as "Problem 1971A5" in the 1985–2000 MAA volume. The answer was confirmed by brute force over all $1 \le b < a < 80$: $(11, 8)$ is the only pair with exactly 35 non-attainable positive scores including 58.

Fields: number theory (Frobenius / Sylvester coin problem) **Position:** A5 — the formula is the whole problem; with it the rest is a factor table **Answer:** $a = 11$, $b = 8$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Attainable scores are the numbers $xa + yb$ with $x, y \ge 0$. You start experimenting. Suppose $b = 2$: then every even number is attainable and the odd ones below a are not, so there are $(a-1)/2$ gaps; 35 gaps gives $a = 71$. But 58 is even, so $58 = 29 \cdot 2$ is attainable. Discard. Try $b = 3$: the gaps are the numbers below $2a$ that are not reachable... you start listing residues mod 3 and it becomes a table. Try $b = 5$, $b = 7$, each a fresh table, each time counting gaps by hand, and each time the count of 35 is a long tally prone to off-by-one errors. After forty minutes you may have the answer by exhaustion, with no certainty that you have not missed a pair with large b. The problem has quietly punished zombie enumeration: it gives you a *count*, and the only efficient use of a count is a formula.

The model chain. Call $m \ge 0$ *attainable* if $m = xa + yb$ with integers $x, y \ge 0$.

Step 1: $\gcd(a, b) = 1$. If $d = \gcd(a, b) > 1$, every attainable score is a multiple of d, so infinitely many scores are non-attainable. Thirty-five is finite, so $d = 1$. Also $b \ge 2$, since $b = 1$ makes everything attainable.

Step 2: the count is $\frac{1}{2}(a-1)(b-1)$. Let $N = ab - a - b$. Because $\gcd(a, b) = 1$, every integer m has a unique expression $m = xa + yb$ with $0 \le x \le b - 1$ (choose $x \equiv ma^{-1} \pmod b$; then $y = (m - xa)/b$ is determined). Call this the *standard form*.

Claim: m is attainable iff $y \ge 0$ in standard form. If $y \ge 0$, done. Conversely, any representation $m = x'a + y'b$ with $x', y' \ge 0$ can be shifted by $(x', y') \mapsto (x' - b, y' + a)$ until $0 \le x' \le b - 1$; the shifts only raise y', and the result is the standard form, so $y \ge y' \ge 0$.

Now for any integer m, with $m = xa + yb$ in standard form,

$$N - m = (b - 1 - x)a + (-1 - y)b, \qquad 0 \le b - 1 - x \le b - 1,$$

which is the standard form of $N - m$. So $N - m$ is attainable iff $-1 - y \ge 0$ iff $y < 0$ iff m is *not* attainable. Exactly one of m, $N - m$ is attainable.

Consequences: negative integers are not attainable, so every $m > N$ is attainable (its partner $N - m$ is negative). On $\{0, 1, \ldots, N\}$ the map $m \mapsto N - m$ pairs each attainable score with a non-attainable one, so exactly half of these $N + 1$ numbers are non-attainable:

$$\#\{\text{non-attainable}\} = \frac{N+1}{2} = \frac{(a-1)(b-1)}{2}.$$

($N + 1 = (a-1)(b-1)$ is even because a, b coprime cannot both be even, and if either is odd the corresponding factor is even.)

Step 3: the factor table. $(a-1)(b-1) = 70$ with $a > b \geq 2$: $(a-1, b-1) \in \{(70,1), (35,2), (14,5), (10,7)\}$, i.e. $(a,b) \in \{(71,2), (36,3), (15,6), (11,8)\}$. Coprimality removes $(36,3)$ and $(15,6)$. For $(71,2)$, $58 = 0 \cdot 71 + 29 \cdot 2$ is attainable. For $(11,8)$: $58 - 11x$ for $x = 0, \ldots, 5$ is $58, 47, 36, 25, 14, 3$, none divisible by 8, so 58 is not attainable. Hence $a = 11$, $b = 8$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from "attainable means $xa + yb$ with $x, y \geq 0$". At the next line the natural attempt fixes a small b and counts gaps; the model asks for the number of gaps as a function of a and b, because the data given is a number of gaps.

> *They have told me how many scores are missing, not which. A count is only useful if I have a formula for the count, so the problem is really: how many numbers are not of the form $xa + yb$?*

> *To count the missing numbers, I pair each missing one with a present one. The largest missing number is $ab - a - b$, so try $m \leftrightarrow ab - a - b - m$.*

> *The second clue, "58 is not attainable", is there to kill a spurious candidate from the factor table. Use it last, as a filter.*

Layer 3 THE WHY

Why is the formula the natural target? Read the problem as data. Two facts are given: "thirty-five" and "58". Two unknowns are asked for. A single number like 35 cannot pin down two unknowns unless it is a known function of them that factors well; a single example like 58 cannot identify anything unless the candidates are already few. So the design of the problem announces its solution: turn 35 into an equation in a, b, get a short list, and use 58 to choose. If you enumerate by b instead, you are ignoring the only information that makes the answer unique — that is why the enumeration feels endless.

Now why the pairing $m \leftrightarrow ab - a - b - m$, which the official solution does not use? The book instead counts lattice points in a rectangle and halves by the diagonal. Both arguments rest on one axiom-level fact: since $\gcd(a, b) = 1$, the multiples $0, a, 2a, \ldots, (b-1)a$ hit every residue class mod b exactly once. Everything follows from that. In the residue class of xa (with $0 \leq x \leq b - 1$) the attainable numbers are exactly $xa,\ xa + b,\ xa + 2b, \ldots$ — an arithmetic progression starting at xa. So each residue class has a first attainable member xa, and below it the missing members $xa - b,\ xa - 2b, \ldots$ down to 0. The number missing in that class is $\lfloor xa/b \rfloor$, and the total is $\sum_{x=1}^{b-1} \lfloor xa/b \rfloor$, a sum you can evaluate by pairing x with $b - x$: the two arguments add to exactly a, neither is an integer (because $b \nmid xa$ for $0 < x < b$), so the two floors add to $a - 1$, and the $b - 1$ terms total $\frac{1}{2}(a-1)(b-1)$. That is the same symmetry as $m \leftrightarrow N - m$, seen from the side. The pairing I used is the cleanest form of it: reflecting the standard form $(x, y) \mapsto (b - 1 - x, -1 - y)$ is the point reflection of the lattice through the centre of the "staircase" of missing values. The number $N = ab - a - b$ is not pulled from a hat — it is the value where that reflection sends $m = 0$, i.e. the largest gap, $(b-1)a - b$, the last member of the last residue class to start.

So the chain of "of course" is: coprime ⇒ one representative per residue ⇒ gaps are counted class by class ⇒ the classes pair up ⇒ half of $0..N$ is missing.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "thirty-five non-attainable scores". Write down: the number of non-representable values of $xa + yb$ is $\frac{1}{2}(a-1)(b-1)$ when $\gcd(a,b) = 1$, so $(a-1)(b-1) = 70$.

2. The cue suggests the Sylvester count, proved by residues mod b and the pairing $m \leftrightarrow ab - a - b - m$, followed by a divisor table.

3. Hesitation occurs at "must a, b be coprime?" Reinforce: a common factor $d > 1$ makes every non-multiple of d unattainable, infinitely many. The second hesitation is whether 0 counts as a score; it does not matter, since 0 is attainable (zero plays) and is not among the missing ones either way.

4. The trap is $(a, b) = (71, 2)$, which satisfies the count and fails the 58 condition; the second trap is forgetting coprimality and keeping $(36, 3)$ or $(15, 6)$.

5. When a problem gives you the *number* of objects with a property and asks for parameters, the intended route is a closed-form count in those parameters. Stop generating examples and derive the count.

6. To count what is missing from a set with a symmetry, find an involution that swaps "in" and "out"; then the count is half of the range the involution acts on. For representations by coprime a, b, work in standard form with $0 \le x \le b - 1$ and reflect.

▷ **Cue:** a count of non-attainable totals from two step sizes a, b ⟶ **Move:** use $\frac{1}{2}(a-1)(b-1)$ (coprime case) and factor

▷ **Cue:** "infinitely many" versus "finitely many" gaps in $xa + yb$ ⟶ **Move:** finitely many forces $\gcd(a, b) = 1$

▷ **Cue:** deciding whether $m = xa + yb$ with $x, y \ge 0$ ⟶ **Move:** reduce to standard form $0 \le x \le b - 1$; attainable iff $y \ge 0$

▷ **Cue:** counting elements of $\{0, \dots, N\}$ that fail a property ⟶ **Move:** look for an involution $m \mapsto N - m$ swapping success and failure

▷ **Cue:** two clues, one numeric and one example ⟶ **Move:** turn the numeric clue into an equation giving a short list, then filter with the example

1971 A6

1971 A6. Let c be a real number such that n^c is an integer for every positive integer n. Show that c is a non-negative integer.

Caveat. The MAA OCR prints the exponent flat as "nc"; the statement is restored as n^c. Identity check: the book's solution under 1971 A-6 applies the mean value theorem to x^c and then a k-th difference version to $f(x) = x^c$, concluding $c = k - 1$; this matches the restored statement. The generalised mean value theorem the book quotes without proof is proved below.

