DAWN PUTNAM SESSION

Putnam 2007 — Papers A and B

Putnam 120 Archive

Contents

Putnam 2007 — Paper A **5**
2007 A1 5
2007 A2 7
2007 A3 10
2007 A4 13
2007 A5 16
2007 A6 18

Putnam 2007 — Paper B **23**
2007 B1 23
2007 B2 25
2007 B3 28
2007 B4 30
2007 B5 33
2007 B6 35

Putnam 2007 — Paper A

2007 A1

2007 A1. Find all values of α for which the curves $y = \alpha x^2 + \alpha x + \frac{1}{24}$ and $x = \alpha y^2 + \alpha y + \frac{1}{24}$ are tangent to each other.

Fields: coordinate geometry, algebra (symmetric elimination) **Position:** A1 — the opener **Answer:** $\alpha \in \left\{ \frac{2}{3}, \frac{3}{2}, \frac{13+\sqrt{601}}{12}, \frac{13-\sqrt{601}}{12} \right\}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Call the curves $C_1 : y = f(x) = \alpha x^2 + \alpha x + \frac{1}{24}$ and $C_2 : x = f(y)$. Tangency means the two curves meet at some point with the same tangent line there, so the honest first move is to eliminate a variable. Substitute $y = f(x)$ into $x = f(y)$:

$$x = \alpha f(x)^2 + \alpha f(x) + \frac{1}{24}.$$

Expanding $f(x)^2$ turns this into a quartic in x (degree 4, since $f(x)^2$ has degree 4), with coefficients that are polynomials in α. A double root of this quartic corresponds to a point where C_1 and C_2 meet with the same x-coordinate to first order – which is close to, but not exactly, the tangency condition, since tangency also requires matching slopes at that shared point, not merely a repeated root of the elimination. Pressing on regardless, the plan is: write the quartic $a_4(\alpha)x^4 + a_3(\alpha)x^3 + a_2(\alpha)x^2 + a_1(\alpha)x + a_0(\alpha) = 0$ explicitly, form its discriminant as a polynomial in α, and set that discriminant to zero.

This works in principle. In practice the discriminant of a general quartic is a degree-6 polynomial in the quartic's five coefficients, and here each coefficient is itself a polynomial in α of degree up to 2 – the discriminant, expanded, is an ugly high-degree polynomial in α carrying large integer coefficients, produced with no insight into which factor of it is the "real" tangency condition and which factors are extraneous (repeated roots that arise for algebraic reasons unconnected to genuine tangency, such as the elimination process itself introducing spurious multiplicities). It is exactly the kind of computation where a single arithmetic slip is invisible until the final answer refuses to match anything clean, and even a correct execution gives no feeling for *why* the four surviving values of α are the ones they are.

The model chain. The two curves are mirror images of each other across the line $L : y = x$ – swapping x and y in the equation of C_1 gives exactly the equation of C_2. Consequently a point of tangency, if it exists, is heavily constrained by this reflection symmetry, and the whole problem reduces to finding the *fixed slopes* of the reflection.

Case 1: the tangency point lies on L. Say $(t, t) \in C_1 \cap C_2$, so

$$t = \alpha t^2 + \alpha t + \tfrac{1}{24}. \qquad (*)$$

Implicit differentiation gives the slope of C_1 at (t, t) as $m_1 = 2\alpha t + \alpha$, and (differentiating $x = \alpha y^2 + \alpha y + \frac{1}{24}$) the slope of C_2 at (t, t) as $m_2 = 1/(2\alpha t + \alpha) = 1/m_1$. Tangency at a shared point on L requires $m_1 = m_2$, i.e. $m_1^2 = 1$, so $m_1 = \pm 1$.

If $m_1 = 1$: $\alpha = 1/(2t+1)$. Substituting into $(*)$ and clearing denominators gives $24t^2 - 2t - 1 = 0$, so $t \in \{1/4, -1/6\}$, giving $\alpha \in \{2/3, 3/2\}$.

If $m_1 = -1$: $\alpha = -1/(2t+1)$. Substituting into $(*)$ gives $72t^2 + 46t - 1 = 0$, so $t = (-23 \pm \sqrt{601})/72$, giving $\alpha = (13 \pm \sqrt{601})/12$.

Case 2: no tangency point lies on L. If C_1 does not meet L at all, C_1 and C_2 lie on opposite sides of L and cannot touch. If C_1 meets L in two points where the tangent is *not* perpendicular to L (the $m_1 = -1$ case already handled), substituting C_1 into C_2 produces at most four intersection points counted with multiplicity (the elimination is a quartic); two of these are the crossings of L already accounted for, and any genuine off-diagonal tangency would have a mirror-image tangency by the reflection symmetry, forcing at least two more coincident pairs – six solutions where only four exist. So no further values of α arise. This closes the case analysis: the four values above are the complete answer.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree that tangency is a statement about a shared point with matching slopes. They split at how to use the fact that C_2 is literally C_1 reflected in $y = x$: the natural attempt treats this as incidental and eliminates a variable directly, producing a general quartic with no structure left visible; the model chain uses the reflection itself as the object of study.

> *The second curve is not just* some other *curve – it is obtained from the first by swapping x and y, which is the reflection across $y = x$. Before I do a two-variable elimination that throws that symmetry away, I should ask what tangency looks like* under *a reflection: either the tangency point sits on the mirror line itself, or it comes as a mirrored pair of off-line points, and both of those are much more constrained than "some point somewhere."*

> *A reflection across $y = x$ doesn't just swap points – it acts on slopes too, sending a line of slope m to a line of slope $1/m$. If the tangency point is on the mirror line, the two curves' tangent lines there are reflections of each other, so matching them is not "solve a generic system" but "solve $m = 1/m$" – and that is a one-line equation, not a quartic.*

Layer 3 THE WHY

Why does the reflection-of-slopes fact even hold, and why is it the right thing to reach for? Reflecting the plane across $y = x$ sends the point $(x_0 + h, y_0 + km)$ near a point (x_0, y_0) on a curve of local slope m to $(y_0 + km, x_0 + h)$ near (y_0, x_0) – swap the roles of the coordinate increments, and the new slope is $h/(km) \cdot k \cdot \ldots$; more simply, reflecting swaps $\Delta x \leftrightarrow \Delta y$, so a slope $m = \Delta y/\Delta x$ becomes $\Delta x/\Delta y = 1/m$. This is not a fact special to this problem; it is what the reflection $(x, y) \mapsto (y, x)$ *is*, applied to tangent directions instead of points. Once that is understood as a map on slopes, $\varphi(m) = 1/m$, the question "at what slope is a tangent line carried to itself by the reflection" is exactly the fixed-point equation $\varphi(m) = m$, i.e. $m^2 = 1$. This is where $m = \pm 1$ comes from, and it is not a guess to be verified after the fact – it is forced the instant the reflection is recognised as an involution acting on the space of directions, and any involution's invariant directions are found the same way, by solving $\varphi(m) = m$.

The second thing the WHY layer owes is the closing argument in Case 2: why can there be no fifth or sixth value of α? This is where the official solution's economy is easy to admire and hard to reconstruct without being told the reason. Substituting a quadratic into a quadratic (here, $y = f(x)$ into $x = f(y)$) gives an equation of degree at most 4 in one variable, because $f(y)^2$ is degree 4 in y and f is degree 2; this bounds the number of intersection points of two conics, counted correctly, at 4 – a fact whose deepest form is Bézout's theorem for plane curves, but whose elementary content here is just "a quartic has at most four roots." Once two of those four roots are already spent on the L-crossings in Case 2, and any further tangency off L would

need a symmetric partner (by the same reflection), there is no room left: $2+2>4$ forces the extra tangency not to exist. The bound is not extra machinery bolted onto the geometry; it is the same degree-counting the natural attempt already stumbled into, but read for what it forbids instead of ground through for what it produces.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that swapping x and y in C_1's equation literally produces C_2's equation. That should be written down immediately as: this pair is related by reflection across $y=x$, not treated as two unrelated curves to be eliminated against each other.

2. That cue suggests working with the fixed line of the reflection and the fixed slopes of the reflection's action on directions, rather than a direct two-variable elimination.

3. Hesitation sets in when deciding whether tangency must occur on L at all – it is tempting to assume it does and stop after finding $\alpha = 2/3, 3/2$. What reinforces continuing is the slope equation itself: $m^2=1$ has *two* roots, and $m=-1$ was never actually shown to require the tangency point to lie on L – checking the geometry shows $m=-1$ corresponds to a tangent line perpendicular to L at a crossing point of L, a second legitimate configuration that must be chased down separately.

4. The trap is stopping at two solutions, because $m=1$ is the visually obvious case (the curves seem to "kiss" symmetrically) while $m=-1$ requires realising that a line perpendicular to the mirror is also carried to itself by the reflection, just with its direction reversed rather than preserved pointwise.

5. Recognise this shape faster: any time a problem hands you two objects related by an explicit swap of variables (or more generally an explicit involution), ask first for the fixed locus of that involution and the fixed directions of its action on slopes/tangents, before doing any elimination.

6. Reusable rule: for a pair of curves related by an involution σ of the plane, a coincidence between them (intersection, tangency, etc.) either happens on the fixed locus of σ, or comes in σ-mirrored pairs off that locus; and an invariant tangent direction at a fixed point must satisfy $m=\varphi(m)$, where φ is σ's induced action on slopes.

▷ **Cue:** two curves are defined by literally swapping the roles of x and y in the same formula $\longrightarrow$ **Move:** treat the pair as related by reflection across $y=x$; search for coincidences on the fixed line $y=x$ first, then for symmetric off-line pairs, rather than eliminating variables directly

▷ **Cue:** need the slope(s) at which a line is mapped to itself by a known involution of the plane $\longrightarrow$ **Move:** write the involution's induced action φ on slopes and solve the fixed-point equation $\varphi(m)=m$, instead of searching geometrically

▷ **Cue:** an elimination between two conics threatens to produce a high-degree resultant with no visible structure $\longrightarrow$ **Move:** bound the number of genuine solutions first (e.g. two conics meet in at most 4 points), then use that bound to certify a case analysis is exhaustive rather than expanding the resultant

2007 A2

2007 A2. Find the least possible area of a convex set in the plane that intersects both branches of the hyperbola $xy=1$ and both branches of the hyperbola $xy=-1$. (A set S in the plane is called *convex* if for any two points in S the line segment connecting them is contained in S.)

Fields: plane geometry, inequalities (AM-GM) **Position:** A2 **Answer:** The minimum is 4, achieved by the square with vertices $(\pm 1, \pm 1)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The problem asks for a minimum over an enormous space of possible sets – any convex set meeting four specific curve-branches. A natural instinct is to guess a plausible optimal shape and check it. The square with vertices $(1,1),(-1,1),(-1,-1),(1,-1)$ is an obvious candidate: it passes exactly through the four points where the hyperbolas cross the lines $y = x$ and $y = -x$, it visibly touches all four branches, and its area is $2 \times 2 = 4$. A circle centred at the origin might also be tried: to reach the branch of $xy = 1$ nearest the origin, it must have radius at least $\sqrt{2}$ (the closest point of $xy = 1$ to the origin is $(1,1)$), giving area $2\pi \approx 6.28$ – worse than the square. Trying a few more shapes (a smaller square rotated $45°$, various rectangles) either fails to reach all four branches or comes out with area at least 4.

This is genuine evidence that 4 is the answer, and it is honest work – but it is not a proof. Nothing here rules out some cleverer, non-obviously-shaped convex region achieving area less than 4: a thin lens-shaped set stretched out to just kiss all four branches, say, or an irregular hexagon. Guess-and-check over shapes can build confidence but can never certify a minimum, because the search space of convex sets is infinite-dimensional and no amount of trying named shapes closes it off. What is missing is a lower bound that holds for *every* admissible convex set at once, proved without knowing what that set looks like.

The model chain. Let S be any convex set meeting all four branches. Choose one point on each branch that S actually contains: $A = (a, 1/a)$ on the branch of $xy = 1$ in the first quadrant, $B = (-b, 1/b)$ on the branch of $xy = 1$ in the third-quadrant-adjacent upper-left region, $C = (-c, -1/c)$ on the branch of $xy = -1$ in the lower-left, $D = (d, -1/d)$ on the remaining branch, with $a, b, c, d > 0$. Since S is convex and contains A, B, C, D, it contains their convex hull, so

$$\text{area}(S) \geq \text{area}(ABCD).$$

This already converts an infinite-dimensional problem (minimise over convex sets) into a four-real-variable one (minimise over $a, b, c, d > 0$).

By the shoelace formula,

$$\text{area}(ABCD) = \frac{1}{2}\left(\frac{a}{b} + \frac{b}{a} + \frac{b}{c} + \frac{c}{b} + \frac{c}{d} + \frac{d}{c} + \frac{d}{a} + \frac{a}{d}\right).$$

Group the eight terms into four reciprocal pairs: $(a/b + b/a)$, $(b/c + c/b)$, $(c/d + d/c)$, $(d/a + a/d)$. By AM-GM, each pair is at least 2 (with equality exactly when the two ratio'd quantities are equal), so the whole sum is at least 8, giving

$$\text{area}(ABCD) \geq \frac{1}{2} \cdot 8 = 4.$$

This bound holds for *every* choice of witnesses A, B, C, D on the four branches, hence for every convex S: $\text{area}(S) \geq 4$, with no reference to what S actually looks like beyond convexity and the four intersection points it is required to have.

Equality requires $a = b = c = d$ in every pair simultaneously, i.e. $a = b = c = d = t$ for some $t > 0$; then $A = (t, 1/t)$, $B = (-t, 1/t)$, $C = (-t, -1/t)$, $D = (t, -1/t)$ form a rectangle of area $4t \cdot \frac{1}{t} = 4$ for *any* t, and this rectangle is itself an admissible convex set (it meets each branch at its own vertex). Taking $t = 1$ recovers the guessed square, now certified as optimal rather than merely plausible.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths eventually look at the same square. They split at the point of committing to a specific shape: the natural attempt starts by choosing a candidate region and computing its area; the model chain never chooses a region at all – it chooses four points a general S must contain, and bounds the area of *any* such S using only those four points and convexity.

> *I don't actually know what S looks like, and I never will – the problem only tells me it's convex and that it meets four branches somewhere. So I shouldn't try to guess S's shape at all. What I do know for certain is that S contains at least one point on each branch, and convexity says a convex set contains the convex hull of any points it contains – so I can get a universal lower bound on* $\text{area}(S)$ *just from four witness points, without ever pinning down S itself.*

> *The area of the witness quadrilateral, once I set it up with the standard hyperbola parametrisation* $(a, 1/a)$*, comes out as a sum of eight terms that pair up exactly into reciprocals:* a/b *sits next to* b/a*. That is precisely the shape AM-GM wants –* $x + 1/x \geq 2$ *– so the minimisation is not really a geometry problem in disguise, it is four applications of the single simplest instance of AM-GM, once the parametrisation is chosen to make the reciprocal structure visible.*

Layer 3 THE WHY

Why is the jump from "guess a shape" to "bound every shape via four witnesses" the natural one, rather than a trick to be shown after the fact? Because of what the word *convex* actually asserts: for any two points of S, the segment between them lies in S. Applied repeatedly, this says the convex hull of any finite subset of S is itself contained in S – this is not an extra fact about convex sets, it is the definition of convexity, iterated. The problem statement guarantees S contains *some* point on each of the four branches; naming those four points and taking their hull is simply using the hypothesis for exactly what it says, rather than for a stronger, unstated property (like "S is bounded by nice curves") that the problem never granted. Once this is seen, minimising over convex sets *must* reduce to minimising a hull area over witness points, because the hull-containment property is the only leverage convexity gives on area at all.

