DAWN PUTNAM SESSION

Putnam 1969 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1969 — Paper A **5**
1969 A1 5
1969 A2 8
1969 A3 11
1969 A4 13
1969 A5 17
1969 A6 20

Putnam 1969 — Paper B **23**
1969 B1 23
1969 B2 25
1969 B3 28
1969 B4 32
1969 B5 35
1969 B6 37

Putnam 1969 — Paper A

1969 A1

1969 A1. Let $f(x, y)$ be a polynomial with real coefficients in the real variables x and y defined over the entire x-y plane. What are the possibilities for the range of $f(x, y)$?

Caveat. The MAA 1965–1984 OCR renders this statement as "Letf(x, y) be a polynomial with real coefficients in the real variables x andy defined over the entire x-y plane. What are the possibilities for the range of j(x, y)?" – a run-together "Let f", a run-together "x and y", and the OCR's characteristic f/j font confusion in "$j(x, y)$". These are typographic artefacts only; there is no mathematical ambiguity, since the book's own solution consistently uses $f(x, y)$ throughout. The reconstruction above corrects only the spacing and the f/j substitution.

Fields: Real analysis, topology of the plane, polynomials in two variables **Position:** A1 (opener)
Answer: Exactly four possibilities: a single point, a semi-infinite interval with its endpoint (closed), a semi-infinite interval without its endpoint (open), or all of $\mathbb{R}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The natural first move on seeing "polynomial in two variables, what does its range look like" is to start generating examples and sorting them by degree. Degree 1: $f = x$ has range $\mathbb{R}$; only constants fail to hit everything. Degree 2: $f = x^2 + y^2$ has range $[0, \infty)$, $f = x^2 - y^2$ has range $\mathbb{R}$ (it hits every value: fix $y = 0$, vary x), $f = x^2$ has range $[0, \infty)$ again. Already the reflex is to classify by the type of the quadratic form: positive definite gives a closed ray, indefinite gives all of $\mathbb{R}$, and so on. This is correct as far as it goes, and it quickly produces most of the answer by pattern-matching: single point, closed ray, everything. But two things go wrong if this is pursued as *the* method. First, it never proves there is nothing else – an exam script that lists three or four example-types and stops has demonstrated existence, not exhausted the possibilities, and Putnam graders mark on completeness. Second, it offers no route at all to the open ray (an infimum that is approached but never attained), because every low-degree homogeneous example either attains its extreme value (positive definite forms do) or has none (indefinite forms range over all of $\mathbb{R}$). The case-by-case route can spend the whole clock generating more and more degree-3, degree-4 examples looking for the missing type, with no guarantee it is looking in the right place.

The model chain. Since f is a polynomial, it is continuous on $\mathbb{R}^2$, and $\mathbb{R}^2$ is connected (indeed path-connected: any two points are joined by a line segment). The continuous image of a connected set is connected, so the range of f is a connected subset of $\mathbb{R}$ – that is, an interval, possibly degenerate to a point, possibly unbounded on one or both sides. This already rules out, for instance, a range that is two disjoint closed intervals.

Now suppose the range is bounded both above and below. For each fixed real k, consider the restriction to the line through the origin of slope k: $g_k(x) = f(x, kx)$, a polynomial in the single

variable x. Its values are a subset of the (bounded) range of f, so g_k is a bounded polynomial in one variable. A single-variable polynomial of positive degree is unbounded (its leading term dominates as $x \to \pm\infty$), so boundedness forces g_k to have degree 0, i.e. to be constant, for every k. Together with the vertical line $x = 0$ (where $f(0, y)$ must likewise be constant by the same argument, or by taking $k \to \infty$ and using continuity), this means f is constant on every line through the origin, hence constant everywhere with value $f(0, 0)$.

So: if the range is bounded on both sides, it is a single point. Otherwise it is unbounded on at least one side, and being an interval it is a ray or all of $\mathbb{R}$, and a ray may or may not include its finite endpoint. This yields exactly the four types. Examples: (i) $f \equiv 0$ gives a point; (ii) $f = x^2$ gives $[0, \infty)$, a closed ray; (iv) $f = x$ gives $\mathbb{R}$. For (iii), an open ray, take

$$f(x, y) = (xy - 1)^2 + x^2.$$

This is a sum of squares, so $f \geq 0$. For fixed $x \neq 0$, minimising over y by setting $y = 1/x$ kills the first square, leaving $f = x^2$; letting $x \to 0$ (through nonzero values) drives $f \to 0$, yet $f = 0$ would require simultaneously $xy = 1$ and $x = 0$, which is impossible. So 0 is an infimum that is never attained, while f is clearly unbounded above (fix $y = 0$, let $x \to \infty$: $f = 1 + x^2 \to \infty$). The range is $(0, \infty)$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both readings start from the same question – "what are the possibilities for the range?" – but split at the very first decision: what kind of question is this. The natural attempt treats it as an algebra question about the coefficients and degree of f, to be answered by generating and sorting examples. The model treats it as a topology question about the map $f : \mathbb{R}^2 \to \mathbb{R}$ on a connected domain, to be answered by pinning down the shape of the range *before* any example is built.

> *I do not need to know anything about this specific f to know its range is an interval – continuity on a connected domain is the only fact that matters here, so I should establish that constraint first and only afterwards go looking for which intervals are actually achievable.*

> *If I restrict f to any single straight line through the origin, the two-variable problem collapses to single-variable calculus, where I already know boundedness is impossible for anything of positive degree – so testing every line through one point is enough to force the bounded case down to a single value, without ever describing f's coefficients in general.*

- ▷ **Cue:** a question asks for every possible shape of the range/image of a continuous (in particular polynomial) function on a domain known to be connected $\longrightarrow$ **Move:** first invoke continuity plus connectedness of the domain to pin down that the range must be an interval or a point; only then classify the remaining cases (bounded/unbounded, attained/not attained) and search for examples realising each one
- ▷ **Cue:** need an explicit function whose infimum (or supremum) is approached but never attained $\longrightarrow$ **Move:** build it as a sum of two squares of conditions that are each individually satisfiable but never simultaneously satisfiable (e.g. $xy = 1$ and $x = 0$), so one variable can be tuned to kill one square while the other square is driven to 0 in a limit that the first condition forbids from ever being reached exactly

Layer 3 THE WHY

The first move – reach for connectedness before reaching for examples – is licensed by nothing specific to this f at all. A polynomial in x, y is a finite sum of products of the continuous coordinate functions and constants, hence continuous; $\mathbb{R}^2$ is path-connected because any two points are joined by the straight-line segment between them, on which f restricts to a continuous

(indeed polynomial) function of one real parameter, so by the Intermediate Value Theorem f hits every value between any two values it takes on that segment. That is the entire content of "continuous image of a connected set is connected" specialised to this setting, and it costs nothing to invoke because it uses no property of f beyond being continuous. This is why it belongs first: it is free information, obtained before a single coefficient of f has been examined, and it already collapses the answer from "some arbitrary subset of $\mathbb{R}$" to "an interval or a point."

