DAWN PUTNAM SESSION

Putnam 1982 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1982 — Paper A **5**
1982 A1 5
1982 A2 7
1982 A3 10
1982 A4 13
1982 A5 16
1982 A6 18

Putnam 1982 — Paper B **23**
1982 B1 23
1982 B2 25
1982 B3 28
1982 B4 31
1982 B5 34
1982 B6 36

Putnam 1982 — Paper A

1982 A1

1982 A1. Let V be the region in the cartesian plane consisting of all points (x, y) satisfying the simultaneous conditions $|x| \leq y \leq |x| + 3$ and $y \leq 4$. Find the centroid $(\bar{x}, \bar{y})$ of V.

Fields: coordinate geometry, centroids of plane regions **Position:** A1 — the opener **Answer:** The centroid is $(\bar{x}, \bar{y}) = (0, 13/5)$. V is the large triangle with vertices $(0,0), (-4,4), (4,4)$ (bounded by $y = |x|$ and $y = 4$) with the small similar triangle at its apex, vertices $(0,3), (-1,4), (1,4)$ (where $y = |x| + 3$ overruns $y = 4$), removed.

Layer 1 THE WHAT — surface analysis

The natural attempt. You read "find the centroid" and reach for the definition: $\bar{x} = \frac{1}{A}\iint_V x\,dA$, $\bar{y} = \frac{1}{A}\iint_V y\,dA$. To use it you need V as a y-range in x, so you sketch the three lines $y = |x|$, $y = |x| + 3$, $y = 4$ and find, correctly, that the shape of the vertical strip changes at $x = \pm 1$: for $|x| \leq 1$ the strip runs from $y = |x|$ up to $y = |x| + 3$ (the constraint $y \leq 4$ is slack there, since $|x| + 3 \leq 4$); for $1 \leq |x| \leq 4$ the upper line $y = |x| + 3$ has already passed 4, so the strip is capped and runs from $y = |x|$ up to $y = 4$; for $|x| > 4$ the strip is empty.

You now have a genuine plan, and it is correct, but it is four separate integrals (by symmetry you can do $x \geq 0$ and double, which is some relief): area over $[0,1]$, area over $[1,4]$, and the same split again for the first moment $\iint y\,dA$ (you do not need $\iint x\,dA$ once you notice V is symmetric under $x \mapsto -x$, but you have to notice that first, and it costs a moment of doubt about whether the two pieces really do share that symmetry). Each piece is an easy polynomial integral — $\int_0^1 3\,dx$, $\int_1^4 (4 - x)\,dx$ for area; $\int_0^1 \left[(x+3)^2 - x^2\right]/2\,dx$, $\int_1^4 (16 - x^2)/2\,dx$ for the y-moment — but there are four of them, each with its own limits, and a single sign slip or a dropped factor of 2 in any one of the four throws off the final fraction. Carried through carefully it gives area 15, y-moment 39, hence $\bar{y} = 39/15 = 13/5$: the right answer, purchased with eight lines of arithmetic across two case splits.

The model chain. Look at the three boundary lines again, but ask what shapes they bound rather than what strip they cut. $y = |x|$ together with $y = 4$ bounds a single triangle U with vertices $(0,0)$, $(-4,4)$, $(4,4)$ (the lines meet the cap at $x = \pm 4$). $y = |x| + 3$ together with $y = 4$ bounds a smaller, similarly-shaped triangle T with vertices $(0,3)$, $(-1,4)$, $(1,4)$ (the lines meet the cap at $x = \pm 1$). Because $|x| + 3 \geq |x|$ always, T sits inside U, and V — the set with $|x| \leq y$ and $y \leq |x| + 3$ and $y \leq 4$ — is exactly U with T removed: T and V overlap only on their shared boundary and $T \cup V = U$.

The centroid of a triangle is the average of its three vertices — no integration required. So

$$G_U = \left(\tfrac{0-4+4}{3}, \tfrac{0+4+4}{3}\right) = (0, \tfrac{8}{3}), \qquad G_T = \left(\tfrac{0-1+1}{3}, \tfrac{3+4+4}{3}\right) = (0, \tfrac{11}{3}).$$

Areas, likewise without integration: U has base 8 (from $x = -4$ to $x = 4$ at $y = 4$) and height 4, area 16; T has base 2 and height 1, area 1. Hence $[V] = 16 - 1 = 15$.

Centroids combine like point masses located at their own centroids, weighted by area, whenever the region is an area-disjoint union: $\text{area}(U) \cdot G_U = \text{area}(T) \cdot G_T + \text{area}(V) \cdot G_V$. By symmetry $\bar{x} = 0$ immediately. For $\bar{y}$:

$$16 \cdot \tfrac{8}{3} = 1 \cdot \tfrac{11}{3} + 15 \cdot \bar{y} \implies \tfrac{128}{3} - \tfrac{11}{3} = 15\bar{y} \implies \bar{y} = \frac{117/3}{15} = \frac{39}{15} = \frac{13}{5}.$$

No integral is evaluated anywhere in this chain.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that the strip shape changes at $|x| = 1$ and vanishes at $|x| = 4$ — that much is forced by reading the inequalities. They split at the question asked immediately after the sketch is drawn. The natural attempt asks "what is y as a function of x on each piece?" and heads for $\iint x\,dA$, $\iint y\,dA$. The model asks "what *shape* is each piece, as a whole object?" and notices that U, T, V are each a single triangle or triangle-minus-triangle, so their centroids are already known formulas.

> *Every boundary curve here is a straight line. A region bounded entirely by straight lines is a polygon, and a polygon can be built out of triangles whose centroids I already know how to write down — I do not need calculus for this at all.*

> *$y = |x| + 3$ is $y = |x|$ shifted up by 3; both are capped by the same line $y = 4$. Two parallel copies of the same V-shape capped by the same horizontal line give two similar triangles, one nested inside the other, and my region is their difference.*

> *Centroid $\times$ area is a moment, and moments add over disjoint pieces because the integral defining them is additive. If $T \cup V = U$ with no overlapping area, then $\text{area}(U)\,G_U = \text{area}(T)\,G_T + \text{area}(V)\,G_V$ exactly — I can solve this single vector equation for G_V instead of computing G_V's integral directly.*

Layer 3 THE WHY

Why is it legitimate to treat T and V as point masses at their own centroids and simply add? Go back to the definition: $\bar{y}_U = \frac{1}{[U]} \iint_U y\,dA$. Since $U = T \cup V$ with the two pieces overlapping only on a set of area zero (their shared boundary segments), the integral over U splits exactly as $\iint_U y\,dA = \iint_T y\,dA + \iint_V y\,dA$ — this is not an approximation, it is additivity of the (Riemann or Lebesgue) integral over a partition of the domain. Multiplying through by the respective areas, $[U]\bar{y}_U = [T]\bar{y}_T + [V]\bar{y}_V$. This identity holds for *any* decomposition of a region into two area-disjoint pieces, whatever their shape; the reason triangles are special is only that their individual centroids are already closed-form (the average of the vertices), not that the additivity itself needs triangles.

Why is the centroid of a triangle the average of its vertices, with no integral needed once you accept that fact? Slice the triangle into thin strips parallel to one side. Each strip's own centroid lies at its midpoint, and as the strip width shrinks to zero the locus of these midpoints is exactly the median from the opposite vertex. So the triangle's centroid lies on every median, hence at their common intersection, which elementary geometry places at the point dividing each median 2 : 1 from the vertex — and that point is $\frac{1}{3}(P_1 + P_2 + P_3)$. This fact is itself provable by the same additivity argument one level down (strips are point masses at their midpoints), so nothing here is asserted without appeal to the same underlying principle: moments add.

Why does the problem hand you two similar triangles rather than one awkward shape? Because $y = |x|$ and $y = |x| + 3$ are literally the same function translated vertically, and a

vertical translation of a V-shape, re-capped by the same horizontal line $y = 4$, produces a scaled copy of the same triangle sitting at the same apex $x = 0$. The problem-setter's phrasing "$|x| \le y \le |x| + 3$ and $y \le 4$" is already telling you, before you draw anything, that you are looking at an outer shape minus an inner shape of the same family. Reading the inequality structurally — as "two nested regions of the same shape, capped identically" — is what removes the need for calculus; reading it pointwise — as "y ranges over this interval for this x" — is what forces you into it. Centroids are additive with respect to area no matter which reading you take; only one reading lets you skip the integral altogether.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that every bounding curve in the problem — $y = |x|$, $y = |x| + 3$, $y = 4$ — is piecewise linear. Write down immediately: this region is a polygon (in fact a union/difference of triangles), not a general curved region.

2. That cue suggests decomposing V into triangles with known centroids and combining them by the area-weighted lever-arm identity $[U]G_U = [T]G_T + [V]G_V$, rather than integrating.

3. Hesitation occurs right after the sketch, choosing between "set up $\iint x\,dA$" and "name the shapes." It is resolved by asking what $y = |x| + 3$ is *as a curve relative to* $y = |x|$ — a vertical shift of the same function — which flags the two-similar-triangles structure before any integral is written.

4. The trap is doing the direct double integral: it is not wrong, but it requires two case splits (area and y-moment, each over $[0, 1]$ and $[1, 4]$), four separate antiderivatives, and a correct symmetry argument to avoid computing $\bar{x}$'s integral — eight opportunities to drop a factor of 2 or a limit of integration.

5. Recognisable faster next time by treating "$|x| \le y \le |x| + 3$" as a template: two translates of an absolute-value graph, capped by a line, always bound a triangle-minus-triangle, and both triangles' vertices can be read off directly from where each pair of bounding lines meets.

6. Reusable decision rule: before integrating for a centroid, check whether the region decomposes (by addition or subtraction) into pieces — triangles, rectangles, circular sectors — whose centroids are already closed-form. If so, use $[\text{whole}]\, G_{\text{whole}} = \sum[\text{piece}]\, G_{\text{piece}}$ and solve for the unknown piece; reserve direct integration for regions with genuinely curved boundaries.

▷ **Cue:** a region bounded entirely by straight lines $\longrightarrow$ **Move:** treat it as a polygon built from triangles with known vertex-average centroids, not as a generic $\iint dA$ problem

▷ **Cue:** two bounding curves that are vertical translates of the same function, both capped by the same line $\longrightarrow$ **Move:** recognise nested similar shapes and express the target region as outer minus inner

▷ **Cue:** need the centroid of a region that is one known shape minus another known shape $\longrightarrow$ **Move:** use the area-weighted identity $[\text{whole}]\, G_{\text{whole}} = [\text{piece}_1]\, G_{\text{piece}_1} + [\text{piece}_2]\, G_{\text{piece}_2}$ and solve for the unknown centroid

▷ **Cue:** the region is symmetric under $x \mapsto -x$ $\longrightarrow$ **Move:** $\bar{x} = 0$ immediately, without computing $\iint x\,dA$

1982 A2

1982 A2. For positive real x, let $B_n(x) = 1^x + 2^x + \cdots + n^x$. Prove or disprove the convergence of

$$\sum_{n=2}^{\infty} \frac{B_n(\log_n 2)}{(n \log_2 n)^2}.$$

Fields: real analysis, series convergence, comparison test **Position:** A2 **Answer:** The series converges. Since $x = \log_n 2 > 0$, each term of $B_n(x)$ is at most the largest one, n^x, so $B_n(\log_n 2) \le n \cdot n^{\log_n 2} = 2n$, and the series is dominated by $\sum 2/\left(n(\log_2 n)^2\right)$, which converges by the integral test.

Layer 1 THE WHAT — surface analysis

The natural attempt. The exponent $\log_n 2$ is unusual — it depends on n and shrinks to 0 as $n \to \infty$, since $\log_n 2 = \ln 2 / \ln n$. Your instinct with a sum of n terms of the form k^x is to estimate it the way you estimate any such sum: compare it to an integral. You write

$$B_n(x) = \sum_{k=1}^{n} k^x \approx \int_1^n t^x \, dt = \frac{n^{x+1} - 1}{x+1} \qquad (x \neq -1),$$

and substitute $x = \log_n 2$, so $x + 1 = 1 + \dfrac{\ln 2}{\ln n} \to 1$ as $n \to \infty$, and $n^{x+1} = n \cdot n^x = n \cdot n^{\log_n 2}$. At this point you need the fact $n^{\log_n 2} = 2$ anyway — so you already have to stop and justify it before your asymptotic formula means anything — and having got it, your integral approximation gives $B_n(\log_n 2) \approx \dfrac{2n - 1}{1 + \ln 2 / \ln n} \sim 2n$.

This is correct, and it is a stronger statement than the problem needs: you now know $B_n(\log_n 2)$ is asymptotically $2n$, not merely that it is at most some multiple of n. But getting there rigorously means controlling the error in "$\sum \approx \int$" (the integral test error bound, or comparing B_n to both $\int_1^n$ and $\int_0^n$, or invoking Euler–Maclaurin), and separately controlling the error in $x + 1 \to 1$ as n varies. None of that machinery is wrong, but the problem only asked you to *prove or disprove convergence*, a yes/no threshold question, and you have built an asymptotic equality to answer it.

The model chain. Since $x = \log_n 2$ is positive and $t \mapsto t^x$ is increasing for $t > 0$, every term in $B_n(x) = 1^x + 2^x + \cdots + n^x$ satisfies $k^x \le n^x$ for $1 \le k \le n$. There are n such terms, so

$$B_n(\log_n 2) \le n \cdot n^{\log_n 2}.$$

By the definition of $\log_n 2$ as the exponent to which n must be raised to give 2, $n^{\log_n 2} = 2$ exactly, for every $n \ge 2$. Hence

$$0 < \frac{B_n(\log_n 2)}{(n \log_2 n)^2} \le \frac{2n}{(n \log_2 n)^2} = \frac{2}{n(\log_2 n)^2}.$$

It remains to check that $\sum_{n=2}^{\infty} \dfrac{2}{n(\log_2 n)^2}$ converges. Substituting $u = \log_2 n$ (so $du = \dfrac{dn}{n \ln 2}$), $\int_2^\infty \dfrac{dn}{n(\log_2 n)^2} = \ln 2 \int_1^\infty \dfrac{du}{u^2}$, which converges. By the integral test the comparison series converges, and since $0 \le \dfrac{B_n(\log_n 2)}{(n \log_2 n)^2} \le \dfrac{2}{n(\log_2 n)^2}$ for all $n \ge 2$, the comparison test gives convergence of the original series.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same identity, $n^{\log_n 2} = 2$, and both must confront it. They split immediately after: the natural attempt uses that identity as one ingredient inside an *asymptotic* formula for $B_n(\log_n 2)$, aiming to know the sum's behaviour precisely; the model uses it as the value of the single largest term inside a one-line *inequality*, aiming only to know the sum's behaviour well enough.

> *I am being asked to prove or disprove convergence — a threshold, not a value. I do not need to know what $B_n(\log_n 2)$ is; I need a true inequality that is good enough for the comparison test. The crudest bound available — n copies of the biggest term — might already be exactly good enough, and it costs nothing to check before reaching for asymptotics.*

> *$n^{\log_n 2}$ is not something to approximate: it is 2, by the definition of $\log_n 2$ as the inverse of $n^{(\cdot)}$. Evaluate that piece to a constant the instant I see it, before doing anything else with the exponent.*

Layer 3 THE WHY

Why is the crude bound "n terms, each at most the largest" the right tool here rather than a compromise? Because the comparison test asks only for an inequality $0 \le a_n \le b_n$ with $\sum b_n$ known to converge — it is indifferent to how loose that inequality is, provided it is true and provided $\sum b_n$ still converges. Reaching for the tightest possible estimate of $B_n(x)$ is solving a harder problem than the one asked. The discipline the problem is testing is exactly the ability to notice when a question is a threshold question and to stop optimising once a sufficient bound is in hand. This is not laziness dressed up as a principle: an asymptotic estimate for $B_n(x)$ still has to be turned into an inequality before it can be fed to a comparison test, so the crude bound was always going to be needed in some form — the natural attempt just builds a stronger, unnecessary intermediate result on the way to it.