Fields: analysis (finite differences, mean value theorem), number theory flavour **Position:** A6 — short, but only once you know that finite differences of integers are integers and that high derivatives of x^c decay; otherwise it looks like number theory and stalls **Answer:** Only $c \in \{0, 1, 2, \dots\}$; the proof shows the k-th difference of n^c, $k = \lfloor c \rfloor + 1$, is an integer tending to 0, hence eventually 0, which forces c to be an integer.

Layer 1 THE WHAT — surface analysis

The natural attempt. You treat it as number theory. First, 2^c is an integer; if $c < 0$ then $0 < 2^c < 1$, impossible. So $c \ge 0$. Good.

Next you try to show c is rational. $2^c = p$ and $3^c = q$ are integers, so $c = \log p/\log 2 = \log q/\log 3$. Then $p^{\log 3} = q^{\log 2}$... you start trying unique factorisation: if $c = r/s$ in lowest terms then $2^r = p^s$, so $2^{r/s}$ being an integer forces $s \mid r$, so c is an integer. That argument is correct *if* c is rational. But nothing so far shows c is rational, and showing that $2^c, 3^c \in \mathbb{Z}$ forces $c \in \mathbb{Q}$ is a transcendence question (it is essentially the six exponentials theorem territory). You have hit a real wall, not a gap in effort: the arithmetic of individual values n^c is too weak. The hypothesis you have not used is "for *every* n" — infinitely many values, all at once.

The model chain. Let $f(x) = x^c$ on $x > 0$.

$c \geq 0$. 2^c is an integer; if $c < 0$ then $0 < 2^c < 1$. So $c \geq 0$.

Finite differences. For a function g define $\Delta g(x) = g(x+1) - g(x)$ and $\Delta^k = \Delta(\Delta^{k-1})$. By induction on k,

$$\Delta^k f(u) = \sum_{i=0}^{k} (-1)^{k-i} \binom{k}{i} f(u+i),$$

so for a positive integer u, $\Delta^k f(u)$ is an integer combination of the integers $u^c, (u+1)^c, \ldots, (u+k)^c$: it is an integer.

Integral formula. If g is C^1 then $\Delta g(x) = \int_0^1 g'(x+t)\,dt$. Applying this k times (note Δ commutes with differentiation and with integration in a shift parameter),

$$\Delta^k f(u) = \int_0^1 \cdots \int_0^1 f^{(k)}(u + t_1 + \cdots + t_k)\,dt_1 \cdots dt_k.$$

(This is the book's "extension of the mean value theorem": since $f^{(k)}$ is continuous on $[u, u+k]$, the average equals $f^{(k)}(\xi)$ for some $\xi \in (u, u+k)$.)

Choose k. Let $k = \lfloor c \rfloor + 1$, so $k - 1 \leq c < k$ and $k \geq 1$. Then

$$f^{(k)}(x) = P\,x^{c-k}, \qquad P = c(c-1)\cdots(c-k+1),$$

with $c - k < 0$. On $[u, u+k]$ we have $|f^{(k)}(x)| \leq |P|\,u^{c-k}$, so

$$|\Delta^k f(u)| \leq |P|\,u^{c-k} \longrightarrow 0 \qquad (u \to \infty).$$

An integer of absolute value less than 1 is 0, so $\Delta^k f(u) = 0$ for all large integers u.

Conclude. If $P \neq 0$ then $f^{(k)}(x) = Px^{c-k}$ has constant nonzero sign on $(0, \infty)$, so the integral above is nonzero for every u — contradiction. Hence $P = c(c-1)\cdots(c-k+1) = 0$, so $c \in \{0, 1, \ldots, k-1\}$; with $k - 1 \leq c$, $c = k - 1$ is a non-negative integer.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains get $c \geq 0$ from $n = 2$. At the next step the natural attempt looks at single values $(2^c, 3^c)$ and asks what their arithmetic says about c; the model looks at a fixed integer combination of *consecutive* values, $\Delta^k f(u)$, which is still an integer but is controlled by calculus.

> *Individual values tell me nothing I can prove. What do I get for free from "integer at every n"? Sums and differences of those values are integers too. And differences of x^c are controlled by derivatives.*

> *An integer that I can show is small must be zero. So I want an integer combination of the n^c that tends to 0 — and x^c gets flatter under repeated differentiation once the exponent goes negative.*

> *How many differences? Enough that $c - k < 0$: take $k = \lfloor c \rfloor + 1$.*

Layer 3 THE WHY

Why is the finite difference the natural object, not a trick? Start from what the hypothesis actually gives. "$n^c \in \mathbb{Z}$ for every n" is a statement about a function on the integers taking values in a discrete set. The only operations guaranteed to preserve integrality are integer linear combinations. So the candidate tools are integer combinations of values of f, and the simplest family of those, built only from consecutive values, is the difference operator Δ and its powers. That is not an invention; it is the complete list of cheap things you can do with integer-valued sequences.

Next: why should differences be small? Because of the defining property of the derivative. $\Delta f(u) = f(u+1) - f(u)$ is the discrete derivative, and by the mean value theorem it equals $f'(\xi)$. For $0 < c < 1$ this already ends the problem: $f'(x) = cx^{c-1} \to 0$, so $(u+1)^c - u^c$ is a positive integer smaller than 1 for large u, which is absurd. That special case (the book does it first) is the whole idea in miniature. For $1 < c < 2$, f' grows, but $f'' = c(c-1)x^{c-2} \to 0$; so use the second difference. In general, each difference lowers the exponent by one, and after $\lfloor c \rfloor + 1$ steps it is negative. The choice of k is forced by the arithmetic of exponents, not guessed.

Then: why does the argument not kill integer c? It must not, since n^2 is an integer. For $c = 2$, $k = 3$, and $f''' \equiv 0$: the third difference of n^2 is identically 0 — the integer 0 that the argument forces is actually attained. The product $P = c(c-1)\cdots(c-k+1)$ is exactly the switch between the two cases: it vanishes precisely at $c = 0, 1, \ldots, k-1$. So the proof is not "show a contradiction", it is "show $P = 0$", and $P = 0$ *is* the conclusion.

Where does the book's generalised mean value theorem come from? From iterating $\Delta g(x) = \int_0^1 g'(x+t)\,dt$, which is the fundamental theorem of calculus. The k-fold integral is an average of $f^{(k)}$ over a box of shifts, and an average of a continuous function lies between its minimum and maximum. That is all the book means by "there is a ξ with $f^{(k)}(\xi) = \Delta^k f(a)$"; the statement is quoted there without proof, and the integral form is the proof.

The reusable content: the arithmetic of a function is often reachable only through its analysis. "Integer" is a discreteness condition; "tends to 0" is an analytic one; together they force exact vanishing. The same mechanism drives the standard proof that e is irrational (if $e = p/q$ then $q!\,(e - \sum_{j \le q} 1/j!)$ is a positive integer smaller than 1) and Liouville's construction of transcendental numbers: build an integer, bound it analytically, conclude it is 0.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "integer for *every* positive integer n". Write down: every integer combination $\sum_i \lambda_i f(u+i)$, in particular $\Delta^k f(u)$, is an integer.

2. The cue suggests finite differences, the mean value theorem (or integral form of Δ), and the principle "an integer of absolute value < 1 is 0".

3. Hesitation occurs at the choice of k and at the last step: after showing $\Delta^k f(u) = 0$ eventually, how does that give $c \in \mathbb{Z}$? Reinforce by writing $f^{(k)} = Px^{c-k}$: a nonzero multiple of a positive function cannot average to 0, so $P = 0$.

4. The trap is number theory: trying to show c rational from $2^c, 3^c \in \mathbb{Z}$. That is a transcendence problem and not provable with competition tools. A smaller trap is forgetting the case $c < 0$, where x^c decays but the difference argument would still need care; $n = 2$ disposes of it at once.

5. When a real-analytic function takes integer values at all integers, reach for differences and derivatives at once, not for prime factorisations. The sign is "every n" rather than "some n".

6. Integer + small $\Rightarrow$ zero. Build an integer-valued quantity from the hypothesis (a finite difference, a determinant, a norm), bound it by analysis, and conclude it vanishes; then read off what vanishing means.

▷ **Cue:** a real function takes integer values at every positive integer $\longrightarrow$ **Move:** its finite differences $\Delta^k f(u)$ are integers

▷ **Cue:** an integer-valued quantity that can be made smaller than 1 $\longrightarrow$ **Move:** it is 0

- ▷ **Cue:** $f(x) = x^c$ with non-integer c $\longrightarrow$ **Move:** differentiate $\lfloor c \rfloor + 1$ times to get a negative power, which tends to 0
- ▷ **Cue:** need to bound $\Delta^k f$ $\longrightarrow$ **Move:** $\Delta^k f(u) = \int_{[0,1]^k} f^{(k)}(u + t_1 + \cdots + t_k)\, dt$, the mean value theorem for differences
- ▷ **Cue:** temptation to prove an exponent is rational from a couple of values $\longrightarrow$ **Move:** stop: use infinitely many values through analysis instead

Putnam 1971 — Paper B

1971 B1

1971 B1. Let S be a set and let $\circ$ be a binary operation on S satisfying the two laws $x \circ x = x$ for all x in S, and $(x \circ y) \circ z = (y \circ z) \circ x$ for all x, y, z in S. Show that $\circ$ is associative and commutative.

Fields: abstract algebra, binary operations, identities in magmas **Position:** B1 (opener) **Answer:** No numerical answer; the operation is proved commutative and associative directly from the two given laws.