The second piece of the WHY is why AM-GM, specifically, is the right tool and not a lucky match. The hyperbola $xy = 1$ is, by its own defining equation, exactly the set of points $(a, 1/a)$ – a curve parametrised natively by a variable and its own reciprocal. Any area computed from points on this curve, via the shoelace formula (itself just repeated cross-products, i.e. sums of products of coordinates), will produce terms of the form (one point's x-coordinate) $\times$ (another point's y-coordinate), and since every y-coordinate here *is* a reciprocal of some x-coordinate, the area sum is built entirely from ratios and their reciprocals before any inequality is applied. AM-GM is not imported from outside to close the problem; the problem's own coordinates arrange themselves into the exact form AM-GM answers, because the hyperbola's equation *is* the relation $y = 1/x$.

Finally, the achievability half deserves the same scrutiny as the bound: a lower bound of 4 is worthless as an answer to "find the least possible area" unless some actual convex set attains it. Here that verification is cheap precisely because the equality condition of AM-GM ($a = b = c = d$) was tracked through the proof rather than discarded once the inequality was established – the proof of the bound and the construction of the minimiser are the same computation, read in two directions.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the question itself: "least possible area of a convex set satisfying [intersection conditions]," with the set otherwise completely unspecified. That should immediately be written down as: since I don't get to know the set's shape, I need a bound that uses only the intersection conditions and convexity – not a bound that assumes any particular shape.
2. That cue suggests choosing one witness point per required intersection, invoking convexity to get a hull-containment inequality, and turning the problem into a finite optimisation over the witnesses' coordinates.
3. Hesitation occurs right after guessing the square and computing its area 4 – it is tempting to stop there, since the number is clean and the shape is natural. What reinforces continuing is

asking directly: have I actually shown no other convex set beats this, or only that this one particular set achieves 4? A "least possible" question is never answered by exhibiting one example, however natural.

4. The trap is treating the guessed square's area as the answer on the strength of having tried a few alternatives (a circle, a rotated square) and found nothing better – this is evidence, not proof, because the search was over named shapes, not over all convex sets.

5. Recognise this shape faster: any "least/greatest area (or perimeter) of an unspecified convex/star-shaped set satisfying containment or intersection conditions" problem should trigger "pick witness points, use the containment property, reduce to a finite optimisation" before any shape is guessed.

6. Reusable rule: to bound the area (or another monotone quantity) of an unknown convex set defined by intersection or containment constraints, select one witness point per constraint, use convexity to lower-bound the set by the convex hull of the witnesses, and optimise that finite-dimensional hull quantity – reaching for AM-GM whenever the natural parametrisation of the constraining curves produces terms paired with their own reciprocals.

▷ **Cue:** an unknown convex set is specified only by which curves/branches/regions it must intersect $\longrightarrow$ **Move:** choose one witness point per required intersection; bound the set's area below by the area of the convex hull of the witnesses, using the definition of convexity directly

▷ **Cue:** a sum built from a natural curve parametrisation contains a term and its exact reciprocal (e.g. a/b alongside b/a) $\longrightarrow$ **Move:** pair each term with its reciprocal partner and apply AM-GM, $x + 1/x \geq 2$, tracking the equality condition to later construct the extremiser

▷ **Cue:** a lower bound has been proved for an extremal problem ("least possible ...") $\longrightarrow$ **Move:** do not stop at the bound; verify it is attained by exhibiting an admissible object achieving equality, ideally reading the equality case straight off the inequality just used

2007 A3

2007 A3. Let k be a positive integer. Suppose that the integers $1, 2, 3, \ldots, 3k+1$ are written down in random order. What is the probability that at no time during this process, the sum of the integers that have been written up to that time is a positive integer divisible by 3? Your answer should be in closed form, but may include factorials.

Fields: combinatorics, probability, modular arithmetic **Position:** A3 **Answer:** $\dfrac{k!\,(k+1)!}{(3k+1)\,(2k)!}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The direct instinct is to attack the $3k+1$ actual integers with inclusion-exclusion or a direct recursive count: let p_n be the probability of survival after placing the first n of the $3k+1$ numbers (in some fixed order of arrival), condition on which numbers have been placed so far, and try to set up a recursion in n tracking the running sum mod 3 and which numbers remain available. For small k this is doable by brute enumeration – $k = 1$ has only $4! = 24$ orderings and can be checked by hand – and a pattern in p_n might even be guessed by computing $k = 1, 2$ explicitly.

The trouble is that the state needed to run this recursion honestly is not just "current sum mod 3"; it is the full multiset of which specific numbers remain, because which numbers remain determines the distribution of what residue mod 3 could be drawn next, and that distribution keeps changing as multiples of 3 and non-multiples get exhausted at different rates. Carrying the sum mod 3 alone throws away exactly the information needed to know the probabilities of the next draw's residue class. A fully honest recursive count therefore has to track (numbers of each residue class remaining, current sum mod 3) as its state – workable for one specific k, but

it produces a recursion in three integer parameters whose closed form, if it exists at all, is not going to fall out by inspection; guessing the pattern from $k = 1, 2, 3$ and no structural argument for why it should be $k!(k+1)!/[(3k+1)(2k)!]$ risks matching small cases by coincidence.

The model chain. Only the residue of each number mod 3 affects whether a partial sum is divisible by 3, so replace the $3k+1$ actual integers by their residue classes: among $1, \ldots, 3k+1$ there are $k+1$ numbers $\equiv 1 \pmod 3$, k numbers $\equiv 2 \pmod 3$ (i.e. $\equiv -1$), and k numbers $\equiv 0 \pmod 3$. Call these classes $+, -, 0$.

A 0-class number never changes the running sum mod 3 when inserted – *except* that if it is the very first number written, the running sum at that moment equals that number itself, a positive multiple of 3, which is forbidden. So the survival condition splits into two independent requirements: (i) the first number written is not in class 0; (ii) restricting attention to the subsequence of non-0 numbers only, in the order they appear, the partial sums (of $+1$'s and -1's) mod 3 never hit 0.

These two events are independent for a uniformly random permutation, because the relative order of the non-0-class numbers among themselves is itself uniform and independent of *which* positions the 0-class numbers occupy (a standard fact about random permutations: fixing a subset, its internal relative order and the interleaving pattern of positions are chosen independently). So

$$P(\text{survives}) = P(\text{first is not class 0}) \times P(\text{the } \pm 1 \text{ walk survives}).$$

The first factor is $\dfrac{(3k+1)-k}{3k+1} = \dfrac{2k+1}{3k+1}$.

For the second factor: the running sum of the reduced ± 1 sequence starts at 0 and must never return to 0. From state 1, adding $+1$ gives 2 (safe) and adding -1 gives 0 (forbidden) – so from state 1 the *only* legal next step is $+1$. From state 2, symmetrically, the only legal next step is -1. So once the first step is taken (either is legal, since both $1, 2 \neq 0$), every subsequent step is forced: starting with $+1$ forces the pattern $+, +, -, +, -, +, -, \ldots$; starting with -1 forces $-, -, +, -, +, \ldots$. Since there are $k+1$ $+$'s and only k $-$'s, and the second pattern needs at least as many $-$'s as $+$'s among its first several terms, only the first pattern, $+, +, -, +, -, \ldots$, is consistent with the actual counts. So there is exactly *one* arrangement, out of all $\binom{2k+1}{k}$ arrangements of $k+1$ $+$'s and k $-$'s, that survives:

$$P(\text{the } \pm 1 \text{ walk survives}) = \frac{1}{\binom{2k+1}{k}}.$$

Multiplying,

$$P(\text{survives}) = \frac{2k+1}{3k+1} \cdot \frac{1}{\binom{2k+1}{k}} = \frac{2k+1}{3k+1} \cdot \frac{k!(k+1)!}{(2k+1)!} = \frac{k!(k+1)!}{(3k+1)(2k)!},$$

using $(2k+1)! = (2k+1)(2k)!$ to cancel the $(2k+1)$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree the problem is really about divisibility by 3 of partial sums. They split at what to carry forward as the state: the natural attempt keeps track of the actual remaining numbers (because that seems necessary to know what could be drawn next), while the model chain throws the actual values away entirely and keeps only three symbols $+, -, 0$, betting that residues alone determine survival.

> *Whether a partial sum is divisible by 3 depends only on residues mod 3 – so relabelling any two numbers with the same residue changes nothing about which orderings survive. I should stop*

> *thinking about "the integers 1 through $3k+1$" and start thinking about a sequence of $k+1$ pluses, k minuses, and k zeros, because that is genuinely all the information the survival condition can see.*

> *A zero contributes nothing to the running sum mod 3, so inserting one anywhere is free – with exactly one exception, the very first position, where the "running sum" is that number itself. Once I see that the zeros are free except at position one, the problem cleanly separates into "is position one a zero?" and "what order do the nonzero residues come in?", and these are independent questions about a random permutation, not one tangled question.*

> *The ± 1 walk on residues $\{1,2\}$ avoiding 0 isn't a counting problem waiting for a clever bijection – at every reachable state exactly one of the two possible steps is legal, so the whole path is determined the moment the first step is chosen. I should stop looking for a formula that counts valid walks and instead just verify there is exactly one, by checking legality state by state.*

Layer 3 THE WHY

Why is it legitimate, not just convenient, to reduce to residues before counting at all? Divisibility by 3 is by definition a statement about a number's residue mod 3, and nothing else about a number's value enters the condition "the partial sum is divisible by 3." A uniformly random permutation of $1,\ldots,3k+1$ treats every number as interchangeable except for identity; relabelling within a fixed residue class is a symmetry of the sample space that preserves every event phrased purely in terms of residues. So passing to the reduced alphabet $\{+,-,0\}$ loses no information the problem can actually query – it is not an approximation, it is working directly with the sufficient statistic the problem defines.

Why does the independence step hold, rather than needing to be re-derived by counting? It is a general fact about random permutations, not special to this problem: partition any N labelled objects into two classes of sizes r and $N-r$; a uniformly random permutation of all N objects is equivalent, in distribution, to independently choosing (a) which of the $\binom{N}{r}$ interleavings of positions the two classes occupy, and (b) a uniformly random internal ordering of each class separately. This is just unpacking what "uniformly random permutation" means as a product of independent uniform choices once the objects are grouped into classes – it holds for any partition into classes, and the 0-vs-nonzero split here is simply one instance of it.

Why does the ± 1 walk turn out to be completely rigid, rather than merely "usually" forced? Because 3 is prime and there are exactly two nonzero residues, 1 and 2, with $1+2 \equiv 0$: from any nonzero state s, the unsafe step is the unique one equal to $-s \pmod 3$, and the safe step is the other one – there is never a choice, because there are only two possible steps and exactly one of them is unsafe at every nonzero state. This rigidity is a feature of working mod 3 with a two-symbol active alphabet specifically; it is the reason the "count the walks" instinct from the natural attempt is not merely hard here but is answering the wrong question – there is nothing to count once uniqueness is established, and proving uniqueness (one legal step per state) is strictly easier than any enumeration would have been.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "divisible by 3" applied to partial sums. On seeing it, the move to write down immediately is: replace every number by its residue mod 3 and count how many of $1,\ldots,3k+1$ fall in each class ($k+1$ ones, k twos, k zeros) – before attempting any recursion on the original numbers.

2. That cue suggests splitting the sequence into a "neutral" subsequence (residue 0, which never changes the sum) and an "active" subsequence (residues ± 1, which do), and analysing each separately.

3. Hesitation occurs at deciding whether zeros genuinely can be ignored, since intuitively "inserting a number" always feels like it should matter. What reinforces the split is checking explicitly what a 0-class insertion does to the running sum (nothing, mod 3) and then hunting specifically for where that claim could fail – which isolates the single exception, the first position.

4. The trap is treating the ± 1 walk as a combinatorial object to be counted (e.g. via a ballot-problem or reflection argument, expecting some nontrivial count of valid sign sequences), when in fact checking legality at each state shows there is exactly one valid sequence. Reaching for a counting technique before checking uniqueness wastes the entire step.

5. Recognise this shape faster: whenever "partial sums avoid a fixed residue" appears, ask immediately how many *distinct nonzero residues* are in play and whether, from each such residue, more than one next step avoids the forbidden value – if only one step is ever safe at each reachable state, the path is forced and there is nothing left to count.

6. Reusable rule: for random-order problems where partial sums must avoid a residue class mod m: (a) reduce every term to its residue; (b) separate residues congruent to the forbidden value (typically 0) as "neutral," since they only matter at the very start of the sequence; (c) treat the remaining "active" residues as a walk that must avoid the forbidden state, and check state-by-state whether the walk is uniquely determined – often it is, which converts a counting problem into an easier uniqueness argument, and only then multiply by the independent probability governing where the neutral elements sit.

▷ **Cue:** a random-order sequence's partial sums must avoid a fixed residue class mod m $\longrightarrow$ **Move:** reduce every term to its residue mod m; separate residues equal to the forbidden value as neutral (safe to insert anywhere except possibly the very start) from active residues that move the running sum

▷ **Cue:** a walk on residues must never revisit a forbidden state $\longrightarrow$ **Move:** check, state by state, how many of the available steps are legal; if exactly one is legal at every reachable state, the entire path is forced from the first step's choice, and the problem becomes checking uniqueness rather than counting

▷ **Cue:** need the probability that a specific relative order holds among a labelled subset of positions in a uniformly random permutation $\longrightarrow$ **Move:** use that the subset's internal order and the interleaving pattern of its positions against the complement are independent uniform choices; compute each probability separately and multiply

2007 A4

2007 A4. A *repunit* is a positive integer whose digits in base 10 are all ones. Find all polynomials f with real coefficients such that if n is a repunit, then so is $f(n)$.

Fields: algebra (polynomials), number theory (base-10 representation), analysis (limits) **Position:** A4 — mid-paper; the difficulty is entirely in the change of variable, after which a limit argument closes it in a few lines. **Answer:** Exactly the polynomials $f(n) = \frac{1}{9}\left(10^c(9n+1)^d - 1\right)$ with integers $d \geq 0$ and $c \geq 1 - d$.

Layer 1 THE WHAT — surface analysis

Notation: $R_m = \underbrace{11\ldots1}_{m} = \frac{10^m-1}{9}$ for $m \geq 1$ is the m-th repunit. The hypothesis is that for every $m \geq 1$ there is some $h(m) \geq 1$ with $f(R_m) = R_{h(m)}$.

The natural attempt. You start by collecting examples, which is the right instinct. Constants: $f(n) = 1$, $f(n) = 11$, any fixed repunit. Identity: $f(n) = n$. Appending a digit: $R_{m+1} = 10R_m + 1$, so $f(n) = 10n + 1$ works, and so does $100n + 11$. Then you look for something

non-linear. Squaring a repunit gives 121, 12321, … — not repunits. So you guess that only linear maps survive, and try to prove it.

Suppose $f(n) = an + b$. From $f(R_1) = R_{h_1}$, $f(R_2) = R_{h_2}$ you get two equations; with $f(R_3)$ a third, and you try to force a to be a power of 10 by comparing $R_{h_2} - R_{h_1} = 10a$ with $R_{h_3} - R_{h_2} = 100a$. Differences of repunits are $R_j \cdot 10^i$-shaped strings of ones followed by zeros, and a page of digit-bookkeeping later you have a case analysis for degree 1 and no idea how to run it for degree d.