Why does $n^{\log_n 2} = 2$ collapse so cleanly, and why is noticing it the real content of the problem? The logarithm $\log_n 2$ is defined as the unique real number x with $n^x = 2$; this is not a computation to perform, it is the definition being unwound. The problem-setter chose the exponent $\log_n 2$ specifically so that this identity would trigger: any exponent that made n^x into an unpleasant function of n would have buried the problem in genuine analysis, but this one is engineered to hand you a clean constant, 2, for free. Recognising $a^{\log_a b} = b$ on sight is the single moment where the problem is decided; everything after it (bounding the sum, applying the integral test) is standard machinery.

Why does the resulting series, $\sum 2/(n(\log_2 n)^2)$, need the integral test rather than a simpler comparison, and why should its convergence be recognisable on sight? $1/n$ alone diverges (the harmonic series), so the extra factor $1/(\log_2 n)^2$ has to be doing real work, and it is exactly the boundary case: $\sum 1/n$ diverges, $\sum 1/(n \log n)$ diverges (barely), $\sum 1/(n(\log n)^2)$ converges (barely). This trio is the standard demonstration that the integral test can distinguish rates the direct comparison test cannot, via $\int du/u = \ln u \to \infty$ against $\int du/u^2 = -1/u \to$ finite after the substitution $u = \log n$. Once the crude bound on B_n produces exactly this shape, no further estimation is needed — the series was engineered to land precisely on the convergent side of that boundary, which is also why any *sharper* estimate of B_n than "$\le 2n$" would have been wasted effort: the comparison series it produces would still just be $\sum c/(n(\log_2 n)^2)$ for some constant c, convergent regardless of the exact value of c.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the exponent $\log_n 2$: it is a function of n that tends to 0 and, more importantly, satisfies $n^{\log_n 2} = 2$ by definition. Write that identity down before manipulating B_n at all.
2. That cue suggests bounding the whole sum $B_n(x)$ by n copies of its largest term n^x (valid since $x > 0$ makes t^x increasing), turning $B_n(\log_n 2)$ into the clean bound $2n$, and then applying the comparison test.
3. Hesitation occurs at exactly the point where the crude bound is available: is "n times the largest term" too wasteful to actually prove convergence? It is resolved by computing where the bound lands — $2/(n(\log_2 n)^2)$ — and recognising that series as the standard convergent boundary case via the integral test, i.e. checking the bound is sufficient before discarding it as too crude.

4. The trap is treating a yes/no convergence question as a request for an asymptotic formula, and spending the effort of an integral approximation (with its error terms) on a quantity that only needed to be bounded above.

5. Recognisable faster next time by scanning any exponent of the form $\log_a b$ appearing in a competition sum and evaluating $a^{\log_a b} = b$ immediately, before touching the rest of the expression.

6. Reusable decision rule: when a problem asks only whether a series converges, look for the cheapest true bound that produces a series with recognisable convergence behaviour (geometric, p-series, or the $1/(n(\log n)^p)$ family), rather than for the tightest asymptotic — and specifically, a positive increasing sequence's partial sum of n terms is always at most n times its largest term, a bound worth trying first.

▷ **Cue:** an exponent of the form $\log_a b$ appears in the problem $\longrightarrow$ **Move:** evaluate $a^{\log_a b} = b$ immediately; it is the definition of the logarithm, not something to estimate

▷ **Cue:** need to bound a sum of n positive, increasing-in-index terms $\sum_{k=1}^{n} k^x$ with $x > 0$ $\longrightarrow$ **Move:** bound crudely by n times the largest term, $n \cdot n^x$, before attempting an asymptotic

▷ **Cue:** a series is asked only for convergence, not a value $\longrightarrow$ **Move:** seek the cheapest sufficient comparison bound rather than the tightest asymptotic estimate

▷ **Cue:** a series of the shape $\sum 1/(n(\log n)^p)$ $\longrightarrow$ **Move:** substitute $u = \log n$ and apply the integral test: converges iff $p > 1$

1982 A3

1982 A3. Evaluate

$$\int_0^\infty \frac{\arctan(\pi x) - \arctan x}{x}\, dx.$$

Fields: real analysis, integral calculus (Frullani-type integrals, Fubini/Leibniz) **Position:** A3 **Answer:** The value is $\frac{\pi}{2}\ln\pi$. The difference $\arctan(\pi x) - \arctan x$ is rewritten as an integral over an auxiliary parameter, the order of integration is swapped, and each remaining integral is the elementary $\int_0^\infty dx/(1 + a^2x^2) = \pi/(2a)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The integrand is a difference of two things divided by x, and the reflex from linearity of the integral is to split it:

$$\int_0^\infty \frac{\arctan(\pi x) - \arctan x}{x}\, dx \stackrel{?}{=} \int_0^\infty \frac{\arctan(\pi x)}{x}\, dx - \int_0^\infty \frac{\arctan x}{x}\, dx.$$

Near $x = 0$ each piece is fine — $\arctan(cx)/x \to c$ as $x \to 0$, no singularity. The trouble is at infinity: $\arctan(cx) \to \pi/2$ for any $c > 0$, so $\arctan(cx)/x \sim (\pi/2)/x$ as $x \to \infty$, and $\int^\infty dx/x$ diverges (logarithmically). Both separated integrals diverge. You have turned a convergent integral into "$\infty - \infty$", and linearity of the integral does not apply to a difference of two divergent integrals — you cannot split and then reunite them.

Recognising this, you might try integration by parts on the original integrand, or a substitution $x \to 1/x$ to relate the pieces near 0 and near ∞ to each other, hoping the two logarithmic divergences visibly cancel. This can be made to work, but it requires you to keep re-deriving, by hand, exactly the fact that the two tails cancel at the same rate — the very fact a smarter setup would build in from the start rather than fight for at the end. You are treating the two arctangents as separate objects, when the problem handed them to you as one.

The model chain. Keep the two arctangents coupled. For fixed $x > 0$, the function $u \mapsto \arctan(ux)$ has derivative $\dfrac{x}{1+u^2x^2}$ in u, so by the Fundamental Theorem of Calculus,

$$\arctan(\pi x) - \arctan(x) = \int_1^\pi \frac{x}{1+u^2x^2}\,du.$$

Substitute this into the original integral:

$$I = \int_0^\infty \frac{1}{x}\left(\int_1^\pi \frac{x}{1+u^2x^2}\,du\right)dx = \int_0^\infty \int_1^\pi \frac{1}{1+u^2x^2}\,du\,dx.$$

The integrand $\dfrac{1}{1+u^2x^2}$ is positive on $x \geq 0$, $u \in [1,\pi]$, so Tonelli's theorem permits swapping the order of integration freely:

$$I = \int_1^\pi \left(\int_0^\infty \frac{dx}{1+u^2x^2}\right)du.$$

The inner integral is elementary: substituting $t = ux$ gives $\int_0^\infty \dfrac{dx}{1+u^2x^2} = \dfrac{1}{u}\int_0^\infty \dfrac{dt}{1+t^2} = \dfrac{1}{u}\cdot\dfrac{\pi}{2} = \dfrac{\pi}{2u}$. Hence

$$I = \int_1^\pi \frac{\pi}{2u}\,du = \frac{\pi}{2}\big[\ln u\big]_1^\pi = \frac{\pi}{2}\ln\pi.$$

(An equally valid route: set $f(a) = \int_0^\infty \frac{\arctan(ax)-\arctan x}{x}\,dx$ for $a > 0$ and differentiate under the integral sign, $f'(a) = \int_0^\infty \frac{x}{x(1+a^2x^2)}\,dx = \int_0^\infty \frac{dx}{1+a^2x^2} = \dfrac{\pi}{2a}$; then $f(a) = \frac{\pi}{2}\ln a + C$ with $f(1) = 0$ forcing $C = 0$, so $f(\pi) = \frac{\pi}{2}\ln\pi$. This is the same identity, differentiated rather than integrated into existence.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree the two arctangents cannot simply be pulled apart. They split at what to do instead. The natural attempt keeps trying to evaluate $\arctan(\pi x)$ and $\arctan(x)$ as functions of x alone and hunt for a way their infinite tails cancel after the fact. The model chain stops treating π and 1 as two fixed constants and instead reads them as the two endpoints of a parameter u that the difference is secretly an integral over: $\arctan(\pi x) - \arctan(x) = \int_1^\pi \partial_u[\arctan(ux)]\,du$.

> *This is a difference of the same function evaluated at two different scalings of x. Splitting it by linearity throws away the coupling between the two terms that is the entire reason the combination converges. I should keep them coupled — and the Fundamental Theorem of Calculus, run backwards, turns a difference $g(\pi) - g(1)$ into an integral $\int_1^\pi g'(u)\,du$, which is exactly a way of keeping two values of a function coupled inside one expression.*

> *Once the difference is an integral over u, I now have an iterated integral $\int_0^\infty \int_1^\pi \frac{1}{1+u^2x^2}\,du\,dx$ with a positive integrand. Swapping the order costs nothing (Tonelli) and turns the hard variable of integration into the easy one: integrating over x first for fixed u is the standard $\int_0^\infty dx/(1+a^2x^2) = \pi/(2a)$, whereas integrating over u first for fixed x is exactly the original difficult difference.*

Layer 3 THE WHY

Why does splitting by linearity fail here specifically, rather than just "happening" to fail? Both $\arctan(\pi x)/x$ and $\arctan(x)/x$ behave like $(\pi/2)/x$ as $x \to \infty$ — the *same* leading behaviour, because arctan tends to the same limit $\pi/2$ regardless of the scaling constant inside it.

A logarithmic divergence, $\int^{\infty} dx/x$, is the marginal case: neither integrable nor obviously not, and here it appears with an identical coefficient in both halves. The two divergences are not merely both present, they are equal, which is exactly the situation in which they must cancel in the difference but cannot be shown to cancel by treating the two integrals as independent objects with their own limits — a limit of a difference is not the difference of two (nonexistent) limits. The only way to see a cancellation of equal, divergent things is to never let them become fully separate quantities; the FTC rewrite is precisely a way of doing this, because it replaces "two values of arctan" with "one integral of arctan'", and $\arctan'(t) = 1/(1+t^2)$ decays fast enough ($\sim 1/t^2$) that the resulting double integral is absolutely convergent and Fubini/Tonelli applies without any special pleading.

Why is $\int_1^{\pi} \partial_u[\arctan(ux)]\, du$ the natural object to write down, rather than a trick pulled from nowhere? It is the Fundamental Theorem of Calculus applied to the function $u \mapsto \arctan(ux)$ on $[1, \pi]$: $g(\pi) - g(1) = \int_1^{\pi} g'(u)\, du$ for any differentiable g, is not a device invented for this problem, it is the definition of the derivative turned into a statement about accumulation. What makes it the *right* move here, rather than a valid-but-useless identity, is that differentiating $\arctan(ux)$ in u removes the arctangent entirely and replaces it with the rational function $x/(1+u^2x^2)$ — a strictly easier object to integrate in x. The problem's two constants, π and 1, were never meant to be evaluated as two separate numbers; they were the two endpoints of the parameter the whole expression is differentiating with respect to. This is the same insight, differently packaged, as differentiating under the integral sign in the alternate solution: whenever a hard-to-evaluate quantity is naturally indexed by a parameter (here, the scaling constant π vs. 1), introduce or exploit that parameter as a variable of integration or differentiation, because the derivative of arctan is elementary even though arctan itself, divided by x and integrated to infinity, is not.

Why is $\int_0^{\infty} dx/(1+a^2x^2) = \pi/(2a)$ the correct final ingredient, and not a coincidence that it appears twice (once directly, once via $f'(a)$ in the alternate route)? It is, again, just $\arctan'(t) = 1/(1+t^2)$ together with $\arctan(\infty) - \arctan(0) = \pi/2 - 0 = \pi/2$, rescaled by the substitution $t = ax$. The entire problem, both routes, is the single fact "arctan is the antiderivative of $1/(1+t^2)$, with known limits at 0 and ∞" used twice: once to build the difference of arctangents as an integral, once to evaluate the integral that results. Seeing that repetition is what makes the two official routes recognisably the same idea rather than two unrelated tricks.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape $\dfrac{f(ax) - f(bx)}{x}$ integrated over $(0, \infty)$, with $f = \arctan$, $a = \pi$, $b = 1$: a difference of the same function at two scales, divided by x. Write down immediately that this is a Frullani-integral shape and that splitting by linearity will fail.

2. That cue suggests either invoking Frullani's theorem directly ($\int_0^{\infty} \frac{f(ax)-f(bx)}{x}\, dx = (f(0) - f(\infty))\ln(b/a)$ under mild hypotheses), or manufacturing the same coupling by hand via the Fundamental Theorem of Calculus in a parameter, followed by Fubini/Tonelli — or, equivalently, differentiating under the integral sign.

3. Hesitation occurs right after linearity looks tempting: does the difference actually converge, or does it just look like it should? It is resolved by checking the behaviour of each piece at $x \to \infty$ *before* splitting — both are $\sim (\pi/2)/x$, an identical divergent rate, which is the signal that the two pieces must stay coupled.

4. The trap is splitting $\int_0^{\infty} [f(ax) - f(bx)]/x\, dx$ into two separate integrals by linearity; each is individually infinite, and linearity of the integral does not license splitting a convergent integral into two divergent ones and reuniting them later.

5. Recognisable faster next time by keeping the named pattern in memory: $\dfrac{f(ax) - f(bx)}{x}$ integrated over $(0, \infty)$ is a Frullani integral, and its value is $(f(0) - f(\infty))\ln(b/a)$ whenever

$f(0)$ and $f(\infty)$ exist and f is suitably continuous — check the hypotheses, then quote the value, rather than re-deriving from scratch every time.

6. Reusable decision rule: when an integrand is a difference of the same function evaluated at two different linear scalings of the variable, divided by that variable, do not split by linearity; either apply Frullani's theorem or introduce a parameter (via the Fundamental Theorem of Calculus or Leibniz's rule) so the two values never exist as separate, individually divergent objects.