Layer 1 THE WHAT — surface analysis

The natural attempt. The problem asks for associativity and commutativity, and the axiom $(x \circ y) \circ z = (y \circ z) \circ x$ visibly rearranges a triple product, so the natural first move is to attack associativity head-on: take $(x \circ y) \circ z$ and try to manipulate it into $x \circ (y \circ z)$ using the axiom directly. Applying the axiom to $(x \circ y) \circ z$ just gives $(y \circ z) \circ x$ — a cyclic shift, not the target shape. Trying the axiom again on $(y \circ z) \circ x$ gives $(z \circ x) \circ y$, and again gives back $(x \circ y) \circ z$. Three applications return to the start: the axiom alone generates the cyclic group of order 3 on $\{x, y, z\}$ and nothing else. At this point a student has a true fact — $(x \circ y) \circ z = (y \circ z) \circ x = (z \circ x) \circ y$ — but no route from any of these three expressions to $x \circ (y \circ z)$, because that expression's outermost grouping is different in shape from all three cyclic forms. The idempotent law $x \circ x = x$ sits unused off to the side, feeling like a throwaway fact for closing edge cases, not like the engine of the proof. Time gets spent substituting concrete relabelings of x, y, z into the axiom, hoping something eventually collides with $x \circ (y \circ z)$, without a principle for which substitution to try next.

The model chain. Do not aim at associativity first, even though it is asked first in some phrasings of this style of problem — aim at commutativity, and manufacture a situation where the idempotent law can fire. Idempotence only ever applies to a *repeated* argument, $a \circ a$, so the productive move is to feed a repeated argument into the cyclic axiom. Take $a = x \circ y$ and observe

$$x \circ y = (x \circ y) \circ (x \circ y)$$

by idempotence applied to $a = x \circ y$. Now expand the right-hand side using the axiom with $(A, B, C) = (x \circ y,\ x,\ y)$, since $(A \circ B) \circ C = (x \circ y) \circ (x \circ y)$ when $B \circ C = x \circ y$:

$$(x \circ y) \circ (x \circ y) = \big[(x \circ y) \circ x\big] \circ y,$$

so $x \circ y = \big[(x \circ y) \circ x\big] \circ y$. Apply the axiom once more, this time to the inner product $(x \circ y) \circ x$, with $(A, B, C) = (x, y, x)$:

$$(x \circ y) \circ x = (y \circ x) \circ x.$$

Substituting back, $x \circ y = \big[(y \circ x) \circ x\big] \circ y$. Apply the axiom a third time to the inner product $(y \circ x) \circ x$, with $(A, B, C) = (y, x, x)$:

$$(y \circ x) \circ x = (x \circ x) \circ y = x \circ y$$

using idempotence $x \circ x = x$ in the last step. Substituting back once more,

$$x \circ y = (x \circ y) \circ y.$$

Apply the axiom a fourth and final time, to $(x \circ y) \circ y$, with $(A, B, C) = (x, y, y)$:

$$(x \circ y) \circ y = (y \circ y) \circ x = y \circ x$$

using idempotence again. Chaining the last two displays,

$$x \circ y = (x \circ y) \circ y = y \circ x.$$

Commutativity is proved. Associativity now falls out of the original axiom for free: since $\circ$ is commutative, $(y \circ z) \circ x = x \circ (y \circ z)$, so

$$(x \circ y) \circ z = (y \circ z) \circ x = x \circ (y \circ z).$$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains have access to the same two facts and both notice, after one or two substitutions, that the axiom applied directly to $(x \circ y) \circ z$ only cycles x, y, z and never reaches $x \circ (y \circ z)$. They split at that point of stall: the natural attempt keeps trying fresh relabelings of the three original variables x, y, z, treating the axiom as the only tool in play. The model abandons the three free variables and instead substitutes the *compound* expression $x \circ y$ itself into the axiom's slots, specifically so that idempotence — the one law that can shrink an expression rather than merely rearrange it — is forced to become applicable.

> *The axiom only ever cycles a product of three things; it can never shrink one. Idempotence is the only law here that shrinks something, from $a \circ a$ down to a, but it needs a repeated argument to fire — so instead of feeding x, y, z into the axiom, I should feed $x \circ y$ itself in, twice, and let the axiom manufacture the repetition idempotence needs.*

Layer 3 THE WHY

The axiom $(x \circ y) \circ z = (y \circ z) \circ x$ is, on its own, a purely combinatorial statement: for any three elements of S under any labelling, the two ways of forming this particular triple product related by a cyclic shift agree. Read at face value it never reduces the number of distinct symbols in play — every application replaces one triple product with another triple product of the same three underlying elements, permuted. A proof strategy built only on this axiom is, structurally, a search inside the cyclic group of order 3 acting on (x, y, z), and $x \circ (y \circ z)$ (the target shape for associativity) is not a member of that orbit when x, y, z are kept as three independent free variables — there is no cyclic relabelling of x, y, z that turns $(x \circ y) \circ z$ into $x \circ (y \circ z)$ as an uninterpreted string, because the latter groups y and z together first, while the axiom only ever regroups by moving one whole factor to the far end. This is why chasing associativity first, using only the axiom, is a search with no terminating condition: the tool being used cannot change the "shape" of the expression, only its labelling.

Idempotence is different in kind: $x \circ x = x$ is a collapsing law, the only rule in the system that reduces two occurrences of an object to one. But it is inert unless its hypothesis — a doubled argument — is already present. The generative move is therefore to treat $x \circ y$ not as an answer being computed but as a raw element of S in its own right, call it a, and observe that $a \circ a = a$ is available for free, for any $a \in S$, including $a = x \circ y$. Writing $a \circ a$ as $(x \circ y) \circ (x \circ y)$ and then expanding via the axiom is what converts a static fact ($a \circ a = a$) into a chain of equalities involving x and y separately. Every subsequent application of the axiom in the model chain exists solely to peel that expanded form back down, using idempotence again at each of the

two points where a genuine repetition ($x \circ x$, then $y \circ y$) has been engineered to appear. The lesson generalises beyond this problem: when an axiom set mixes a shape-preserving relational law with a shape-reducing law, the relational law's only productive use is to manufacture the reducing law's precondition, not to be iterated on the original variables in the hope of a direct hit.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the axiom system contains two laws of different character: one relational law that only rearranges a triple product ($(x \circ y) \circ z = (y \circ z) \circ x$), and one collapsing law that shrinks a doubled argument ($x \circ x = x$). On seeing both, write down: the collapsing law needs a repeated argument to fire, so the goal is to manufacture that repetition using the relational law.

2. That cue suggests substituting a compound expression, here $x \circ y$, into the relational axiom's slots in place of a bare variable, specifically so that $(x \circ y) \circ (x \circ y)$ appears and idempotence becomes applicable to it.

3. Hesitation sets in right after the axiom is applied once or twice to the bare variables x, y, z and only produces cyclic permutations of the same triple product. Reinforce against it by asking explicitly: does this application shrink the expression or only relabel it? A relabelling-only move cannot, by itself, close a proof that needs a genuinely new relationship between $x \circ y$ and $y \circ x$.

4. The trap is that the axiom's visual form, $(x \circ y) \circ z = (y \circ z) \circ x$, resembles a statement about associativity and invites a direct assault on associativity first, even though the problem's own harder content — commutativity — is what must be established before associativity becomes a one-line consequence of the axiom.

5. Recognise this faster next time by treating idempotent-type laws ($x \circ x = x$, or more generally absorbing/fixed-point laws) as a standing invitation: whenever one appears alongside a relational axiom, ask what compound expression, substituted into the relational axiom, would create the doubled argument the idempotent law needs.

6. Decision rule: to prove structural properties (commutativity, associativity) of an operation from a small axiom set that includes an idempotent-type law, do not iterate the relational axiom on the original free variables alone; substitute a compound expression built from those variables into the relational axiom so that the idempotent law is forced to apply, and only then read off the resulting identity between the variables.

1971 B2

1971 B2. Let $F(x)$ be a real valued function defined for all real x except for $x = 0$ and $x = 1$ and satisfying the functional equation $F(x) + F\left(\dfrac{x-1}{x}\right) = 1 + x$. Find all functions $F(x)$ satisfying these conditions.

Fields: functional equations, Möbius transformations, finite-order substitution **Position:** B2 **Answer:** $F(x) = \dfrac{x^3 - x^2 - 1}{2x(x-1)}$, and this is the unique solution.

Layer 1 THE WHAT — surface analysis

The natural attempt. One equation is written down and it contains two unknown values, $F(x)$ and $F\left(\frac{x-1}{x}\right)$, so the honest first instinct is to try to generate a second equation by substituting $\frac{x-1}{x}$ in place of x everywhere in the functional equation. That produces

$$F\left(\frac{x-1}{x}\right) + F\left(\frac{\frac{x-1}{x} - 1}{\frac{x-1}{x}}\right) = 1 + \frac{x-1}{x}.$$

The new inner expression simplifies (correctly, with some fraction-within-fraction care) to $\frac{-1}{x-1} = \frac{1}{1-x}$, so this second equation reads $F\left(\frac{x-1}{x}\right) + F\left(\frac{1}{1-x}\right) = 1 + \frac{x-1}{x}$. Progress feels real: two equations now exist. But counting unknowns honestly, there are three distinct function values in play — $F(x)$, $F\left(\frac{x-1}{x}\right)$, and this new $F\left(\frac{1}{1-x}\right)$ — and only two equations. The natural next move is to substitute $\frac{1}{1-x}$ in turn, expecting a fourth unknown value to appear and the chain to run on forever, each substitution introducing one more unknown function value than it resolves. Some solvers stop here, convinced the system is underdetermined without extra input (such as guessing F is a rational function of a particular degree and matching coefficients), which does eventually work but commits to a large, unmotivated ansatz rather than a proof that the equation pins down F uniquely.