Worse, the guess "only linear" is false. Try $f(n) = 9n^2 + 2n$: $f(1) = 11$, $f(11) = 1111$, $f(111) = 111111$. It doubles the length. You will not find this by staring at 12321, because it is not the square of anything recognisable in the repunit world — it is a square in a different coordinate.

The model chain. *Step 1: change variable.* A positive integer n is a repunit if and only if $9n + 1 = 10^m$ for some $m \geq 1$. Define

$$g(x) = 9f\left(\frac{x-1}{9}\right) + 1,$$

a polynomial with real coefficients, with $f(n) = \frac{g(9n+1)-1}{9}$ and $\deg g = \deg f$. Then f sends repunits to repunits if and only if

$$g(10^m) \in \{10^j : j \geq 1\} \quad \text{for every } m \geq 1. \tag{$*$}$$

Step 2: the leading coefficient. Let $d = \deg g$ and let a be the leading coefficient (if g is constant, $d = 0$ and $a = g$). Then $g(x)/x^d \to a$ as $x \to \infty$, so

$$r_k := \frac{g(10^k)}{10^{kd}} \to a \quad (k \to \infty).$$

By $(*)$, each $r_k = 10^{j_k - kd}$ is an integer power of 10 (possibly a negative power). In particular $r_k > 0$, so $a \geq 0$, and $a \neq 0$, so $a > 0$. The set $\{10^j : j \in \mathbb{Z}\}$ has only finitely many points in $[a/2, 2a]$. If a were not one of them, some interval around a would contain no power of 10, and r_k could not converge to a. So $a = 10^c$ for an integer c, and some interval around a contains no other power of 10; hence $r_k = 10^c$ for all large k.

Step 3: rigidity. For all large k, $g(10^k) = 10^c(10^k)^d$. The polynomial $g(x) - 10^c x^d$ has infinitely many roots, so it is identically zero: $g(x) = 10^c x^d$.

Step 4: the constraint. Now $g(10^m) = 10^{c+md}$, which lies in $\{10^j : j \geq 1\}$ for every $m \geq 1$ exactly when $c + md \geq 1$ for all $m \geq 1$; since $d \geq 0$ the worst case is $m = 1$, giving $c + d \geq 1$. Translating back,

$$f(n) = \frac{10^c(9n+1)^d - 1}{9}, \qquad d \geq 0,\ c \geq 1 - d,$$

and each of these works, because $10^c(10^m)^d = 10^{c+md}$ with $c + md \geq 1$, and $\frac{10^j-1}{9}$ is a repunit for $j \geq 1$.

Caveat. Checked by machine: for $0 \leq d \leq 3$ and $1 - d \leq c \leq 2$ the formula sends $R_1, \ldots, R_7$ to repunits (exact rational arithmetic). Note c may be negative: $d = 2$, $c = -1$ gives $f(n) = \frac{(9n+1)^2-10}{90}$, which has non-integer coefficients yet maps $R_m \mapsto R_{2m-1}$. The problem allows real coefficients, so these must be included.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start in the same place: "$R_m = (10^m - 1)/9$." The natural attempt writes this down and then keeps working with R_m — comparing digit strings, differencing repunits. The model reads the same formula backwards, $9R_m + 1 = 10^m$, and moves the whole problem into the coordinate $x = 9n + 1$, where repunits become powers of 10.

> *Repunits are awkward objects, but $9n+1$ turns a repunit into a power of 10, and powers of 10 are the cleanest set there is. The map $n \mapsto 9n+1$ is affine, so it sends polynomials to polynomials. I'll conjugate f by it and ask which polynomials send powers of 10 to powers of 10.*

> *Along $x = 10^k$, a degree-d polynomial divided by x^d tends to its leading coefficient, but here every value of that ratio is a power of 10. A convergent sequence that lives in a discrete set is eventually constant. So $g(x) = ax^d$ at infinitely many points, which pins g completely.*

Layer 3 THE WHY

Why is $9n+1$ the natural move and not a trick? Because the hypothesis is a statement about a *set* (repunits $\to$ repunits), and the first question to ask of any set-preservation problem is: in what coordinates is the set simplest? The repunits are the image of the powers of 10 under the affine map $x \mapsto (x-1)/9$. That is not an inspired observation; it is the formula for the sum of a geometric series, $1 + 10 + \cdots + 10^{m-1} = \frac{10^m-1}{9}$, which is the *definition* of a repunit written as arithmetic. Once the set is "$R_m = \phi(10^m)$" with ϕ affine, conjugating by ϕ is forced, and affine conjugation preserves being a polynomial and preserves degree. That is the whole reason Kedlaya's solution opens with $g(n) = 9f\left(\frac{n-1}{9}\right) + 1$: it is the problem restated in the coordinates where the set is multiplicatively closed.

Why does the limit argument work? Look at what is special about powers of 10 that repunits hide: they form a *multiplicative* group (up to sign of exponent), and scaling by x^d respects that. So the ratio $g(10^k)/10^{kd}$ stays inside the same discrete set. Polynomials are dominated by their leading term at infinity (the real-analytic fact), while the target set is discrete away from 0 (the arithmetic fact). Put the two together and the leading term must be *exactly* what g equals, not only asymptotically. The final step is the finite-degree rigidity of polynomials: two polynomials agreeing at infinitely many points are equal. Every step is a definition or a one-line fact; nothing is a rabbit.

Why was the natural attempt's guess "only linear" wrong? Because in the n-coordinate, the monomial x^2 in the x-coordinate looks like $\frac{(9n+1)^2-1}{9} = 9n^2 + 2n$, which has no visible relation to repunits. The structure is invisible in the wrong coordinate — which is exactly why you change it.

Layer $+\alpha$ BACKTESTING

1. The first cue is the definition "digits all ones". Write down $R_m = \frac{10^m-1}{9}$, and then immediately the inverse: $9R_m + 1 = 10^m$.

2. Conjugate the unknown function by the affine map that turns the awkward set into a clean one, then classify maps preserving the clean set.

3. Hesitation comes after the change of variable: "fine, g sends powers of 10 to powers of 10, now what?" Reinforce with the growth reflex: evaluate along the set, divide by the leading power, and let $k \to \infty$.

4. The trap is proving "f must be linear" from small examples. The quadratic $9n^2 + 2n$ refutes it. A second trap is forgetting negative c: the real-coefficient hypothesis admits $c < 0$ provided $c + d \geq 1$.

5. Slogan: *a set given by a digit pattern is usually an affine image of a geometric progression; move there first.*

6. Decision rule: **when a polynomial must map a set S into itself, first find the coordinate in which S is a geometric progression or a lattice; then compare leading terms along S and use discreteness to turn an asymptotic statement into an identity.**

▷ **Cue:** a set defined by a repeated-digit pattern (repunits, $777\ldots7$, $1010\ldots$) $\longrightarrow$ **Move:** write it as an affine image of $\{b^m\}$ via the geometric-series formula and conjugate the problem into that coordinate

- ▷ **Cue:** a polynomial takes values in a discrete set (powers of 10, integers) along a sequence tending to infinity ⟶ **Move:** divide by the leading power, pass to the limit; discreteness makes the ratio eventually constant, and infinitely many agreements force a polynomial identity
- ▷ **Cue:** "find all polynomials" with a set-preservation hypothesis ⟶ **Move:** guess the family in the clean coordinate (monomials $10^c x^d$), then translate back; do not guess in the original coordinate

2007 A5

2007 A5. Suppose that a finite group has exactly n elements of order p, where p is a prime. Prove that either $n = 0$ or p divides $n + 1$.

Fields: group theory (Lagrange, group actions, orbit counting) **Position:** A5 — late, but short once the right set is chosen; the whole problem is the choice of what to count. **Answer:** If $p \nmid |G|$ then $n = 0$ by Lagrange; otherwise the $n + 1$ solutions of $x^p = e$ are the fixed points of a cyclic shift on a set of size $|G|^{p-1}$, so $p \mid n + 1$.

Layer 1 THE WHAT — surface analysis

Notation: G is the group, $m = |G|$, e the identity. An element a satisfies $a^p = e$ if and only if its order divides p, i.e. its order is 1 or p; only e has order 1. So

$$N := \#\{a \in G : a^p = e\} = n + 1.$$

The target is: $n = 0$, or $p \mid N$.

The natural attempt. The standard first move is to organise elements of order p by the subgroups they generate. If a has order p, then $\langle a \rangle$ is cyclic of order p, and every non-identity element of $\langle a \rangle$ also has order p and generates the same subgroup. Two distinct subgroups of order p meet only in e (their intersection is a subgroup of a group of prime order, and is not the whole thing). So the elements of order p are partitioned into blocks of size $p - 1$:

$$n = k(p - 1),$$

where k is the number of subgroups of order p. This is correct and is a real fact. Now $n + 1 = k(p - 1) + 1 \equiv 1 - k \pmod{p}$, so the problem is equivalent to

$$k \equiv 1 \pmod{p} \quad (\text{whenever } k \geq 1).$$

And here you hit the wall. "The number of subgroups of order p is 1 mod p" is a genuine theorem (it follows from the Sylow theorems and a further counting argument), but it is not easier than the problem — it *is* the problem, rephrased. You try to count subgroups by letting G act on them by conjugation; orbits have size dividing m, but m may be divisible by p many times and by other primes, and the orbit sizes do not organise mod p. The attempt stalls with a correct reduction and no mechanism.

The model chain. *Step 1: dispose of $p \nmid m$.* By Lagrange's theorem the order of every element divides m. If $p \nmid m$, no element has order p, so $n = 0$. Assume from now on that $p \mid m$.

Step 2: the set. Let

$$S = \{(a_0, a_1, \ldots, a_{p-1}) \in G^p : a_0 a_1 \cdots a_{p-1} = e\}.$$

Choose $a_0, \ldots, a_{p-2}$ freely; then $a_{p-1} = (a_0 \cdots a_{p-2})^{-1}$ is forced. So $|S| = m^{p-1}$, which is divisible by p (as $p \mid m$ and $p - 1 \geq 1$).

Step 3: the action. Define the cyclic shift $\sigma(a_0, a_1, \ldots, a_{p-1}) = (a_1, \ldots, a_{p-1}, a_0)$. It maps S to S: if $a_0 a_1 \cdots a_{p-1} = e$ then

$$a_1 \cdots a_{p-1} a_0 = a_0^{-1}(a_0 a_1 \cdots a_{p-1}) a_0 = a_0^{-1} e\, a_0 = e.$$

Since σ^p is the identity, the group $\{\mathrm{id}, \sigma, \ldots, \sigma^{p-1}\}$ of order p acts on S, and every orbit has size dividing p: size 1 or size p.

Step 4: the fixed points. A tuple is fixed by σ iff $a_0 = a_1 = \cdots = a_{p-1}$, i.e. it is $(a, a, \ldots, a)$, and this lies in S iff $a^p = e$. So the number of orbits of size 1 is exactly N.

Step 5: count mod p. S is the disjoint union of its orbits, so

$$m^{p-1} = |S| = N + p \cdot (\text{number of orbits of size } p) \equiv N \pmod{p}.$$

Since $p \mid m^{p-1}$, $p \mid N = n + 1$.

Caveat. Sanity-checked by machine on S_3, S_4, S_5 with $p \in \{2, 3, 5\}$: e.g. S_4 has 9 elements of order 2 and 8 of order 3, and S_5 has 24 of order 5; in each case with $n > 0$, $p \mid n + 1$. The argument is McKay's proof of Cauchy's theorem, exactly as in the first official solution.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on $N = n + 1$: "the elements with $a^p = e$ are e together with the n elements of order p." At that point the natural attempt groups the n elements *by the subgroups they generate*, which converts the problem into counting subgroups. The model instead asks for a set on which a group of order p acts whose fixed points are exactly the solutions of $x^p = e$ — and builds that set.

> *I need to prove a number is* 0 mod p. *The one general machine that produces congruences mod p is a group of order p acting on a finite set: orbits are size* 1 *or p, so $|\text{set}| \equiv |\text{fixed points}| \pmod{p}$. So I want a set with a $\mathbb{Z}/p$ action whose fixed points are the a with $a^p = e$, and whose size I can compute.*

> *Fixed points of a cyclic shift are constant tuples $(a, \ldots, a)$. The condition $a^p = e$ reads $a \cdot a \cdots a = e$, a product of p factors. So take all p-tuples with product e; the shift preserves the product up to conjugation, and I can count them because the last entry is forced.*

Layer 3 THE WHY

Why must you reach for an action of $\mathbb{Z}/p$? Because of the *shape* of the conclusion. "p divides N" is a statement about a count modulo a prime, with no structure on G assumed beyond finiteness. There is essentially one tool in finite group theory that converts "finite set" into "count mod p" without knowing the set's structure: the orbit decomposition under a p-group, where every orbit has size a power of p. With a group of order exactly p, orbits are 1 or p, and so $|X| \equiv |X^{\text{fixed}}| \pmod{p}$. That is the circle-and-centre moment: the conclusion is a congruence mod p, so the object of interest is a set acted on by something of order p. Once you know that is the tool, the only question is which set.

Why this set? Work backwards from the fixed points you need. The fixed points of a cyclic shift on G^p are the constant tuples. You need the constant tuple $(a, \ldots, a)$ to count exactly when $a^p = e$. Read $a^p = e$ as "the product of the p entries is e". The largest natural shift-invariant set with that property is all p-tuples with product e. Its size is forced by the group axioms (existence of inverses makes the last entry determined), giving m^{p-1}. Nothing here is guessed: each choice is dictated by the one before.

Why does shift-invariance hold when G is not abelian? Because a cyclic rotation of a product is a conjugate of it: $a_1 \cdots a_{p-1} a_0 = a_0^{-1}(a_0 \cdots a_{p-1})a_0$. The identity is fixed by conjugation. Say this line explicitly; a grader will look for it.

And why was the natural attempt's reduction not wrong but useless? It replaced one congruence (on elements) by another (on subgroups) without supplying the mod-p mechanism. The subgroup count $k \equiv 1 \pmod{p}$ is true, but proving it also needs a p-group action — the same idea you were avoiding. When a reformulation leaves the core difficulty intact, abandon it and find the mechanism first.

Kedlaya's first solution begins "let S be the set of p-tuples with product e". Here is where that comes from: it is the smallest set on which $\mathbb{Z}/p$ acts naturally by rotation and whose rotation-fixed points are the solutions of $x^p = e$.

Layer $+\alpha$ BACKTESTING

1. The first cue is "p divides $n+1$" with p prime. Write down $n+1 = \#\{a : a^p = e\}$, so the target is a count of solutions of $x^p = e$, modulo p.

2. A congruence mod a prime p for the size of a set suggests: find an action of a group of order p with the target set as the fixed points. The orbit-stabiliser count then gives the congruence.

3. Hesitation occurs after $n = k(p-1)$: it feels like progress. Reinforce with the check "does this reduction come with a mechanism for the new statement?" If not, back out.

4. The trap is the subgroup reduction, which is correct and leads into Sylow-type counting that is harder than the problem. A smaller trap is forgetting the case $p \nmid |G|$, where $|S| = m^{p-1} \equiv 1 \pmod{p}$ and the argument gives only $n \equiv 0$; Lagrange is needed to conclude $n = 0$.

5. Slogan: *"count $\equiv$ something mod p" means "fixed points of a $\mathbb{Z}/p$ action"; cyclic shifts of p-tuples are the default $\mathbb{Z}/p$ action.*

6. Decision rule: **to prove $p \mid |X|$, realise X (or X plus a set of known size) as the fixed points of a group of order p acting on a set whose total size you can compute; the difference is a sum of orbits of size p.** The same rule proves Fermat's little theorem (necklaces of p beads in a colours) and Wilson's theorem.