The second move – why bounded-on-both-sides forces *constant* rather than merely some closed interval $[a, b]$ – needs one further fact, and it is a fact about single-variable polynomials, not about f: a real polynomial of degree $d \geq 1$ in one variable is unbounded, because its leading term $c\,x^d$ dominates every lower-order term as $|x| \to \infty$ and so $|g(x)| \to \infty$. Restricting f to a line through the origin is precisely the device that converts the two-variable problem into this one-variable fact, because $f(x, kx)$ is, term by term, still a polynomial in x alone. If the range of f is bounded, every such restriction is a bounded polynomial in x, hence (by the fact just stated) of degree 0: a constant. Since this holds for every direction k through the origin – and the constants must agree as k varies continuously, by continuity of f itself at the origin – f is forced to be the single constant $f(0, 0)$ on the whole plane. The archive's official solution states "the polynomial $f(x, kx)$ is a constant for each value of k" as a single line with no justification; this is exactly where that line comes from.

The third move – constructing $(xy - 1)^2 + x^2$ rather than guessing at an example – follows from asking directly what property an "approached but never attained infimum" example needs: two conditions, each achievable on its own, that cannot be achieved together. $xy = 1$ is achievable for any $x \neq 0$; $x = 0$ is achievable; the two together are impossible, since $x = 0$ forces $xy = 0 \neq 1$. Writing f as the sum of the squared deviations from these two conditions guarantees $f \geq 0$ with equality precisely when both hold, which never happens – yet by fixing x small and choosing $y = 1/x$ to zero out the first square exactly, the second square x^2 can be driven arbitrarily close to 0. This is not a lucky guess; it is the direct algebraic translation of "pick two mutually exclusive but individually satisfiable conditions" into a nonnegative polynomial.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "possibilities for the range" of a polynomial defined on the whole plane. On seeing it, write down immediately: this is a question about the image of a continuous map on a connected domain, so the range is an interval (or a point) before any specific f is considered.

2. That cue suggests classifying by boundedness (bounded on neither side, one side, or both), and for the bounded-on-both-sides case, testing the restriction of f to lines through a fixed point, since that converts a two-variable question into single-variable polynomial behaviour that is already fully understood.

3. The hesitation comes at exactly the step "bounded implies constant." It is tempting to think boundedness on both sides could still leave room for a nontrivial closed interval $[a, b]$. It is reinforced away by remembering that a positive-degree single-variable polynomial is never bounded, so every line-restriction of a bounded f must be literally constant, and continuity forces all those constants to agree.

4. The trap is trying to classify the range by enumerating polynomials degree by degree (linear, quadratic forms, cubic, ...) and hoping the pattern of examples found so far is exhaustive. This can consume the whole allotted time without ever proving completeness, and it gives no natural route to the hardest case, the open ray.

5. To recognise this faster: any question of the form "what can the image/range of a continuous function on a connected domain look like" should trigger connectedness first, boundedness-via-restriction second, and example-hunting only third and last.

6. Decision rule: for the range of a continuous (in particular polynomial) function on a connected

domain, first fix the topological shape of the range using continuity plus connectedness; then resolve boundedness questions by restricting to a lower-dimensional subset (a line, a ray) where a known one-variable fact (growth of polynomials, IVT) applies directly; only then construct explicit examples, and build any "infimum not attained" example as a sum of squares of two individually satisfiable but jointly unsatisfiable conditions.

1969 A2

1969 A2. Let D_n be the determinant of order n of which the element in the ith row and jth column is the absolute value of the difference of i and j. Show that D_n is equal to $(-1)^{n-1}(n-1)2^{n-2}$.

Caveat. The MAA 1965–1984 OCR truncates the problem statement itself: it reads "Let D_n be the determinant of order n of which the element in the ith row and the ith column is the absolute value of the difference of i and j. Show that D_n is equal to" and then breaks to problem A-3 without ever printing the formula (also note the OCR's own "ith row and the ith column" should read "ith row and jth column", an obvious slip since the entry is described as a function of i and j). The formula was reconstructed from the book's solution, which states the answer (itself OCR-garbled as "Dn = (-t)n-1(n -l)zn-2") and derives it via row/column reduction to a triangular minor with diagonal entries $1, 2, 2, \ldots, 2$. I independently verified $D_n = (-1)^{n-1}(n-1)2^{n-2}$ by direct computation for $n = 1, 2, 3, 4$ (values $0, -1, 4, -12$) before writing this analysis, so the reconstruction is confirmed by the mathematics, not merely by trusting the OCR of the solution line.

Fields: Linear algebra (determinants), structured/kernel matrices **Position:** A2 **Answer:** $D_n = (-1)^{n-1}(n-1)2^{n-2}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The determinant is given concretely – $D_n = \det(|i-j|)_{i,j=1}^{n}$ – so the obvious first move is to just compute it: cofactor-expand along the first row or column of the matrix as given. For $n = 1, 2, 3$ this is entirely manageable by hand: $D_1 = 0$, $D_2 = \det\begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix} = -1$, and $D_3 = \det\begin{pmatrix} 0 & 1 & 2 \\ 1 & 0 & 1 \\ 2 & 1 & 0 \end{pmatrix} = 4$. A student comfortable with small determinants might push to $n = 4$ as well and, seeing $0, -1, 4, -12$, try to spot a pattern (ratios $-1, -4, -3$; magnitudes $0, 1, 4, 12$) and guess a closed form by curve-fitting rather than derivation. This is a real risk: with only four data points and two free parameters in a guess like $c \cdot k^n \cdot n$, it is easy to lock onto a plausible-looking but wrong formula, or to get the right shape by luck without any argument that it continues to hold for $n = 5, 6, \ldots$. Worse, direct cofactor expansion of the *un-simplified* $n \times n$ matrix of absolute values does not visibly simplify as n grows: each cofactor is itself a full $(n-1) \times (n-1)$ determinant of absolute-value entries with no special structure exposed, so the work grows combinatorially and no inductive step presents itself. Guessing-then-verifying-by-induction is possible in principle but only works once the correct closed form is already in hand – which is exactly what direct expansion struggles to produce.

The model chain. Rather than expanding the matrix as given, first perform elementary column and row operations to expose structure, since determinants are unchanged by adding a multiple of one row (or column) to another.

Subtract column 1 from every other column. The (i, j) entry for $j \geq 2$ becomes $|i-j| - |i-1|$.

Then add row 1 to every other row. Carrying this out (verified directly for $n = 4$: the matrix $(|i-j|)$ becomes, after these two operations,

$$\begin{pmatrix} 0 & 1 & 2 & 3 \\ 1 & 0 & 2 & 4 \\ 2 & 0 & 0 & 2 \\ 3 & 0 & 0 & 0 \end{pmatrix}$$

for $n = 4$), the last row becomes all zeros except for the value $n-1$ in the first column. This holds for general n by the same mechanism: the operations concentrate the effect of the "corner" entry into the bottom-left position.

Expanding the determinant along this last row, only the entry $(n, 1) = n-1$ survives, with cofactor sign $(-1)^{n+1}$:

$$D_n = (-1)^{n+1}(n-1) \cdot M,$$

where M is the minor obtained by deleting row n and column 1 from the transformed matrix. Direct computation (again verified at $n = 4$: deleting row 4 and column 1 from the matrix above leaves $\begin{pmatrix} 1 & 2 & 3 \\ 0 & 2 & 4 \\ 0 & 0 & 2 \end{pmatrix}$) shows this minor is upper triangular with diagonal entries $1, 2, 2, \ldots, 2$ ($n-2$ copies of 2 after the leading 1), so

$$M = 1 \cdot 2^{n-2} = 2^{n-2}.$$

Since $(-1)^{n+1} = (-1)^{n-1}$,

$$D_n = (-1)^{n-1}(n-1)2^{n-2}.$$

(An alternate route the book also records: derive a two-term recursion $D_{n+1} = -\frac{2n}{n-1} D_n$ for $n > 1$ directly from row operations on D_{n+1}, then iterate from $D_2 = -1$; this reaches the same formula without ever exhibiting the triangular minor, but at the cost of not explaining *why* the recursion has that particular ratio.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both readings begin at the same matrix $(|i-j|)$. They split at the very next decision: whether to expand it as it stands, or to spend a move transforming it first. The natural attempt treats "compute the determinant" as synonymous with "expand it," immediately. The model treats the entries $|i-j|$ as data with exploitable structure and asks what row/column operations that structure suggests, before any expansion is attempted.