The model chain. Before generating a third equation, check whether the chain of substitutions actually closes. Define $\varphi(x) = \frac{x-1}{x}$ and compute its own iterates directly, in closed form, rather than assuming they keep producing new expressions:

$$\varphi(x) = \frac{x-1}{x}, \qquad \varphi(\varphi(x)) = \frac{1}{1-x}, \qquad \varphi(\varphi(\varphi(x))) = x.$$

The third iterate returns to x: φ has order exactly 3 under composition. So there are only three distinct expressions in the whole substitution chain, not infinitely many: write $x_1 = x$, $x_2 = \varphi(x) = \frac{x-1}{x}$, $x_3 = \varphi(x_2) = \frac{1}{1-x}$, with $\varphi(x_3) = x_1$ closing the loop. The functional equation, applied at x_1, x_2, and x_3 in turn, gives exactly three equations in exactly three unknowns:

$$F(x_1) + F(x_2) = 1 + x_1, \qquad F(x_2) + F(x_3) = 1 + x_2, \qquad F(x_3) + F(x_1) = 1 + x_3.$$

This is a closed linear system. Take the first equation, minus the second, plus the third; on the left every $F(x_2)$ and $F(x_3)$ cancels in pairs, leaving $2F(x_1)$:

$$2F(x_1) = 1 + x_1 - x_2 + x_3.$$

Substituting $x_2 = \frac{x-1}{x}$ and $x_3 = \frac{1}{1-x}$ and combining over the common denominator $x(x-1)$ gives, after expansion,

$$x_1 - x_2 + x_3 = \frac{x^3 - 2x^2 + x - 1}{x^2 - x},$$

so

$$F(x) = \frac{1}{2}\left(1 + \frac{x^3 - 2x^2 + x - 1}{x^2 - x}\right) = \frac{x^3 - x^2 - 1}{2x(x-1)}.$$

Sufficiency does not need to be checked by raw substitution and simplification. The formula for F was derived from the same alternating-sum pattern at every point of the cycle, so it obeys $2F(x_1) = 1 + x_1 - x_2 + x_3$ and, by the identical argument applied starting from x_2, also $2F(x_2) = 1 + x_2 - x_3 + x_1$. Adding these two displays,

$$F(x_1) + F(x_2) = \frac{(1 + x_1 - x_2 + x_3) + (1 + x_2 - x_3 + x_1)}{2} = \frac{2 + 2x_1}{2} = 1 + x_1,$$

which is exactly the original functional equation at $x = x_1$. So F as derived automatically satisfies the equation, at every point where x_1, x_2, x_3 are all defined — which the domain restriction $x \neq 0, 1$ guarantees, since one checks $x_2, x_3 \notin \{0, 1\}$ whenever $x_1 \notin \{0, 1\}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same second equation, involving $F(x_1)$, $F(x_2)$, and a new expression $F(x_3)$ with $x_3 = \frac{1}{1-x}$. They split immediately after that: the natural attempt treats x_3 as one more unresolved substitution and prepares to substitute it in turn, expecting the chain to keep growing. The model instead stops substituting and computes $\varphi(x_3)$ in closed form before doing anything else, checking whether it equals x_1.

> *I keep getting a new function value every time I substitute — but the substitution itself is a fixed rational map, $\varphi(x) = \frac{x-1}{x}$, and a fixed map applied repeatedly to a real number either wanders forever or returns to where it started after finitely many steps. Before assuming this chain is infinite, I should just compute $\varphi(\varphi(\varphi(x)))$ directly and see which one it is.*

Layer 3 THE WHY

The functional equation has the shape $F(x) + F(\varphi(x)) = g(x)$ for a fixed map φ and fixed g. Such an equation pins down F uniquely, by the alternating-sum method used above, precisely when the orbit of x under φ is finite — that is, when φ has finite order as a map. If φ had infinite order (its iterates $x, \varphi(x), \varphi^2(x), \ldots$ all distinct forever), the same substitution strategy would generate an infinite triangular system with no way to close it into a fixed number of equations in a fixed number of unknowns, and F would not be determined by this single functional equation without extra assumptions (continuity, boundedness, or the like). The domain restriction in the problem, $x \neq 0, 1$, is not an arbitrary technicality to be checked at the end; it is exactly the condition that keeps every iterate of φ defined, because $\varphi(x) = \frac{x-1}{x}$ is undefined at $x = 0$ and $\varphi(x) = 1$ has no solution (so 1 is never hit, but φ itself needs $x \neq 0$ to be defined, and the problem also excludes 1 so that F is total on the same set the equation needs). The fact that $\varphi^3 = \mathrm{id}$ is a computation, not a coincidence to be marvelled at: φ is a Möbius transformation, and Möbius transformations of the real (or complex) line correspond to 2×2 matrices up to scalar, composition corresponds to matrix multiplication, and a matrix of finite order forces the corresponding map to be of the same finite order. Recognising $\varphi(x) = \frac{x-1}{x} = 1 - \frac{1}{x}$ as a Möbius map, rather than as an ad hoc formula, is what should trigger the instinct to check its order before trusting that the substitution chain runs forever. Once the order is known to be 3, the alternating sum $E_1 - E_2 + E_3$ is the correct combination for exactly the reason that a cyclic system of three equations, each summing two of three unknowns, has that combination as its unique way to isolate one unknown with coefficient 2 and cancel the other two — the same linear-algebra fact that makes $2F(x_1) = 1 + x_1 - x_2 + x_3$ solvable at all.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the functional equation relates F at x to F at a specific rational substitution, $\frac{x-1}{x}$. On seeing it, write down: identify this substitution as a Möbius map and check whether it has finite order under composition before assuming the substitution chain is infinite.
2. That cue suggests computing the map's iterates in closed form — here $\varphi(x)$, $\varphi(\varphi(x))$, $\varphi(\varphi(\varphi(x)))$ — to find the orbit size, then writing down one copy of the functional equation at each point of the orbit.
3. Hesitation sets in right after the second equation is derived and a third, apparently new, function value appears. Reinforce against it by refusing to substitute a fourth time until the third iterate has actually been computed and compared with x.
4. The trap is assuming that "one new unknown per substitution" continues indefinitely just because it held for the first two substitutions; a finite-order map produces exactly that pattern for a bounded number of steps and then closes the loop, and mistaking a closing system for an open one leads to abandoning a solvable problem as underdetermined.
5. Recognise this faster next time by treating any substitution of the form $x \mapsto \frac{ax+b}{cx+d}$ inside a functional equation as a Möbius transformation on sight, and reflexively computing its first few iterates to find its order before setting up any system of equations.
6. Decision rule: for $F(x) + F(\varphi(x)) = g(x)$ with φ a fixed map, find the order n of φ, write the n cyclically-substituted equations, and solve the resulting $n \times n$ linear system; for odd n the alternating sum (add, subtract, add, ...) around the cycle isolates $2F(x)$ directly.

1971 B3

1971 B3. Two cars travel around a track at equal and constant speeds, each completing a lap every hour. From a common starting point, the first starts at time $t = 0$ and the second at an arbitrary later time $t = T > 0$. Prove that there is a total period of exactly one hour during the motion in which the first has completed twice as many laps as the second.

Fields: real analysis, floor function, interval decomposition **Position:** B3 **Answer:** The total measure of the set of times at which car 1's completed-lap count is twice car 2's is exactly 1 hour, realised as two sub-intervals of lengths $1 - \delta$ and δ, where δ is the fractional part of T.

Layer 1 THE WHAT — surface analysis

The natural attempt. Car 1, moving from $t = 0$, has completed $\lfloor t \rfloor$ laps by time t. Car 2, moving from $t = T$, has completed $\lfloor t - T \rfloor$ laps by time $t \geq T$. The condition to satisfy is $\lfloor t \rfloor = 2\lfloor t - T \rfloor$, and the natural instinct is to treat this as an equation to solve for t directly, perhaps by writing $t = 2(t - T) +$ (fractional correction) and trying to isolate t, or by sketching the graphs of $\lfloor t \rfloor$ and $2\lfloor t - T \rfloor$ as staircases and hunting visually for where they coincide. Both staircases jump at integers, but the two families of jump points — the integers themselves for car 1, and T plus the integers for car 2 — are offset by T, which is an arbitrary real number, not a nice one. A student who tries to track the coincidence globally, lap by lap across the whole time axis, ends up doing an unbounded case analysis: for each integer k, is there a sub-interval where the condition holds, and if the pattern looks like it might repeat every couple of hours, does that mean infinitely many hour-long qualifying intervals exist, contradicting the problem's claim of "a total period of exactly one hour"? Without a way to organise the offset T against the unit grid, the analysis sprawls: separate cases for every value of $\lfloor T \rfloor$, unclear whether the total length one hour is a coincidence of a particular T or a fact true for every T, and no visible reason the qualifying set should be finite in total measure rather than periodic and infinite.