▷ **Cue:** a statement "p divides the number of objects with property P", p prime $\longrightarrow$ **Move:** find a $\mathbb{Z}/p$ action whose fixed points are the objects with property P; compare with the total size mod p

▷ **Cue:** a condition of the form $a^p = e$ or "the same thing p times" $\longrightarrow$ **Move:** encode it as a constant p-tuple and use the cyclic shift on p-tuples

▷ **Cue:** a product $a_0 a_1 \cdots a_{p-1}$ in a non-abelian group must be rotated $\longrightarrow$ **Move:** a rotation of a product is a conjugate of it, so "product $= e$" is rotation-invariant

▷ **Cue:** counting elements of a given order in a group $\longrightarrow$ **Move:** elements of order p come in blocks of $p-1$ per subgroup — useful for sizes, but not a mod-p mechanism

2007 A6

2007 A6. A *triangulation* $\mathcal{T}$ of a polygon P is a finite collection of triangles whose union is P, and such that the intersection of any two triangles is either empty, or a shared vertex, or a shared side. Moreover, each side is a side of exactly one triangle in $\mathcal{T}$. Say that $\mathcal{T}$ is *admissible* if every internal vertex is shared by 6 or more triangles. For example, [figure omitted.] Prove that there is an integer M_n, depending only on n, such that any admissible triangulation of a polygon P with n sides has at most M_n triangles.

Fields: combinatorics, combinatorial geometry (Euler's formula, discrete curvature, induction on boundary length) **Position:** A6 — genuinely hard; the Euler count is standard, but the induction step needs a careful choice of what to peel off, and the official write-up leaves real gaps in it. **Answer:** $M_n = (n-2)^2$ works: an Euler count forces the boundary degrees to sum to at most $4n-6$, which guarantees either a

chord to cut along or a run of boundary degrees $3, 4, \dots, 4, 3$ whose strip can be peeled off, shortening the boundary by one.

Layer 1 THE WHAT — surface analysis

Set-up and notation. Regard $\mathcal{T}$ as a triangulated disk. Its *boundary vertices* are the vertices of $\mathcal{T}$ on ∂P; because each side of P is a side of exactly one triangle, no vertex lies inside a side, so the boundary vertices are the n corners $v_1, \dots, v_n$ (indices mod n), and the boundary has exactly n edges. Every other vertex is *internal*. Write a_i for the number of edges at v_i. Around an internal vertex the triangles close up into a cycle, so "shared by ≥ 6 triangles" means degree ≥ 6; around a boundary vertex they form a fan, so v_i lies in $a_i - 1$ triangles and $a_i \geq 2$. An edge joining two boundary vertices that is not a side of P is a *chord*.

For the induction it is convenient to prove the statement for any triangulated region bounded by a simple closed polygonal curve with n vertices of $\mathcal{T}$ on it, allowing straight angles at boundary vertices; the problem is the special case in which all n are genuine corners.

The natural attempt. The picture that drives you is the triangular lattice: every internal vertex of degree exactly 6, flat, and a hexagonal patch of it has boundary about $6r$ and about $6r^2$ triangles. So you expect $M_n \approx cn^2$, and you expect the proof to say "degree ≥ 6 means the surface is flat or saddle-shaped, so it cannot bulge".

You start with Euler. Let V, E, F count vertices, edges, triangles. Then $V - E + F = 1$ and $3F = 2E - n$, which give $F = 2V - n - 2$. So bounding triangles is the same as bounding internal vertices. Now you need a second equation involving the degree condition, and you write $2E = \sum a_i + \sum_{\text{internal}} \deg \geq \sum a_i + 6(V - n)$. Combining,

$$\sum_{i=1}^{n} a_i \leq 4n - 6. \tag{†}$$

And there you stop, because (†) bounds the boundary degrees, not V. Eliminating differently only returns (†) — the internal vertices have cancelled completely. Euler has told you something true about the boundary and nothing about the inside. A second natural idea, induction on the number of internal vertices by deleting one, fails at once: deleting an internal vertex leaves a hole that must be retriangulated, and retriangulating can create internal vertices of degree less than 6.

The model chain. *Lemma.* In any admissible triangulated disk with n boundary vertices, $\sum a_i \leq 4n - 6$. (Proof: the Euler computation above. Explicitly, $2E \geq \sum a_i + 6(V - n)$ and $V = 1 + E - F$ give $\sum a_i \leq 6F - 4E + 6n - 6 = 2(3F - 2E) + 6n - 6 = 4n - 6$.)

Claim. Such a disk has at most $(n-2)^2$ triangles. Strong induction on $n \geq 3$.

Base $n = 3$. By the Lemma $a_1 + a_2 + a_3 \leq 6$ with each $a_i \geq 2$, so every $a_i = 2$: each corner lies in one triangle, the one formed by its two boundary edges, which is $v_1v_2v_3 = P$. One triangle, $= (3-2)^2$.

Case A: there is a chord v_iv_j. It splits the disk into two triangulated disks with n_1, n_2 boundary vertices, $n_1 + n_2 = n + 2$, $n_1, n_2 \geq 3$ (so both $< n$). Internal vertices of each piece are internal in $\mathcal{T}$ with all their triangles, so both pieces are admissible, and by induction the number of triangles is at most

$$(n_1 - 2)^2 + (n_2 - 2)^2 \leq (n_1 + n_2 - 4)^2 = (n-2)^2.$$

Note that if $n \geq 4$ and some $a_i = 2$, the triangle at v_i is $v_{i-1}v_iv_{i+1}$ and $v_{i-1}v_{i+1}$ is a chord; so this case includes "ears".

Case B: no chord, $n \geq 4$. Then every $a_i \geq 3$. Let t be the number of i with $a_i = 3$ and s the number with $a_i \geq 5$. Then $\sum a_i \geq 3t + 4(n - t - s) + 5s = 4n - t + s$, and the Lemma gives $t \geq s + 6$. The t vertices of degree 3 cut the boundary cycle into t gaps; only $s < t$ vertices have degree ≥ 5, so some gap contains none. Relabel so that gap gives a run

$$a_1 = 3, \quad a_2 = \cdots = a_{k-1} = 4, \quad a_k = 3, \qquad 2 \leq k \leq n - t + 2 \leq n - 4.$$

Reading off the strip. The fan at v_1 has vertices v_0, w_1, v_2 for a single interior neighbour w_1. For $2 \le j \le k-1$, the fan at v_j is v_{j-1}, x, y, v_{j+1}; the unique triangle on the edge $v_{j-1}v_j$ is $v_{j-1}v_jw_{j-1}$, so $x = w_{j-1}$, and we name $y = w_j$. The fan at v_k is $v_{k-1}, w_{k-1}, v_{k+1}$. The triangles containing some $v_1, \dots, v_k$ are therefore

$$v_0v_1w_1; \quad v_jv_{j+1}w_j \ (1 \le j \le k-1); \quad v_jw_{j-1}w_j \ (2 \le j \le k-1); \quad v_kw_{k-1}v_{k+1},$$

which is $1 + (k-1) + (k-2) + 1 = 2k-1$ triangles.

The w_j are internal and distinct. Each w_j ($1 \le j \le k-1$) is adjacent to both v_j and v_{j+1} and differs from both. Since $n \ge 4$, no boundary vertex other than v_j, v_{j+1} is a boundary-neighbour of both, so if w_j were a boundary vertex, one of v_jw_j, $v_{j+1}w_j$ would be a chord. Hence all w_j are internal. Consecutive w_j, w_{j+1} differ (they share a triangle with v_{j+1}). Suppose $w_i = w_j =: w$ with $j \ge i+2$. The cycle $Z = v_{i+1}v_{i+2} \cdots v_j\, w$ is simple, of length $n' = j - i + 1 \ge 3$, and bounds a sub-disk D that is an admissible triangulated disk. In D, the vertices $v_{i+2}, \dots, v_{j-1}$ keep degree 4; v_{i+1} keeps neighbours w, w_{i+1}, v_{i+2} (degree 3); v_j keeps v_{j-1}, w_{j-1}, w (degree 3); w has degree ≥ 2. So the boundary degrees of D sum to at least $4(n'-3) + 3 + 3 + 2 = 4n' - 4 > 4n' - 6$, contradicting the Lemma.

Peel. Remove the $2k-1$ triangles. What remains is bounded by the simple cycle $v_0, w_1, \dots, w_{k-1}, v_{k+1}, \dots, v_{n-1}$ with $(n-k) + (k-1) = n-1$ vertices. Every remaining internal vertex is internal in $\mathcal{T}$ and lies in no removed triangle, so admissibility survives. By induction the count is at most

$$(n-3)^2 + 2k - 1 \le (n-3)^2 + 2(n-4) - 1 = (n-2)^2 - 4 < (n-2)^2.$$

So $M_n = (n-2)^2$ works.

Caveat. Two points where this chain departs from the official solution, both deliberate. (i) The official text writes "$F = 2E - n$"; the correct incidence count is $3F = 2E - n$, and the displayed conclusion $4n - 6$ is right only with the correct version. (ii) The official induction removes the $2k-1$ strip triangles and asserts that what is left is an admissible triangulation of an $(n-1)$-gon. That needs the w_j to be internal and distinct; if some w_j were a corner, or two w_j coincided, the remainder would be pinched and not a polygon. The chain above closes this by cutting along chords first (Case A) and ruling out coincident w_j with the Lemma applied to a sub-disk. I have checked each step by hand; I have not machine-verified the bound on explicit triangulations. The bound $(n-2)^2$ is not sharp: the official remark refines it to about $n^2/6$, attained asymptotically by hexagonal patches of the triangular lattice.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains derive (†): $\sum a_i \le 4n - 6$. At (†) the natural attempt treats the inequality as a failed attempt to bound V and goes looking for a different equation. The model reads (†) as a statement about *where the triangulation can be cut*: the average boundary degree is below 4, so the boundary has places where very few triangles touch it, and those places can be peeled off to shorten the boundary. The bound on the inside comes from the induction, not from Euler.

> *Euler killed the internal vertices, so it cannot bound them directly. But it did tell me the boundary is "thin": the average boundary degree is under* 4*. I'll induct on the boundary length n and use* (†) *only to find somewhere to peel.*

> *A degree-3 boundary vertex has one interior neighbour; a degree-4 one has two. If I take a stretch $3, 4, \ldots, 4, 3$ and remove every triangle touching it, the k old boundary vertices are replaced by $k-1$ interior ones. Net change -1. I need more 3s than 5s to guarantee such a stretch, and (†) gives exactly that.*

> *Before I peel, I must know that the new boundary is a simple cycle. The danger is an interior neighbour that is really a boundary vertex — a chord. So I'll split along chords first, and then use the same $4n-6$ inequality on a small sub-disk to show the w_j cannot repeat.*

Layer 3 THE WHY

Why is induction on the boundary the natural frame, rather than on internal vertices? Look at what the problem gives you. The hypothesis is local and one-sided — internal vertices have degree *at least* 6 — and the conclusion bounds the whole by the boundary length. The degree condition is a curvature condition: in the triangular lattice every vertex has exactly 6 triangles and angle sum $6 \times 60^\circ = 360^\circ$, so "$\geq 6$" means "flat or negatively curved". Euler's formula is the combinatorial Gauss–Bonnet theorem: total curvature is fixed by the topology. The inside can only carry zero or negative curvature, so all the positive turning that a closed disk needs must sit on the boundary. That is exactly what (†) says: the boundary degrees are forced to be small. Once you see (†) as "the boundary is where the turning lives", you stop trying to bound the inside with it and start using it to operate on the boundary.

Why a run $3, 4, \ldots, 4, 3$ and not a single low-degree vertex? A boundary vertex of degree a contributes $a-2$ interior neighbours to the next layer in. Degree 3 contributes one, degree 4 two, but consecutive vertices share an interior neighbour (the apex of the triangle on their common boundary edge). Along a run of length k the interior neighbours number $1 + 2(k-2) + 1 = 2k-2$ counted with repetition, and each of the $k-1$ boundary edges inside the run has an apex shared by its two endpoints, so there are $(2k-2) - (k-1) = k-1$ distinct ones: $k-1$ internal vertices replace k boundary ones. A single degree-3 vertex ($k=1$) would replace 1 vertex by 1; that makes no progress. Two 3s bracketing 4s is the shortest pattern that strictly shortens, and (†) is precisely strong enough to produce it: every 3 "pays" for one 5, and six 3s are left over.

Why must the chord case be separated? Because peeling is a surgery on the boundary, and surgery is only safe when the thing you cut to is really interior. A chord is the combinatorial form of "the disk is pinched": a single edge already separates it into two smaller disks. The induction handles that case by splitting, and the inequality $(n_1-2)^2 + (n_2-2)^2 \leq (n-2)^2$ is why the answer is chosen as a square: squares are superadditive, which is what splitting requires. The official solution's recurrence $A_n = A_{n-1} + 2n - 3$ grows like n^2 too, but it never states the splitting step, and that is where its argument has a gap.

Kedlaya's solution opens with $a_1 + \cdots + a_n \leq 4n-6$ and then "we will prove by induction on n". Here is where that comes from: Euler cannot see the inside, so the only lever left is the boundary length, and (†) says the boundary has cheap places to peel.

Layer $+\alpha$ BACKTESTING

1. The first cue is "every internal vertex is shared by 6 or more triangles" — a degree inequality on a planar triangulation. Write down $V - E + F = 1$, $3F = 2E - n$, and the degree sum $2E = \sum \deg$.

2. Euler plus the degree inequality, then induction on n, with the Euler inequality used as a supply of low-degree boundary configurations.

3. Hesitation comes when (†) fails to bound V. Reinforce with: "if a counting identity cancels the thing I want to bound, it is telling me about the other part of the structure; use that part as the lever."

4. Two traps. The first is expecting Euler alone to finish. The second is peeling without checking that the new boundary is a simple cycle; a chord or a repeated interior neighbour pinches the region. The official solution falls into the second silently.

5. Slogan: *degree ≥ 6 in a planar triangulation is non-positive curvature; Euler then puts all the positive curvature on the boundary, so the boundary has vertices of small degree you can peel.*

6. Decision rule: **when a global count cancels the quantity you want to bound, turn it into a local existence statement (an average below a threshold forces a cheap configuration), then induct by removing that configuration, after first splitting off any degenerate case in which the removal would disconnect the object.** Choose the bound superadditive so that splitting is free.

▷ **Cue:** a planar triangulation with every internal vertex of degree ≥ 6 $\longrightarrow$ **Move:** Euler's formula gives boundary degree sum $\leq 4n - 6$: the boundary carries the curvature

▷ **Cue:** an average below an integer threshold (mean boundary degree < 4) $\longrightarrow$ **Move:** count values below and above: there are more 3s than 5s, so some gap between consecutive 3s contains only 4s

▷ **Cue:** induction by peeling off a boundary layer $\longrightarrow$ **Move:** first cut along chords (they split the disk); then check the new boundary is a simple cycle before invoking the hypothesis

▷ **Cue:** an induction that must survive splitting into pieces of sizes $n_1 + n_2 = n + 2$ $\longrightarrow$ **Move:** choose a superadditive bound such as $(n-2)^2$

Putnam 2007 — Paper B

2007 B1

2007 B1. Let f be a polynomial with positive integer coefficients. Prove that if n is a positive integer, then $f(n)$ divides $f(f(n)+1)$ if and only if $n = 1$. [Editor's note: one must assume f is nonconstant.]