> *The entries of this matrix are not n^2 independent numbers – they are one fixed function, $|i-j|$, evaluated at every pair (i, j). A function like $|x|$ has a single kink and is otherwise perfectly linear; if I take a difference across the matrix the way I would take a derivative of $|x|$, the smooth linear parts should cancel and only something sparse, concentrated near the kink, should survive. I should manipulate the matrix to expose that before doing any cofactor expansion at all.*

▷ **Cue:** a determinant's entries are given by one fixed structured kernel function of the two indices (e.g. $|i-j|$, $\min(i, j)$, $\gcd(i, j)$), rather than by unrelated numbers $\longrightarrow$ **Move:** take row or column differences (a discrete analogue of differentiating the kernel) to convert the structure into a sparse, near-triangular matrix, and only then expand the determinant, instead of expanding the original matrix directly

Layer 3 THE WHY

Why should differencing the matrix do anything useful at all? Because $|i-j|$, viewed as a function of i for fixed j, is piecewise linear: it decreases with slope -1 for $i < j$ and increases

with slope $+1$ for $i > j$, with a single kink exactly at $i = j$. This is the discrete analogue of the fact that the second derivative of $|x|$ (in the distributional sense) is $2\delta(x)$, a spike concentrated at the single point $x = 0$: taking one finite difference of a piecewise-linear function removes its smooth linear pieces and leaves a step, and a second finite difference of that step leaves a spike. The two operations performed – subtracting column 1 from every other column, then adding row 1 to every other row – are exactly two finite-difference operations of this kind, applied to a matrix built from a kinked kernel. That is why they concentrate almost the entire matrix into a sparse, near-triangular shape: the operations are doing to the matrix what differentiation does to $|x|$, turning a globally-spread-out linear structure into something local. This is the piece of insight the archive's official solution omits entirely – it states the two operations and their effect as a fait accompli, with no hint of why row/column differencing of *this particular* matrix would be the right thing to try, rather than any of the many other elementary operations available.

A second, smaller piece of WHY: once the matrix is sparsified, the minor left after deleting row n and column 1 is upper triangular *because* the differencing operations were applied in a fixed order (all columns relative to column 1, then all rows relative to row 1), which respects the natural ordering $1, 2, \ldots, n$ of the indices – the same ordering that makes $|i - j|$ kink at $i = j$ in the first place. Triangularity is not a coincidence; it is what happens when a structure that is "ordered" in one variable is differenced in a way that is consistent with that same order.

Finally, why does no case analysis on n (odd vs. even, etc.) enter anywhere? Because the sign $(-1)^{n+1} = (-1)^{n-1}$ and the magnitude $(n-1)2^{n-2}$ both fall out of the single cofactor expansion along the sparsified last row – the formula is uniform in n precisely because the row/column operations used are themselves uniform (defined the same way for every n), unlike a direct expansion of the original matrix, whose recursive structure (if any) is not visible without first performing exactly this reduction.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the matrix entries $|i - j|$ come from one fixed kernel function of the indices, not from n^2 unrelated numbers. On seeing it, write down: this structure should be exploited by row/column operations before any expansion.

2. That cue suggests taking column differences (relative to column 1) and then row sums (relative to row 1) – a discrete analogue of differentiating the kinked function $|i - j|$ – to try to sparsify the matrix before computing the determinant.

3. The hesitation occurs right after performing the operations, when the resulting matrix (worked out explicitly, e.g. at $n = 4$) is not obviously triangular at a glance and it is tempting to abandon the approach and fall back to raw numerical expansion. It is reinforced away by actually carrying the reduction through by hand for one concrete size and checking that a triangular minor genuinely emerges in the corner, confirming the mechanism generalises rather than trusting it on faith.

4. The trap is guessing the closed form from the first three or four values $(0, -1, 4, -12)$ without a mechanism, then either failing to prove it for general n or, worse, mis-fitting a plausible but incorrect formula from too little data (a two-parameter guess needs more than two or three confirmed values to be trustworthy).

5. To recognise this faster: any determinant whose (i, j) entry is a single structured kernel function of i and j – $|i - j|$, $\min(i, j)$, $\gcd(i, j)$, and similar – should trigger "try row/column differencing to expose sparsity" as the first move, before cofactor expansion of the raw matrix is attempted.

6. Decision rule: for a determinant built from a fixed structured kernel of the indices, always attempt discrete differencing (subtracting a reference row or column) to convert the kernel's structure into sparsity before expanding; and never trust a closed-form guess fitted only to the first few small cases without either a derivation or verification on strictly more data points than the guess has free parameters.

1969 A3

1969 A3. Let P be a non-self-intersecting closed polygon with n sides. Let its vertices be $P_1, P_2, \ldots, P_n$. Let m other points, $Q_1, Q_2, \ldots, Q_m$, interior to P be given. Let the figure be triangulated. This means that certain pairs of the $(n+m)$ points $P_1, \ldots, P_n, Q_1, \ldots, Q_m$ are connected by line segments such that (i) the resulting figure consists exclusively of a set T of triangles, (ii) if two different triangles in T have more than a vertex in common then they have exactly a side in common, and (iii) the set of vertices of the triangles in T is precisely the set of $(n+m)$ points $P_1, \ldots, P_n, Q_1, \ldots, Q_m$. How many triangles are in T?

Caveat. The MAA OCR renders the point labels as "P1, Ps, · · · , P,." and "Q11 Q2, · · · , Q,.", with digit/letter substitutions (2 → "s", n → ",.", a doubled "1" for the comma after Q_1) typical of this scan. There is no ambiguity: the surrounding sentence explicitly names n sides and m interior points, so the intended labels are $P_1, \ldots, P_n$ and $Q_1, \ldots, Q_m$, as reconstructed above. The book's own solution confirms the reading – it uses exactly n and m in the answer $t = 2m + n - 2$.

Fields: Combinatorial/plane geometry (triangulations, Euler's formula) **Position:** A3 **Answer:** $t = n + 2m - 2$ triangles.

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong first instinct is to draw small cases and look for the pattern directly. A triangle with no interior points ($n = 3, m = 0$) is already one triangle: $t = 1$. A quadrilateral with no interior points ($n = 4, m = 0$), triangulated by one diagonal, gives $t = 2$. A triangle with one interior point joined to all three vertices ($n = 3, m = 1$) splits into three smaller triangles: $t = 3$. Tabulating $(n, m, t) = (3, 0, 1), (4, 0, 2), (3, 1, 3)$ against the candidate formula $t = n + 2m - 2$ gives $1, 2, 3$ – a match in all three cases, so the formula is plausible. But this is pattern-matching from three data points to a two-parameter family, and it proves nothing about a general triangulation with an arbitrary mix of n and m, built by an arbitrary (and unknown) sequence of cuts. A more careful student might try to prove the formula by induction: add the points $Q_1, \ldots, Q_m$ one at a time, and track how t changes at each step. Adding a single new interior point strictly inside an existing triangle and connecting it to that triangle's three corners removes 1 triangle and adds 3, a net change of $+2$ – consistent with the "$+2m$" term. This is a genuine proof strategy, but it requires justifying that *every* valid triangulation can actually be built this way, one interior point at a time, landing inside a single existing triangle at each step and never, say, needing to also insert new boundary structure or re-triangulate a neighbourhood in a more complicated way. That justification is nontrivial to state cleanly, and a script that assumes it without comment is assuming exactly the fact that would need to be proved.