▷ **Cue:** integrand of the shape $[f(ax) - f(bx)]/x$ over $(0, \infty)$ $\longrightarrow$ **Move:** recognise the Frullani pattern; do not split by linearity, since each half typically diverges at the same rate

▷ **Cue:** a difference of the same function at two fixed constants, $g(\pi) - g(1)$, sits inside a harder integral $\longrightarrow$ **Move:** rewrite as $\int_1^\pi g'(u)\,du$ (Fundamental Theorem of Calculus) to trade the difference for an integral over an auxiliary parameter

▷ **Cue:** a positive double integral results from introducing an auxiliary parameter $\longrightarrow$ **Move:** swap the order of integration (Fubini/Tonelli) to make the inner integral elementary

▷ **Cue:** a definite integral depends on a constant that is hard to plug in directly $\longrightarrow$ **Move:** introduce that constant as a parameter a, differentiate under the integral sign to get an elementary $f'(a)$, then integrate back and fix the constant using an easy value of a

1982 A4

1982 A4. Assume that the system of simultaneous differential equations

$$y' = -z^3, \qquad z' = y^3$$

with the initial conditions $y(0) = 1$, $z(0) = 0$ has a unique solution $y = f(x)$, $z = g(x)$ defined for all real x. Prove that there exists a positive constant L such that for all real x,

$$f(x + L) = f(x), \qquad g(x + L) = g(x).$$

Fields: differential equations (autonomous systems, conserved quantities), real analysis **Position:** A4 — mid-paper; the idea (a conserved quantity plus "it goes round") arrives quickly, and the marks are in making "it goes round" rigorous **Answer:** The point (f, g) runs round the curve $y^4 + z^4 = 1$ anticlockwise with angular speed between $1/\sqrt{2}$ and 1, so it returns to $(1, 0)$ at some time $L > 0$, and uniqueness makes the motion L-periodic. In fact $L = \int_0^{2\pi} \frac{2\,d\theta}{\sqrt{3+\cos 4\theta}} = \frac{\Gamma(1/4)^2}{\sqrt{\pi}} \approx 7.4163$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You recognise the shape. If the cubes were first powers, $y' = -z$, $z' = y$ would give $y = \cos x$, $z = \sin x$, period 2π. So you look for an explicit solution: differentiate again, $y'' = -3z^2z' = -3z^2y^3$, and you are now staring at a non-linear second-order equation with no standard method attached. You try separating: $dz/dy = -y^3/z^3$, so $z^3\,dz = -y^3\,dy$ and $y^4 + z^4 = 1$. Good — the orbit is known. But the orbit is not the solution: you still need x as a function of position, $x = \int dz/y^3 = \int dz/(1 - z^4)^{3/4}$, which is not an elementary function.

So you write, as the official solution does: "the point moves on the closed curve $y^4 + z^4 = 1$ with speed at least some positive constant, so it goes all the way round in finite time L, and then it repeats." That is the right picture. But notice what it quietly assumes. Speed bounded below means the *arc length* travelled grows at least linearly. It does not by itself say the point makes progress *round* the curve rather than going back and forth on an arc; speed is the size of the velocity, not its direction. You need a monotone coordinate on the curve, and the speed argument does not supply one. That is the wall: the picture is right and the sentence proving it is missing.

The model chain. *Step 1 (conserved quantity).* Along the solution,

$$\frac{d}{dx}(f^4+g^4) = 4f^3f' + 4g^3g' = 4f^3(-g^3) + 4g^3f^3 = 0,$$

so $f^4+g^4 \equiv f(0)^4+g(0)^4 = 1$. In particular $(f,g) \neq (0,0)$ for every x.

Step 2 (an angle that only increases). Write $\rho^2 = f^2+g^2 > 0$. Because (f,g) is a differentiable curve avoiding the origin, it has a continuous (indeed differentiable) polar angle $\theta(x)$ with $\theta(0)=0$, $f=\rho\cos\theta$, $g=\rho\sin\theta$, and

$$\theta' = \frac{fg'-gf'}{f^2+g^2} = \frac{f\cdot f^3 - g\cdot(-g^3)}{\rho^2} = \frac{f^4+g^4}{\rho^2} = \frac{1}{\rho^2}.$$

Now bound ρ^2 on the curve. Since $(f^2+g^2)^2 = f^4+g^4+2f^2g^2 \geq 1$, we get $\rho^2 \geq 1$; and since $2f^2g^2 \leq f^4+g^4$, we get $(f^2+g^2)^2 \leq 2$, so $\rho^2 \leq \sqrt{2}$. Hence

$$\tfrac{1}{\sqrt{2}} \leq \theta'(x) \leq 1 \qquad \text{for all } x.$$

Step 3 (it comes back). θ is continuous, $\theta(0)=0$, and $\theta(x) \geq x/\sqrt{2}$ for $x \geq 0$, so by the intermediate value theorem there is $L>0$ with $\theta(L) = 2\pi$ (unique, as θ is strictly increasing; and $L \geq 2\pi$ because $\theta' \leq 1$). At $x=L$ the point $(f(L),g(L))$ lies on the positive y-axis direction, i.e. $g(L)=0$, $f(L)>0$, and $f(L)^4=1$. So $(f(L),g(L)) = (1,0)$.

Step 4 (uniqueness does the rest). Define $F(x)=f(x+L)$, $G(x)=g(x+L)$. The system does not contain x explicitly, so $F'=-G^3$, $G'=F^3$, and $F(0)=1$, $G(0)=0$. By the assumed uniqueness, $F=f$ and $G=g$, i.e. $f(x+L)=f(x)$, $g(x+L)=g(x)$ for all real x. □

The period, as a check. On the curve, $\rho^{-4} = \cos^4\theta + \sin^4\theta = \frac{1}{4}(3+\cos 4\theta)$, so $\theta' = \frac{1}{2}\sqrt{3+\cos 4\theta}$ is a function of θ alone and $L = \int_0^{2\pi} 2\,d\theta/\sqrt{3+\cos 4\theta}$. Alternatively the rotation $(y,z)\mapsto(-z,y)$ maps solutions to solutions, so each quarter-turn takes the same time $\int_0^1 dz/(1-z^4)^{3/4} = \frac{1}{4}B(\frac{1}{4},\frac{1}{4})$, giving $L = B(\frac{1}{4},\frac{1}{4}) = \Gamma(\frac{1}{4})^2/\sqrt{\pi} \approx 7.4163$. Both expressions agree numerically, and a numerical integration of the system returns to $(1,0)$ at this L.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains find $y^4+z^4=1$. They split at the sentence "so it goes all the way round". The natural attempt justifies it by speed, which measures how fast, not which way. The model replaces speed by the polar angle and computes θ', which is where the problem turns out to be designed: $\theta' = (y^4+z^4)/\rho^2$, and the numerator is the conserved quantity.

> *Going round a closed curve means some angle increases by 2π. So I want an angle, not a speed. The angle of a point about the origin has derivative $(yz'-zy')/(y^2+z^2)$, and substituting the equations gives $(y^4+z^4)/(y^2+z^2)$ — the numerator is my conserved quantity, so $\theta' = 1/\rho^2 > 0$.*

> *Once it is back at the starting point, I do not need to know anything else about the motion. The equations do not contain x, so the shifted solution solves the same initial value problem, and the problem has handed me uniqueness precisely so that I can say "therefore they are equal".*

Layer 3 THE WHY

Why look for a conserved quantity first? Because the question asks for periodicity, and periodicity of an autonomous system is a statement about the orbit being a closed curve traversed repeatedly. The first thing to find is the curve. For a planar system $y'=P$, $z'=Q$, the curve is a level set of some H with $H_yP + H_zQ = 0$; here $P=-z^3$, $Q=y^3$ pair off against $H_y = 4y^3$, $H_z = 4z^3$ in exactly the way $y'=-z$, $z'=y$ pairs against y^2+z^2. You are

not guessing: you are asking "which H makes $H_y \cdot (-z^3) + H_z \cdot y^3$ vanish?", and the answer $H_y \propto y^3$, $H_z \propto z^3$ is forced.

Why the angle, specifically? A closed curve that surrounds the origin is parametrised by the polar angle (it meets each ray from the origin once, because on the ray at angle θ the equation $\rho^4(\cos^4\theta + \sin^4\theta) = 1$ has exactly one positive root). So "the point goes round" is literally "θ increases by 2π", and that is a one-variable statement about a real function — exactly what the intermediate value theorem handles. The speed argument tries to prove a statement about position using the size of velocity; the angle argument proves it using the component of velocity that measures progress round the curve. The model is not smarter than the natural attempt; it is aimed at the right quantity.

Why is θ' positive, and not by luck? The linear model $y' = -z$, $z' = y$ is rotation: velocity perpendicular to position, pointing anticlockwise. Here the velocity is $(-z^3, y^3)$, and its cross product with the position (y, z) is $y \cdot y^3 - z \cdot (-z^3) = y^4 + z^4$. The cubes preserve signs, so the velocity still has the rotational orientation; the cross product is a sum of fourth powers and cannot vanish off the origin. That is why the conserved quantity reappears in θ': the same sum of fourth powers measures both "how far from the origin" and "how much rotation".

Why the hypothesis about uniqueness? The official solution's last line, "As the speed depends only on y and z (and not on x), the motion is periodic", is uniqueness in disguise. The problem states uniqueness because the right-hand side is a polynomial (so locally Lipschitz, so Picard–Lindelöf applies) and the setters did not want to examine that. Use it by name. Returning to the start is not the same as being periodic unless the future is determined by the present, and uniqueness is exactly that sentence.

Layer $+\alpha$ BACKTESTING

1. The cue is an autonomous planar system whose right-hand sides are "rotation-shaped" (y' involves $-z$, z' involves y). Write down: find H with $\dot{H} = 0$; here $H = y^4 + z^4$.

2. Conserved quantity for the orbit; polar angle for progress along it; intermediate value theorem for the return time; uniqueness for periodicity.

3. Hesitation comes at "so it goes round". Reinforce it away by asking "what real-valued function increases by exactly 2π on one lap?" and computing its derivative from the equations.

4. The trap is trying to solve the system explicitly, which leads to non-elementary integrals. The subtler trap is accepting "speed bounded below" as a proof of circulation.

5. Any time a problem says "periodic" about a solution of an autonomous equation, the proof is: show the state returns to its initial value, then invoke uniqueness. Look for the return via a monotone angle or other monotone coordinate on the orbit.

6. Reusable rule: to prove a trajectory circulates, compute the derivative of its polar angle, $\theta' = (yz' - zy')/(y^2 + z^2)$, and bound it away from zero; to prove periodicity of an autonomous system, show one return to the initial state and cite uniqueness.

▷ **Cue:** planar system $y' = -P(z)$, $z' = Q(y)$ with P, Q odd $\longrightarrow$ **Move:** conserved quantity $\int Q\,dy + \int P\,dz$; here $y^4 + z^4$

▷ **Cue:** need to show a point goes round a closed curve about the origin $\longrightarrow$ **Move:** polar angle $\theta' = (yz' - zy')/(y^2 + z^2)$; show it is bounded below by a positive constant

▷ **Cue:** "prove the solution is periodic" for an autonomous system with uniqueness given $\longrightarrow$ **Move:** show the state returns to its initial value once; translate in time and invoke uniqueness

▷ **Cue:** quantity $y^2 + z^2$ on the curve $y^4 + z^4 = 1$ $\longrightarrow$ **Move:** $1 \le (y^2 + z^2)^2 \le 2$ from $0 \le 2y^2z^2 \le y^4 + z^4$

▷ **Cue:** explicit solution requires a non-elementary integral $\longrightarrow$ **Move:** stop solving; prove the qualitative property from invariants and monotonicity

1982 A5

1982 A5. Let a, b, c, and d be positive integers and

$$r = 1 - \frac{a}{b} - \frac{c}{d}.$$

Given that $a + c \leq 1982$ and $r > 0$, prove that

$$r > \frac{1}{1983^3}.$$

Fields: number theory (Diophantine approximation, integer gaps), inequalities **Position:** A5 — late-paper by position, but short once the right case split is seen; the danger is spending the time on the wrong extremal problem **Answer:** Since $r = (bd - ad - bc)/(bd)$ has a positive integer numerator, $r \geq 1/(bd)$. Order so that $b \leq d$. If $b \geq 1983$ then $r \geq 1 - \frac{a+c}{1983} \geq \frac{1}{1983}$. If $b \leq 1982$, the smallest admissible d is at most $bc + 1$, so $r \geq \frac{1}{b(bc+1)} \geq \frac{1}{1982(1982 \cdot 1981 + 1)} > \frac{1}{1983^3}$. The exponent 3 is the right order: $a = 1321$, $b = 1322$, $c = 661$, $d = 661 \cdot 1322 + 1$ gives $r \approx 6.75/1983^3$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You see a fraction to be bounded below and reach for a common denominator:

$$r = \frac{bd - ad - bc}{bd}.$$

The numerator is a positive integer, so $r \geq 1/(bd)$. So far so good. Now you want $bd < 1983^3$, and you realise b and d are not bounded by anything in the hypotheses. Only $a + c$ is. You try to bound b and d in terms of a and c and cannot: a/b can be tiny with b enormous.

The second instinct is to find the extremal configuration. Make $a/b + c/d$ as close to 1 as possible from below with $a + c \leq 1982$. You guess a/b close to 1, say $a = b - 1$, then choose c/d just under $1/b$. That gives $r = \frac{1}{b(cb+1)}$, and optimising b^2c subject to $(b - 1) + c = 1982$ is a calculus exercise with answer $b \approx \frac{2}{3} \cdot 1982$. You now have a candidate minimum around $\frac{27}{4 \cdot 1982^3}$. But you have not proved it is the minimum; you have only evaluated one family. The extremal route produces the right *size* and no proof. That is the wall.

The model chain. Put $N = bd - ad - bc$, an integer; $r = N/(bd)$ and $r > 0$ give $N \geq 1$, so

$$r \geq \frac{1}{bd}. \tag{1}$$

The hypotheses are symmetric under $(a, b) \leftrightarrow (c, d)$, so we may assume $b \leq d$.

Case 1: $b \geq 1983$. Then $d \geq b \geq 1983$, so $\frac{a}{b} + \frac{c}{d} \leq \frac{a+c}{1983} \leq \frac{1982}{1983}$ and $r \geq \frac{1}{1983} > \frac{1}{1983^3}$.

Case 2: $b \leq 1982$. Here b is small but d may be huge. Fix a, b, c. Since $r > 0$ we have $a < b$, and $r > 0$ reads

$$\frac{c}{d} < \frac{b - a}{b} \iff d > \frac{bc}{b - a}.$$

Let $d_0 = \lfloor bc/(b - a) \rfloor + 1$, the least integer satisfying this. As a function of d, $r(d) = \frac{b-a}{b} - \frac{c}{d}$ is increasing, and our d is admissible, so $d \geq d_0$ and $r(d) \geq r(d_0)$. Applying (1) with d_0 in place of d (legitimate, since $r(d_0) > 0$),

$$r \geq r(d_0) \geq \frac{1}{bd_0}.$$

Now $b - a \geq 1$ gives $d_0 \leq \frac{bc}{b-a} + 1 \leq bc + 1$. Also $a \geq 1$ gives $c \leq 1981$. Therefore

$$bd_0 \leq b(bc + 1) \leq 1982\,(1982 \cdot 1981 + 1) < 1982 \cdot 1983^2 < 1983^3,$$

using $1982 \cdot 1981 + 1 = 1982^2 - 1981 < 1983^2$. Hence $r \geq 1/(bd_0) > 1/1983^3$.

In both cases $r > 1/1983^3$. □

Sanity checks. Replace 1982 by a small n and 1983 by $n+1$: an exhaustive search for $n \leq 8$ finds the minimum of $r(n+1)^3$ between 4.5 and 6.4, always above 1. The family $a = k$, $b = k+1$, $d = c(k+1)+1$ with $k + c = 1982$ gives $r = \frac{1}{(k+1)(c(k+1)+1)}$, smallest at $k = 1321$, $c = 661$, where $r \cdot 1983^3 \approx 6.75$. So the bound is off by a constant factor only.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach $r \geq 1/(bd)$. They split at the next line. The natural attempt tries to bound b and d simultaneously, or to find the global minimiser. The model notices that the bound (1) is only useful when bd is small, that the crude bound $r \geq 1 - \frac{a+c}{\min(b,d)}$ is useful when both denominators are large, and that the only remaining case has one small denominator, where the other can be pushed down to its least admissible value.

> *The integer-gap bound $1/(bd)$ is strong when the denominators are small and useless when they are large. When both are large the fractions themselves are small, and $a + c \leq 1982$ finishes it. So split on the size of the denominators, and the only case left has one small and one large.*

> *With a, b, c fixed, r is increasing in d. So the worst d is the smallest one that keeps $r > 0$, and that d is at most $bc + 1$. I do not need the extremal configuration; I need an upper bound on the worst denominator.*

Layer 3 THE WHY

Why start with the common denominator? Because the problem is a lower bound on a positive rational with integer data, and the only universal lower bound on a positive rational is "numerator at least 1". Every Diophantine-approximation argument begins here: $p/q > 0$ with p an integer means $p/q \geq 1/q$. Everything else is controlling q.