The model chain. Do not track the two staircases against the whole time axis at once. Fix attention on a single, generic unit interval where car 1's lap count $\lfloor t \rfloor$ is constant, and express T relative to that same unit grid. Write $T = k + \delta$ with $k = \lfloor T \rfloor$ an integer and $\delta \in [0, 1)$ the fractional part. For t in the interval $[m, m + 1)$ with integer $m \geq k$, write $t = m + e$ with $e \in [0, 1)$, so $\lfloor t \rfloor = m$ and

$$t - T = (m - k) + (e - \delta).$$

This single quantity $e - \delta$ decides everything: if $e \geq \delta$ then $\lfloor t - T \rfloor = m - k$, and if $e < \delta$ then $\lfloor t - T \rfloor = m - k - 1$. The condition $\lfloor t \rfloor = 2\lfloor t - T \rfloor$ splits into exactly two linear cases:

$$e \geq \delta: \quad m = 2(m - k) \implies m = 2k, \qquad e < \delta: \quad m = 2(m - k - 1) \implies m = 2k + 2.$$

Each case pins down a *unique* value of m — not a family recurring over every unit interval — because the equation $m = 2(m - k)$ has exactly one integer solution, $m = 2k$, and likewise $m = 2(m - k - 1)$ has exactly one solution, $m = 2k + 2$. So the qualifying set of times is confined to two specific unit intervals, not scattered periodically across infinitely many:

$$m = 2k: \ e \in [\delta, 1) \implies t \in [2k + \delta, 2k + 1), \quad \text{length } 1 - \delta,$$

$$m = 2k + 2: \ e \in [0, \delta) \implies t \in [2k + 2, 2k + 2 + \delta), \quad \text{length } \delta.$$

Both intervals lie in $t \geq T$ (since $2k + \delta \geq k + \delta = T$ as $k \geq 0$, and similarly for the second), so both are legitimate parts of the motion after car 2 has started. Their total length is

$$(1 - \delta) + \delta = 1,$$

exactly one hour, for every value of T, with no case analysis on the size of $\lfloor T \rfloor$ required beyond fixing the single symbol k.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree the condition to analyse is $\lfloor t \rfloor = 2\lfloor t - T \rfloor$, and both recognise that the two floor functions jump at different, T-shifted sets of points. They split at the decision of what to hold fixed while comparing them: the natural attempt tries to compare the two staircases across the whole time axis at once, generating one case per integer $k = \lfloor T \rfloor$ or per lap of car 1, with no natural stopping point. The model instead fixes a single generic unit interval $[m, m+1)$ for t, decomposes T into $k + \delta$ against that same unit grid, and asks only where inside that one interval the local position e crosses the single threshold δ.

> *Both floor functions only change value at integers, so their difference in behaviour can only ever be decided by comparing fractional parts against a single threshold — not by tracking the full history of jumps. If I fix t inside one unit interval $[m, m+1)$, then $\lfloor t \rfloor = m$ is constant there, and the only thing that can vary is whether $t - T$ has crossed its own next integer yet; that crossing happens exactly when t's local position e passes T's fractional part δ. So I should split the unit interval at $e = \delta$ and handle the two pieces separately, instead of trying to track both staircases together.*

Layer 3 THE WHY

The floor function $\lfloor x \rfloor$ is piecewise constant, with its only discontinuities at the integers; consequently, comparing $\lfloor t \rfloor$ with $\lfloor t - T \rfloor$ for a fixed shift T is entirely a question of where the two integer grids — the integers themselves, and the integers shifted by T — interleave. Shifting the argument of a floor function by T is the same, locally, as shifting the integer grid by $-T \pmod 1$, i.e. by $-\delta$; within any single unit interval of t, this shift is felt as a single crossing point at $e = \delta$, and nowhere else, because within one interval there is only one place a new integer boundary (belonging to the shifted grid) can fall. This is the structural fact that makes fixing a generic unit interval, rather than the whole time axis, the right unit of analysis: it isolates the problem to its essential local content, a single inequality $e \gtrless \delta$, and turns an apparently unbounded case analysis over all integers k into a symbol, k, that never needs to be split into cases at all — it cancels out of both solved equations $m = 2(m - k)$ and $m = 2(m - k - 1)$, leaving $m = 2k$ and $m = 2k + 2$ as the only two admissible interval indices regardless of what k actually is.

The finiteness of the total qualifying period is not an extra fact to be checked afterward; it follows from the same structural observation. Because $m = 2(m - k)$ is a genuinely linear equation in the single unknown m with k fixed, it has exactly one solution, not one solution per period — so the phrase "a total period of exactly one hour" in the problem statement is a promise that the analysis, once localised correctly, produces a finite, computable answer rather than an infinite recurring one. Recognising that $\lfloor t \rfloor = 2\lfloor t - T \rfloor$ is linear in m once e is fixed relative to δ is precisely what guarantees the qualifying set cannot recur: any attempt to search "lap by lap" without this localisation risks mistaking the two genuinely distinct, non-recurring solutions ($m = 2k$ and $m = 2k + 2$) for the first two terms of an infinite periodic family.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is an equation comparing $\lfloor t \rfloor$ with $\lfloor t - T \rfloor$ for a fixed real shift T. On seeing it, write down: floor functions are piecewise constant and only compare meaningfully via the fractional part of the shift, so decompose $T = k + \delta$ before doing anything else.

2. That cue suggests fixing a single generic unit interval $[m, m+1)$ for the variable, writing the local position as $t = m + e$, and comparing e against δ rather than tracking the global behaviour of both floor functions at once.

3. Hesitation sets in right when the offset T is arbitrary and it is tempting to case-split on the size of $\lfloor T \rfloor$ or to track which "lap number" of car 1 corresponds to which of car 2. Reinforce against it by fixing m as a free integer symbol and solving for it algebraically, letting $k = \lfloor T \rfloor$ remain a symbol rather than a case.

4. The trap is assuming the qualifying condition recurs periodically — once per some fixed cycle length — and therefore that the total qualifying time must be summed over infinitely many equal-length intervals, when in fact the equation $m = 2(m - k)$ pins down a single value of m and the analogous equation for the other branch pins down exactly one more, giving a finite, two-piece answer.
5. Recognise this faster next time by treating any comparison of $\lfloor f(t) \rfloor$ against $\lfloor g(t) \rfloor$, where f and g differ by a constant shift, as a signal to decompose that constant into integer and fractional parts and localise the analysis to one unit interval before attempting anything global.
6. Decision rule: to solve conditions of the form $\lfloor t \rfloor = c \cdot \lfloor t - T \rfloor$ (or similar floor-function comparisons under a constant shift), write $T = k + \delta$, fix $t = m + e$ within a generic unit interval, split at $e = \delta$, solve each resulting linear equation in m for its unique admissible integer solution, and sum the resulting sub-interval lengths — the symbol k should cancel rather than force a case split.

1971 B4

1971 B4. A "spherical ellipse" with foci A, B on a given sphere is defined as the set of all points P on the sphere such that $\overset{\frown}{PA} + \overset{\frown}{PB} =$ constant. Here $\overset{\frown}{PA}$ denotes the shortest distance on the sphere between P and A. Determine the entire class of real spherical ellipses which are circles.

Fields: spherical geometry; symmetry arguments; spherical trigonometry (spherical Pythagoras) **Position:** B4 — middle of the paper; the answer is guessable, the proof of the negative half is where the marks are **Answer:** Exactly the great circles, and only with constant sum π (on a unit sphere). A great circle Γ is the spherical ellipse $\overset{\frown}{PA} + \overset{\frown}{PB} = \pi$ precisely when B and the antipode $-A$ are mirror images in the plane of Γ (equivalently: A, B lie on one great circle perpendicular to Γ, on the same side of Γ and at equal distances from it). For any other constant a genuine (non-degenerate) spherical ellipse is never a circle.

Caveat. Degenerate cases are excluded, as in the official solution. Take the sphere to have radius 1, write $2a$ for the constant, and assume $0 < \overset{\frown}{AB} < \pi$ and $\overset{\frown}{AB} < 2a < 2\pi - \overset{\frown}{AB}$. If $A = B$ is allowed, every circle is a (degenerate) "spherical ellipse" (a circle of spherical radius a about A). If A, B are antipodal, the sum is π for every P and the locus is the whole sphere. At the endpoints $2a = \overset{\frown}{AB}$ or $2a = 2\pi - \overset{\frown}{AB}$ the locus is an arc. The book OCR drops the arc accents from the statement; they are restored here from the book's own solution, which uses $\overset{\frown}{PA}$ throughout.

Layer 1 THE WHAT — surface analysis

The natural attempt. The first instinct is to borrow from the plane: a planar ellipse is a circle only when the foci coincide. So you try to prove that a spherical ellipse with distinct foci is never a circle. You set up coordinates, $A = (\cos c, -\sin c, 0)$ and $B = (\cos c, \sin c, 0)$, so $\overset{\frown}{AB} = 2c$, and write the condition as

$$\arccos(P \cdot A) + \arccos(P \cdot B) = 2a.$$

Then you try to eliminate the arccosines: take cosines of both sides, isolate the square-root term, square. After two rounds of squaring you have a quartic-looking relation in x, y, z with $x^2 + y^2 + z^2 = 1$, and you have to decide whether it can be the intersection of the sphere with a plane. The algebra is heavy, and the squaring has introduced extraneous branches. There is a

worse problem: the planar analogy is false. Test the constant $2a = \pi$ with one numerical point off the equator and the sum comes out to exactly π again, on a whole great circle. The "never a circle" conjecture is wrong, and nobody who stays inside the algebra ever sees why.