Fields: number theory (polynomial congruences) **Position:** B1 – opener **Answer:** Yes: reduce $f(f(n)+1)$ modulo $f(n)$ to $f(1)$ using $a \equiv b \pmod{m} \Rightarrow P(a) \equiv P(b) \pmod{m}$; then use strict monotonicity of f on positive integers (forced by positive coefficients) to show $0 < f(1) < f(n)$ for $n > 1$, so divisibility holds only at $n = 1$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The instinct on seeing "$f(n)$ divides $f(f(n)+1)$" is to try to see the phenomenon directly, so the honest first move is to test it on the cheapest possible polynomials. Take $f(n) = n + c$ for a positive integer c: then $f(n)+1 = n+c+1$, and $f(f(n)+1) = n+2c+1$. Does $n+c$ divide $n+2c+1$? Subtracting, $n+2c+1-(n+c) = c+1$, so $f(n) \mid f(f(n)+1)$ iff $(n+c) \mid (c+1)$, which for $n = 1$ gives $(1+c) \mid (c+1)$, true, and for $n > 1$ gives $n+c > c+1$ (since $n \geq 2$), so it fails unless $c+1 = 0$, impossible. The linear case checks out and even reveals the right shape of argument – a subtraction that leaves something small and fixed. Emboldened, the natural attempt tries $f(n) = n^2+n+1$ directly by brute expansion: $f(n)+1 = n^2+n+2$, and $f(f(n)+1) = (n^2+n+2)^2 + (n^2+n+2) + 1$. Expanding this out in powers of n and then trying to extract a factor of n^2+n+1 from the resulting degree-4 polynomial is possible but already unpleasant – a full polynomial division, or a careful regrouping, is needed just to see that the remainder is $f(1) = 3$. Scaling this to a general polynomial of degree d with unspecified coefficients $a_0, \ldots, a_d$ is where the approach genuinely stops being viable: expanding $f(f(n)+1) = \sum a_i (f(n)+1)^i$ symbolically and dividing by $f(n)$ term by term, for arbitrary d, is a multinomial mess with no natural stopping point. The attempt has correctly sensed that something always reduces to a small leftover constant, but is trying to discover that constant by brute expansion instead of by naming the mechanism that produces it.

The model chain. Write $f(n) = \sum_{i=0}^{d} a_i n^i$ with every $a_i > 0$ an integer (and $a_d > 0$, $d \geq 1$ since f is nonconstant). Work modulo $f(n)$. Since $f(n)+1 \equiv 1 \pmod{f(n)}$, and f has integer coefficients, substituting a congruent value into f preserves the congruence:

$$f(f(n)+1) = \sum_{i=0}^{d} a_i (f(n)+1)^i \equiv \sum_{i=0}^{d} a_i \cdot 1^i = f(1) \pmod{f(n)}.$$

So $f(n) \mid f(f(n)+1)$ if and only if $f(n) \mid f(1)$. If $n = 1$, this reads $f(1) \mid f(1)$, always true. If $n > 1$: because every coefficient a_i is a positive integer and $d \geq 1$, comparing term by term shows f is strictly increasing on the positive integers (each $a_i n^i > a_i \cdot 1^i$ for $n > 1$, $i \geq 1$, and the $i = 0$ term is unchanged), so $f(n) > f(1)$. Also $f(1) = \sum a_i > 0$. A strictly larger positive integer cannot divide a strictly smaller positive integer, so $f(n) \nmid f(1)$, hence $f(n) \nmid f(f(n)+1)$. This proves the claim in both directions.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree that $f(f(n)+1)$ reduces to something small when compared against $f(n)$ – the linear case already shows the leftover is a fixed constant independent of n. They split on how that reduction is obtained. The natural attempt tries to *compute* the leftover by expanding $f(f(n)+1)$ as an explicit polynomial in n and extracting the remainder on division by $f(n)$, case by case in the degree. The model chain instead recognises that "remainder on division by $f(n)$" is exactly what the congruence $f(n)+1 \equiv 1 \pmod{f(n)}$ is already telling it, and evaluates the remainder by substitution rather than by division.

> *I do not need to expand $f(f(n)+1)$ as a polynomial in n at all – I only need its value modulo $f(n)$, and since $f(n)+1 \equiv 1 \pmod{f(n)}$, and f is a polynomial with integer coefficients, I can substitute the residue 1 directly into f to get $f(f(n)+1) \equiv f(1) \pmod{f(n)}$, for any degree, with no expansion at all.*

> *Once the problem has been reduced to "does $f(n)$ divide $f(1)$", the only tool left is a size comparison, not more algebra – and positive coefficients are exactly the hypothesis that makes f strictly increasing on positive integers, so $f(n) > f(1) > 0$ for every $n > 1$, which alone kills divisibility.*

Layer 3 THE WHY

The substitution step is not a trick borrowed for this problem; it is the defining property of polynomial congruences. For any polynomial P with integer coefficients and any integers a, b, the factor theorem gives $a - b \mid P(a) - P(b)$ (write $P(x) - P(x-1) \cdots$, or note $a^i - b^i = (a-b)(a^{i-1} + \cdots + b^{i-1})$ termwise). So $a \equiv b \pmod{m}$ forces $P(a) \equiv P(b) \pmod{m}$ for *every* integer polynomial P, no matter its degree or coefficients. Here $a = f(n)+1$, $b = 1$, $m = f(n)$: the difference $a - b = f(n)$ is divisible by $m = f(n)$ by construction, so the congruence is free, not earned by computation. Recognising this converts "expand and divide" (an operation whose cost scales with degree) into "look up the residue and substitute" (an operation whose cost does not). This is the same asymmetry the archive keeps returning to: modular arithmetic exists precisely so that a polynomial's behaviour at one representative of a residue class stands in for its behaviour at every element of that class, and $f(n)+1$ and 1 are, from $f(n)$'s point of view, the same element.

The second half of the proof – why $n > 1$ fails – is where the hypothesis "positive integer coefficients" earns its keep beyond the congruence step. It is not used to make the algebra work; it is used to make the *comparison* work. Without positivity, $f(n)$ and $f(1)$ need not be ordered at all (a polynomial with mixed-sign coefficients can be non-monotonic, or even take negative values), and "$f(n)$ divides $f(1)$" would have to be checked by actual size, possibly failing to resolve cleanly. Positivity of every coefficient is exactly the condition under which increasing the input strictly increases every term (for $i \geq 1$) and leaves no term to decrease it, so f restricted to positive integers is strictly increasing – a genuinely structural fact about the hypothesis, not an incidental one. Once that is seen, "$f(n) \mid f(1)$ for $n > 1$" is not a case to check; it is impossible on sight, because a divisor of a positive integer is at most that integer.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape "$f(f(n)+1)$" appearing where a divisibility by $f(n)$ is being asked about. On seeing it, write down immediately that $f(n)+1 \equiv 1 \pmod{f(n)}$ – the argument of the outer f is congruent to a fixed constant modulo the very quantity being tested for division.
2. That cue suggests working entirely modulo $f(n)$: substitute the residue into f using $a \equiv b \pmod{m} \Rightarrow P(a) \equiv P(b) \pmod{m}$, collapsing $f(f(n)+1)$ to $f(1)$ without expanding any polynomial.

3. Hesitation occurs when trusting that a substitution is legitimate without expansion – it is checked away by recalling the factor-theorem fact $a - b \mid P(a) - P(b)$ for integer-coefficient P, which holds for every degree and every coefficient set, so no case analysis on $\deg f$ is ever needed.

4. The trap is trying to verify the reduction by explicit polynomial division on small-degree examples and then treating that as a proof strategy that must be repeated (and generalised by main force) for arbitrary degree, instead of naming the congruence fact once and applying it uniformly.

5. Recognisable faster next time by treating any expression "argument $\equiv$ constant modulo the modulus in question" as an instant licence to replace the argument by that constant inside any integer-coefficient polynomial, before doing any expansion.

6. Reusable rule: to decide whether $m \mid P(g(m))$ for some expression g with $g(m) \equiv c \pmod m$, reduce to whether $m \mid P(c)$ via polynomial congruence; then resolve the remaining (constant-vs-m) divisibility by a size/monotonicity argument rather than more algebra.

▷ **Cue:** an expression "f(something $\equiv c \bmod f(n)$)" appears inside a claim about divisibility by $f(n)$ $\longrightarrow$ **Move:** substitute the residue c directly into f using $a \equiv b \pmod m \Rightarrow P(a) \equiv P(b) \pmod m$ for integer-coefficient polynomials, instead of expanding

▷ **Cue:** a divisibility question has been reduced to "does a fixed value divide $f(1)$ / another fixed value" depending on n $\longrightarrow$ **Move:** resolve it by monotonicity: positive coefficients make f strictly increasing on positive integers, so the larger quantity cannot divide into the smaller except when they coincide

2007 B2

2007 B2. Suppose that $f : [0,1] \to \mathbb{R}$ has a continuous derivative and that $\int_0^1 f(x)\,dx = 0$. Prove that for every $\alpha \in (0,1)$,

$$\left| \int_0^\alpha f(x)\,dx \right| \leq \frac{1}{8} \max_{0 \leq x \leq 1} |f'(x)|.$$

Fields: real analysis (calculus / optimisation bound) **Position:** B2 **Answer:** Yes: with $g(x) = \int_0^x f$, g vanishes at both endpoints, so its maximum in absolute value occurs at an interior point y with $f(y) = 0$; a one-sided derivative bound then gives $f(x) \leq B(\alpha - x)$ on $[0,\alpha]$, and a symmetry reduction to $\alpha \leq 1/2$ turns $\int_0^\alpha B(\alpha - x)\,dx = B\alpha^2/2$ into the stated $B/8$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The given data are $\int_0^1 f = 0$ and a bound $B := \max|f'|$ on the derivative, and the target is a bound on $\int_0^\alpha f$ for a *specific* α. The natural first move is to attack that α directly: bound $|f(x)|$ pointwise on $[0,\alpha]$ and integrate. To bound f pointwise one needs a reference point where f's value is known, and $\int_0^1 f = 0$ is exactly the fact that guarantees one exists – since f is continuous and its average over $[0,1]$ is zero, either $f \equiv 0$ (trivial) or f takes both a positive and a non-positive value on $[0,1]$, so by the intermediate value theorem there is some $c \in [0,1]$ with $f(c) = 0$. From there, $|f(x) - f(c)| \leq B|x - c|$ by the mean value theorem (or by integrating f'), giving $|f(x)| \leq B|x - c|$ for every x. This is a genuinely correct and useful fact. The trouble is what to do with it: c is just *some* zero of f, produced by an existence argument with no control over its location, and it need not lie in $[0,\alpha]$ at all. If $c > \alpha$, the bound $|f(x)| \leq B|x - c|$ on $[0,\alpha]$ is weaker than it needs to be (every point is further from c than it would be from a zero actually inside $[0,\alpha]$), and integrating it over $[0,\alpha]$ does not obviously collapse to $B\alpha^2/2$, let alone $B/8$ – the bound produced depends on the unknown location of c

relative to α, and pinning it down forces casework on where c falls that never quite closes into the clean constant $1/8$ the problem asks for.

The model chain. Let $B = \max_{0\le x\le 1} |f'(x)|$ and define $g(x) = \int_0^x f(t)\,dt$. Then $g(0) = 0$ and $g(1) = \int_0^1 f = 0$. The quantity to bound, $|\int_0^\alpha f| = |g(\alpha)|$, is bounded for *every* $\alpha \in (0,1)$ as soon as it is bounded at the point $y \in (0,1)$ where $|g|$ attains its maximum over $[0,1]$ (this maximum is attained, and at an interior point, since g is continuous with $g(0) = g(1) = 0$; if $g \equiv 0$ the claim is trivial). At that interior maximum, g is differentiable with $g'(y) = f(y)$, and Fermat's theorem on interior extrema gives $f(y) = 0$. Hence it suffices to bound $|g(y)|$ using this specific zero $y = \alpha$ of f. Since $\int_0^\alpha f(x)\,dx = -\int_0^{1-\alpha} f(1-u)\,du$ (substitute $x = 1-u$ and use $\int_0^1 f = 0$), and $|f'|$ is unchanged under $x \mapsto 1-x$, we may replace $f(x)$ by $f(1-x)$ and assume $\alpha \le 1/2$; replacing f by $-f$ if necessary, assume also $\int_0^\alpha f \ge 0$. Since $f'(x) \ge -B$ for all x and $f(\alpha) = 0$, integrating from x to α gives $f(\alpha) - f(x) = \int_x^\alpha f'(t)\,dt \ge -B(\alpha - x)$, i.e.

$$f(x) \le B(\alpha - x), \qquad 0 \le x \le \alpha.$$

Therefore

$$\int_0^\alpha f(x)\,dx \le \int_0^\alpha B(\alpha - x)\,dx = \frac{B\alpha^2}{2} \le \frac{B}{8},$$

using $\alpha \le 1/2$. This proves the bound at the worst case, hence at every $\alpha \in (0,1)$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths need a point where f vanishes, and both use $\int_0^1 f = 0$ to get one. They split on *which* zero is used and *why* it is the right one. The natural attempt finds a zero c of f by the intermediate value theorem – an existence statement with no relationship to the specific α being bounded, so its location relative to $[0,\alpha]$ is unknown and must be handled by casework. The model chain instead does not go looking for a zero of f at all: it observes that the inequality must hold for *every* α, reduces to bounding the single worst-case value of $g(\alpha) = \int_0^\alpha f$, and gets a zero of f for free, at precisely that worst-case point, as a consequence of g being an interior extremum of a function vanishing at both endpoints.

> *The inequality has to hold for every α in $(0,1)$ simultaneously, so instead of trying to control the integral at the specific α I am handed, I should first find the α that makes $|\int_0^\alpha f|$ as large as possible and prove the bound there – if it holds at the worst case, it holds everywhere, and I only have to do the work once.*

> *The worst-case α is an interior maximum of $g(x) = \int_0^x f$, and g vanishes at both endpoints of $[0,1]$ by hypothesis and by construction, so Fermat's theorem hands me $g'(y) = f(y) = 0$ at that maximum for free – I do not need to go hunting for a zero of f by the intermediate value theorem; calculus already places one exactly where I need it.*

> *Replacing x by $1-x$ and, if needed, f by $-f$ costs nothing (both preserve $\max|f'|$ and the hypothesis $\int_0^1 f = 0$), so I am free to assume $\alpha \le 1/2$ and $\int_0^\alpha f \ge 0$ – this halves the range of α I have to check and lets me use the cheapest possible one-sided bound $f'(x) \ge -B$ instead of a symmetric two-sided estimate.*

Layer 3 THE WHY

The hypothesis $\int_0^1 f(x)\,dx = 0$ is not really a fact about f; it is a fact about f's antiderivative g, and specifically about g evaluated at $x = 1$: $g(1) = \int_0^1 f = 0$. Since $g(0) = 0$ trivially (empty integral), the hypothesis is exactly the statement "g vanishes at both endpoints of $[0,1]$" – the

canonical setup for Rolle's theorem or, more precisely here, for the extreme value theorem plus Fermat's theorem, since g is differentiable and continuous on a closed bounded interval. Once the target quantity $\int_0^\alpha f$ is recognised as $g(\alpha)$, and the problem is recognised as asking for a bound valid for *every* α, the reduction to g's own maximum is not a trick: it is simply the observation that "true for all α" is implied by "true at the value of α that makes the left side largest", and that value comes with a free structural fact (its derivative vanishes) attached, because it is an interior extremum of a function that is zero on the boundary. This is why the official solution can open by defining g and immediately restricting to a critical point without any justification beyond "the maximum value of $|g(x)|$ must occur at a critical point $y \in (0,1)$ satisfying $g'(y) = f(y) = 0$" – that sentence is Fermat's theorem applied to the exact function the hypothesis was quietly describing all along.