The model chain. Two independent global counts settle the question without ever describing how the triangulation was built.

Angle sum. Every triangle contributes an angle sum of exactly π, so summing over all t triangles in T, the grand total of every angle in the figure is πt.

Now compute the same total a different way, vertex by vertex. Condition (ii) says triangles meet only along a full shared edge or at a single shared vertex – the triangulation tiles the interior of P exactly, with no gaps and no overlaps. So at each interior point Q_j, the triangles meeting there fill out a complete turn around that point: their angles at Q_j sum to exactly 2π. There are m such points, contributing $2\pi m$ in total. At each boundary vertex P_i, the triangles meeting there fill out exactly the interior angle of the polygon P at P_i – nothing more, since there is no "outside" region past the polygon boundary being tiled. Summing over all n boundary vertices gives the sum of interior angles of an n-gon, which is $(n-2)\pi$.

Equating the two counts of the same total:

$$\pi t = 2\pi m + (n-2)\pi \quad \Longrightarrow \quad t = 2m + n - 2.$$

Euler's formula (independent check). View the triangulation as a planar graph with $V = n + m$ vertices, together with one additional unbounded outer face outside P, so $F = t + 1$ faces in total. Euler's formula for a connected planar graph gives

$$V - E + F = 2.$$

Count edges by summing, over every face (including the outer one), the number of edges bounding it: each edge borders exactly two faces, so this sum is $2E$. Each of the t triangular faces contributes 3 edges; the single outer face is bounded by the n sides of P, contributing n. So

$$3t + n = 2E.$$

From Euler's formula, $E = V + F - 2 = (n+m) + (t+1) - 2 = n + m + t - 1$. Substituting,

$$3t + n = 2(n + m + t - 1) = 2n + 2m + 2t - 2 \quad \Longrightarrow \quad t + n = 2n + 2m - 2 \quad \Longrightarrow \quad t = n + 2m - 2.$$

Both routes agree: $t = n + 2m - 2$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both readings ask the same question – how many triangles – but split at the decision of what kind of quantity to compute. The natural attempt tries to *construct* the count, by building the triangulation up one point (or one cut) at a time and tracking local changes at each step. The model instead looks for a *global invariant* – a total (the sum of all angles, or the Euler characteristic) that every valid triangulation of this shape must produce, regardless of the particular sequence of line segments used to build it.

> *I do not need to know how this specific triangulation was assembled – I need a quantity that is forced to come out the same no matter which valid triangulation someone drew, and "every angle in the whole figure, summed two different ways," is exactly such a quantity, because condition (ii) guarantees the triangles tile the interior exactly, with nothing left over and nothing double-counted.*

▷ **Cue:** count the number of pieces (triangles, faces, regions) in a planar subdivision that exactly tiles a fixed shape, where the specific combinatorial layout of the subdivision is unknown or arbitrary $\longrightarrow$ **Move:** find a globally conserved quantity (total angle measure, at π per triangle; or the Euler characteristic $V - E + F = 2$) and compute it two different ways – once by summing per piece, once by summing per vertex or via the topological identity – then equate and solve, rather than inducting on how the subdivision was assembled

Layer 3 THE WHY

Why does a global invariant even exist here, rather than the count genuinely depending on the specific triangulation chosen? Because the triangulation is required, by condition (ii), to tile the interior of P exactly: no two triangles overlap in area, and none leaves a gap, since any two triangles that share more than a vertex must share a full edge. That exactness is precisely what licenses the claim that the angles meeting at any single point account for *all* of the angle at that point and *no more* – which is what makes "total angle around each point" a fixed, known number independent of how the triangulation was drawn. At an interior point, that fixed number is a full turn, 2π, which is not a fact about triangulations at all but a basic fact about angles surrounding a point in the Euclidean plane. At a boundary vertex P_i, the fixed number is the interior angle of the polygon there, because the tiling stops exactly at the polygon's boundary – there is no region beyond it for the triangles to also fill. The archive's official solution states

"the sum of all the angles is πt... and it is also $2\pi m + (n-2)\pi$" as a single line; the reason the second expression is forced, rather than merely plausible, is this exactness condition (ii), which is doing all the work and is easy to read past.

The Euler's-formula route earns its own WHY: $V - E + F = 2$ is a topological invariant of any connected planar graph (with the plane compactified by one point at infinity, giving the outer face somewhere to live), true regardless of how the graph is drawn – it does not know or care that the graph in question happens to be a triangulation. What turns this generic topological fact into the specific triangle count is a second, independent counting of the same edges: summing "edges per face" over every face counts each edge exactly twice, because every edge in a planar embedding borders exactly two faces (an edge with only one adjacent face would mean the graph fails to separate the plane into distinct regions there). The triangular faces contribute 3 each, and the single outer face contributes exactly n, because its boundary is precisely the original n-gon. Neither counting equation alone determines t; it is the intersection of a purely topological fact ($V - E + F = 2$) with a purely combinatorial fact (every edge is shared by exactly two faces) that pins the answer down, and this is why the method needs no knowledge of the specific triangulation used – both facts hold for every one of them.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the question asks for a count (of triangles) that ought to be the same no matter which specific triangulation was drawn. On seeing it, write down: look for an invariant, not a construction.
2. That cue suggests one of two standard invariants for planar subdivisions: the total angle sum of all pieces, or Euler's formula $V - E + F = 2$; either one converts the counting question into solving a linear equation in n, m, t.
3. The hesitation occurs when deciding what each type of vertex contributes to the angle sum – whether a boundary vertex P_i contributes a full 2π like an interior point, or only its own polygon interior angle. It is reinforced away by checking one tiny concrete example by hand (a triangle with a single interior point) and confirming the angle sum matches πt only when boundary vertices contribute their polygon angle, not 2π.
4. The trap is assuming every vertex, boundary or interior, contributes 2π, which silently over-counts and produces the wrong formula; a second trap is trying to induct on "add one point at a time" without first confirming every valid triangulation is actually reachable that way.
5. To recognise this faster: any request to count the pieces of an exact tiling of a fixed region – no gaps, no overlaps – should immediately trigger "sum a conserved local quantity two different ways," before any attempt to construct or induct on the specific dissection.
6. Decision rule: to count the pieces of any planar dissection that exactly tiles a region, find a globally conserved quantity (total angle measure, at a fixed amount per piece; or the Euler characteristic) and evaluate it in two independent ways – once locally per piece, once globally per vertex or via the topological identity – then equate and solve the resulting linear relation, entirely avoiding the need to describe or induct on the dissection's specific combinatorial structure.

1969 A4

1969 A4. Show that

$$\int_0^1 x^x \, dx = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n^n}.$$

(The integrand is taken to be 1 at $x = 0$.)