Why does controlling $q = bd$ require cases? Look at what the hypothesis actually restricts. It restricts numerators, $a + c \leq 1982$, not denominators. So there are two regimes, and the problem's number 1983 is the boundary between them. If a denominator exceeds 1982, its fraction is at most $\frac{\text{numerator}}{1983}$, and the numerators share a budget of 1982; if both denominators exceed 1982, the two fractions together fall short of 1 by at least $1/1983$ — no integer arithmetic needed. If a denominator is at most 1982, it is small enough to feed the integer-gap bound. The number $1983 = 1982 + 1$ in the conclusion is a signal that exactly this split is intended.

Why is the mixed case not hopeless, given that d is unbounded? Because an unbounded d makes c/d small, which makes r *larger*. The danger is not large d; it is d just large enough. This is the monotonicity observation: for fixed a, b, c, r increases with d, so the adversary chooses the least d with $c/d < (b-a)/b$. That least value is a floor plus one, $\lfloor bc/(b-a) \rfloor + 1$, and dividing by $b - a \geq 1$ can only help, so it is at most $bc + 1$. Now both denominators of the worst case are bounded in terms of the data: $b \leq 1982$ and $d_0 \leq 1982 \cdot 1981 + 1$. The cubic exponent is the product $b \cdot b \cdot c$.

The official solution does exactly this but compresses the reason into "the d that minimizes r for fixed a, b, c is $1 + \lfloor bc/(b-a) \rfloor$". Here is where that comes from: r is monotone in the one unbounded variable, so its worst value sits at the boundary of the constraint $r > 0$.

Why is the extremal approach a trap here? Finding the true minimum of r is a genuinely harder problem (it involves the structure of Farey neighbours) and the question does not ask for it. The bound actually proved in Case 2, $1/(b(bc+1))$, throws away information at every step ($b \leq 1982$, $c \leq 1981$, $b - a \geq 1$ all used at their worst simultaneously, which no single configuration achieves), and it still clears $1/1983^3$. A statement that survives that much crudeness is an invitation to be crude.

Layer $+\alpha$ BACKTESTING

1. The cue is "positive rational built from integers, prove it exceeds a tiny explicit number". Write down the common denominator and "numerator ≥ 1".

2. Integer-gap bound $r \geq 1/(\text{denominator})$, combined with a case split on the size of the denominators, and monotonicity in the unbounded variable to reduce to its least admissible value.

3. Hesitation occurs when you notice b, d are unbounded. Reinforce it away by asking "which way does r move when the unbounded variable grows?" — if it helps you, the worst case is the smallest admissible value.

4. The trap is trying to find the exact minimiser. It gives the right order of magnitude and no proof, and it eats the time.

5. When a hypothesis bounds numerators but the conclusion needs bounded denominators, split on whether denominators exceed the numerator budget; the threshold is usually visible in the numbers of the statement (1982 versus 1983).

6. Reusable rule: to bound a positive quantity below, find the variable it is monotone in, push that variable to the edge of the feasible region, and bound the quantity there; a constant-factor-loose bound needs only the boundary, never the global extremum.

▷ **Cue:** positive rational p/q with integer p $\longrightarrow$ **Move:** $p/q \geq 1/q$; now bound q

▷ **Cue:** hypothesis bounds numerators, conclusion needs denominators bounded $\longrightarrow$ **Move:** split on whether denominators exceed the numerator budget; large denominators make fractions small directly

▷ **Cue:** quantity to bound below is monotone in an unbounded variable $\longrightarrow$ **Move:** the worst case is the least admissible value of that variable; compute it as a floor plus one

▷ **Cue:** constant $N+1$ in the conclusion and N in the hypothesis $\longrightarrow$ **Move:** expect a case split at denominator N versus $N+1$

▷ **Cue:** "prove $r > \varepsilon$" with ε much smaller than typical $\longrightarrow$ **Move:** do not seek the exact minimum; crude case bounds suffice

1982 A6

1982 A6. Let σ be a bijection of the positive integers, that is, a one-to-one function from $\{1, 2, 3, \dots\}$ onto itself. Let $x_1, x_2, x_3, \dots$ be a sequence of real numbers with the following three properties:

(i) $|x_n|$ is a strictly decreasing function of n;

(ii) $|\sigma(n) - n| \cdot |x_n| \to 0$ as $n \to \infty$;

(iii) $\lim_{n\to\infty} \sum_{k=1}^{n} x_k = 1$.

Prove or disprove that these conditions imply that

$$\lim_{n\to\infty} \sum_{k=1}^{n} x_{\sigma(k)} = 1.$$

Fields: real analysis (conditionally convergent series, rearrangements), construction of counterexamples **Position:** A6 — hardest slot; the answer is a counterexample, and building it needs two independent design decisions (the size of the terms and the shape of the permutation) **Answer:** False. Take $y_n = \frac{1}{(n+1)\ln(n+1)}$, $g = \sum_{n\geq 1}(-1)^{n+1} y_n > 0$, $x_n = (-1)^{n+1} y_n / g$, and let σ act on consecutive blocks of length $4a_i$ by listing the $2a_i$ odd (positive) positions of the block first and then the $2a_i$ even (negative) positions. Then $|\sigma(n) - n| < n$, so (ii) holds because $n y_n \to 0$; and with the a_i large, the rearranged

partial sums exceed the original ones by more than 1 infinitely often, so the rearranged series does not converge to 1 (in fact it diverges).

Layer 1 THE WHAT — surface analysis

The natural attempt. Condition (ii) says the permutation moves n by an amount that is small compared with $1/|x_n|$. The natural heuristic: the rearranged partial sum $D(n) = \sum_{k\le n} x_{\sigma(k)}$ differs from $C(n) = \sum_{k\le n} x_k$ only by the terms that have been moved across the cut at n; each moved term travels a distance $|\sigma(k) - k|$ and has size about $|x_k|$; so the discrepancy is like "displacement $\times$ size", which (ii) sends to 0. So you set out to prove the statement. You write

$$D(n) - C(n) = \sum_{m\in A_n} x_m - \sum_{m\in B_n} x_m, \quad A_n = \sigma(\{1,\dots,n\}) \setminus \{1,\dots,n\},\ B_n = \{1,\dots,n\} \setminus \sigma(\{1,\dots,n\}),$$

and try to show both sums are small. You need $|A_n|$ small, or the terms in it small. Condition (ii) gives neither directly: it controls $|\sigma(k) - k|$ *weighted by* $|x_k|$, the size of the term at the step k, not the size of the term $x_{\sigma(k)}$ actually being moved. A step k where x_k is tiny may fetch a term $x_{\sigma(k)}$ from far away, and the terms left behind in B_n can be many and not tiny. The heuristic does not become a proof, and the gap is exactly where a counterexample can live.

The model chain. *The terms.* Let $y_n = \frac{1}{(n+1)\ln(n+1)}$ for $n \ge 1$. Then y_n is strictly decreasing with limit 0, so by the alternating series test $\sum(-1)^{n+1} y_n$ converges to some g with $g \ge y_1 - y_2 > 0$. Put $x_n = (-1)^{n+1} y_n / g$. Then (i) holds, (iii) holds, $x_n > 0$ for odd n and $x_n < 0$ for even n. Also $\sum y_n = \infty$ by the integral test ($\int \frac{dt}{t\ln t} = \ln\ln t$), and $n y_n = \frac{n}{(n+1)\ln(n+1)} \to 0$.

The permutation. Choose positive integers $a_0, a_1, \dots$ (specified below), set $b_0 = 0$, $b_{i+1} = b_i + 4a_i$, and on the block $b_i < n \le b_{i+1}$ define

$$\sigma(n) = \begin{cases} 2n - 1 - b_i, & b_i < n \le b_i + 2a_i, \\ 2n - b_{i+1}, & b_i + 2a_i < n \le b_{i+1}. \end{cases}$$

Writing $n = b_i + j$: for $1 \le j \le 2a_i$ we get $\sigma(n) = b_i + 2j - 1$, running through the odd offsets $1, 3, \dots, 4a_i - 1$; for $2a_i < j \le 4a_i$ we get $\sigma(n) = b_i + 2(j - 2a_i)$, running through the even offsets $2, 4, \dots, 4a_i$. So σ permutes each block and is a bijection of the positive integers. Since every b_i is even, odd offsets are odd positions: in each block σ takes all $2a_i$ positive terms first, then all $2a_i$ negative ones.

Condition (ii). In the first half of a block, $\sigma(n) - n = n - 1 - b_i \in [0, n)$; in the second half, $n - \sigma(n) = b_{i+1} - n \in [0, 2a_i)$ and $n > b_i + 2a_i \ge 2a_i$. So $|\sigma(n) - n| < n$ always, and $|\sigma(n) - n|\,|x_n| < n y_n / g \to 0$.

The discrepancy. Let $C(n) = \sum_{k\le n} x_k$, $D(n) = \sum_{k\le n} x_{\sigma(k)}$. Since σ permutes each complete block, $D(b_i) = C(b_i)$. At $n = b_i + 2a_i$, D has added the $2a_i$ positive terms of the block, while C has added the first $2a_i$ terms, a_i positive and a_i negative. Hence

$$D(b_i + 2a_i) - C(b_i + 2a_i) = \frac{1}{g}\sum_{j=1}^{a_i} y_{b_i+2j} + \frac{1}{g}\sum_{k=1}^{a_i} y_{b_i+2a_i+2k-1}, \tag{A}$$

the first sum being the negatives C took and D did not, the second the positives D took and C did not. Both sums are positive. Since y is decreasing, $2\sum_{j\ge1} y_{b+2j} \ge \sum_{m>b+1} y_m = \infty$ for any fixed b, so we may choose each a_i (recursively, once b_i is known) with $\sum_{j=1}^{a_i} y_{b_i+2j} > g$. Then

$$D(b_i + 2a_i) > C(b_i + 2a_i) + 1 \qquad \text{for every } i.$$

As $i \to \infty$, $C(b_i + 2a_i) \to 1$, so $D(n) > 1.9$ for infinitely many n, while $D(b_i) = C(b_i) \to 1$. Thus $D(n)$ does not converge at all, and in particular not to 1. The assertion is false. □

Numerical check. With $g \approx 0.5264$ the first two blocks need $a_0 = 4$ and $a_1 = 1974$; the construction gives $D(8) - C(8) \approx 1.29$ and $D(3964) - C(3964) \approx 1.08$, while $D = C$ at the block ends 16 and 7912, and $|\sigma(n) - n||x_n|$ stays below 0.23 on the second half of the range and tends to 0.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write $D(n) - C(n)$ as "terms fetched early minus terms left behind". They split on how to read condition (ii). The natural attempt reads it as "each moved term is small times its displacement" and tries to prove the statement. The model reads it literally: it weights the displacement of step k by $|x_k|$, the size of the term that *would* have been used at step k, and in a slowly decaying series a term far down the line is much smaller than one near the start of a long block.

> *If $\sum |x_n|$ converged, every rearrangement would give the same sum. So any counterexample is conditionally convergent, and I should use the standard Riemann-rearrangement move: take positives early. The question is only whether I can do that while keeping $|\sigma(n) - n|\,|x_n| \to 0$.*

> *Any rearrangement that takes positives early in long blocks moves terms by distances comparable to n. So I need $n|x_n| \to 0$ while $\sum |x_n|$ still diverges, so that the delayed negatives in a block can add up to more than 1. That is exactly $|x_n| \approx 1/(n \ln n)$.*

Layer 3 THE WHY

Why is the answer "disprove"? Consider what would make it true. If σ moved every index by a bounded amount M, then $D(n)$ and $C(n)$ would differ by at most $2M$ terms all with index near n, so the difference tends to 0. Condition (ii) is a weakening of bounded displacement: it allows the displacement to grow as long as it is $o(1/|x_n|)$. For $x_n \approx 1/n$ that is $o(n)$, and a sum over a mismatch set of size $o(n)$ of terms about $1/n$ is still $o(1)$. The trouble is that the terms in the mismatch set need not be about $1/n$: they can be terms from the start of a long block, which are much larger, and the condition never looks at them. This is the Riemann rearrangement theorem's mechanism — conditional convergence means the positive part and the negative part both diverge, and delaying one part lets the other run away — slowed down to the speed (ii) permits.

Why $1/(n \ln n)$ specifically? Two requirements pull in opposite directions. The delayed negatives in a block must sum to more than a fixed amount, so $\sum |x_n|$ must diverge: $|x_n|$ cannot decay like $n^{-1-\varepsilon}$. The permutation moves indices near the middle of a block by about the block length, which is comparable to n, so (ii) needs $n|x_n| \to 0$: $|x_n|$ cannot be as large as $1/n$. The canonical sequence that is $o(1/n)$ with a divergent sum is $1/(n \ln n)$ — the integral test puts it on the boundary, with $\sum_{n \le N} \approx \ln \ln N$. Once both requirements are written down, the choice is forced, not inspired.

Why blocks of length $4a_i$ with a_i huge? The divergence of $\sum 1/(n \ln n)$ is very slow: to gain a fixed amount g from the even terms after b_i you need $\ln(b_i + 2a_i)$ to be about $e^{2g} \ln b_i$, so a_i is a large power of b_i. The blocks must be long enough to let the slow divergence accumulate. Making the block a permutation of itself guarantees that $D = C$ at every block end, which also shows 1 is a limit point of D, so D diverges.

Why block "positives first" rather than a cleverer rearrangement? It is the simplest shape that makes $D - C$ explicit. Each block is split into halves by a single parity rule, $\sigma(n)$ is a linear formula in each half, and the discrepancy (A) is a sum of positive terms with no cancellation to track. The official solution gives exactly this construction; its unstated design constraint is the pair of opposing requirements above.

Layer $+\alpha$ BACKTESTING

1. The cue is a "prove or disprove" about rearranging a series under a condition that weights displacement by a term size. Write down: "absolutely convergent series rearrange freely, so a counterexample must be conditionally convergent".

2. Riemann rearrangement: in long blocks, take the positive terms before the negative ones; tune the decay of the terms so that the displacement condition still holds.

3. Hesitation comes from the "displacement times size" heuristic, which suggests the statement is true. Reinforce it away by reading the condition literally: which term's size multiplies the displacement, and is it the term that moves?

4. The trap is trying to prove the statement, or taking $x_n = (-1)^{n+1}/n$ as the test sequence: with it, block rearrangements violate (ii), and one wrongly concludes that no counterexample exists.

5. When two constraints pull a sequence in opposite directions ("sum diverges" but "$n\cdot$term $\to 0$"), write both down and pick the canonical sequence on the boundary: $1/(n \ln n)$, or $1/(n \ln n \ln\ln n)$ for a finer boundary.

6. Reusable rule: to refute an "implies convergence" statement, identify the mechanism the hypothesis fails to control (here, the size of the moved term), build the simplest block construction exploiting it, and choose the growth of block lengths after the fact to make the error exceed a fixed amount.