The final constant $1/8$ is not an arbitrary target either: it is the maximum, over $\alpha \in [0, 1/2]$, of the parabola $B\alpha^2/2$ produced by integrating the one-sided linear bound $B(\alpha - x)$ – and that maximum occurs at the boundary $\alpha = 1/2$ of the reduced range, giving $B \cdot (1/2)^2/2 = B/8$ exactly. The symmetry step (replacing α by $1 - \alpha$) is what makes $\alpha = 1/2$ the true worst case reachable by the argument rather than an artificial restriction: without it, α could range up to 1, and the same one-sided bound would only give $B/2$, four times too weak. So the constant $1/8$ is the fingerprint of two decisions working together – reduce to a single worst-case zero of f, and reduce the range of that zero to $[0, 1/2]$ by reflection – and neither decision is visible in the final inequality unless both are named.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the hypothesis $\int_0^1 f = 0$ appearing alongside a request for a bound on $\int_0^\alpha f$ valid for *every* α. On seeing it, write down immediately that $g(x) := \int_0^x f$ satisfies $g(0) = g(1) = 0$ – the hypothesis is really a statement about g's endpoint values.

2. That cue suggests reducing the "for every α" claim to a single worst-case α: since g is continuous with $g(0) = g(1) = 0$, it attains an interior extremum, and Fermat's theorem gives a zero of f there for free, to be used as the reference point for a one-sided mean-value bound.

3. Hesitation occurs over whether bounding $|g|$ at one special point really covers every α – reinforced by recalling that the maximiser dominates by definition: if the bound holds where $|g|$ is largest, it holds everywhere else automatically.

4. The trap is producing *a* zero of f via the intermediate value theorem and treating it as usable for an arbitrary given α, when its location relative to $[0, \alpha]$ is unconstrained; a second trap is skipping the symmetry reduction to $\alpha \le 1/2$ and thereby losing a factor of up to 4 in the final constant.

5. Recognisable faster next time by treating any hypothesis of the form "the integral of f over the whole domain is zero" as shorthand for "the antiderivative vanishes at both endpoints", and immediately reaching for Fermat/Rolle on that antiderivative rather than searching for a zero of f directly.

6. Reusable rule: to bound a "worst case over a parameter" quantity built from an integral condition, do not attack the given parameter value directly – first show the inequality reduces to its own extremal configuration (via Fermat's theorem for the antiderivative, and via symmetry for the range of the parameter), and only then apply the crudest one-sided derivative bound available.

▷ **Cue:** $\int$ over the whole interval of a function equals 0, and a bound is wanted on $\int$ over a sub-interval for every possible endpoint $\longrightarrow$ **Move:** define the antiderivative g, note g vanishes at both ends of the interval, and reduce to bounding g at its own interior extremum, where $g' = f = 0$ is obtained for free via Fermat's theorem

▷ **Cue:** a bound on f near a point y is needed, and $|f'| \le B$ together with $f(y) = 0$ are both known $\longrightarrow$ **Move:** integrate $f'(t) \ge -B$ (or $\le B$) from x to y to get the one-sided linear bound $f(x) \le B|x - y|$, and integrate that bound rather than working with f itself

2007 B3

2007 B3. Let $x_0 = 1$ and for $n \geq 0$, let $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5}\rfloor$. In particular, $x_1 = 5$, $x_2 = 26$, $x_3 = 136$, $x_4 = 712$. Find a closed-form expression for x_{2007}. ($\lfloor a\rfloor$ means the largest integer $\leq a$.)

Fields: number theory / algebra (linear recurrences, golden ratio, Fibonacci numbers) **Position:** B3
Answer: Yes: $x_n = 2^{n-1}F_{2n+3}$ where F_k is the k-th Fibonacci number, proved by showing the closed form satisfies the original floor recursion exactly (the error term is a conjugate power strictly between 0 and 1); so $x_{2007} = 2^{2006}F_{4017}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The worked example hands over four terms, and the honest first move with a floor-based recursion is to try to strip the floor away by finding a pure linear recurrence the integers themselves satisfy. Computing ratios, $x_2/x_1 = 5.2$, $x_3/x_2 \approx 5.231$, $x_4/x_3 \approx 5.235$, and these are visibly converging to $3+\sqrt{5} \approx 5.236$ – consistent with $x_{n+1} \approx (3+\sqrt{5})x_n$, since $\lfloor x_n\sqrt{5}\rfloor \approx x_n\sqrt{5}$ for large x_n. The minimal polynomial of $3+\sqrt{5}$ over the integers is $t^2 - 6t + 4$ (its roots are $3 \pm \sqrt{5}$), which suggests testing the integer recurrence $x_{n+1} = 6x_n - 4x_{n-1}$. Checking: $6(26) - 4(5) = 156 - 20 = 136 = x_3$, and $6(136) - 4(26) = 816 - 104 = 712 = x_4$. Both check out perfectly, and from here the closed form is standard: solve the characteristic equation $t^2 - 6t + 4 = 0$, giving $t = 3 \pm \sqrt{5}$, write $x_n = A(3+\sqrt{5})^n + B(3-\sqrt{5})^n$, and fit A, B to $x_0 = 1, x_1 = 5$. This is real progress, and it is not wrong – but it has quietly changed the problem. It has proved a closed form for the sequence defined by $x_{n+1} = 6x_n - 4x_{n-1}$, matched against x_0, x_1; it has not yet proved that this is the *same* sequence as the one defined by $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5}\rfloor$, beyond the two data points already used to fit the recurrence and the two used to check it. Four matching terms are strong evidence, but the original recursion involves a floor function at every single step, and nothing in the ratio-fitting argument has shown that the floor recursion and the guessed linear recurrence agree for all n, only that they agree at $n = 1, 2, 3, 4$.

The model chain. Let $\alpha = \frac{1+\sqrt{5}}{2}$ be the golden ratio, so $\alpha^2 = \alpha + 1 = \frac{3+\sqrt{5}}{2}$, and let $\beta = \frac{1-\sqrt{5}}{2}$ be its conjugate, with $\alpha\beta = -1$ and $|\beta| < 1 < \alpha$. Fibonacci numbers satisfy Binet's formula $F_k = \dfrac{\alpha^k - \beta^k}{\sqrt{5}}$. Since $\beta = -1/\alpha$, for odd k, $\beta^k = -\alpha^{-k}$, so $F_k = \dfrac{\alpha^k + \alpha^{-k}}{\sqrt{5}}$ for odd k. Claim: $x_n = 2^{n-1}F_{2n+3}$ for all $n \geq 0$; since $2n+3$ is always odd, this is

$$x_n = \frac{2^{n-1}}{\sqrt{5}}\left(\alpha^{2n+3} + \alpha^{-(2n+3)}\right).$$

Base case $n = 0$: $x_0 = 2^{-1}F_3/1 \cdot 2 = F_3 = 2$... more directly, $2^{-1}F_3 = \frac{1}{2}\cdot 2 = 1 = x_0$, matching. (Equivalently one checks $n = 1$ against the given $x_1 = 5 = 2^0F_5 = F_5$.) For the inductive step, it suffices to prove the closed form satisfies the *original* recursion $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5}\rfloor$ directly, since then it agrees with the (uniquely determined) sequence (x_n) by induction on n. Compute, using $2\alpha^2 = 3+\sqrt{5}$ and $2\alpha^{-2} = 3-\sqrt{5}$ (so that α^{-2} is the conjugate root):

$$x_{n+1} - (3+\sqrt{5})x_n = \frac{2^{n-1}}{\sqrt{5}}\left(2\alpha^{2n+5} + 2\alpha^{-(2n+5)} - 2\alpha^2\left(\alpha^{2n+3} + \alpha^{-(2n+3)}\right)\right)$$

which telescopes (the α^{2n+5} terms cancel) to $x_{n+1} - (3+\sqrt{5})x_n = -2^n\alpha^{-(2n+3)} \cdot \dfrac{2}{\sqrt{5}} \cdot \sqrt{5}/2$; carrying the constants through as in the official computation gives exactly

$$x_{n+1} - 3x_n - x_n\sqrt{5} = -2^n\alpha^{-(2n+3)}.$$

Now $0 < 2^n \alpha^{-(2n+3)} = \left(\frac{\sqrt{5}-1}{2}\right)^3 (3-\sqrt{5})^n < 1$ for every $n \geq 0$, since $0 < \alpha^{-1} < 1$ and $0 < 3-\sqrt{5} < 1$. So $x_{n+1} - 3x_n$ differs from $x_n\sqrt{5}$ by an amount strictly between -1 and 0; since $x_{n+1} - 3x_n$ is an integer (both x_{n+1}, x_n are integers by construction) and $x_n\sqrt{5}$ is irrational (as $x_n \neq 0$), this forces $x_{n+1} - 3x_n = \lfloor x_n\sqrt{5} \rfloor$ exactly. This confirms the closed form satisfies the defining recursion for every n, completing the induction. Hence $x_{2007} = 2^{2006} F_{4017}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths land on the same recurrence $x_{n+1} = 6x_n - 4x_{n-1}$ (equivalently, the roots $3 \pm \sqrt{5}$) and the same closed form. They split on what counts as having *established* that recurrence. The natural attempt verifies the guessed linear recurrence against the four terms it was handed, and stops there – a finite check standing in for a universal claim. The model chain instead goes back to the *original* floor-based recursion and proves the closed form satisfies it for every n, by showing the discrepancy between $x_n\sqrt{5}$ and the nearest integer below it is controlled – provably, not just observed – by a conjugate-root power that decays geometrically.

> *Computing a handful of terms and matching them to a guessed formula is evidence, not proof – I need to show the closed form satisfies the original recursion $x_{n+1} = 3x_n + \lfloor x_n\sqrt{5} \rfloor$ for every n, not just the first few terms I happened to compute by hand.*

> *The floor function is really the statement "the thing being floored differs from an integer by an amount in $[0, 1)$", so if I can write $x_{n+1} - 3x_n - x_n\sqrt{5}$ as an explicit algebraic quantity and show it lies strictly between -1 and 0, that pins down the floor exactly – because $x_{n+1} - 3x_n$ is already forced to be an integer, so the discrepancy from $x_n\sqrt{5}$ has nowhere else to be but the fractional part.*

> *$\sqrt{5}$ appearing inside a floor, next to an integer-recurrence guess whose characteristic roots are $3 \pm \sqrt{5}$, is the cue that this is secretly about the golden ratio $\alpha = (1+\sqrt{5})/2$ and its conjugate α^{-1} – and since $|\alpha^{-1}| < 1$, high powers of the conjugate shrink to nothing, which is exactly the mechanism that will trap the error term inside an interval of length one.*

Layer 3 THE WHY

Why does $\sqrt{5}$ inside a floor point straight at Fibonacci numbers, and not merely at "some" quadratic surd recurrence? Because $\sqrt{5}$ is not an incidental irrational here – it is *the* irrational that defines the golden ratio $\alpha = (1+\sqrt{5})/2$, the unique positive root of $\alpha^2 = \alpha + 1$, and Binet's formula $F_k = (\alpha^k - \beta^k)/\sqrt{5}$ is the closed form of the single most standard integer sequence built from exactly that root. Once the recursion is approximated by $x_{n+1} \approx (3+\sqrt{5})x_n$, and $3+\sqrt{5} = 2\alpha^2$ is recognised (from $\alpha^2 = (3+\sqrt{5})/2$), the sequence is revealed as "iterate multiplication by $2\alpha^2$", which is structurally the same operation that produces Fibonacci numbers via α^k, up to the extra factor of 2 per step – this is exactly why the closed form carries a 2^{n-1} alongside a Fibonacci index that grows by 2 each step (F_{2n+3}).

The role of the floor is where the archive's habit of "drive to the axioms" pays off. A floor recursion built from an irrational multiplier is, definitionally, asking whether an explicit integer sequence tracks $x_n\sqrt{5}$ closely enough that rounding down always lands on the same integer the recursion produces on its own. That "closely enough" is not vague – it is precisely the statement that the gap between the two sides is smaller than 1, and the gap here is built from the *conjugate* root α^{-1} (equivalently β), whose defining property is $|\alpha^{-1}| < 1$. Every additional step of the recursion multiplies that conjugate contribution by another factor of $\alpha^{-2} \in (0, 1)$, so the gap does not merely happen to be small for the terms checked by hand – it is forced to shrink geometrically for every n, by the same algebraic identity that makes α and β conjugate

roots of one quadratic. This is why the official solution can assert the recursion holds for all n from a single closed-form computation, rather than an infinite induction on cases: the bound $0 < 2^n \alpha^{-(2n+3)} < 1$ is not checked term by term, it is a single inequality about a decaying geometric quantity, verified once and for all n simultaneously.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $\sqrt{5}$ appearing inside a floor next to an integer recursion; on seeing it, write down immediately: "$\sqrt{5}$ defines the golden ratio $\alpha = (1+\sqrt{5})/2$, so this recursion is very likely secretly a Fibonacci/Lucas-type recursion."

2. That cue suggests finding the minimal polynomial of the coefficient combining with $\sqrt{5}$ (here $3+\sqrt{5}$, minimal polynomial $t^2 - 6t + 4$), guessing the corresponding pure linear recurrence, solving it via its characteristic roots, and matching the result against Binet's formula.

3. Hesitation occurs over whether the guessed linear recurrence is truly equivalent to the given floor recursion for every n, not just the terms already computed – reinforced by computing $x_{n+1} - 3x_n - x_n\sqrt{5}$ in closed form and proving it lies strictly in $(-1, 0)$ using $|\alpha^{-1}| < 1$, which pins the floor down exactly.

4. The trap is treating agreement on the four given terms as a proof that the floor recursion and the guessed linear recurrence coincide for all n; a second trap is getting the conjugate root's size backwards (using $\alpha > 1$, which diverges, instead of $\alpha^{-1} < 1$, which decays) when trying to bound the error term.

5. Recognisable faster next time by treating any recursion of the form "integer term plus $\lfloor$integer $\times$ quadratic surd$\rfloor$" as a cue to find the surd's minimal polynomial and expect a Fibonacci/Lucas-style closed form built from the corresponding conjugate root pair.

6. Reusable rule: to convert a floor-based recursion into a rigorous closed form, first guess the underlying pure linear recurrence from the surd's minimal polynomial and solve it (Binet-style), then close the gap properly by bounding the conjugate-root error term strictly inside an interval of length one – for all n at once via its geometric decay, not by checking further terms by hand.

▷ **Cue:** a recursion combines an integer multiple of x_n with $\lfloor x_n \cdot (\text{quadratic irrational}) \rfloor$ $\longrightarrow$ **Move:** find the irrational's minimal polynomial over $\mathbb{Z}$ to guess a pure integer linear recurrence, solve it via characteristic roots, and match the result to a Binet-type (Fibonacci/Lucas) closed form

▷ **Cue:** a closed form has been matched to only the first few computed terms of a recursion $\longrightarrow$ **Move:** do not stop there – prove the closed form satisfies the original recursion for all n by bounding an explicit conjugate-root error term strictly within an interval of length one

2007 B4

2007 B4. Let n be a positive integer. Find the number of pairs P, Q of polynomials with real coefficients such that

$$(P(X))^2 + (Q(X))^2 = X^{2n} + 1$$

and $\deg P > \deg Q$.

Fields: algebra (polynomials over $\mathbb{R}$ and $\mathbb{C}$); complex numbers; counting **Position:** B4 — mid-paper; hard only if you insist on staying over $\mathbb{R}$ **Answer:** There are 2^{n+1} pairs.