Caveat. The MAA OCR of the problem statement drops the displayed formula entirely: it reads "A-4. Show that" followed directly by "(The integrand is taken to be 1 at x =0.)". The identity above is reconstructed, not transcribed. The reconstruction is secured three ways. (a) The book's index lists 1969 A-4 under both "Infinite series" and "Definite integrals". (b) The book's own solution for A-4 writes $\int_0^1 x^x\,dx = \sum_{m\ge 0} \frac{1}{m!}\int_0^1 x^m(\log x)^m\,dx$, evaluates the m-th integral as $(-1)^m m!\,(m+1)^{-m-1}$, and says "replace $m+1$ by n", which produces exactly $\sum_{n\ge 1}(-1)^{n+1}n^{-n}$; the parenthetical about the value at $x=0$ only makes sense for the integrand x^x. (c) Both sides agree numerically to 30 digits (0.783430510712134407…). The typography of the book's display (whether it prints $(-1)^{n+1}n^{-n}$ or $\frac{(-1)^{n+1}}{n^n}$) cannot be recovered from the scan; the mathematics is not in doubt.

Fields: Real analysis (power series, term-by-term integration, uniform convergence); integral calculus
Position: A4 **Answer:** Write $x^x = e^{x\log x} = \sum_{m\ge 0}\frac{(x\log x)^m}{m!}$, integrate term by term (legal by uniform convergence, since $|x\log x| \le 1/e$ on $[0,1]$), and use $\int_0^1 x^m(\log x)^m\,dx = \frac{(-1)^m m!}{(m+1)^{m+1}}$; the $m!$ cancels and $n = m+1$ gives $\sum_{n\ge 1}(-1)^{n+1}n^{-n}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You see a definite integral and your A-level reflex fires: find an antiderivative. So you differentiate x^x to see what family it lives in. Writing $y = x^x$, $\log y = x\log x$, so $y'/y = 1+\log x$ and

$$\frac{d}{dx}x^x = x^x(1+\log x).$$

That is the derivative, not the integral, and nothing in it rearranges into "something whose derivative is x^x". You try integration by parts with $u = x^x$, $dv = dx$:

$$\int_0^1 x^x\,dx = \big[x\cdot x^x\big]_0^1 - \int_0^1 x\cdot x^x(1+\log x)\,dx = 1 - \int_0^1 x^{x+1}(1+\log x)\,dx,$$

and the new integrand is worse than the old one. Each further round of parts raises the power of x and adds another $\log x$ factor. This is the wall, and it is a genuine wall: x^x has no elementary antiderivative, so no amount of substitution or parts will close the integral in finite form. The zombie route consumes the exam here.

A better-read student looks at the right-hand side, sees an infinite series, and tries to manufacture a series from the integral by Riemann sums: $\frac{1}{N}\sum_{k=1}^{N}(k/N)^{k/N}$. That produces an infinite series in the limit, but the terms $(k/N)^{k/N}$ bear no resemblance to n^{-n}, and the book's own commentary records that this route "apparently doesn't work". A third try is to expand x^x in powers of x about 0. That fails immediately: $x^x = 1 + x\log x + \cdots$, and $x\log x$ is not a power of x, so x^x has no Taylor series at 0 at all.

The model chain. **Step 1: rewrite by definition.** For $x>0$, x^x means $e^{x\log x}$. Put $g(x) = x\log x$ for $0<x\le 1$ and $g(0)=0$. Since $x\log x \to 0$ as $x\to 0^+$, g is continuous on $[0,1]$, and $e^{g(0)} = 1$ matches the convention in the problem. So the integrand is $e^{g(x)}$ on all of $[0,1]$.

Step 2: bound the exponent. $g'(x) = 1+\log x$ vanishes only at $x = 1/e$, where $g(1/e) = -1/e$; also $g(0) = g(1) = 0$ and $g \le 0$ on $[0,1]$. Hence $|g(x)| \le 1/e$ on $[0,1]$.

Step 3: expand and interchange. By the exponential series,

$$x^x = \sum_{m=0}^{\infty}\frac{g(x)^m}{m!}.$$

Each term is continuous on $[0,1]$ and $\left|g(x)^m/m!\right| \le e^{-m}/m!$, with $\sum_m e^{-m}/m! = e^{1/e} < \infty$. By the Weierstrass M-test the series converges uniformly on $[0,1]$, and a uniformly convergent

series of continuous functions on a closed bounded interval may be integrated term by term:

$$\int_0^1 x^x\,dx = \sum_{m=0}^{\infty} \frac{1}{m!} \int_0^1 x^m (\log x)^m\,dx.$$

Step 4: evaluate each integral. Substitute $x = e^{-t}$, so $dx = -e^{-t}\,dt$, $\log x = -t$, and x running from 0 to 1 corresponds to t running from ∞ to 0:

$$\int_0^1 x^m (\log x)^m\,dx = (-1)^m \int_0^\infty t^m e^{-(m+1)t}\,dt.$$

With $s = (m+1)t$ this is $(-1)^m (m+1)^{-(m+1)} \int_0^\infty s^m e^{-s}\,ds = \dfrac{(-1)^m\, m!}{(m+1)^{m+1}}$, using $\int_0^\infty s^m e^{-s}\,ds = m!$ (induction on m via integration by parts; the boundary terms $s^m e^{-s}$ vanish at 0 for $m \geq 1$ and at ∞ always).

Step 5: assemble. The $m!$ cancels:

$$\int_0^1 x^x\,dx = \sum_{m=0}^{\infty} \frac{(-1)^m}{(m+1)^{m+1}} = \sum_{n=1}^{\infty} \frac{(-1)^{n+1}}{n^n},$$

with $n = m+1$. ■

The book's variant of Step 4. It keeps x and integrates by parts: with $F(m,k) = \int_0^1 x^m (\log x)^k\,dx$, taking $u = (\log x)^k$, $dv = x^m\,dx$ gives $F(m,k) = -\frac{k}{m+1} F(m,k-1)$ for $k \geq 1$ (the boundary term $\frac{x^{m+1}}{m+1}(\log x)^k$ is 0 at $x = 1$ and tends to 0 as $x \to 0^+$). Iterating from $F(m,0) = \frac{1}{m+1}$ gives $F(m,m) = (-1)^m m!\,(m+1)^{-m-1}$, which agrees with Step 4.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths start by treating x^x as the object to be integrated. The split is at the first line: the natural attempt asks "what is the antiderivative of x^x?", while the model asks "what *is* x^x?" and writes $x^x = e^{x \log x}$ before doing any calculus. Once that line is on the page, the exponential series is the only tool in reach, and the rest is bookkeeping. The second, smaller split comes at Step 3: the natural instinct is to swap $\int$ and $\sum$ without comment, and the model stops to pay for the swap with the bound $|x \log x| \leq 1/e$.