The model chain. Place A, B on the equator, symmetric about the point $C = (1, 0, 0)$, as above. Let σ be reflection in the equatorial plane $z = 0$.

Step 1: reduce to $2a \le \pi$. Write $A' = -A$. For every P, $\overset{\frown}{PA} + \overset{\frown}{PA'} = \pi$, because A, P, $-A$ lie on a great circle and the two arcs make up half of it. Hence

$$\overset{\frown}{PA} + \overset{\frown}{PB} = 2a \iff \overset{\frown}{PA'} + \overset{\frown}{PB'} = 2\pi - 2a,$$

and $\overset{\frown}{A'B'} = \overset{\frown}{AB}$. So an ellipse with constant $2a > \pi$ is also an ellipse with constant $2\pi - 2a < \pi$.

Step 2: $2a < \pi$ *gives no circle.* Let E be the ellipse. On the equator, with P at longitude θ, the sum equals $2c$ for $|\theta| \le c$ and grows linearly to $2\pi - 2c$ at $\theta = \pi$. So E meets the equator in exactly two points, the vertices $V_{1,2}$ at longitude $\pm a$. Since σ fixes A and B and is an isometry, $\sigma(E) = E$.

Suppose E is a circle Γ, lying in a plane Π. Then $\sigma(\Pi) = \Pi$ (a circle determines its plane), so the normal of Π is vertical or horizontal. A horizontal plane $z = h$ is σ-invariant only for $h = 0$, which would make Γ the equator; but Γ meets the equator in only two points. So Π is vertical. It then contains V_1, V_2 (the only points of Γ on the equator) and the vertical direction. That forces $\Pi = \{x = \cos a\}$, and so

$$\Gamma = \{P : \ P \cdot C = \cos a\} = \{P : \ \overset{\frown}{PC} = a\}.$$

Take $M = (\cos a, 0, \sin a) \in \Gamma$. The spherical triangle MCA has a right angle at C (meridian meets equator), with legs $\overset{\frown}{MC} = a$ and $\overset{\frown}{CA} = c$. Spherical Pythagoras (or the dot product $M \cdot A = \cos a \cos c$) gives

$$\cos \overset{\frown}{MA} = \cos a \cos c < \cos a,$$

using $c > 0$ and $\cos a > 0$ (because $a < \pi/2$). So $\overset{\frown}{MA} > a$, likewise $\overset{\frown}{MB} > a$, and $\overset{\frown}{MA} + \overset{\frown}{MB} > 2a$. So $M \notin E$, a contradiction.

Step 3: $2a = \pi$ *gives a great circle.* Now

$$\overset{\frown}{PA} + \overset{\frown}{PB} = \pi \iff \overset{\frown}{PB} = \pi - \overset{\frown}{PA} = P(\overset{\frown}{-A}).$$

So E is the set of points equidistant from B and $-A$. That is the perpendicular-bisector great circle of B and $-A$ (the intersection of the sphere with the plane through O perpendicular to $B - (-A) = A + B$). Here it is the meridian $x = 0$ through $V_{1,2} = (0, \pm 1, 0)$.

Step 4: the whole class. Conversely, given any great circle Γ and any point B that is neither on Γ nor one of its two poles (the poles would give $A = B$, the excluded degenerate case), let B^* be the mirror image of B in the plane of Γ and put $A = -B^*$. Then Γ is the ellipse with foci A, B and constant π. The spherical ellipses that are circles are exactly the great circles, each arising in this way. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same definition. They split at the first sentence after it. The natural attempt writes $\arccos(P \cdot A) + \arccos(P \cdot B) = 2a$ and starts squaring. The model asks what is *different* about a sphere, and the answer is antipodes: $\overset{\frown}{PA} + P(\overset{\frown}{-A}) = \pi$ for every P. That single identity does three jobs: it reduces $2a > \pi$ to $2a < \pi$; it turns the case $2a = \pi$ into a perpendicular-bisector problem; and it tells you which constant to test.

> *On a sphere, distance to A and distance to $-A$ always add to π. So a sum of two distances equal to π is really a difference: $PB = P(-A)$. Equidistant from two points is a great circle. So the planar intuition "ellipse circle means equal foci" must fail on the sphere, and it fails at the constant π.*

> *For the other constants, I will not fight the equation. The ellipse is symmetric in the plane of A, B and O, so if it is a circle I know exactly which circle: the one with the two vertices as a diameter. Then I test one point, the top of that circle.*

- ▷ **Cue:** a locus on a sphere defined by distances to given points $\longrightarrow$ **Move:** bring in the antipodes: $d(P,X) + d(P,-X) = \pi$ converts sums to differences and large constants to small ones
- ▷ **Cue:** "is this locus a circle / a line?" for a locus with a visible symmetry $\longrightarrow$ **Move:** use the symmetry to pin down the only candidate circle, then test one well-chosen point on it
- ▷ **Cue:** a sum of distances equal to a quantity that is the distance to a reflected point $\longrightarrow$ **Move:** rewrite as "equidistant from two points" and recognise a perpendicular bisector

Layer 3 THE WHY

Why antipodes? Go back to what a spherical distance is. $\overset{\frown}{PA}$ is the angle $\angle POA$ at the centre. The angle $\angle POA$ and the angle $\angle PO(-A)$ are supplementary, because A, O, $-A$ are collinear. That is a fact with no planar analogue: in the plane there is no point whose distance from P is "π minus" the distance to A. Any argument that sees the difference between a sphere and a plane is likely to use it, and this problem is built around it. Once you have written $\overset{\frown}{PA} = \pi - P\overset{\frown}{(-A)}$ you hold a device that converts the sum in the definition into a difference. The special constant is the one where the difference is zero, and zero difference means equidistance, which is the most familiar locus there is.

Why is testing the top point M natural, rather than a trick? The definition of a circle is "all points at fixed distance from a centre". Symmetry has already told us the centre must be C and the radius a (the vertices are on E and are at distance a from C). So every point of Γ is at distance a from C. The ellipse condition compares $\overset{\frown}{MA} + \overset{\frown}{MB}$ with $2a = 2\overset{\frown}{MC}$. For the vertices, the two arcs to A and B run along the equator and "bend back" through C, so equality holds. For the point furthest from the equator, both arcs are the hypotenuses of right triangles with leg a, and a hypotenuse is longer than a leg. That is the whole content, and it is the spherical version of "the ellipse is fatter at the ends than in the middle". The one place the sphere could spoil it is $\cos a \leq 0$: when $a = \pi/2$ the hypotenuse equals the leg ($\cos \overset{\frown}{MA} = 0 \cdot \cos c = 0$), which is exactly the exceptional great-circle case. So the proof tells you where the exception lives.

The official solution states "Γ would have to be symmetric with respect to the equatorial plane, thus lie in a plane perpendicular to it". That line hides a case: a plane parallel to the equator is also invariant only if it is the equator itself, and that is ruled out by counting the equatorial points of E. The model chain above writes it out.

Layer $+\alpha$ BACKTESTING

1. The cue is "shortest distance on the sphere". Write down $\overset{\frown}{PA} + P\overset{\frown}{(-A)} = \pi$ immediately.
2. It suggests splitting by the size of the constant: reduce $2a > \pi$ to $2a < \pi$ via antipodal foci, and treat $2a = \pi$ as an equidistance locus.
3. Hesitation comes at "if it were a circle, which circle?". Reinforce it away with the symmetry of the configuration under reflection in the plane of A, B, O: a symmetric circle through the two vertices has them as a diameter.
4. The trap is the planar analogy, which predicts "never, unless $A = B$" and misses the great circles. A second trap is squaring the arccos equation, which buries the geometry.

5. Whenever a problem moves a planar notion (ellipse, bisector, parallel) to a sphere, ask what the antipodal map does to it before computing anything.

6. Reusable decision rule: *to decide whether a symmetric locus is a given kind of curve, let the symmetry name the unique candidate, then test the single point where the candidate and the locus are most likely to disagree (the point furthest from where they were forced to agree).*

1971 B5

1971 B5. Show that the graphs in the x-y plane of all solutions of the system of differential equations

$$x'' + y' + 6x = 0, \qquad y'' - x' + 6y = 0 \qquad (' = d/dt)$$

which satisfy $x'(0) = y'(0) = 0$ are hypocycloids, and find the radius of the fixed circle and the two possible values of the radius of the rolling circle for each such solution. (A hypocycloid is the path described by a fixed point on the circumference of a circle which rolls on the inside of a given fixed circle.)

Fields: differential equations (linear, constant coefficients); complex numbers; plane curves **Position:** B5 — late in the paper, but routine once the system is packed into one complex equation **Answer:** Every solution is $z = x + iy = A\,(2e^{3it} + 3e^{-2it})$ with $A \in \mathbb{C}$. With $\rho = 5|A| = \sqrt{x(0)^2 + y(0)^2}$, the fixed circle has radius ρ and the rolling circle has radius $\frac{2}{5}\rho$ or $\frac{3}{5}\rho$.

Caveat. The zero solution ($A = 0$, i.e. $x(0) = y(0) = 0$) traces a single point, a degenerate hypocycloid with all radii 0. The statement "all solutions are hypocycloids" is meant for the nonzero ones; the official solution passes over this silently.