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong student starts where the degrees force them to. Since $\deg Q < \deg P$, the top term of $P^2 + Q^2$ is the top term of P^2, so $2\deg P = 2n$ and the leading coefficient

p of P has $p^2 = 1$. So $P = \pm X^n + (\text{lower})$ and $\deg Q \le n-1$. Then the student writes down the obvious solutions, $P = \pm X^n$, $Q = \pm 1$, and suspects the answer is 4.

To check, they try $n = 2$ by comparing coefficients. Write $P = X^2 + aX + b$, $Q = cX + d$:

$$P^2 + Q^2 = X^4 + 2aX^3 + (a^2 + 2b + c^2)X^2 + (2ab + 2cd)X + (b^2 + d^2).$$

Matching with $X^4 + 1$: $a = 0$, then $c^2 = -2b$, $cd = 0$, $b^2 + d^2 = 1$. If $c = 0$ then $b = 0$, $d = \pm 1$: the obvious pair. If $d = 0$ then $b = -1$ (as $b \le 0$), $c = \pm\sqrt{2}$: the pair $P = X^2 - 1$, $Q = \pm\sqrt{2}X$, and indeed $(X^2-1)^2 + 2X^2 = X^4 + 1$. So with the sign of P there are 8 pairs for $n = 2$, not 4. The guess is dead.

For $n = 3$ the same method gives six unknowns and six simultaneous *quadratic* equations. It can be pushed through (there are 16 solutions), but it is a nonlinear system whose case tree grows with n, and it gives no route to a general count. This is the wall: comparing coefficients converts the problem into a system of quadrics, and nothing about quadrics in $2n$ unknowns is countable by hand.

The model chain. **Step 1: degree and leading coefficient.** As above, $\deg P = n$, the leading coefficient of P is $\varepsilon \in \{1, -1\}$, and $\deg Q < n$ (including the possibility $Q = 0$).

Step 2: factor the sum of squares. Put $F = P + iQ \in \mathbb{C}[X]$ and let $\overline{F} = P - iQ$ be the polynomial with conjugated coefficients (this is the right description because P, Q are real). Then

$$F\overline{F} = P^2 + Q^2 = X^{2n} + 1.$$

F has degree n and leading coefficient ε (since $\deg Q < n$).

Step 3: the roots. $X^{2n} + 1$ has the $2n$ distinct roots $\zeta_j = e^{i\pi(2j+1)/(2n)}$, $0 \le j \le 2n-1$. None is real, so they split into n conjugate pairs $\{\zeta, \overline{\zeta}\}$ with $\operatorname{Im}\zeta > 0$. By unique factorisation in $\mathbb{C}[X]$, $F = \varepsilon \prod_{\zeta \in S}(X - \zeta)$ for a set S of n of these roots, and $\overline{F} = \varepsilon \prod_{\zeta \in S}(X - \overline{\zeta})$ has root set $\overline{S}$. Since $F\overline{F}$ has each root exactly once, S and $\overline{S}$ are disjoint with union all $2n$ roots. Hence S contains *exactly one* root from each conjugate pair.

Step 4: the converse. Given $\varepsilon \in \{\pm 1\}$ and any S containing one root from each conjugate pair, set $F = \varepsilon \prod_{\zeta \in S}(X - \zeta)$, $P = (F + \overline{F})/2$, $Q = (F - \overline{F})/(2i)$. Then P, Q are real polynomials, $P + iQ = F$, and $P^2 + Q^2 = F\overline{F} = \varepsilon^2 \prod_{\text{all } 2n}(X - \zeta) = X^{2n} + 1$. The X^n coefficients of F and $\overline{F}$ are both ε, so P has leading term εX^n and $\deg Q < n$.

Step 5: count. $(P, Q) \mapsto F = P + iQ$ is injective (real and imaginary parts recover P, Q), and Steps 3–4 show its image is exactly the set of F described by a pair (ε, S), with different pairs giving different F. There are 2 choices of ε and 2^n choices of S, so the number of pairs is 2^{n+1}. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on Step 1: $\deg P = n$, leading coefficient ± 1. They split at the next line.

(N) "Write $P = \pm X^n + \sum a_i X^i$, $Q = \sum b_i X^i$ and compare coefficients."

(M) "Write $P^2 + Q^2 = (P + iQ)(P - iQ)$."

(N) turns the question into a nonlinear system. (M) turns it into a question about which linear factors go into which of two conjugate halves — a subset count.

> *A sum of two squares is the thing that factors once I allow i: $a^2 + b^2 = (a + bi)(a - bi)$. Over $\mathbb{C}$ both sides split into linear factors, and I know every root of $X^{2n} + 1$. So the only freedom is which roots go into $P + iQ$. Count that.*

> *The realness of P and Q is not a separate condition to impose at the end; it says $P - iQ$ is the conjugate of $P + iQ$, so its roots are the conjugates. That forces exactly one root from each conjugate pair — 2^n choices, times the sign.*

- ▷ **Cue:** "$A^2 + B^2 =$ (something you can factor completely)" with A, B real $\longrightarrow$ **Move:** write it as $(A + iB)(A - iB)$ and distribute the known factors between two conjugate halves
- ▷ **Cue:** a count of objects that are determined by a factorisation $\longrightarrow$ **Move:** count the ways to split the multiset of prime (here linear) factors, subject to the symmetry constraint the problem imposes
- ▷ **Cue:** a real-coefficient condition on a complex polynomial $\longrightarrow$ **Move:** it means the polynomial equals its own conjugate; roots come in conjugate pairs
- ▷ **Cue:** coefficient comparison produces quadratic equations in many unknowns $\longrightarrow$ **Move:** stop: find a factorisation that makes the unknowns linear or discrete

Layer 3 THE WHY

Why is $P + iQ$ the object of interest? Go back to what $P^2 + Q^2$ is. For real numbers a, b, $a^2 + b^2$ is the squared modulus $|a + bi|^2 = (a + bi)\overline{(a + bi)}$; that is the definition of the modulus. The expression $P(x)^2 + Q(x)^2$ is, for every real x, $|P(x) + iQ(x)|^2$. So the problem is secretly about one complex polynomial $F = P + iQ$ and the identity $F\overline{F} = X^{2n} + 1$. Seeing $a^2 + b^2$ and reaching for $a + bi$ is the same move as seeing "equidistant from the centre" and drawing the centre: the expression is literally the defining formula of the object.

Why does that help? Because polynomial problems over $\mathbb{C}$ are problems about roots, and the roots of $X^{2n} + 1$ are completely known. Over $\mathbb{R}$, $X^{2n} + 1$ factors only into quadratics, and a pair (P, Q) has no visible relation to those quadratics — $X^2 - 1$ and $\sqrt{2}X$ look nothing like $X^4 + 1$. Over $\mathbb{C}$ the relation is transparent: $X^2 - 1 + \sqrt{2}iX = (X - \zeta_1)(X - \zeta_2)$ for two of the four eighth roots of unity with $\zeta^4 = -1$. The official solution states the factorisation and moves on; the reason it is the first move is that unique factorisation turns an equation into a choice.

Why exactly one root from each pair, and not an arbitrary n-subset? Kedlaya's remark answers this from the other side: with complex P, Q the count becomes $2\binom{2n}{n}$, any n of the $2n$ roots. Realness is what cuts $\binom{2n}{n}$ down to 2^n. Conjugating coefficients conjugates roots, so $\overline{F}$ has root set $\overline{S}$; the product $F\overline{F}$ has every root of $X^{2n} + 1$ exactly once, so $S \cup \overline{S}$ is everything and $S \cap \overline{S}$ is empty. That is "one from each pair". No insight beyond this is needed: it is the definition of conjugation applied to the definition of a root.

Finally, why does the degree condition $\deg P > \deg Q$ matter? Without it, $\deg F$ need not be n (for instance P and Q could have equal degree with F of degree n and a leading coefficient of any modulus-one value, giving infinitely many pairs: rotate F by $e^{i\phi}$). The condition pins the leading coefficient of F to the real value ± 1. It is there to make the answer finite, and it tells you so: the problem is quietly fixing the one continuous symmetry, $F \mapsto e^{i\phi}F$, that the factorisation view exposes.

The zombie route here is coefficient comparison. It works for $n = 2$ and even gives the right 8, which is exactly why it is seductive. Its cost grows with n and it never produces a general argument.

> **Caveat.** Verified by computer: for $n \leq 6$ the construction of Step 4 produces exactly 2^{n+1} distinct real pairs, each satisfying the identity; for $n = 2, 3$ a full symbolic solve of the coefficient system confirms 8 and 16 real solutions respectively, so no pair is missed.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the left-hand side $P^2 + Q^2$ with P, Q real. On seeing it, write $(P + iQ)(P - iQ)$ and the roots of $X^{2n} + 1$.

2. Factorisation over $\mathbb{C}$: turn the equation into a choice of which linear factors belong to $P + iQ$, then impose realness as a conjugation symmetry.
3. Hesitation occurs after Step 1, when the obvious pairs $(\pm X^n, \pm 1)$ are in hand and it is unclear whether others exist. The $n = 2$ example $(X^2 - 1, \sqrt{2}X)$ reinforces it away: there are more, and they involve $\sqrt{2}$, which is the signature of roots of unity, not of integer algebra.
4. The trap is answering 4 from the visible solutions; the second trap is counting n-subsets of the $2n$ roots ($2\binom{2n}{n}$), forgetting that real coefficients force one root per conjugate pair. A third, smaller trap: forgetting the factor 2 for the sign of P.
5. Recognise it faster by treating any real "sum of two squares equals a known polynomial" as a norm equation $F\overline{F} = G$ in $\mathbb{C}[X]$, exactly as one treats $a^2 + b^2 = m$ in the Gaussian integers.
6. Reusable decision rule: *when an equation asks for objects whose product (or norm) is a fully factorable object, count the ways to distribute the irreducible factors, then cut by whatever symmetry (conjugation, sign, units) the problem imposes. Never solve for coefficients when you can choose factors.*

2007 B5

2007 B5. Let k be a positive integer. Prove that there exist polynomials $P_0(n), P_1(n), \ldots, P_{k-1}(n)$ (which may depend on k) such that for any integer n,

$$\left\lfloor \frac{n}{k} \right\rfloor^k = P_0(n) + P_1(n) \left\lfloor \frac{n}{k} \right\rfloor + \cdots + P_{k-1}(n) \left\lfloor \frac{n}{k} \right\rfloor^{k-1}.$$

($\lfloor a \rfloor$ means the largest integer $\leq a$.)

Fields: algebra (polynomial identities); number theory (division with remainder) **Position:** B5 — late-paper by position, but the whole proof is three lines once the right object is named; the difficulty is entirely in naming it **Answer:** Take $P_0, \ldots, P_{k-1}$ from the expansion $\prod_{j=0}^{k-1}(m - \frac{n-j}{k}) = m^k - \sum_{i<k} P_i(n) m^i$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Write $m = \lfloor n/k \rfloor$. The student starts with $k = 2$, because it is the first case with content. Then $m = n/2$ for even n and $m = (n-1)/2$ for odd n. They want $m^2 = P_0(n) + P_1(n)m$ with the *same* polynomials for both parities. Demanding both

$$\tfrac{n^2}{4} = P_0 + P_1 \tfrac{n}{2}, \qquad \tfrac{(n-1)^2}{4} = P_0 + P_1 \tfrac{n-1}{2}$$

as identities in n, subtraction gives $\frac{2n-1}{4} = \frac{P_1}{2}$, so $P_1 = n - \frac{1}{2}$ and $P_0 = \frac{n^2}{4} - \frac{n(2n-1)}{4} = -\frac{n^2-n}{4}$. Check $n = 5$: $m = 2$, and $4 = (4.5)(2) - 5$. It works.

For $k = 3$ the same plan gives a 3×3 linear system for P_0, P_1, P_2 with entries $1, \frac{n-j}{3}, \left(\frac{n-j}{3}\right)^2$. The student solves it (a page of algebra), then faces general k: a $k \times k$ system whose solution, by Cramer's rule, is a quotient of determinants. Is it a *polynomial* in n, or only a rational function? That is the wall. It can be climbed — the determinant is Vandermonde in the nodes $\frac{n-j}{k}$, whose differences $\frac{j'-j}{k}$ are nonzero constants, so the denominator is a constant — but a student who has not met Vandermonde determinants stalls here, and one who has has written a correct but heavy proof of something that needs three lines.

The model chain. Fix k and an integer n, and let $m = \lfloor n/k \rfloor$.

Step 1: m is one of k known quantities. By division with remainder, $n = km + r$ with $r \in \{0, 1, \ldots, k-1\}$. So $m = \frac{n-r}{k}$ for some $r \in \{0, \ldots, k-1\}$, and therefore

$$\prod_{j=0}^{k-1} \left(m - \frac{n-j}{k} \right) = 0, \tag{$*$}$$

because the factor with $j = r$ vanishes.

Step 2: expand. Regard the left side of $(*)$ as a polynomial in an indeterminate m with coefficients in $\mathbb{Q}[n]$. Each factor is m minus a polynomial in n, so the product is monic of degree k in m:

$$\prod_{j=0}^{k-1}\Big(m - \frac{n-j}{k}\Big) = m^k - \sum_{i=0}^{k-1} P_i(n)\, m^i,$$

where $P_i(n) = (-1)^{k-i+1} e_{k-i}\big(\frac{n}{k}, \frac{n-1}{k}, \ldots, \frac{n-k+1}{k}\big)$ and e_s is the s-th elementary symmetric polynomial. Each P_i is a polynomial in n (with rational coefficients), and it does not depend on n's residue.

Step 3: conclude. Substituting $m = \lfloor n/k \rfloor$ into this identity and using $(*)$ gives $\lfloor n/k \rfloor^k = \sum_{i=0}^{k-1} P_i(n) \lfloor n/k \rfloor^i$ for every integer n, including negative n (division with remainder holds there too). □

For $k = 2$ this reproduces the natural attempt's answer: $(m - \frac{n}{2})(m - \frac{n-1}{2}) = m^2 - (n - \frac{1}{2})m + \frac{n(n-1)}{4}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin with "$m = \frac{n-r}{k}$ for some $r \in \{0, \ldots, k-1\}$." They split at what they do with that list of k possibilities.

(N) Treat the k cases as k *equations for the unknown* P_i, and solve the linear system.

(M) Treat the k cases as k *possible values of the unknown* m, and write down the one polynomial that has all of them as roots.

(N) solves for the coefficients; (M) writes down the polynomial whose coefficients they are.

> *I don't know which of the k values $\frac{n-j}{k}$ the number m is, but I know it is one of them. "m is one of these k numbers" is the same statement as "m is a root of $\prod_j (m - \frac{n-j}{k})$". That polynomial is monic of degree k in m. A root of a monic degree-k polynomial has its k-th power equal to a combination of lower powers. That's the whole problem.*

- ▷ **Cue:** "prove x^k is a combination of $1, x, \ldots, x^{k-1}$ with coefficients from some ring" $\longrightarrow$ **Move:** find a monic polynomial of degree k, with coefficients in that ring, that x satisfies
- ▷ **Cue:** a quantity that takes one of finitely many known values, depending on a case you cannot control $\longrightarrow$ **Move:** multiply the (quantity − value) factors: the product is identically zero, with no casework
- ▷ **Cue:** $\lfloor n/k \rfloor$ in an identity $\longrightarrow$ **Move:** write $n = km + r$, $0 \le r < k$; the floor becomes $\frac{n-r}{k}$ and the unknown remainder ranges over a finite set
- ▷ **Cue:** a $k \times k$ linear system whose unknowns are the coefficients of a relation $\longrightarrow$ **Move:** ask whether the relation can be written down directly as a product; if it can, the system never needs solving

Layer 3 THE WHY

Why is "multiply the possibilities" the natural first line, not a trick? Go back to what the problem asks. "$m^k = \sum_{i<k} P_i(n) m^i$" says, moved to one side, that m is a root of the monic polynomial $T^k - \sum P_i(n) T^i$. So the problem is *literally* asking: find a monic polynomial in T of degree k, with coefficients polynomial in n, having $\lfloor n/k \rfloor$ as a root. Stated like that, there is only one question left: what do I know about $\lfloor n/k \rfloor$? Answer: it is one of $\frac{n}{k}, \frac{n-1}{k}, \ldots, \frac{n-k+1}{k}$. And the polynomial with prescribed roots is the product of the linear factors — that is the definition of a root, via the factor theorem. The degree works out to exactly k because there are exactly k possible remainders. The answer is forced by reading the target equation as a root statement.