▷ **Cue:** rearrangement of a series preserves the sum — prove or disprove $\longrightarrow$ **Move:** counterexamples must be conditionally convergent; use positives-first blocks

▷ **Cue:** need $\sum a_n = \infty$ but $na_n \to 0$ $\longrightarrow$ **Move:** $a_n = 1/(n \ln n)$; integral test gives $\ln\ln N$ growth

▷ **Cue:** condition weights displacement by $|x_n|$ at step n $\longrightarrow$ **Move:** check whether the term actually moved is $x_{\sigma(n)}$; the condition may not see it

▷ **Cue:** permutation defined on blocks $(b_i, b_{i+1}]$ $\longrightarrow$ **Move:** partial sums agree at block ends; compute the discrepancy at mid-block only

▷ **Cue:** slowly divergent series must exceed a fixed bound inside a block $\longrightarrow$ **Move:** choose block lengths recursively after the previous blocks are fixed

Putnam 1982 — Paper B

1982 B1

1982 B1. Let M be the midpoint of side BC of a general $\triangle ABC$. Using the smallest possible n, describe a method for cutting $\triangle AMB$ into n triangles which can be reassembled to form a triangle congruent to $\triangle AMC$.

Fields: plane geometry (dissection via rigid motion, congruence) **Position:** B1 – opening problem of the second session **Answer:** $n = 2$. Cut $\triangle AMB$ along the segment MD, where D is the midpoint of AB; a 180° rotation of $\triangle DBM$ about D reassembles the two pieces into a triangle congruent to $\triangle AMC$.

Layer 1 THE WHAT — surface analysis

The natural attempt. $\triangle AMB$ and $\triangle AMC$ have equal area – same base length $BM = MC$, same height from A – so your first thought is the Bolyai–Gerwien theorem: any two polygons of equal area are equidecomposable, full stop, so some n certainly works. That settles existence but not the actual question, which asks for the *smallest* n and a *method*. You check $n = 1$ first: does $\triangle AMB$ already sit congruent to $\triangle AMC$ without cutting at all? The two triangles share the side AM and the equal pair $BM = MC$, so they fail to be congruent only if the remaining sides differ, i.e. $AB \neq AC$ – true in a general (scalene) triangle, so $n = 1$ is out. Past that you have no obvious next move: Bolyai–Gerwien's standard proof routes both polygons through a common rectangle, and doing that honestly for two triangles usually costs three or four pieces, not two. You might start drawing altitudes and dropping perpendiculars to build that rectangle, only to notice you have committed to a construction that will not obviously beat $n = 3$, let alone reach 2, and you have not used the one piece of extra structure the problem handed you for free: M is a midpoint.

The model chain. Use the midpoint that is already sitting unused: let D be the midpoint of AB, and cut $\triangle AMB$ along MD into $\triangle ADM$ and $\triangle DBM$.

Rotate $\triangle DBM$ by 180° about D. This is an isometry, so it preserves the triangle's shape and size. Work in vectors with D as a temporary origin for the map: a 180° rotation about D sends any point P to $2D - P$. Since $D = \frac{1}{2}(A + B)$, it sends $B \mapsto 2D - B = A + B - B = A$, and it sends $M \mapsto M' := 2D - M = A + B - M$.

So the rotated copy of $\triangle DBM$ is $\triangle DAM'$, and it shares the full edge AD with $\triangle ADM$ (both triangles have this edge, and the rotation fixes D while sending $B \to A$, so the edge DB of the first piece lands exactly on the edge DA of the second). The two triangles fit together along AD without gap or overlap, so their union is bounded by AM, AM', and the pair of segments DM, DM'. These two segments are collinear – because M and $M' = 2D - M$ are, by construction, reflections of each other through D, so D is their midpoint and M, D, M' lie on one line. Hence $DM \cup DM'$ is a single straight edge MM', and the reassembled figure is exactly the triangle AMM', with D an interior point of side MM', not a vertex.

Now identify this triangle. Using $M = \frac{1}{2}(B + C)$:

$$M' - M = (A + B - M) - M = A + B - 2M = A + B - (B + C) = A - C,$$

so MM' has the same length and direction as CA: $|MM'| = |CA|$. Also

$$M' - A = (A + B - M) - A = B - M,$$

so $|AM'| = |BM| = |MC|$, using M the midpoint of BC. Finally AM is common to both triangles AMM' and AMC. So $\triangle AMM'$ has side lengths $\{AM,\ MC,\ CA\}$ under the correspondence $A \leftrightarrow A,\ M \leftrightarrow M,\ M' \leftrightarrow C$ – exactly the side lengths of $\triangle AMC$, matched vertex for vertex. By SSS, $\triangle AMM' \cong \triangle AMC$.

So $n = 2$ suffices, and it fails only if $n = 1$ can be forced to work, which the argument above rules out for a general (non-isosceles) triangle. Hence $n = 2$ is smallest.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree that $n = 1$ is impossible in general, by comparing the side triples $\{AM, BM, AB\}$ and $\{AM, MC, AC\}$. They split immediately after that, at the question of *which extra structure to invoke for* $n \geq 2$: the natural attempt reaches for Bolyai–Gerwien, an existence theorem that answers "can some dissection be found" but is silent on how few pieces it takes and generically needs several. The model instead asks what isometry is already lying in the figure for free, and answers: M is a midpoint, so is the natural second cut point, D, the midpoint of AB – because a midpoint is exactly the fixed point of the 180° rotation that swaps its two endpoints.

> *I already have one midpoint, M, doing nothing for me yet. If I want an isometry that relates the two halves of my cut for free, I should cut through a second midpoint, D, because rotating* 180° *about a midpoint is the one motion that swaps its endpoints exactly – $B \to A$ – without me having to construct anything.*

> *Once I rotate $\triangle DBM$ about D, I don't need to guess whether the two pieces glue into a triangle – I should check whether M and its image M' end up on a line through D, because if D is their midpoint, the two collinear half-edges DM and DM' merge into one straight side and the two pieces really do form a triangle, not a quadrilateral.*

Layer 3 THE WHY

Why does a midpoint license this rotation for free, rather than counting as a discovered trick? A midpoint of a segment is, by definition, the unique point equidistant from both ends and lying between them – which is precisely the fixed point of the 180° rotation exchanging the two endpoints. That rotation is not something you construct; it is already latent in the phrase "let D be the midpoint of AB" the moment you write it down. So whenever a dissection or transformation problem hands you a midpoint, the first isometry to try costs nothing: rotate 180° about it and see what swaps. Here it swaps B and A, which is exactly the swap the problem wants, since we are relating a triangle built on B ($\triangle AMB$) to one built on C ($\triangle AMC$). The genuinely non-obvious step is realising there is a *second* midpoint available – M was given, but D has to be introduced – and the reason to introduce it at the midpoint of AB specifically, rather than anywhere else on AB, is that only the midpoint makes the rotation an isometry that fixes the cut point and swaps the vertices you need swapped.

Why does this beat Bolyai–Gerwien for minimality? Because Bolyai–Gerwien is an existence statement: it guarantees a common refinement into finitely many polygonal pieces exists whenever areas agree, but the standard proof (route both shapes through a rectangle) treats the two polygons as arbitrary and discards any structure specific to this pair. Here the two triangles are not arbitrary equal-area polygons – they are two halves of the same triangle, split by a cevian to a midpoint – and that shared ancestry is exactly the extra structure that collapses a generic multi-piece argument down to a single cut and a single rotation. The reusable principle:

an existence theorem answers whether a construction is possible; it should never be the tool you reach for once you also need the construction to be minimal. Minimality comes from structural symmetry already present in the problem, and midpoints, centres, and points of tangency are the standard places that symmetry hides.

The official solution states the rotation and the resulting congruence without proving it lands correctly (it says the pieces can be "placed alongside" one another to form $\triangle AMC$, and stops there). The vector computation above is where the archive supplies what is skipped: it verifies both that M, D, M' are collinear (so the union really is a triangle) and that the three side lengths match $\triangle AMC$'s exactly, rather than merely asserting a picture.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "M is the midpoint of BC". On seeing it, write down: a midpoint is the fixed point of the 180° rotation swapping its endpoints – so that rotation is available for free.

2. It suggests looking for a *second* midpoint to cut through, so that the same free-rotation trick relates the two triangles being compared, rather than reaching immediately for a general equidecomposition theorem.

3. Hesitation occurs right after ruling out $n = 1$, when the honest next thought is "some dissection must exist" (Bolyai–Gerwien) and it is tempting to stop there. It is reinforced away by asking: does the problem's specific structure (a midpoint cevian) let me beat the generic bound, rather than merely meet it?

4. The trap is treating a hard existence theorem as if it answered a minimality question. A second, smaller trap: asserting the rotated pieces "obviously" form a triangle without checking that the two half-edges are collinear – they could instead form a quadrilateral if the midpoint were placed anywhere else.

5. Whenever a dissection, folding, or congruence problem hands you a midpoint (or centre, or point of tangency), immediately ask what 180° rotation (or reflection) that point is the fixed point of, and whether that motion swaps the two objects the problem is relating.

6. Reusable rule: for a minimum-pieces dissection between two shapes that share structural ancestry (halves of one figure, images under a known symmetry), look for a free isometry fixed at an already-given special point before invoking a general existence theorem; verify collinearity/gluing by vectors, not by picture.

▷ **Cue:** a midpoint of a segment is given or constructible in a dissection/congruence problem $\longrightarrow$ **Move:** rotate 180° about that midpoint; it is the unique isometry fixing that point and swapping the segment's two endpoints, and costs nothing to invoke

▷ **Cue:** two polygons are known to have equal area and a minimal-piece dissection is wanted $\longrightarrow$ **Move:** do not reach for Bolyai–Gerwien (it proves existence only); look for structural symmetry already in the figure (a shared cevian, a common midpoint) that forces a small dissection directly

▷ **Cue:** after cutting and rotating a piece, you must check the pieces reassemble into a single triangle rather than a quadrilateral $\longrightarrow$ **Move:** verify by vectors that the two candidate half-edges are collinear, i.e. that the cut point is the midpoint of the two images, not merely that the picture looks right

1982 B2

1982 B2. Let $A(x, y)$ denote the number of points (m, n) in the plane with integer coordi-

nates m and n satisfying $m^2+n^2 \le x^2+y^2$. Let $g = \sum_{k=0}^{\infty} e^{-k^2}$. Express

$$\int_{-\infty}^{\infty}\int_{-\infty}^{\infty} A(x,y)\, e^{-x^2-y^2}\, dx\, dy$$

as a polynomial in g.

Fields: analysis (lattice-point counting, Gaussian integrals, Tonelli's theorem) **Position:** B2 – middle problem of the second session **Answer:** $\pi(2g-1)^2$.

Layer 1 THE WHAT — surface analysis

The natural attempt. $A(x,y)$ is the Gauss circle-counting function: the number of integer lattice points inside or on the circle of radius $r=\sqrt{x^2+y^2}$. Your instinct is to get a handle on it before integrating, so you reach for the standard estimate $A(x,y)=\pi r^2+E(r)$, where $E(r)$ is an error term – and you already half-remember that bounding $E(r)$ well is a genuinely hard, partly open problem in number theory (the Gauss circle problem). Even granting the crude bound $|E(r)|=O(r)$, you would need to integrate $\big(\pi r^2+E(r)\big)e^{-r^2}$ over the whole plane and show the error term's contribution vanishes or is controllable – a serious estimate, not a five-minute one, and overkill for a problem that gives you the tidy constant g up front as a hint that an exact answer, not an asymptotic one, is wanted. You are trying to integrate the counting function's *closed form*, and it does not have a clean one; you are stuck holding an approximation on a competition that wants an identity.

The model chain. Do not approximate $A(x,y)$; re-read its definition. "The number of lattice points (m,n) with $m^2+n^2 \le x^2+y^2$" is literally the sum of an indicator over all integer pairs:

$$A(x,y) = \sum_{m\in\mathbb{Z}}\sum_{n\in\mathbb{Z}} \mathbf{1}\big[m^2+n^2 \le x^2+y^2\big]\,.$$

Substitute this into the integral and swap the (countable) sum with the (double) integral. Every term is nonnegative, so Tonelli's theorem licenses the swap unconditionally – no absolute-convergence check is needed:

$$I := \iint_{\mathbb{R}^2} A(x,y)e^{-x^2-y^2}\,dx\,dy = \sum_{m,n\in\mathbb{Z}} \iint_{\{x^2+y^2\ge m^2+n^2\}} e^{-x^2-y^2}\,dx\,dy.$$

Each inner integral is now an honest, individually computable Gaussian integral over the exterior of a disk of radius $r_{mn}=\sqrt{m^2+n^2}$. In polar coordinates,

$$\int_0^{2\pi}\int_{r_{mn}}^{\infty} e^{-r^2}\, r\, dr\, d\theta = 2\pi\left[-\tfrac{1}{2}e^{-r^2}\right]_{r_{mn}}^{\infty} = \pi e^{-r_{mn}^2} = \pi e^{-m^2}e^{-n^2}.$$

So

$$I = \sum_{m,n\in\mathbb{Z}} \pi e^{-m^2}e^{-n^2} = \pi\left(\sum_{m\in\mathbb{Z}} e^{-m^2}\right)\left(\sum_{n\in\mathbb{Z}} e^{-n^2}\right) = \pi\left(\sum_{k\in\mathbb{Z}} e^{-k^2}\right)^2.$$

Finally convert the sum over *all* integers to g, which sums only $k\ge 0$: split off $k=0$ and pair k with $-k$ for $k\ge 1$,

$$\sum_{k\in\mathbb{Z}} e^{-k^2} = e^0 + 2\sum_{k=1}^{\infty} e^{-k^2} = 1+2(g-1) = 2g-1.$$

Hence

$$I = \pi(2g-1)^2.$$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree that $A(x,y)$ is the lattice-point counting function and that the target integral is $\iint A(x,y)e^{-x^2-y^2}\,dx\,dy$. They split at the moment $A(x,y)$ must actually be integrated: the natural attempt treats A as an unknown function of $r = \sqrt{x^2+y^2}$ to be estimated (Gauss circle asymptotics), while the model chain expands $A(x,y)$ back into the sum of indicators that its own definition already is, and integrates *that* instead.

> *$A(x,y)$ is not really a single function I need to approximate – the problem* defined *it as a count, and a count of points satisfying a condition is, by definition, a sum of indicator functions over those points. I should undo the collapse from "sum of indicators" to "counting function" before I integrate, not after.*

> *Every term in that sum is nonnegative, so I don't need to justify swapping $\sum$ and $\iint$ with an epsilon argument – Tonelli's theorem gives it to me for free whenever the summand is nonnegative, regardless of whether the total is finite. That means I can push the sum outside the integral first and worry about convergence of the final series afterward, not before.*

Layer 3 THE WHY

Why is unfolding $A(x,y)$ back into indicators the right move, rather than a trick spotted in hindsight? Because it costs nothing and loses nothing: the problem's own definition of $A(x,y)$ *is* that sum, in words rather than symbols. Collapsing "count the lattice points satisfying P" into a single opaque function A is a convenience for stating the problem, not a structural simplification – and once you need to integrate against that count, the convenience becomes an obstruction, because closed forms for lattice-point counts inside disks are exactly the object the Gauss circle problem is still open about. The rule is general: whenever an integral or sum involves a counting function of the form "the number of objects satisfying a condition", write it back out as a sum of indicators before doing calculus on it, because the resulting swap of sum and integral very often turns one hard, irregular object into infinitely many easy, regular ones.

Why is the swap itself free, and not something requiring a separate justification you have to search your memory for? Because the summand $\mathbf{1}[m^2+n^2 \le x^2+y^2]\,e^{-x^2-y^2}$ is nonnegative everywhere. Tonelli's theorem states that for nonnegative measurable functions, sum and integral (or any two integrals) may be interchanged without checking absolute convergence first – convergence of the rearranged expression is a consequence, not a precondition. This is worth internalising as a standing habit: the moment you see you are about to swap the order of two integrals or a sum and an integral, check nonnegativity first; if it holds, the swap needs no further defence, and you can go straight to computing the new, hopefully simpler, iterated expression.