Layer 1 THE WHAT — surface analysis

The natural attempt. The natural move is the textbook one for a coupled linear system: eliminate. From the first equation, $y' = -x'' - 6x$. Differentiate the second equation once and substitute $y'' = (-x'' - 6x)' = -x''' - 6x'$, $y''' = -x^{(4)} - 6x''$ and $y' = -x'' - 6x$:

$$y''' - x'' + 6y' = 0 \implies -x^{(4)} - 6x'' - x'' - 6x'' - 36x = 0 \implies x^{(4)} + 13x'' + 36x = 0.$$

The characteristic polynomial $r^4 + 13r^2 + 36 = (r^2 + 4)(r^2 + 9)$ gives $x = a_1 \cos 2t + b_1 \sin 2t + a_2 \cos 3t + b_2 \sin 3t$. Then you must recover y from $y' = -x'' - 6x$ and an integration constant, impose $x'(0) = y'(0) = 0$, and fix the constant by returning to one of the original equations (the fourth-order equation has more solutions than the system). That is four real constants, two conditions, a consistency check, and a pair of real formulas for $X(t)$, $Y(t)$ which you then have to recognise as a hypocycloid. It works. It costs half an hour, and the recognition step at the end is where students stall, because the real form hides the rolling.

The model chain. Put $z = x + iy$. Then

$$z'' = x'' + iy'' = (-y' - 6x) + i(x' - 6y) = i(x' + iy') - 6(x + iy) = iz' - 6z,$$

so the system is the single equation $z'' - iz' + 6z = 0$ (equivalent to it, by taking real and imaginary parts). Its characteristic equation $r^2 - ir + 6 = 0$ has roots

$$r = \frac{i \pm \sqrt{-1 - 24}}{2} = \frac{i \pm 5i}{2} \in \{3i,\, -2i\},$$

so $z = c_1 e^{3it} + c_2 e^{-2it}$. The condition $z'(0) = 0$ reads $3ic_1 - 2ic_2 = 0$, so $c_1 = 2A$, $c_2 = 3A$ for some $A \in \mathbb{C}$:

$$z(t) = A\big(2e^{3it} + 3e^{-2it}\big).$$

Write $A = Re^{i\alpha}$ with $R > 0$. Multiplying by $e^{i\alpha}$ is a rotation of the plane, so the graph is congruent to that of $Z(t) = 2Re^{3it} + 3Re^{-2it}$.

The rolling-circle formula. Let a circle of radius r roll inside a fixed circle of radius $F > r$ centred at 0, and let the traced point start at F. When the contact point is $Fe^{i\varphi}$, the rolling circle's centre is $(F-r)e^{i\varphi}$. Write the traced point as $(F-r)e^{i\varphi} + re^{i\beta}$. No slipping means the arc of the fixed circle already covered, $F\varphi$, equals the arc of the rolling circle from the traced point to the contact point, $r(\varphi - \beta)$. So $\beta = -\frac{F-r}{r}\varphi$ and

$$w(\varphi) = (F-r)e^{i\varphi} + re^{-i\frac{F-r}{r}\varphi}. \tag{H}$$

Matching. Take $F = 5R$, $r = 3R$: (H) is $2Re^{i\varphi} + 3Re^{-\frac{2}{3}i\varphi}$, and $\varphi = 3t$ gives exactly $Z(t)$. Take $F = 5R$, $r = 2R$: (H) is $3Re^{i\varphi} + 2Re^{-\frac{3}{2}i\varphi}$, and $\varphi = -2t$ gives $3Re^{-2it} + 2Re^{3it} = Z(t)$. Both reparametrisations are linear in t, so they trace the same set as t runs over $\mathbb{R}$.

So every nonzero solution traces a hypocycloid. (Its cusps occur where $|Z| = 5R$, that is where $e^{5it} = 1$, so it has five cusps.) The fixed circle has radius $5R = 5|A| = |z(0)|$, and the rolling circle has radius $2R$ or $3R$, that is $\frac{2}{5}|z(0)|$ or $\frac{3}{5}|z(0)|$. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains read the same two equations. The natural attempt sees "system" and eliminates a variable. The model sees the pattern (x'', y'') paired with $(-y', x')$: that is multiplication of the velocity by i. The split is at the very first line: eliminate, or combine into z.

> *The first-derivative terms are $+y'$ in the x equation and $-x'$ in the y equation. That swap with a sign change is a rotation by 90° of the vector (x', y'), which is multiplication by i. So I put $z = x + iy$ and the system becomes one equation.*

> *The answer is a sum of two circular motions, $2e^{3it}$ and $3e^{-2it}$. A rolling circle is also a sum of two circular motions: the centre going round, and the point going round the centre. I only need to match speeds and radii.*

- ▷ **Cue:** a planar system whose coupling terms are $(+y', -x')$ or $(-y, +x)$ with equal coefficients on the diagonal ⟶ **Move:** set $z = x + iy$; the coupling is multiplication by $\pm i$ and the system becomes one complex ODE
- ▷ **Cue:** a curve of the form $pe^{imt} + qe^{-int}$ (opposite senses) ⟶ **Move:** hypocycloid: fixed radius $p + q$, rolling radius q or p (after matching angular speeds); same senses would give an epicycloid
- ▷ **Cue:** "the two possible values" of a rolling radius ⟶ **Move:** expect the double-generation symmetry $r \leftrightarrow F - r$

Layer 3 THE WHY

Why is $z = x + iy$ the right move rather than a trick? Look at what the equations say. Write $\mathbf{v} = (x', y')$. The terms $+y'$ and $-x'$ form the vector $(y', -x')$, which is $\mathbf{v}$ rotated by $-90°$. So the system reads $\mathbf{r}'' = J\mathbf{v} - 6\mathbf{r}$ for a rotation J: a particle on a spring ($-6\mathbf{r}$) feeling a force always perpendicular to its velocity (a magnetic field, physically). Rotation by 90° in the plane is, by definition of complex multiplication, multiplication by i. So the problem is written in complex numbers already; z is its native coordinate, and elimination is translating it into a worse language and back. The same equation with a 2×2 real matrix would need eigenvectors; the complex form is the eigen-decomposition done in advance.

Why is a hypocycloid a sum of two exponentials? Go to the definition. The traced point is "centre of the rolling circle" plus "offset from the centre". The centre moves on a circle of radius $F - r$ at a steady rate, so it is $(F-r)e^{i\varphi}$. The offset has length r and rotates at a steady rate, so

it is $re^{i\beta}$. Rolling without slipping is the one condition linking the two rates, and it forces the offset to rotate the other way (the circle rolls clockwise as its centre moves anticlockwise). So a hypocycloid is exactly "two uniform circular motions in opposite senses whose radii and speeds satisfy one relation". The ODE solution has two opposite-sense motions, e^{3it} and e^{-2it}. The relation in (H) is that the ratio of angular speeds equals the ratio of radii with sides swapped: speed ratio $\frac{F-r}{r}$ against radii $F-r$ and r. Check it here: $|3|\cdot 2R = |-2|\cdot 3R$. That identity, $\text{speed}_1 \cdot \text{radius}_1 = \text{speed}_2 \cdot \text{radius}_2$, says the two contributions to the velocity have equal length, and it is precisely what the initial condition $z'(0) = 0$ produced ($3i\cdot 2A = 2i\cdot 3A$). The cusps are the moments when the two velocity contributions cancel. That is why the condition $x'(0) = y'(0) = 0$ is in the problem: without it you get a general "trochoid", not a hypocycloid.

Why two rolling radii? In (H) the two terms play symmetric roles once you allow a reparametrisation, so either term can be "the centre" and the other "the offset". That swaps r with $F-r$ and keeps F. This is the classical double generation of hypocycloids, and the problem's phrase "the two possible values" is announcing it.

Layer $+\alpha$ BACKTESTING

1. The first cue is the antisymmetric coupling $+y'$, $-x'$ with equal diagonal terms. Write $z = x + iy$.
2. A constant-coefficient complex linear ODE: characteristic roots, then match to the rolling-circle parametrisation.
3. Hesitation comes at the recognition step: "is $2Re^{3it} + 3Re^{-2it}$ a hypocycloid?". Reinforce it away by deriving (H) from the definition once, in two lines, rather than trying to recall it.
4. The trap is elimination to a fourth-order equation, which doubles the work and introduces spurious solutions that must be removed. A second trap is forgetting the second rolling radius, or the zero solution.
5. When two real equations are "each other rotated", treat them as one complex equation. When a curve is a sum of uniform circular motions, read it as a rolling-circle curve.
6. Reusable decision rule: *before eliminating in a coupled system, check whether the coupling matrix is a rotation (or a scalar plus a rotation); if it is, use the complex coordinate and the system collapses to one equation.*

1971 B6

1971 B6. Let $\delta(x)$ be the greatest odd divisor of the positive integer x. Show that

$$\left| \sum_{n=1}^{x} \frac{\delta(n)}{n} - \frac{2x}{3} \right| < 1,$$

for all positive integers x.

Fields: number theory (2-adic structure); recurrences; induction **Position:** B6 — last problem of the paper, but short once the even/odd split is made **Answer:** Proof. In fact $F(x) = \sum_{n\le x} \delta(n)/n - \frac{2x}{3}$ satisfies $0 < F(x) < \frac{2}{3}$; explicitly, if $x = \sum_j b_j 2^j$ in binary then $F(x) = \frac{1}{3}\sum_j b_j 2^{-j}$.