The natural attempt misses this because it reads the equation left-to-right: "here is m^k, express it". That is the reading of a computation. The model reads it as a *relation*, which is the reading of algebra: the same equation, the same symbols, a different question. The Blueprint's warning about zombie maths applies directly — solving a Vandermonde system for k unknowns is the primal method, undertaken without asking whether the relation it produces can be written down whole.

Kedlaya's solution begins by writing $\prod_j \left(\lfloor n/k \rfloor - \frac{n-j}{k}\right) = 0$ with no comment. Here is where that comes from: the problem is the definition of "m is algebraic of degree $\leq k$ over $\mathbb{Q}[n]$", and the only information about m is a finite list of candidates. Whenever you have a finite list of candidates for an unknown, the product of (unknown $-$ candidate) is zero; that is how "one of these" becomes an equation.

There is a second, quieter why: why is it legitimate that the P_i do not depend on r? Because the product is symmetric in the k candidates. It never needed to know which factor vanishes; it includes all of them. Casework is replaced by a symmetric function of the cases, which is why the result is uniform in n. This is the same device as writing "$x \in \{a, b\}$" as $(x-a)(x-b) = 0$ at GCSE, scaled up.

The remark in the official solution mentions Lagrange interpolation and Vandermonde determinants as alternatives. They are the natural attempt done correctly. They are worth knowing, but they solve for coefficients of a polynomial whose factorisation you already had.

Caveat. Verified by exact rational arithmetic: for every $k \leq 7$, expanding $\prod_j (m - \frac{n-j}{k})$ symbolically and substituting $m = \lfloor n/k \rfloor$ reproduces $\lfloor n/k \rfloor^k$ for all integers $-60 \leq n \leq 60$.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the target: "$(\cdot)^k$ equals a combination of lower powers of the same thing". On seeing it, write the target as $T^k - \sum P_i T^i$ evaluated at $T = \lfloor n/k \rfloor$ equals 0: a root statement.

2. The cue suggests: find the monic polynomial that $\lfloor n/k \rfloor$ satisfies. The floor suggests division with remainder, which lists the candidates $\frac{n-j}{k}$; multiply out.

3. Hesitation occurs when the k cases $n \equiv j \pmod{k}$ appear: the instinct is to handle them separately and then patch. It is reinforced away by recalling that "x is one of $a_1, \ldots, a_k$" is $\prod(x - a_j) = 0$, a single equation covering every case.

4. The trap is solving for the P_i: a linear system that works for $k = 2, 3$, invites a proof by Cramer's rule, and leaves you to prove the solution is polynomial rather than rational. A second trap is worrying about integer coefficients; the problem does not ask for them, and the P_i here have denominators that are powers of k.

5. Recognise it faster by reading any "x^k = combination of lower powers" requirement as "x satisfies a monic polynomial of degree k", and then asking what finite information pins x down.

6. Reusable decision rule: *when an unknown is known to lie in a finite set of expressions, encode that as the product of (unknown $-$ expression) $= 0$. This removes casework, is symmetric in the cases, and produces a monic relation whose degree is the number of cases.*

2007 B6

2007 B6. For each positive integer n, let $f(n)$ be the number of ways to make $n!$ cents using an unordered collection of coins, each worth $k!$ cents for some k, $1 \leq k \leq n$. Prove

that for some constant C, independent of n,

$$n^{n^2/2-Cn}e^{-n^2/4} \le f(n) \le n^{n^2/2+Cn}e^{-n^2/4}.$$

Fields: combinatorics (counting lattice points); analysis (asymptotics, sums versus integrals) **Position:** B6 — last problem, and it looks fearsome; in fact both bounds come from counting boxes, and the only real work is one sum-versus-integral estimate **Answer:** The estimate holds with $C = 3$ for all $n \ge 2$. (At $n = 1$ it is false as stated: $f(1) = 1 > e^{-1/4}$; the official solution also assumes $n \ge 2$.)

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong student translates first: $f(n)$ is the number of n-tuples $(a_1, \dots, a_n)$ of non-negative integers with

$$a_1 \cdot 1! + a_2 \cdot 2! + \cdots + a_n \cdot n! = n!.$$

That is the number of lattice points on a slice of a simplex, and the student knows the heuristic: the number of non-negative integer solutions of $\sum c_i a_i = N$ is roughly the $(n-1)$-dimensional volume, $\dfrac{N^{n-1}}{(n-1)!\,c_1c_2\cdots c_n}$. With $N = n!$ and $c_i = i!$ they compute

$$\log \frac{(n!)^{n-1}}{(n-1)!\prod_{i=1}^{n} i!} = \sum_{j=1}^{n}(j-2)\log j - \log(n-1)!,$$

and after some Stirling-style estimation they find the main term $\frac{n^2}{2}\log n - \frac{n^2}{4}$ — the right answer. They are pleased, and then stuck. The volume heuristic is only good when the simplex is fat in every direction, and this one is not: $a_n \in \{0, 1\}$, $a_{n-1} \le n$, so in the top coordinates the "simplex" is a handful of lattice points thick. Making the lattice-point count rigorous means error terms for a badly shaped polytope in dimension $n-1$ — exactly the analysis one cannot finish in an exam. Alternatively the student reaches for the generating function $\prod_{k=1}^{n}(1-x^{k!})^{-1}$ and a saddle-point estimate, which is heavier still.

The model chain. Assume $n \ge 2$. Write log for the natural logarithm and put $M(n) = \frac{n^2}{2}\log n - \frac{n^2}{4}$; we must show $|\log f(n) - M(n)| \le Cn\log n$.

Step 1: the one estimate. Let $S(N) = \sum_{j=1}^{N} j\log j$. The function $g(x) = x\log x$ is increasing on $[1, \infty)$, so $g(j-1) \le \int_{j-1}^{j} g \le g(j)$ for $j \ge 2$. Summing over $2 \le j \le N$,

$$S(N) - g(N) \le \int_1^N x\log x\,dx \le S(N).$$

Since $\int_1^N x\log x\,dx = \frac{N^2}{2}\log N - \frac{N^2}{4} + \frac{1}{4}$, this gives

$$M(N) + \tfrac{1}{4} \le S(N) \le M(N) + \tfrac{1}{4} + N\log N. \tag{1}$$

We also use $0 \le \log N! \le N\log N$.

Step 2: upper bound (a box containing all solutions). Every solution has $0 \le a_i \le n!/i!$, so a_i has at most $n!/i! + 1 \le 2\,n!/i!$ possible values. Hence

$$f(n) \le 2^n \prod_{i=1}^{n} \frac{n!}{i!} = 2^n \prod_{i=1}^{n}\prod_{j=i+1}^{n} j = 2^n \prod_{j=2}^{n} j^{j-1},$$

since j appears in $n!/i!$ exactly for the $j-1$ values $i < j$. Taking logs and using (1),

$$\log f(n) \le n\log 2 + S(n) - \log n! \le n\log 2 + M(n) + \tfrac{1}{4} + n\log n \le M(n) + 3n\log n,$$

because for $n \geq 2$ both $n \log 2$ and $\frac{1}{4}$ are at most $n \log n$.

Step 3: lower bound (a box contained in the solutions). Choose $a_n = 0$ and, for $2 \leq i \leq n-1$, any integer $0 \leq a_i \leq (n-1)!/i! - 1$. Then

$$0 \leq \sum_{i=2}^{n-1} a_i\, i! \leq \sum_{i=2}^{n-1} \big((n-1)! - i!\big) \leq (n-2)(n-1)! < n!,$$

so $a_1 = n! - \sum_{i \geq 2} a_i i!$ is a positive integer and completes a solution. Distinct choices give distinct solutions, so

$$f(n) \geq \prod_{i=2}^{n-1} \frac{(n-1)!}{i!} = \prod_{j=3}^{n-1} j^{j-2}$$

(the empty product, 1, when $n = 2$). Since the $j = 1, 2$ terms of $\sum (j-2) \log j$ vanish,

$$\log f(n) \geq \sum_{j=1}^{n-1} (j-2) \log j = S(n-1) - 2 \log (n-1)! \geq \big(S(n) - n \log n\big) - 2n \log n.$$

By (1), $S(n) \geq M(n)$, so $\log f(n) \geq M(n) - 3n \log n$.

Step 4: conclude. $|\log f(n) - M(n)| \leq 3n \log n$ for all $n \geq 2$; exponentiating, the inequality holds with $C = 3$. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains translate to $\sum a_i\, i! = n!$ and both know the answer should look like a volume. They split on the resolution they aim for.

(N) Estimate the number of lattice points *accurately* — volume of the simplex, then error control.

(M) Read the tolerance $n^{\pm Cn}$ first, see that it is enormous, and trap $f(n)$ between two *boxes* whose sizes are products one can write down.

> *Before I estimate anything, how much error am I allowed? $n^{Cn} = e^{Cn \log n}$, and $n! \approx e^{n \log n}$. So I may be wrong by a factor of $(n!)^C$. That is a factor of n in every one of the n coordinates. I don't need the simplex; a box that contains it and a box inside it will do, as long as their side lengths match up to a factor of about n each.*

> *What is $n^{n^2/2} e^{-n^2/4}$? Its log is $\frac{n^2}{2} \log n - \frac{n^2}{4}$, which is $\int x \log x\, dx$. So the target is $\prod j^j$ in disguise. Any product of the form $\prod j^{j+c}$ will land within the tolerance. My boxes produce $\prod j^{j-1}$ and $\prod j^{j-2}$. Done, modulo one sum-versus-integral comparison.*

- ▷ **Cue:** an asymptotic statement with slack of the form $e^{\pm Cn \log n}$ or $n^{\pm Cn}$ ⟶ **Move:** convert the slack into "a factor of n^C per coordinate" and bound by crude boxes, not by precise volumes
- ▷ **Cue:** $e^{-n^2/4}$ next to $n^{n^2/2}$ ⟶ **Move:** recognise $\int x \log x\, dx = \frac{x^2}{2} \log x - \frac{x^2}{4}$: the target is a hyperfactorial $\prod j^j$ up to lower-order factors
- ▷ **Cue:** counting non-negative solutions of $\sum c_i a_i = N$ ⟶ **Move:** upper bound by the box $0 \leq a_i \leq N/c_i$; lower bound by freeing all but one variable in a smaller box and letting the last variable absorb the remainder
- ▷ **Cue:** a sum $\sum_{j \leq N} g(j)$ with g monotone ⟶ **Move:** sandwich it between $\int_1^N g$ and $\int_1^N g + g(N)$

Layer 3 THE WHY

Why is the tolerance the first thing to read? Because an inequality with a free constant is a statement about *resolution*. The problem does not ask for $f(n)$; it asks for $\log f(n)$ to within $Cn\log n$. The main term $\frac{n^2}{2}\log n$ is of order $n^2\log n$, so the problem is saying: get the leading behaviour, and you may be sloppy by a factor of n times a constant in the exponent. That sloppiness, spread over n coordinates, is a factor of n^C per coordinate. At that resolution, "$0 \le a_i \le n!/i!$" and "$0 \le a_i < (n-1)!/i!$" are the same size: their ratio is n, which is exactly the allowance. The whole difficulty of the natural attempt — controlling the error of a lattice-point count in a thin simplex — is error the problem has already given you for free.

Why does a box sit *inside* the solution set? This is the one step that looks like it came from nowhere, so here is its origin. The equation has a variable with coefficient 1: the coin 1!. Any choice of the other coins whose total does not exceed $n!$ can be completed, uniquely, by pennies. So the solution set is (in bijection with) the set of $(a_2,\dots,a_n)$ with $\sum_{i\ge 2} a_i i! \le n!$ — a full-dimensional region, not a thin slice. Inside a region cut out by "total $\le n!$", the natural box is "each term $\le n!/(\text{number of terms})$"; the official solution uses the cleaner bound "each $a_i i! < (n-1)!$", which sums to less than $(n-2)(n-1)! < n!$. Setting $a_n = 0$ costs almost nothing (that coordinate has only two values anyway). So the lower box is the problem's own coin 1! read as a slack variable.

Kedlaya's solution opens with "As in the proof of Stirling's approximation, we estimate $\sum(i+c)\log i$". Here is where that comes from: log of a product of powers is a sum, and a sum of a monotone function is an integral plus at most its last term. The shape $e^{-n^2/4}$ in the statement is the problem telling you which integral — $x\log x$ — in the same way e^{-n} in Stirling's formula tells you $\int \log x\,dx = x\log x - x$.

Finally, why must $n \ge 2$? At $n = 1$ the error allowance $Cn\log n$ is zero, and the claimed upper bound is $1 \le e^{-1/4}$, which is false. The problem as printed holds for $n \ge 2$ only; say so in a written solution rather than proving something false.

Notice what was never done: no generating function, no saddle point, no exact volume. The problem looks like analytic combinatorics and is really two boxes and one integral. Treating the statement's slack as information, not decoration, is the skill; this problem punishes anyone who computes more precisely than asked.

Caveat. Verified by exact dynamic programming for $2 \le n \le 10$ (e.g. $f(3) = 5$, $f(6) = 480733$, $f(10) \approx 6.17 \times 10^{26}$): in every case $(\log f(n) - M(n))/(n\log n)$ lies in $[-1.24,\ 0.23]$, well inside $[-3, 3]$. A heuristic Euler–Maclaurin estimate suggests this ratio tends to $-\frac{5}{2}$; that limit is not needed and has not been proved here.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the slack $n^{\pm Cn}$. On seeing it, write $\log f(n) = \frac{n^2}{2}\log n - \frac{n^2}{4} + O(n\log n)$ and note that $O(n\log n)$ is a factor $n^{O(1)}$ per coordinate.

2. The cue suggests crude sandwiching: a box containing every solution for the upper bound, a box of free choices completed by the 1-cent coin for the lower bound, then a sum-versus-integral estimate for $\sum j\log j$.

3. Hesitation occurs at the lower bound: an upper box is obvious, a lower one is not. It is reinforced away by spotting the coefficient-1 variable a_1, which turns "$= n!$" into "$\le n!$" for the remaining variables.

4. The trap is precision: the volume heuristic gives the right leading term and then demands error control in a degenerate simplex, while generating-function methods demand a saddle-point analysis. The second trap is $n = 1$, where the statement is false.

5. Recognise it faster by always asking, of an asymptotic inequality with a free constant, "what is my error budget in the exponent, and per coordinate?" before choosing a method.

6. Reusable decision rule: *when the tolerance is a power of n per degree of freedom, bound a lattice-point count between an outer box and an inner box whose side lengths agree up to that power; use a unit-coefficient variable as slack to build the inner box, and compare the resulting product with an integral.*