Why does the geometry collapse to a single clean exponential per lattice point? Because $\{x^2+y^2 \ge m^2+n^2\}$ is exactly the exterior of a disk of radius $r_{mn} = \sqrt{m^2+n^2}$, and $e^{-x^2-y^2}$ is radially symmetric – so switching to polar coordinates turns a 2D integral over an annular (here, unbounded) region into a 1D integral in r with the Jacobian factor r, which is precisely what makes $\int re^{-r^2}dr$ elementary via the substitution $u = r^2$. This is the same mechanism as computing $\int_{\mathbb{R}^2} e^{-x^2-y^2}\,dx\,dy = \pi$ in the first place: radial symmetry plus the Jacobian r is what makes Gaussian integrals in the plane closed-form when they are not on the line.

Finally, why does g (summing from $k=0$) turn into $2g-1$ and not $2g$? Because e^{-k^2} is an even function of k, so pairing k with $-k$ for $k \ge 1$ double-counts every term except $k=0$, which has no partner. This is the standard bookkeeping whenever a problem defines a one-sided sum ($k \ge 0$) but the natural object is two-sided ($k \in \mathbb{Z}$): always isolate the unpaired term first.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "the number of points (m, n) … satisfying". On seeing it, write $A(x, y) = \sum_{m,n} \mathbf{1}[\cdot]$ immediately, before doing anything else with A.
2. It suggests swapping the order of summation and integration as the very next step, since the integral is now of a sum, not of an opaque function.
3. Hesitation occurs at the swap itself, from a half-remembered worry that interchanging limits needs justification. It is reinforced away by checking the summand is nonnegative and invoking Tonelli, which needs no further estimate.
4. The trap is treating $A(x, y)$ as a function to be approximated via Gauss circle asymptotics – a real, partly open, number-theoretic estimation problem that this Putnam problem does not require and cannot be finished from A-level machinery in contest time.
5. Whenever an integral or sum involves a phrase of the form "the number of … satisfying …", treat it as an unexpanded sum of indicators and consider swapping the order of operations before doing calculus on the count itself.
6. Reusable rule: a counting function inside an integral is a disguised sum; check the summand's sign, swap by Tonelli if nonnegative, and reduce a 2D radially symmetric integral to a 1D substitution $u = r^2$ per term. Isolate unpaired terms explicitly when converting a one-sided sum to a two-sided one.

▷ **Cue:** an integral or sum contains "the number of … satisfying a condition" $\longrightarrow$ **Move:** rewrite it immediately as a sum of indicator functions before integrating; do not treat it as an opaque function to be estimated
▷ **Cue:** you must swap the order of a sum and an integral (or two integrals) and every term is nonnegative $\longrightarrow$ **Move:** apply Tonelli's theorem; no absolute-convergence check is required, the swap is unconditional
▷ **Cue:** a radially symmetric integrand is integrated over a disk or its exterior in the plane $\longrightarrow$ **Move:** switch to polar coordinates; the Jacobian r turns $\int re^{-r^2}\,dr$ into an elementary substitution $u = r^2$
▷ **Cue:** a problem defines a one-sided sum $\sum_{k\geq 0} f(k)$ for an even function f, but the natural object is two-sided $\sum_{k\in\mathbb{Z}} f(k)$ $\longrightarrow$ **Move:** isolate the unpaired $k = 0$ term and double the rest: $\sum_{\mathbb{Z}} f(k) = 2\big(\sum_{k\geq 0} f(k)\big) - f(0)$

1982 B3

1982 B3. Let p_n be the probability that $c + d$ is a perfect square when the integers c and d are selected independently at random from the set $\{1, 2, 3, \ldots, n\}$. Show that $\lim_{n\to\infty} \left(p_n\sqrt{n}\right)$ exists and express this limit in the form $r(\sqrt{s} - t)$, where s and t are integers and r is a rational number.

Fields: probability / number theory, asymptotic counting (Riemann sums in disguise) **Position:** B3 – closing problem of the opening trio, first appearance of a genuine limit argument in the session **Answer:** $r = \frac{4}{3}$, $s = 2$, $t = 1$: $\lim_{n\to\infty} p_n\sqrt{n} = \frac{4}{3}(\sqrt{2} - 1)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. c, d are chosen independently and uniformly from $\{1, \ldots, n\}$ (with repetition, as ordered pairs), so $p_n = F(n)/n^2$ where $F(n)$ counts ordered pairs (c, d) with $c + d$ a perfect square. You set out to compute $F(n)$ exactly. For each achievable value $c + d = t^2$, the number of ordered pairs in $\{1, \ldots, n\}^2$ with that fixed sum is $t - 1$ once $t \leq n + 1$, and $2n + 1 - t$ once $t > n + 1$ (draw the antidiagonal $c + d = t$ inside the square $[1, n]^2$ and count

lattice points on it). So you write

$$F(n) = \sum_{t=1}^{a(n)} (t^2 - 1) + \sum_{t=a(n)+1}^{b(n)} (2n + 1 - t^2),$$

where $a(n) = \lfloor \sqrt{n+1} \rfloor$ and $b(n) = \lfloor \sqrt{2n} \rfloor$ are the thresholds where t^2 crosses $n+1$ and where t^2 exceeds the maximum possible sum $2n$. This is correct, and with Faulhaber's formula for $\sum t^2$ it is even computable exactly – but the floors $a(n), b(n)$ now sit inside cubic expressions, and extracting $\lim p_n \sqrt{n} = \lim F(n)/n^{3/2}$ from an exact closed form means carrying every $O(1)$ floor correction through the algebra and watching it vanish term by term as $n \to \infty$. It works, but it is exactly the kind of computation where you do three times the bookkeeping the question actually paid for: you were never asked for $F(n)$, only for a limit.

The model chain. The question only wants $\lim_{n\to\infty} p_n \sqrt{n}$, so work asymptotically from the start instead of computing $F(n)$ exactly and taking a limit afterward. Rescale $t = u\sqrt{n}$: as $n \to \infty$, $a(n)/\sqrt{n} = \sqrt{(n+1)/n} \to 1$ and $b(n)/\sqrt{n} = \sqrt{2n/n} \to \sqrt{2}$, and t ranges over $\{1, \ldots, b(n)\}$ in steps of 1, i.e. u ranges over $\{1/\sqrt{n}, 2/\sqrt{n}, \ldots\}$ in steps of $1/\sqrt{n}$. Each of the two pieces of $F(n)$ is now visibly a Riemann sum:

$$\frac{1}{n^{3/2}} \sum_{t=1}^{a(n)} (t^2 - 1) = \sum_{t=1}^{a(n)} \left(\left(\frac{t}{\sqrt{n}} \right)^2 - \frac{1}{n} \right) \frac{1}{\sqrt{n}} \longrightarrow \int_0^1 u^2 \, du = \frac{1}{3},$$

$$\frac{1}{n^{3/2}} \sum_{t=a(n)+1}^{b(n)} (2n + 1 - t^2) = \sum_{t=a(n)+1}^{b(n)} \left(2 - \left(\frac{t}{\sqrt{n}} \right)^2 + O(1/n) \right) \frac{1}{\sqrt{n}} \longrightarrow \int_1^{\sqrt{2}} (2 - u^2) \, du.$$

The second integral: $\displaystyle\int_1^{\sqrt{2}} (2 - u^2) \, du = \left[2u - \frac{u^3}{3} \right]_1^{\sqrt{2}} = \left(2\sqrt{2} - \frac{2\sqrt{2}}{3} \right) - \left(2 - \frac{1}{3} \right) = \frac{4\sqrt{2}}{3} - \frac{5}{3}.$

Adding both pieces,

$$\lim_{n\to\infty} \frac{F(n)}{n^{3/2}} = \frac{1}{3} + \frac{4\sqrt{2}}{3} - \frac{5}{3} = \frac{4\sqrt{2} - 4}{3} = \frac{4}{3}(\sqrt{2} - 1).$$

Since $p_n \sqrt{n} = F(n)\sqrt{n}/n^2 = F(n)/n^{3/2}$, the limit exists and equals $\dfrac{4}{3}(\sqrt{2} - 1)$: $r = 4/3$, $s = 2$, $t = 1$.

The floors $a(n), b(n)$ cause no trouble here precisely because they were only ever going to contribute $O(1/\sqrt{n})$-sized corrections after the rescaling $1/n^{3/2}$ is applied – corrections that vanish in the limit regardless of exactly where $a(n)$ and $b(n)$ round. A Riemann sum is insensitive to whether its last subinterval is slightly short.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree on the exact combinatorial count of ordered pairs for a fixed sum t (that is forced by the problem and identical either way). They split at the line "now compute $\lim_{n\to\infty} F(n)/n^{3/2}$": the natural attempt keeps $F(n)$ in exact closed form (Faulhaber sums with floor-function bounds $a(n), b(n)$) and takes the limit of that closed form as the very last step; the model rescales $t = u\sqrt{n}$ before summing anything, recognises both sums as Riemann sums for integrals over $u \in [0, 1]$ and $u \in [1, \sqrt{2}]$, and computes the limit directly.

> *Since only the limit of $p_n \sqrt{n}$ is wanted, not p_n itself, I don't need an exact formula for $F(n)$ – I need its leading-order behaviour as $n \to \infty$, and any sum of a nice function of t over a range whose length scales like $\sqrt{n}$, once divided by the matching power of n, is a Riemann sum in disguise the moment I rescale $t = u\sqrt{n}$.*

> *The thresholds $a(n) = \lfloor\sqrt{n+1}\rfloor$ and $b(n) = \lfloor\sqrt{2n}\rfloor$ look like they demand careful floor-function algebra, but after dividing by $n^{3/2}$ every $O(1)$ error from a floor is only an $O(1/\sqrt{n})$ correction to the sum, and I am about to let $n \to \infty$ anyway – so I can replace $a(n)/\sqrt{n}$ and $b(n)/\sqrt{n}$ by their limits 1 and $\sqrt{2}$ and not carry the floors through the calculation at all.*

Layer 3 THE WHY

Why is it legitimate – not merely convenient – to abandon the exact formula for $F(n)$ and work with the limit from the start? Because the object the problem actually asks for is a limit, and a limit is precisely the kind of quantity for which only the leading-order (continuum) behaviour of a discrete sum survives; anything the exact formula carries beyond that leading order is, by the time you divide by $n^{3/2}$ and let $n \to \infty$, guaranteed to vanish. Recognising "the target is a limit" as a license to discard lower-order terms *before* computing, rather than after, is a general economy principle: it tells you in advance which parts of a computation are wasted effort, which is exactly the information the exact-then-limit route can only discover after doing the work.

Why does the sum turn into a Riemann sum rather than needing a separate asymptotic tool? Because a Riemann sum is nothing but a sum of $f(u_i)\Delta u$ over a partition with mesh $\Delta u \to 0$, and here $u = t/\sqrt{n}$ literally is such a partition: consecutive integers $t, t+1$ correspond to u-values $1/\sqrt{n}$ apart, which shrinks to 0 exactly as $n \to \infty$. The two pieces of $F(n)$, once divided by $n^{3/2} = n \cdot \sqrt{n}$, are sums of $f(t/\sqrt{n})$ (with $f(u) = u^2$ or $f(u) = 2 - u^2$) times the step size $1/\sqrt{n}$ – the definition of a Riemann sum, not an analogy to one. This is why the answer is a definite integral and not, say, an infinite series: the discreteness in t is an artifact of counting integer sums, and it disappears under rescaling because the range of t itself grows like $\sqrt{n}$.

Why do the floor functions $a(n), b(n)$ contribute nothing to the final answer? Because $\lfloor\sqrt{n+1}\rfloor = \sqrt{n} + O(1)$, so $a(n)/\sqrt{n} = 1 + O(1/\sqrt{n}) \to 1$; the $O(1)$ discrepancy shifts the Riemann sum's endpoint by a vanishing amount and changes the sum itself by at most one extra or missing term of size $O(\sqrt{n})$ (the largest term, at $t \approx a(n)$), which after dividing by $n^{3/2}$ is $O(1/n) \to 0$. This is worth stating as a standing fact: floor-function corrections to a sum's range never survive division by a strictly larger power of n in a Riemann-sum limit, so they may be dropped the moment the rescaling is written down, not verified term by term afterward.

The official solution proves the exact algebraic identity for $F(n)$ first (Faulhaber sums of t^2 evaluated at $a(n)$ and $b(n)$) and only takes $n \to \infty$ at the very end. That order is fully rigorous, but it pays for information – the exact value of $F(n)$ at every finite n – that the question never asked for. The reusable principle: when a problem's final object is a limit or an asymptotic constant, look first for the route that computes that limit directly (Riemann sum, generating function, integral comparison), and reserve the exact finite formula for when the problem also asks for it.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "show that $\lim_{n\to\infty}(p_n\sqrt{n})$ exists". On seeing it, write down: I am being asked for a limit, not an exact value, so I should look for the leading-order behaviour directly rather than computing p_n exactly first.
2. It suggests rescaling the summation variable by the power of n that the final answer is normalised by ($t = u\sqrt{n}$, since the limit is of $p_n\sqrt{n}$), and recognising the resulting sum as a Riemann sum for a definite integral.
3. Hesitation occurs at the floor functions $a(n) = \lfloor\sqrt{n+1}\rfloor$, $b(n) = \lfloor\sqrt{2n}\rfloor$, which look like they demand careful case-by-case tracking. It is reinforced away by checking that any $O(1)$ floor error becomes $O(1/n)$ or smaller after the rescaling that the limit already requires.
4. The trap is "zombie maths": computing $F(n)$ to exact, fully rigorous closed form via Faulhaber's formula and floor-function bookkeeping when the problem only ever asked for an asymptotic constant, spending effort the question does not pay for.

5. Whenever a discrete sum's range grows like n^{α} and the final target is a limit after dividing by a matching power of n, try the substitution that turns the summation index into a variable on $[0, \text{finite}]$ and check whether the result is a Riemann sum before attempting an exact closed form.

6. Reusable rule: if the problem's final object is a limit or asymptotic constant, discard lower-order (floor, $O(1)$) corrections at the start by rescaling into a Riemann sum, rather than carrying an exact discrete formula through to the end and taking the limit last.

▷ **Cue:** a discrete sum's natural range grows proportionally to a power of n, and only $\lim_{n\to\infty}(\text{sum})/n^{\text{matching power}}$ is asked for $\longrightarrow$ **Move:** rescale the index by that power of n and recognise the sum as a Riemann sum for a definite integral, rather than computing the sum in exact closed form first

▷ **Cue:** a sum's bounds are floor functions like $\lfloor\sqrt{n+1}\rfloor$, and the final target is a limit after dividing by a larger power of n $\longrightarrow$ **Move:** replace the floor bound by its continuum limit; the $O(1)$ floor discrepancy vanishes after the rescaling the limit already applies

▷ **Cue:** counting ordered pairs $(c,d) \in \{1,\dots,n\}^2$ with a fixed sum $c+d=t$ $\longrightarrow$ **Move:** the count is $t-1$ for $t \le n+1$ and $2n+1-t$ for $t > n+1$ – the antidiagonal length inside the square, increasing then decreasing

1982 B4

1982 B4. Let $n_1, n_2, \dots, n_s$ be distinct integers such that

$$(n_1+k)(n_2+k)\cdots(n_s+k)$$

is an integral multiple of $n_1 n_2 \cdots n_s$ for every integer k. For each of the following assertions, give a proof or a counterexample:

(a) $|n_i| = 1$ for some i.

(b) If further all n_i are positive, then $\{n_1, n_2, \dots, n_s\} = \{1, 2, \dots, s\}$.