Caveat. The OCR renders the statement as "$|\sum_{n=1}^{x} \delta(n)/n - 2x/3| < 1$" with symbols garbled; the reading above is confirmed by the book's own solution, which defines $S(x) = \sum_{n=1}^{x} \delta(n)/n$ and $F(x) = S(x) - \frac{2}{3}x$. The closed binary formula for F is not in the official solution; it was derived here and checked exactly (rational arithmetic) for all $x < 5000$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The natural start is to compute. Write $S(x) = \sum_{n=1}^{x} \delta(n)/n$. Each term is

$1/2^k$ where 2^k is the exact power of 2 dividing n, since $n = 2^k\delta(n)$. So $S(x)$ counts odd n with weight 1, $n \equiv 2 \pmod 4$ with weight $\frac{1}{2}$, $n \equiv 4 \pmod 8$ with weight $\frac{1}{4}$, and so on. Group by k: the number of $n \le x$ with $2^k \parallel n$ is $\lfloor x/2^k \rfloor - \lfloor x/2^{k+1} \rfloor$, so

$$S(x) = \sum_{k\ge 0} \frac{1}{2^k}\left(\left\lfloor \frac{x}{2^k}\right\rfloor - \left\lfloor \frac{x}{2^{k+1}}\right\rfloor\right).$$

Replacing each floor by its argument gives $\sum_k \frac{x}{2^{k+1}} \cdot \frac{1}{2^k} = x\sum_k 4^{-k}/2 = \frac{2x}{3}$. So the main term is right, and the error is a sum of fractional-part terms $\{x/2^k\}$ with weights. Now you have to bound that error by 1, and a crude bound (each fractional part < 1, weights $1/2^k$) gives roughly 2, not 1. Tightening it means tracking signs of the fractional parts carefully. It can be done, but it is fiddly and easy to get wrong, and on the day this is where the student runs out of time.

The model chain. Split the sum over $n \le 2x$ into even and odd n. For even $n = 2m$, $\delta(2m) = \delta(m)$, so $\delta(2m)/(2m) = \frac{1}{2}\,\delta(m)/m$. For odd n, $\delta(n)/n = 1$. Hence

$$S(2x) = \sum_{m=1}^{x} \frac{\delta(2m)}{2m} + \sum_{m=1}^{x} \frac{\delta(2m-1)}{2m-1} = \tfrac{1}{2}S(x) + x, \qquad S(2x+1) = S(2x) + 1.$$

Put $F(x) = S(x) - \frac{2}{3}x$. Then

$$F(2x) = \tfrac{1}{2}S(x) + x - \tfrac{4}{3}x = \tfrac{1}{2}\big(S(x) - \tfrac{2}{3}x\big) = \tfrac{1}{2}F(x), \qquad F(2x+1) = F(2x) + 1 - \tfrac{2}{3} = \tfrac{1}{2}F(x) + \tfrac{1}{3}.$$

Claim: $0 < F(x) < \frac{2}{3}$ for all $x \ge 1$. For $x = 1$, $F(1) = 1 - \frac{2}{3} = \frac{1}{3}$. Every $x \ge 2$ is $2y$ or $2y+1$ with $1 \le y < x$; if $0 < F(y) < \frac{2}{3}$ then

$$0 < F(2y) = \tfrac{1}{2}F(y) < \tfrac{1}{3}, \qquad \tfrac{1}{3} < F(2y+1) = \tfrac{1}{2}F(y) + \tfrac{1}{3} < \tfrac{2}{3}.$$

By strong induction the claim holds for all x, and in particular $|F(x)| < 1$. □

Unwinding the recursion gives an exact formula. If $x = b_k 2^k + \cdots + b_1 2 + b_0$ in binary ($b_k = 1$), then $F(x) = \frac{1}{3}\big(b_0 + \frac{b_1}{2} + \cdots + \frac{b_k}{2^k}\big)$: the binary digits of x read backwards, as a number in $[2^{-k}, 2)$, divided by 3. The bound $< \frac{2}{3}$ is sharp ($x = 2^k - 1$ gives $\frac{2}{3}(1 - 2^{-k-1})$), and F gets arbitrarily close to 0 ($x = 2^k$ gives $\frac{1}{3}2^{-k}$).

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin with $n = 2^k\delta(n)$. The natural attempt sums over all k at once, getting floor functions. The model peels off only one factor of 2: it splits $\{1, \ldots, 2x\}$ into odds and evens and notices the evens reproduce the original sum at half size. The split is: all dyadic levels at once, or one level and recurse.

> *$\delta(2m) = \delta(m)$. So the even terms up to $2x$ are half of the whole sum up to x. That is a recursion for S in terms of S at half the argument; I don't need floors at all.*

> *The recursion is affine, so subtract its fixed-point line. If S were exactly cx, then $2cx = \frac{1}{2}cx + x$ forces $c = \frac{2}{3}$. The error $F = S - \frac{2}{3}x$ satisfies a recursion with no x in it, and halving shrinks errors, so it stays bounded.*

- ▷ **Cue:** a function of n defined through the power of 2 dividing n (greatest odd divisor, v_2, binary digit sums) ⟶ **Move:** split $n \le 2x$ into even and odd n to get a recursion in $x \mapsto 2x, 2x+1$
- ▷ **Cue:** an affine recursion $S(2x) = \lambda S(x) + \alpha x$ with $|\lambda| < 1$ and a conjectured main term cx ⟶ **Move:** find c from the fixed-point equation, set $F = S - cx$, and bound F by induction on the recursion

▷ **Cue:** a sum over all dyadic levels producing floor functions $\longrightarrow$ **Move:** recurse one level at a time instead of summing all levels

Layer 3 THE WHY

Why is splitting into even and odd the natural move, and not a trick? Go to the definition of δ. The greatest odd divisor is what is left when you strip all factors of 2. Stripping one factor of 2 changes nothing about δ: $\delta(2m) = \delta(m)$. That is the only structural fact the function has, and it is a statement about $n \mapsto 2n$. Any argument that uses it has to compare the sum at x with the sum at $2x$. So a recursion in the doubling map is not a device; it is the definition of δ rewritten as a statement about S. The same happens for every function built from the binary expansion: $v_2(n)$, the binary digit sum $s_2(n)$, the Stern sequence. Each has a rule under $n \mapsto 2n$ and $n \mapsto 2n+1$, and each is controlled by recursions in those two maps.

Why does $\frac{2}{3}$ appear? Heuristically, a random n is odd with probability $\frac{1}{2}$ (term 1), $\equiv 2 \pmod 4$ with probability $\frac{1}{4}$ (term $\frac{1}{2}$), and so on, so the average term is $\sum_{k\geq 0} 2^{-k-1}2^{-k} = \frac{1}{2} \cdot \frac{1}{1-1/4} = \frac{2}{3}$. The recursion finds the same number as the fixed point of $c \mapsto \frac{1}{2}c + \frac{1}{2}$ (divide $S(2x) = \frac{1}{2}S(x) + x$ by $2x$). Subtracting the main term is the move that turns "estimate a growing sum" into "bound a quantity that doesn't grow". The official solution defines $F(x) = S(x) - \frac{2}{3}x$ without comment. That is where it comes from: it is the deviation from the fixed point of the recursion, and deviations are what a contraction ($\lambda = \frac{1}{2}$) keeps small.

Why does induction give $\frac{2}{3}$ rather than a weaker bound? The two maps on F are $F \mapsto \frac{1}{2}F$ and $F \mapsto \frac{1}{2}F + \frac{1}{3}$. Each sends the interval $[0, \frac{2}{3}]$ into itself (onto $[0, \frac{1}{3}]$ and $[\frac{1}{3}, \frac{2}{3}]$ respectively). The interval $[0, \frac{2}{3}]$ is the attractor of those two maps, exactly as $[0, 1]$ is the attractor of $t \mapsto t/2$ and $t \mapsto t/2 + \frac{1}{2}$ (binary expansions). Once you see that, the correct inductive hypothesis is not a guess: it is the invariant interval of the maps. And the starting value $F(1) = \frac{1}{3}$ lies inside it. The problem asks for 1 to give room; the structure gives $\frac{2}{3}$.

Layer $+\alpha$ BACKTESTING

1. The cue is "greatest odd divisor". Write down $\delta(2m) = \delta(m)$ and $\delta(2m+1) = 2m+1$.
2. A recursion for S under $x \mapsto 2x$ and $x \mapsto 2x+1$, then strong induction on x.
3. Hesitation comes at "what should the inductive hypothesis be?". Reinforce it away by computing the invariant interval of the two affine maps $F \mapsto \frac{1}{2}F$, $F \mapsto \frac{1}{2}F + \frac{1}{3}$.
4. The trap is summing over all powers of 2 at once. It produces the right main term and an error term full of floors whose crude bound is about 2, which is not good enough.
5. Any function that "ignores factors of 2" or is defined from binary digits should trigger the doubling recursion. If a sum has a conjectured linear main term, subtract it and look at the recursion for what is left.
6. Reusable decision rule: *for a sum with a self-similar structure under $n \mapsto 2n$, split by parity, derive $S(2x) = \lambda S(x) + (\text{linear})$, subtract the fixed-point line, and prove the error stays in the invariant interval of the resulting affine maps.*