> *The answer is a series, not a closed form, so the question is not asking me to find an antiderivative; it is asking me to turn the integrand into a series and integrate it one term at a time. The integrand is a power whose exponent is a variable, and the only honest definition of that is $e^{x \log x}$, so the series I should use is the exponential one.*

> *Every term I will get is $\int_0^1 x^m (\log x)^m\,dx$. I know that $\log x$ becomes $-t$ under $x = e^{-t}$, which turns this into a Gamma integral $\int_0^\infty t^m e^{-(m+1)t}\,dt = m!/(m+1)^{m+1}$. The $m!$ will cancel the $1/m!$ from the exponential series and leave $(m+1)^{-(m+1)}$. That is the n^{-n} in the target, so I am on the right road.*

> *Before I swap the sum and the integral I need one number: how large can $x \log x$ get on $[0,1]$? It is at most $1/e$ in size, so the series is dominated by $\sum e^{-m}/m!$ and the swap is legal. That single line is where the marks are.*

▷ **Cue:** the integrand is a power with a variable exponent, $f(x)^{h(x)}$ ⟶ **Move:** rewrite it as $e^{h(x) \log f(x)}$ before any other calculus; if an exact value is wanted as a series, expand the exponential

▷ **Cue:** the claimed value of an integral is an infinite series whose terms contain n^{-n} or $k!/a^{k+1}$ $\longrightarrow$ **Move:** look for a term-by-term integration that produces Gamma integrals $\int_0^\infty t^k e^{-at}\,dt = k!/a^{k+1}$, reached from $\int_0^1 x^{a-1}(\log x)^k\,dx$ by $x = e^{-t}$

▷ **Cue:** you are about to swap $\int$ and $\sum$ over a closed bounded interval $\longrightarrow$ **Move:** find a uniform bound M_m on the m-th term with $\sum M_m < \infty$ (Weierstrass M-test) and state it; for an exponential series this means bounding the exponent

Layer 3 THE WHY

Why is $e^{x\log x}$ the first thing to write, rather than a lucky move? Go back to the definition. For a fixed positive base a and a real exponent r, what does a^r mean? For whole-number r it means repeated multiplication. For rational r it means roots. For irrational r, and for any exponent that varies continuously, the only definition that works is $a^r := e^{r\log a}$; that is how the A-level derivative $\frac{d}{dx}a^x = a^x \log a$ is proved in the first place. So when both base and exponent are x, the expression x^x has no meaning other than $e^{x\log x}$. Writing it that way is not a trick. It is reading the definition of the symbol you were given. You used the same definition to differentiate x^x in the natural attempt (taking logs is the same move in disguise). The only change is to use it for integration as well.

Why a series, and why the exponential one? Read the target. The right-hand side is an infinite sum, which tells you the problem-setter does not expect a closed form: the integral has none in elementary terms. An infinite sum from an integral means one of two things, Riemann sums or term-by-term integration of an expanded integrand. Riemann sums give terms that look like values of the integrand at sample points, and $(k/N)^{k/N}$ looks nothing like n^{-n}. So you are looking for an expansion of the integrand. Once you have written $x^x = e^g$, the exponential series is the natural expansion. It also has the one property that makes the swap cheap: it converges absolutely and uniformly wherever its argument is bounded.

Why does n^{-n} come out, and could you have predicted it? The n-th power of n in the denominator is the signature of a scaled Gamma integral: $\int_0^\infty t^k e^{-at}\,dt = k!/a^{k+1}$, and taking $a = k+1$ gives $k!/(k+1)^{k+1}$. So the target is telling you that the m-th term involves t^m and $e^{-(m+1)t}$ together. Under $x = e^{-t}$, the factor $e^{-(m+1)t}$ is $x^m\,dx$ (up to sign), and t^m is $(-\log x)^m$. So the m-th term must be $\int_0^1 x^m(\log x)^m\,dx$, which is $\int (x\log x)^m\,dx$. That is the m-th power of the exponent in $e^{x\log x}$. Worked backwards from the answer, the model solution is forced: the target n^{-n} gives you the Gamma integral, the Gamma integral gives you $(x\log x)^m$, and $(x\log x)^m$ gives you the exponential series. The official solution states "write the integrand as a power of e and use the series expansion for e" without saying where it comes from. This backward reading is where it comes from.

Why does the swap need justification here in particular? Because $\log x$ is unbounded near 0, a reader has a real reason to worry that the terms misbehave at the endpoint. They don't, because x beats $\log x$: $x\log x \to 0$. Finding the actual maximum of $|x\log x|$ ($1/e$, at $x = 1/e$) turns that worry into a single inequality. Here the M-test is not decoration. It answers the one worry the endpoint raises.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the exponent: x appears in both the base and the exponent. On seeing it, write $x^x = e^{x\log x}$ on the first line. The second cue is the right-hand side, an infinite series with n^{-n}. Write next to it "$k!/a^{k+1}$ = Gamma integral".

2. Expand the integrand as a power series (here the exponential series in $x\log x$), integrate term by term after checking uniform convergence, and evaluate each term with the substitution $x = e^{-t}$, which turns $\int_0^1 x^m(\log x)^m\,dx$ into a Gamma integral.

3. Hesitation comes at two points. The first is "x^x has no Taylor series at 0, so series methods are dead". You get past it by expanding in the exponent $x\log x$ rather than in x. The second is the swap of $\int$ and $\sum$ near the singularity of $\log x$. You get past that by computing $\max_{[0,1]}|x\log x| = 1/e$ once and quoting the M-test.

4. There are two traps. The first is spending the exam hunting for an antiderivative of x^x, which does not exist in elementary form: repeated integration by parts only makes the integrand worse. The second is swapping the sum and the integral with no justification, which is exactly the rigour the question is testing once the calculation is routine. A lesser trap is an off-by-one in the reindexing: the m-th term is $(-1)^m/(m+1)^{m+1}$, and the sign becomes $(-1)^{n+1}$ once $n = m+1$.

5. Whenever the claimed value of an integral is an infinite series, stop looking for an antiderivative. Match the shape of the terms to a known integral ($k!/a^{k+1}$ means Gamma, $1/n^2$ means $\int x^{n-1}\log x$, and so on) and work backwards to the expansion of the integrand that would produce those terms.

6. Decision rule: if an integral has no elementary antiderivative, or its answer is presented as a series, expand the integrand in a series that converges uniformly on the interval and integrate term by term. Choose the expansion by working backwards from the form the terms must take, and pay for the interchange with an explicit uniform bound on the terms. A variable exponent always means rewriting as $e^{(\text{exponent})\log(\text{base})}$ first. The same argument gives the companion identity $\int_0^1 x^{-x}\,dx = \sum_{n\geq 1} n^{-n}$.

1969 A5

1969 A5. Let $u(t)$ be a continuous function in the system of differential equations

$$\frac{dx}{dt} = -2y + u(t), \qquad \frac{dy}{dt} = -2x + u(t).$$

Show that, regardless of the choice of $u(t)$, the solution of the system which satisfies $x = x_0$, $y = y_0$ at $t = 0$ will never pass through $(0,0)$ unless $x_0 = y_0$. When $x_0 = y_0$, show that, for any positive value t_0 of t, it is possible to choose $u(t)$ so the solution is at $(0,0)$ when $t = t_0$.

Caveat. The MAA OCR renders the initial values as "x =xo, y =yo", "Xo=Yo" and "to oft"; these are x_0, y_0 and "t_0 of t", reconstructed above. The identity check passes: the book's solution under A-5 subtracts the two equations to get $d(x-y)/dt = 2(x-y)$ (OCR: "x-y-(x0 -yo)e21", i.e. $x - y = (x_0 - y_0)e^{2t}$), which solves exactly this statement. Both the book's control and the alternative constant control given below were checked symbolically.