Fields: number theory; divisibility; polynomials and finite differences **Position:** B4 — mid-to-late paper; (a) is a two-line size argument once you pick the right two values of k, (b) needs one structural fact about polynomials **Answer:** Both assertions are true. (a) follows from $P_0^2 \mid P_1 P_{-1}$ and a size comparison; (b) from $P_0 \mid s!$ via the s-th finite difference. Positivity is needed in (b): $\{-1, 2\}$ satisfies the hypothesis.

Layer 1 THE WHAT — surface analysis

Throughout, write $P_k = (n_1+k)(n_2+k)\cdots(n_s+k)$, so the hypothesis is $P_0 \mid P_k$ for every integer k. First a housekeeping point the problem leaves implicit: no n_i is 0. If one were, $P_0 = 0$, and "P_k is a multiple of 0" forces $P_k = 0$ for every k; but $P_k > 0$ once $k > \max|n_i|$. So $P_0 \neq 0$ and every $n_i \neq 0$.

The natural attempt. For (a), the strong student tests small k. $k = 1$: $P_0 \mid P_1$, i.e. $n_1 \cdots n_s \mid (n_1+1)\cdots(n_s+1)$. Suppose no $|n_i| = 1$. Then each $|n_i| \ge 2$, and one tries to argue that $|n_i+1|$ is "not much bigger" than $|n_i|$, so the quotient $P_1/P_0 = \prod(1+1/n_i)$ can't be an integer. But it can be: with $n = (2,3)$ the quotient is $\frac{3}{2}\cdot\frac{4}{3} = 2$. One value of k is not enough, because the factors $1+1/n_i$ exceed 1 for positive n_i and the product can land on an integer. The student then tries $k = -1$ separately and hits the same wall from the other side: $\prod(1-1/n_i)$ is less than 1 for positive n_i but can be larger for negative ones.

For (b), the natural move is to plug in more and more k and hunt for a contradiction by cases: if $1 \notin \{n_i\}$ then $k = -1$ gives …; if $2 \notin \{n_i\}$ then …. This becomes a case analysis on which integers are missing, with no end in sight, and it does not use the one piece of information that

is special: the divisibility holds for *every* k, which is a statement about the polynomial P_k as a whole.

The model chain. *(a) True.* Use $k = 1$ and $k = -1$ together. From $P_0 \mid P_1$ and $P_0 \mid P_{-1}$,

$$P_0^2 \,\Big|\, P_1 P_{-1} = \prod_{i=1}^{s} (n_i + 1)(n_i - 1) = \prod_{i=1}^{s} (n_i^2 - 1).$$

Suppose, for contradiction, that $|n_i| \neq 1$ for every i. Since also $n_i \neq 0$, each $|n_i| \geq 2$, so $0 < n_i^2 - 1 < n_i^2$. Multiplying,

$$0 < \prod_i (n_i^2 - 1) < \prod_i n_i^2 = P_0^2.$$

A positive integer smaller than P_0^2 cannot be a multiple of P_0^2. Contradiction; so $|n_i| = 1$ for some i.

(b) True. P_k is a monic polynomial in k of degree s. Let Δ be the forward difference, $\Delta f(k) = f(k+1) - f(k)$. Applying Δ to a polynomial of degree d with leading coefficient c gives a polynomial of degree $d-1$ with leading coefficient dc (the k^d terms cancel and $(k+1)^d - k^d = dk^{d-1} + \dots$). Applying it s times to P_k therefore gives the constant $s!$:

$$\Delta^s P_0 = \sum_{j=0}^{s} (-1)^{s-j} \binom{s}{j} P_j = s!.$$

Every P_j on the left is a multiple of P_0 by hypothesis, so $P_0 \mid s!$. All n_i are positive, so $P_0 > 0$ and hence $P_0 \leq s!$.

Now order the n_i: $m_1 < m_2 < \cdots < m_s$. They are distinct positive integers, so $m_1 \geq 1$ and $m_{i+1} \geq m_i + 1$, giving $m_i \geq i$ by induction. Hence $P_0 = \prod m_i \geq \prod i = s!$, with equality only if $m_i = i$ for every i. Combined with $P_0 \leq s!$, equality holds, and $\{n_1, \dots, n_s\} = \{1, \dots, s\}$.

(Conversely $\{1, \dots, s\}$ does satisfy the hypothesis: $P_k = (k+1)\cdots(k+s)$ is a product of s consecutive integers, and $P_k/s! = \binom{k+s}{s}$ for $k \geq 0$, with the negative-k cases either 0 or $\pm$ a binomial coefficient. So (b) is a characterisation, not a vacuous statement.)

Positivity matters in (b). Take $\{n_1, n_2\} = \{-1, 2\}$: $P_0 = -2$ and $P_k = (k-1)(k+2)$; the two factors differ by 3, so exactly one of them is even and P_k is always even. The hypothesis holds, (a) holds ($|-1| = 1$), but the set is not $\{1, 2\}$. (A computer search over $|n_i| \leq 7$, $s \leq 3$ confirms that the only all-positive solutions are $\{1\}, \{1, 2\}, \{1, 2, 3\}$ and that every solution contains ± 1.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. In (a) both chains write down $P_0 \mid P_1$. The natural attempt then studies that single divisibility; the model writes down $P_0 \mid P_{-1}$ *beside it and multiplies*. In (b) the split is at the first line: the natural attempt plugs in particular values of k; the model treats "$P_0 \mid P_k$ for all k" as a statement about a polynomial and asks what integer combinations of its values are forced to be.

> *One value of k makes each factor a bit bigger or a bit smaller depending on the sign of n_i, and I can't control the signs. Pair $k = 1$ with $k = -1$: $(n+1)(n-1) = n^2 - 1$, which is smaller than n^2 whatever the sign of n. Then the product is strictly between 0 and P_0^2, and divisibility is impossible.*

> *"For every integer k" means P_0 divides every value of a degree-s polynomial, so it divides every integer combination of values. Which combination kills everything but a known constant? The s-th difference: it is s! for any monic degree-s polynomial. So $P_0 \mid s!$.*

▷ **Cue:** a product $\prod(n_i + k)$ must be divisible by $\prod n_i$ and the signs of the n_i are unknown $\longrightarrow$ **Move:** use k and $-k$ together: $\prod(n_i^2 - k^2)$ is sign-free and comparable in size with $\prod n_i^2$

▷ **Cue:** "d divides $f(k)$ for every integer k", f a polynomial of degree s with leading coefficient c $\longrightarrow$ **Move:** take the s-th finite difference: $d \mid s!\, c$

▷ **Cue:** a product of s distinct positive integers is bounded above by $s!$ $\longrightarrow$ **Move:** order them; $m_i \geq i$ forces the set to be $\{1, \ldots, s\}$

Layer 3 THE WHY

The official solution begins (a) with "$P_0 \mid P_{-1}$ and $P_0 \mid P_1$ together imply $P_0^2 \mid P_{-1}P_1$". Here is where that comes from.

A divisibility $A \mid B$ between nonzero integers can be refuted in essentially one way without factorising anything: show $0 < |B| < |A|$. That is the definition of divisibility turned into an inequality — a nonzero multiple of A has absolute value at least $|A|$. So the question in (a) is: which P_k is *smaller* than P_0 in absolute value? For a single factor, $|n + k| < |n|$ happens for $k = 1$ when $n < 0$ and for $k = -1$ when $n > 0$. Neither value of k works for every sign pattern. The product $(n + 1)(n - 1)$ removes the dependence on sign, and $|n^2 - 1| < n^2$ for every integer $|n| \geq 2$. Squaring P_0 is the price for multiplying two divisibilities, and it is exactly the price that makes the size comparison sign-free. So "multiply the $k = 1$ and $k = -1$ statements" is what the definition of divisibility asks for, once you know the obstruction is the unknown signs.

For (b), the phrase "for every integer k" is the cue. A condition on all values of a polynomial is a condition on the polynomial, and the cheapest invariant of a polynomial that can be read off from its integer values is its leading coefficient, via finite differences. The reason finite differences are natural here, and derivatives are not, is that divisibility survives only integer linear combinations: $\Delta^s P_0$ is $\sum \pm\binom{s}{j} P_j$, an integer combination of values, so P_0 divides it. And because P_k is monic, the leading coefficient information is exactly $s!$. After that the problem becomes: s distinct positive integers with product at most $s!$. The extremal fact — the i-th smallest positive integer in a set is at least i — ends it.

Why does positivity appear only in (b)? Because the bound $P_0 \geq s!$ needs positive integers; with signs allowed, $\{-1, 2\}$ has $|P_0| = 2 = 2!$ but is not $\{1, 2\}$. The divisibility argument $P_0 \mid s!$ survives, so what fails is only the last step. The counterexample shows the hypothesis is doing work, not decoration.

Layer $+\alpha$ BACKTESTING

1. The cue is "for every integer k". Write down $P_0 \mid P_1$, $P_0 \mid P_{-1}$, and the words "polynomial of degree s in k".

2. For (a): refute divisibility by size, making the size comparison sign-independent by pairing $\pm k$. For (b): integer combinations of values, specifically the s-th difference.

3. Hesitation in (a) comes after $P_0 \mid P_1$ alone gives nothing (e.g. $(2, 3)$ works for $k = 1$). What removes it: ask "what could make the quotient too small to be an integer?" and notice the sign problem. In (b) hesitation comes from case-bashing on missing integers; the reinforcement is that "all k" statements about polynomials are settled by differences.

4. Two traps. First, forgetting that $n_i \neq 0$ needs an argument (it does: $P_0 = 0$ would force $P_k \equiv 0$). Second, assuming (b) also holds with signs, or conversely looking for a counterexample to (b) among positive sets, where none exists.

5. Recognise the shape "a quantity must divide a polynomial's values everywhere": that is always an $s!$-times-leading-coefficient statement. And recognise "divisibility, unknown signs": pair x with $-x$ to get squares.

6. To refute $A \mid B$, show $0 < |B| < |A|$; if the sign of the comparison depends on unknown signs, multiply two instances to make it sign-free. If d divides $f(k)$ for all integers k and $\deg f = s$, then d divides $s!$ times the leading coefficient.

1982 B5

1982 B5. For each $x > e^e$ define a sequence $S_x = u_0, u_1, u_2, \ldots$ recursively as follows: $u_0 = e$, while for $n \geq 0$, u_{n+1} is the logarithm of x to the base u_n. Prove that S_x converges to a number $g(x)$ and that the function g defined in this way is continuous for $x > e^e$.

Fields: real analysis; iteration of a decreasing map; continuity of inverse functions **Position:** B5 — late paper; nothing deep, but a careless convergence argument (a fixed point, then "it converges") scores nothing **Answer:** $u_n \to g(x)$, the unique $g > e$ with $g^g = x$ (equivalently $g \ln g = \ln x$); g is the inverse of the continuous increasing bijection $y \mapsto y^y$ from (e, ∞) onto (e^e, ∞), hence continuous.

Layer 1 THE WHAT — surface analysis

Write $L = \ln x$; the hypothesis $x > e^e$ is $L > e$. Since $\log_u x = \ln x / \ln u$, the recursion is $u_{n+1} = f(u_n)$ with

$$f(u) = \frac{L}{\ln u} \qquad (u > 1).$$

f is continuous and strictly decreasing on $(1, \infty)$.

The natural attempt. The strong student computes the fixed point first: if $u_n \to g$ then $g = L/\ln g$, i.e. $g \ln g = L$, i.e. $g^g = x$. Since $y \mapsto y \ln y$ is increasing on $(1, \infty)$ and passes through e at $y = e$, there is exactly one solution, and $g > e$ because $L > e$. Then they try to prove convergence by contraction: by the mean value theorem, $|u_{n+1} - g| = |f'(\xi_n)|\,|u_n - g|$ with $f'(u) = -L/(u(\ln u)^2)$. At the fixed point $f'(g) = -L/(g(\ln g)^2) = -1/\ln g$, of modulus less than 1 because $g > e$. Encouraging — but $|f'(u)|$ is *not* below 1 everywhere the sequence goes. At $u = e$ it equals $L/e > 1$. So contraction only works once the iterates are near g, and proving that they get near g is the whole problem. The student then computes a few terms ($x = 100$: 2.718, 4.605, 3.015, 4.172, 3.224, ...), sees them oscillate, and is stuck between a local argument that does not start and a numerical picture that does not prove anything.

The model chain. *Step 1: an invariant interval.* Let $I = [e, L]$. For $u \in I$, f decreasing gives $f(L) \leq f(u) \leq f(e) = L$. And $f(L) = L/\ln L > e$: the function $h(L) = L - e \ln L$ has $h(e) = 0$ and $h'(L) = 1 - e/L > 0$ for $L > e$, so $L > e \ln L$. Hence $f(I) \subseteq (e, L] \subseteq I$. Since $u_0 = e \in I$, every $u_n \in I$; in particular every $u_n \geq e > 1$, so the recursion is always defined.

Step 2: monotone subsequences. Let $F = f \circ f$; as a composition of two strictly decreasing maps, F is strictly increasing on I. Now $u_2 \in I$ and $u_0 = e = \min I$, so $u_2 \geq u_0$; in fact $u_2 > e$ by Step 1, so $u_2 > u_0$. Applying F repeatedly, $u_{2n+2} > u_{2n}$ for all n (induction: $u_{2n+2} > u_{2n} \Rightarrow F(u_{2n+2}) > F(u_{2n})$). Likewise $u_1 = L = \max I$, so $u_3 < u_1$ (strictly, as $u_3 \leq L$ and $u_3 = u_1$ would force $u_2 = e$), and u_{2n+1} is strictly decreasing. Both subsequences are monotone and lie in I, so they converge: $u_{2n} \to a$, $u_{2n+1} \to b$, with $a, b \in I$, so $a, b \geq e$.

Step 3: the two limits coincide. Pass to the limit in $u_{2n+1} = f(u_{2n})$ and $u_{2n+2} = f(u_{2n+1})$, using continuity of f on I: $b = f(a)$ and $a = f(b)$. That is, $b \ln a = L = a \ln b$, so

$$\frac{\ln a}{a} = \frac{\ln b}{b}.$$

The function $\varphi(y) = \ln y / y$ has $\varphi'(y) = (1 - \ln y)/y^2 < 0$ for $y > e$, so it is strictly decreasing, hence injective, on $[e, \infty)$. As $a, b \geq e$, $a = b$. Both subsequences have the same limit, so $u_n \to g(x) := a$, and $g = f(g)$ gives $g \ln g = L$, i.e. $g^g = x$, with $g \geq e$; $g = e$ would give $L = e$, so $g > e$.

Step 4: continuity. Let $\psi(y) = y^y = e^{y \ln y}$ on (e, ∞). It is continuous and strictly increasing ($(y \ln y)' = 1 + \ln y > 0$), with $\psi(y) \to e^e$ as $y \to e^+$ and $\psi(y) \to \infty$, so by the intermediate value theorem it is a bijection onto (e^e, ∞). Step 3 shows $g = \psi^{-1}$. A strictly increasing function on an interval whose image is an interval has no jumps, so ψ^{-1} is continuous: if g jumped at x_0,

the values $g(x_0^-) < g(x_0^+)$ would leave a gap in the image (e, ∞) of g, containing at most the single point $g(x_0)$. Hence g is continuous on (e^e, ∞). □

(Checked numerically: $x = 100$ gives $g = 3.59728\ldots$ with $g^g = 100.000$; $x = 10^6$ gives $g = 7.0658\ldots$. Slightly above $x = e^e$ convergence is very slow, as it must be: the fixed-point derivative $-1/\ln g$ tends to -1.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains find the fixed point $g^g = x$. They split at the next line. The natural attempt asks "is f a contraction?" — a question about derivatives. The model asks "where does the sequence live, and what does a decreasing map do to order?" — a question about monotonicity, answered by the interval $[e, L]$ and the increasing map $f \circ f$.