Fields: Ordinary differential equations; linear systems; elementary control **Position:** A5 **Answer:** Along every solution $x - y = (x_0 - y_0)e^{2t}$, which is never zero when $x_0 \neq y_0$; when $x_0 = y_0$ the trajectory stays on the line $x = y$, and the control $u(t) = 2x_0(1 - t/t_0) - x_0/t_0$ (or the constant $u \equiv -2x_0/(e^{2t_0} - 1)$) brings it to the origin at $t = t_0$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You have taken a course in linear ODEs, so you reach for the machinery. Write the system as $\mathbf{r}' = A\mathbf{r} + u(t)\mathbf{1}$ with

$$A = \begin{pmatrix} 0 & -2 \\ -2 & 0 \end{pmatrix}, \qquad \mathbf{1} = \begin{pmatrix} 1 \\ 1 \end{pmatrix}.$$

Or, without matrices: differentiate the first equation and substitute the second, getting $x'' = -2y' + u' = 4x - 2u + u'$, that is, $x'' - 4x = u' - 2u$. This already has a problem: u is only continuous, so u' may not exist. Suppose you ignore that and push on. The homogeneous solutions are e^{2t} and e^{-2t}, and variation of parameters gives $x(t)$ as $C_1e^{2t} + C_2e^{-2t}$ plus a convolution of u against those exponentials. You do the same for y, fit the four constants to

x_0, y_0 and the first-order equations, and now you are supposed to show that $x(t_1) = y(t_1) = 0$ is impossible for every continuous u and every t_1. You are holding two expressions, each containing an unknown integral of an arbitrary function, and asking whether they can vanish together. The integrals do not cancel in any way you can see, because you have not arranged them to. That is the wall. Nothing in the calculation is wrong; it is pointed in the wrong direction. The quantity that settles the question is buried inside two formulas instead of standing alone in one.

The model chain. **Part 1.** Subtract the second equation from the first. The forcing term $u(t)$ appears in both with coefficient 1, so it cancels:

$$\frac{d}{dt}(x - y) = -2y + 2x = 2(x - y).$$

Put $w = x - y$. Then $w' = 2w$, so $\left(we^{-2t}\right)' = (w' - 2w)e^{-2t} = 0$, and we^{-2t} is constant, equal to its value $x_0 - y_0$ at $t = 0$. Hence, for every t at which the solution is defined and whatever continuous u was used,

$$x(t) - y(t) = (x_0 - y_0)e^{2t}.$$

If $x_0 \neq y_0$ the right-hand side is never zero, so $x(t) \neq y(t)$ at every time, and in particular $(x, y) = (0, 0)$ never occurs. Note that this argument does not assume uniqueness of the solution to the full system: any solution satisfies $w' = 2w$, and the integrating-factor step pins w down completely.

Part 2. Now $x_0 = y_0$, so $w \equiv 0$: the trajectory lies on the line $x = y$ for all t. Substituting $y = x$, both equations become the single scalar equation

$$x' = -2x + u(t).$$

The task is now one-dimensional: steer x from x_0 to 0 in time t_0. Choose the path first and read off the control. Take the straight-line path $x(t) = x_0 - at$ with $a = x_0/t_0$, so that $x(t_0) = 0$. Then $x' = -a$, and the equation demands $-a = -2(x_0 - at) + u(t)$, that is,

$$u(t) = 2(x_0 - at) - a = 2x_0\left(1 - \frac{t}{t_0}\right) - \frac{x_0}{t_0},$$

which is continuous (a polynomial in t). With this u, the pair $x(t) = y(t) = x_0(1 - t/t_0)$ satisfies both equations and the initial condition, and is at $(0, 0)$ when $t = t_0$. Because the system is linear with continuous coefficients, its solution is unique, so this is *the* solution the problem refers to. (If $x_0 = y_0 = 0$, $a = 0$ and $u \equiv 0$ keeps the point at the origin throughout.)

An alternative control. Set $s = x + y$; adding the equations gives $s' = -2s + 2u$. The integrating factor e^{2t} gives $s(t_0)e^{2t_0} = 2x_0 + 2\int_0^{t_0} u(\tau)e^{2\tau}\,d\tau$. A constant $u \equiv c$ makes this $2x_0 + c(e^{2t_0} - 1)$, which vanishes for $c = -2x_0/(e^{2t_0} - 1)$. Both controls work; the book uses the first.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same two equations. The split is the very first line: the natural attempt treats the system as a thing to be *solved* (eliminate y, get a second-order equation, variation of parameters). The model treats it as a thing to be *combined* so that the unknown u disappears. At "subtract the two equations" the natural attempt was already differentiating.

> *The claim is "regardless of the choice of u", so whatever proves it cannot depend on u. I don't know u and never will. So before solving anything, I look for a combination of x and y whose equation has no u in it. u enters both equations identically, so $x - y$ is that combination.*

> *Once $x = y$ is forced, I've got one equation in one unknown, $x' = -2x + u$, and I'm the one choosing u. So I don't solve for x given u; I choose the x I want and solve for u.*

- ▷ **Cue:** a system in which an arbitrary or unknown forcing term appears with the same coefficients in every equation, and the claim is "for every choice of the forcing" $\longrightarrow$ **Move:** form the combination of the unknowns that cancels the forcing (here subtract the equations); the resulting equation is autonomous and usually solvable in one line
- ▷ **Cue:** you may choose the input function of an ODE to make the solution hit a target $\longrightarrow$ **Move:** prescribe the trajectory you want (a straight line is the cheapest) and read the input off from the equation, instead of solving the ODE for a general input

Layer 3 THE WHY

Why is subtraction the natural move rather than a lucky one? Go back to what the statement is asking. "Regardless of the choice of $u(t)$" is a universally quantified claim over an infinite-dimensional family of functions. A proof of a statement of the form "for every u, P" has only two honest shapes: either you handle a general u and show its contribution cannot matter, or you find something that u cannot touch at all. The second is always cheaper when it is available. So the question to ask in the arena is: *is there a quantity u cannot touch?* And u enters the system as the vector $u(t)(1,1)$: it pushes the point only along the direction of the line $x = y$. Any linear functional that kills $(1,1)$ – that is, a multiple of $x - y$ – is blind to u. That is not a trick; it is reading off the direction in which the control acts and taking the complementary coordinate.

Now check that the uncontrolled coordinate behaves well. The drift matrix $A = \begin{pmatrix} 0 & -2 \\ -2 & 0 \end{pmatrix}$ has eigenvectors $(1,1)$ with eigenvalue -2 and $(1,-1)$ with eigenvalue 2. So the natural coordinates of this system are exactly $s = x + y$ and $w = x - y$, and in them it decouples completely: $s' = -2s + 2u$, $w' = 2w$. The control acts only on s; the w-coordinate evolves on its own as $w_0 e^{2t}$. An exponential never vanishes, which is Part 1. And if $w_0 = 0$ then w stays 0, which is why Part 2 starts on the line $x = y$ and stays there. The whole problem is the statement "the control direction $(1,1)$ is an eigenvector of A, so the reachable set from a point is confined to a line parallel to it". The origin lies on the line through (x_0, y_0) in direction $(1,1)$ exactly when $x_0 = y_0$.

The book's solution begins "Subtracting the two equations eliminates $u(t)$" and gives no reason for choosing subtraction. The reason is the quantifier plus the shared coefficient: u appears identically in both lines, and the claim has to hold for every u.

For Part 2 the WHY is about direction of computation. In a control question the input is yours to choose, so the unknown is u, not x. Choosing $x(t)$ and reading off $u = x' + 2x$ turns a differential equation into a differentiation. The only constraints are that $x(0) = x_0$, $x(t_0) = 0$, x is C^1, and the resulting u is continuous. Any C^1 path with those endpoints works; the straight line is the cheapest one to write down. That is why the book's ansatz $x = x_0 - at$ is not a guess: it is the simplest member of a family where every member works.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that $u(t)$ appears with coefficient $+1$ in both equations, alongside the phrase "regardless of the choice of $u(t)$". On seeing it, write down $\frac{d}{dt}(x - y) = \ldots$ before anything else.