> *The map is decreasing, so the terms zigzag. A zigzag is two monotone sequences in disguise: $f \circ f$ is increasing, so the even terms move one way and the odd terms the other. I need a box they cannot leave, and the start $u_0 = e$ and first step $u_1 = L$ look like its two ends.*

> *Two limits a, b with $b = f(a)$, $a = f(b)$ is a possible 2-cycle. To kill it I need the equation $a \ln b = b \ln a$ to force $a = b$. That is $\ln a/a = \ln b/b$, and $\ln y/y$ is one-to-one on $[e, \infty)$ — which is exactly why the problem starts at e and demands $x > e^e$.*

- ▷ **Cue:** an iteration $u_{n+1} = f(u_n)$ with f decreasing; numerics oscillate $\longrightarrow$ **Move:** split into even and odd terms; $f \circ f$ is increasing, so each subsequence is monotone
- ▷ **Cue:** you need a monotone sequence to converge $\longrightarrow$ **Move:** find an interval mapped into itself; the first two terms are often its endpoints
- ▷ **Cue:** even and odd limits satisfy $b = f(a)$, $a = f(b)$ $\longrightarrow$ **Move:** rule out a 2-cycle by an injectivity property on the invariant interval
- ▷ **Cue:** the limit is defined implicitly by $\psi(g) = x$ with ψ increasing $\longrightarrow$ **Move:** continuity of g is continuity of an inverse monotone function: no jumps because the image is an interval

Layer 3 THE WHY

The official solution proves the full interleaving $e < u_2 < u_4 < \cdots < u_5 < u_3 < u_1$ by comparing powers $(u_n)^{u_{n+1}}$ with $y^{1/y}$ decreasing. That is correct but longer than it needs to be. Here is where the shorter route comes from.

Start from what the recursion is. $u_{n+1} = L/\ln u_n$ applies a *decreasing* function. A decreasing function reverses order: if $u < v$ then $f(u) > f(v)$. Applying it twice restores order. That single fact from the definition of "decreasing" is why the terms zigzag and why $F = f \circ f$ is the object to study: F is increasing, and iterating an increasing map on an interval produces a monotone sequence as soon as the first step goes in one direction (if $u_2 > u_0$ then $F(u_2) > F(u_0)$, i.e. $u_4 > u_2$, and so on). You don't prove monotonicity of the evens by hand; it is inherited from the first comparison.

Monotone is not enough; you need bounded, and you need the first comparison. Both come from one object, an interval I with $f(I) \subseteq I$. Why $[e, L]$? Because it is $[u_0, u_1]$: once f maps it into itself, u_0 is its bottom (so the evens can only rise) and u_1 its top (so the odds can only fall). The only computation is $f(L) \geq e$, which is $L \geq e \ln L$ — true precisely because $L \geq e$. The hypothesis $x > e^e$ is being used here, and you should see it used.

Then the 2-cycle. Monotone even and odd subsequences can converge to different limits; that is what happens for decreasing maps with a repelling fixed point. The equation of a 2-cycle is $a \ln b = b \ln a$, i.e. $\ln a/a = \ln b/b$, and $\ln y/y$ rises on $(1, e]$ and falls on $[e, \infty)$ — so it is two-to-one overall and one-to-one on $[e, \infty)$. A 2-cycle $a \neq b$ must straddle e. The invariant interval keeps both limits $\geq e$, so no 2-cycle is possible. This is the real reason for the threshold e^e: below it the fixed point lies below e, $|f'(g)| = 1/\ln g > 1$, and the iteration can run away.

(With $L = 2$ the iterates go e, 2, 2.885, 1.887, 3.149, ..., 64.8, 0.479, and the next logarithm is to a base below 1, then negative: the sequence dies.)

So the whole solution is three uses of the word "decreasing": f decreasing gives the zigzag and F increasing; $\ln y/y$ decreasing on $[e,\infty)$ kills the 2-cycle; and y^y increasing gives continuity of the inverse. The contraction attempt fails not because it is wrong but because derivatives answer a local question and the problem asks a global one.

Layer $+\alpha$ BACKTESTING

1. The first cue is "u_{n+1} is the logarithm of x to the base u_n". Write $u_{n+1} = \ln x/\ln u_n$ and note that the map is decreasing in u_n.

2. Decreasing iteration: even/odd subsequences via $f \circ f$, on an invariant interval, then exclude a 2-cycle.

3. Hesitation comes after finding the fixed point, when contraction fails at $u = e$. It is reinforced away by asking for monotonicity instead of a derivative bound.

4. Assuming the limit exists and solving $g = f(g)$ as if that were a proof. Also: proving the evens and odds converge and forgetting they might converge to different numbers.

5. Oscillating numerics plus a decreasing map should trigger "two monotone subsequences" immediately. A strange threshold such as e^e usually marks where an injectivity or monotonicity property changes; look for the function whose turning point it is ($\ln y/y$ at $y = e$).

6. For $u_{n+1} = f(u_n)$ with f decreasing: find $I \ni u_0$ with $f(I) \subseteq I$; then u_{2n} and u_{2n+1} are monotone and converge; they share a limit iff f has no 2-cycle in I. A function defined as the inverse of a strictly monotone continuous function on an interval is continuous.

1982 B6

1982 B6. Let $K(x,y,z)$ denote the area of a triangle whose sides have lengths x, y, and z. For any two triangles with sides a,b,c and a',b',c', respectively, prove that

$$\sqrt{K(a,b,c)} + \sqrt{K(a',b',c')} \le \sqrt{K(a+a',b+b',c+c')},$$

and determine the cases of equality.

Fields: geometry; inequalities; Heron's formula; AM–GM / Hölder **Position:** B6 — the last problem, but short once the right coordinates are chosen; the difficulty is refusing to do geometry **Answer:** The inequality holds, with equality if and only if the two triangles are similar, i.e. (a,b,c) is proportional to (a',b',c').

Layer 1 THE WHAT — surface analysis

First, the right-hand side makes sense: adding $a < b+c$ and $a' < b'+c'$ gives $a+a' < (b+b')+(c+c')$, and similarly for the other two, so $a+a', b+b', c+c'$ are the sides of a (non-degenerate) triangle.

The natural attempt. Areas of triangles with given sides: the strong student reaches for Heron,

$$16K^2 = (a+b+c)(-a+b+c)(a-b+c)(a+b-c),$$

or for $K = \frac{1}{2}bc\sin A$. With Heron, the inequality after squaring becomes

$$K + K' + 2\sqrt{KK'} \le K'',$$

where K'' is the area of the sum triangle, and both K and K'' are square roots of quartics in six variables. Squaring again to clear $\sqrt{KK'}$ leaves a polynomial inequality of degree 8 in six variables. The student starts expanding, perhaps specialises to $a' = a$ to test the case of similar

triangles, and runs out of time. With the sine formula, the angles of the sum triangle have no usable relation to the original angles, and the attempt stalls at once.

This is zombie maths in its purest form: an expansion undertaken before asking what shape the inequality has.

The model chain. *Step 1: better coordinates.* For the triangle (a,b,c) put

$$s=\tfrac{1}{2}(a+b+c),\quad t=s-a,\quad u=s-b,\quad v=s-c.$$

Triangle inequalities say exactly $t,u,v>0$, and Heron's formula reads $K(a,b,c)=\sqrt{stuv}$, so

$$\sqrt{K(a,b,c)}=(stuv)^{1/4}.$$

Define s',t',u',v' likewise for (a',b',c'). Each of s,t,u,v is a *linear* function of (a,b,c), so for the triangle $(a+a',b+b',c+c')$ the corresponding quantities are $s+s'$, $t+t'$, $u+u'$, $v+v'$. The claim becomes

$$(stuv)^{1/4}+(s't'u'v')^{1/4}\le\big((s+s')(t+t')(u+u')(v+v')\big)^{1/4}. \tag{$*$}$$

Step 2: AM–GM, twice. Write $p_1,\dots,p_4$ for s,t,u,v and $q_1,\dots,q_4$ for s',t',u',v'; all are positive. Divide $(*)$ by its right-hand side. The claim is

$$\prod_{i=1}^{4}\Big(\frac{p_i}{p_i+q_i}\Big)^{1/4}+\prod_{i=1}^{4}\Big(\frac{q_i}{p_i+q_i}\Big)^{1/4}\le 1.$$

By AM–GM on four positive numbers, the first product is at most $\frac{1}{4}\sum_i \frac{p_i}{p_i+q_i}$ and the second at most $\frac{1}{4}\sum_i \frac{q_i}{p_i+q_i}$. Adding,

$$\text{LHS}\le\frac{1}{4}\sum_{i=1}^{4}\frac{p_i+q_i}{p_i+q_i}=1.$$

This proves $(*)$ and hence the inequality.

Step 3: equality. Equality in the sum requires equality in both AM–GM steps, i.e. $\frac{p_i}{p_i+q_i}$ is the same for all i (and then so is $\frac{q_i}{p_i+q_i}$). That is, p_i/q_i is constant: $(s,t,u,v)=\lambda(s',t',u',v')$ for some $\lambda>0$. Conversely, if this holds, both sides of $(*)$ equal $(1+\lambda)(s't'u'v')^{1/4}$. Finally, (t,u,v) determines (a,b,c) linearly: $a=u+v$, $b=t+v$, $c=t+u$ (e.g. $u+v=2s-b-c=a$). So $(t,u,v)=\lambda(t',u',v')$ holds if and only if $(a,b,c)=\lambda(a',b',c')$, and then $s=\lambda s'$ automatically. Equality holds exactly when the two triangles are similar (with corresponding sides $a\leftrightarrow a'$, $b\leftrightarrow b'$, $c\leftrightarrow c'$). □

(Check: $(3,4,5)$ and $(6,8,10)$ give $\sqrt{6}+\sqrt{24}=3\sqrt{6}=\sqrt{54}=\sqrt{K(9,12,15)}$. A random search over 2×10^5 pairs of triangles found no violation.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write Heron's formula. The natural attempt keeps it in the variables a,b,c and squares. The model *changes variables to $s,s-a,s-b,s-c$ before doing anything else*, because in those variables the sum triangle is obtained by adding coordinates and $\sqrt{K}$ is a geometric mean.

> *The right-hand side is the same function evaluated at a sum. So the question is whether $\sqrt{K}$ is superadditive, and I want coordinates in which "add the triangles" is "add the coordinates" and $\sqrt{K}$ is a clean function of them. Heron in $s,s-a,s-b,s-c$ does both: the four quantities are linear in the sides, and $\sqrt{K}$ is their fourth root of product — a geometric mean.*

> *"Geometric mean of a sum is at least the sum of geometric means." To prove an inequality of the form $f(p)+f(q)\le f(p+q)$ with f homogeneous of degree 1, divide by $f(p+q)$ and use AM–GM on the normalised ratios; the two AM bounds add to exactly 1.*

- ▷ **Cue:** an inequality $f(X)+f(Y)\le f(X+Y)$ where $X+Y$ is formed coordinate-wise $\longrightarrow$ **Move:** find coordinates in which f is homogeneous of degree 1; then it is a concavity/superadditivity statement
- ▷ **Cue:** Heron's formula appears alongside sums of side lengths $\longrightarrow$ **Move:** use s, $s-a$, $s-b$, $s-c$: linear in the sides, positive exactly for triangles
- ▷ **Cue:** $\prod p_i^{1/n}+\prod q_i^{1/n}\le\prod(p_i+q_i)^{1/n}$ $\longrightarrow$ **Move:** divide by the right side and apply AM–GM to $p_i/(p_i+q_i)$ and to $q_i/(p_i+q_i)$; add
- ▷ **Cue:** equality in a chain of AM–GM steps $\longrightarrow$ **Move:** all the averaged quantities are equal; translate back through the linear change of variables

Layer 3 THE WHY

The official solution proves $(*)$ by first proving the two-variable case $\sqrt{xy}+\sqrt{x'y'}\le\sqrt{(x+x')(y+y')}$ (Cauchy–Schwarz) and applying it twice, pairing st with uv. That works; the one-step AM–GM route above is shorter and generalises to any number of factors. Either way, the solution begins by rewriting Heron in s,t,u,v. Here is where that comes from.

Look at the inequality before touching areas. The right-hand side is $\sqrt{K}$ evaluated at the sum of the arguments. Any inequality of the shape $f(X)+f(Y)\le f(X+Y)$ is a statement about how f interacts with addition, so the first job is to find the variables in which addition is honest. Side lengths already add, but $\sqrt{K}$ is a messy function of them. The quantities $s-a$, $s-b$, $s-c$ are the tangent lengths from the vertices to the incircle; they are linear in the sides, so they add too, and in them $K^2=stuv$ is a monomial. That is the reason, and it's structural: Heron's formula is *already* a product of four linear forms, so the right variables are those four linear forms themselves.

Now read off what $\sqrt{K}$ is: $(stuv)^{1/4}$, the geometric mean of four positive numbers. It is homogeneous of degree 1 (scale the triangle by λ, $\sqrt{K}$ scales by λ) and it is concave. For a function with both properties, $f(X)+f(Y)=2\big(\frac12 f(X)+\frac12 f(Y)\big)\le 2f\big(\frac{X+Y}{2}\big)=f(X+Y)$. So the problem is concavity of the geometric mean, and AM–GM applied to normalised ratios is the elementary proof of that concavity. Homogeneity is also why the check with similar triangles must give equality: if $Y=\lambda X$, both sides are $(1+\lambda)f(X)$. You could predict the equality case before proving anything.

That is also the lesson about the natural attempt. Squaring twice destroys homogeneity-of-degree-one and hides the concavity inside a degree-8 polynomial. The inequality was never about triangles; the triangle is the packaging for a geometric mean.

Caveat. The problem says "triangle", which the archive reads as non-degenerate ($t,u,v>0$). If degenerate triangles (area 0) were admitted, the inequality still holds (by continuity of both sides), but the equality cases enlarge: e.g. $(1,1,2)$ and $(1,2,3)$ are both degenerate with $v=s-c=0$, their sum $(2,3,5)$ is degenerate, and $0+0=0$ although they are not similar. The answer "similar triangles" is correct under the non-degenerate reading, which is the official one.

Layer $+\alpha$ BACKTESTING

1. The cue is the right-hand side $K(a+a',b+b',c+c')$: the same function at a coordinate-wise sum. Write "superadditivity — find coordinates where adding triangles adds coordinates".

2. Heron in $s,s-a,s-b,s-c$, then AM–GM (or Cauchy–Schwarz twice) for superadditivity of the geometric mean.

3. Hesitation comes right after Heron, when the formula in a,b,c looks too heavy to use. It is reinforced away by rewriting Heron as a product of four linear forms and naming them.

4. Squaring to clear roots, which leads to a degree-8 expansion. Also a quiet trap in the equality case: proportional (s, t, u, v) must be translated back to proportional sides, which requires noting the change of variables is invertible.

5. Any time a problem compares $f(X) + f(Y)$ with $f(X + Y)$, ask whether f is homogeneous of degree one; if so the inequality is concavity (or convexity), and the equality case is proportionality.

6. For f positively homogeneous of degree 1 and concave, $f(X) + f(Y) \leq f(X + Y)$, with equality when $Y = \lambda X$. The geometric mean $\left(\prod x_i\right)^{1/n}$ is such an f; prove its superadditivity by applying AM–GM to $x_i/(x_i + y_i)$ and $y_i/(x_i + y_i)$ and adding.