2. It suggests eliminating the forcing by a linear combination, which here decouples the system into its eigen-coordinates $x \pm y$; the u-free coordinate satisfies $w' = 2w$, solved by an integrating factor.

3. Hesitation comes at "should I solve the whole system first?". Remove it by reminding yourself that a universally quantified claim over u is proved most cheaply by a u-free quantity, so look for one first. In Part 2, hesitation comes at "solve $x' = -2x + u$ for general u". You are choosing u, so pick x instead.

4. The trap is eliminating y to get $x'' - 4x = u' - 2u$. It needs u to be differentiable, which the problem does not give you, and it buries the invariant inside two convolution integrals. A smaller trap: citing uniqueness for Part 1 without need. The integrating-factor argument covers every solution.

5. Whenever an arbitrary input enters every equation through the same vector, ask which coordinate that vector is invisible to. Whenever you control an input, ask whether prescribing the output and solving for the input is a differentiation instead of an integration.

6. Decision rule: to prove something "for every choice" of an unknown function or parameter, first search for a combination of the variables from which it cancels, and prove the claim for that combination alone. To show a target is reachable by choosing an input, prescribe the simplest admissible trajectory and back out the input.

1969 A6

1969 A6. Let a sequence $\{x_n\}$ be given, and let $y_n = x_{n-1} + 2x_n$, $n = 2, 3, 4, \cdots$. Suppose that the sequence $\{y_n\}$ converges. Prove that the sequence $\{x_n\}$ also converges.

Caveat. The MAA OCR renders the subscripts as "$\{x, .\}$", "$y, . = x, .1 + 2x, .$", the usual scan corruption of n. The reading $y_n = x_{n-1} + 2x_n$ is confirmed by the book's own solution, which writes $|y_n - y| = |x_{n-1} + 2x_n - 3\xi| = |2(x_n - \xi) + (x_{n-1} - \xi)|$ with $\xi = y/3$; that identity holds only for the coefficients 1 on x_{n-1} and 2 on x_n. No reconstruction beyond the subscripts was needed.

Fields: Real analysis (sequences, linear recurrences, contraction estimates) **Position:** A6 **Answer:** $x_n \to \lim y_n/3$; the proof is a contraction estimate on $e_n = x_n - \frac{1}{3}\lim y_n$.

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong student's first move is to find the candidate limit. If $x_n \to L$, then $y_n \to L + 2L = 3L$, so the only possible limit is $L = y/3$, where $y = \lim y_n$. That part is right, and it is where most scripts stop being right. The next line is usually "take limits in $y_n = x_{n-1} + 2x_n$", which gives $y = 3L$ again – but this uses the conclusion. Limits of x_n may not be taken until x_n is known to converge, and that is the whole problem. Written this way, the proof is circular and scores nothing.

The more careful student then tries the monotone convergence theorem: show $\{x_n\}$ is monotone and bounded. It is neither in general. Take $y_n = 3$ for every n and $x_1 = 0$; then $x_n = (3 - x_{n-1})/2$ gives $0, \frac{3}{2}, \frac{3}{4}, \frac{9}{8}, \ldots$, which oscillates around 1. The monotone route is dead on the first example you try, and a student who has only that tool is now stuck.

A third attempt solves the recurrence outright. From $x_n = \frac{1}{2}y_n - \frac{1}{2}x_{n-1}$, unrolling gives

$$x_n = \left(-\tfrac{1}{2}\right)^{n-1} x_1 + \sum_{k=2}^{n} \tfrac{1}{2}\left(-\tfrac{1}{2}\right)^{n-k} y_k.$$

This is correct, and it can be pushed to a proof, but now you are proving that a weighted alternating average of a convergent sequence converges to the right value – a Toeplitz-type limit with an interchange you must justify by hand. It works; it costs a page; and the error terms are easy to fumble.

The model chain. Let $y = \lim_{n\to\infty} y_n$ and set $L = y/3$. Define the error $e_n = x_n - L$ and the forcing term $d_n = y_n - y$, so $d_n \to 0$. Since $x_{n-1} + 2x_n - 3L = y_n - y$,

$$2e_n + e_{n-1} = d_n, \qquad \text{so} \qquad e_n = \tfrac{1}{2}\big(d_n - e_{n-1}\big), \quad n \geq 2. \tag{†}$$

Taking absolute values,

$$|e_n| \le \tfrac{1}{2}|e_{n-1}| + \delta_n, \qquad \delta_n = \tfrac{1}{2}|d_n| \to 0. \tag{‡}$$

Bounded. A convergent sequence is bounded, so $D = \sup_{n\ge 2} \delta_n$ is finite. Put $B = \max(|e_1|, 2D)$. If $|e_{n-1}| \le B$ then by (‡), $|e_n| \le B/2 + D \le B/2 + B/2 = B$. By induction $|e_n| \le B$ for all n.

Tends to zero. Since $\{|e_n|\}$ is bounded, $M = \limsup_{n\to\infty} |e_n|$ is a finite non-negative number. Take lim sup of both sides of (‡): $\limsup |e_{n-1}| = M$ (shifting the index does not change a lim sup) and $\lim \delta_n = 0$, so

$$M \le \tfrac{1}{2}M + 0, \qquad \text{hence} \qquad M \le 0.$$

So $M = 0$, i.e. $|e_n| \to 0$, and $x_n \to L = \frac{1}{3}\lim y_n$. □

The official solution runs the same inequality with explicit ε: from $|y_n - y| < \varepsilon/2$ for $n > N$ it gets $|x_n - L| < \varepsilon/4 + \frac{1}{2}|x_{n-1} - L|$ and iterates m times to reach $|x_{n+m} - L| < \varepsilon/2 + 2^{-(m+1)}|x_{n-1} - L|$, then takes m large. The lim sup version above is the same argument with the geometric series summed once and for all.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree up to the candidate limit $L = y/3$. They split at the next line. The natural attempt keeps looking at x_n itself – takes limits in it, asks whether it is monotone, or solves for it. The model changes variable to the *error* $e_n = x_n - L$ and asks one question of the relation (†): does it shrink errors or grow them? Everything after (†) is bookkeeping.

> *I know what the limit has to be, so I stop tracking x_n and track $x_n - L$. The recurrence then says the new error is minus half the old error plus something going to zero. Halving wins against a vanishing push – that is the proof, and all I have to do is write it without assuming the limit exists.*

> *Before I trust this, where does the 2 matter? If the coefficients were swapped, $y_n = 2x_{n-1} + x_n$, the error would double each step, and $x_n = (-2)^n$ gives $y_n = 0$. So the proof must use $|-\frac{1}{2}| < 1$ somewhere explicitly, or it is wrong.*

- ▷ **Cue:** a sequence is known only through a linear combination $y_n = ax_{n-1} + bx_n$ of consecutive terms, and y_n converges ⟶ **Move:** find the only possible limit L from $(a+b)L = \lim y_n$, subtract it to get the error recurrence $e_n = -(a/b)e_{n-1} + d_n/b$ with $d_n \to 0$, and check whether $|a/b| < 1$; if so, bound then take lim sup
- ▷ **Cue:** you need to prove convergence and the only tool that comes to mind is "take limits on both sides" ⟶ **Move:** stop: that assumes the conclusion; use the equation only to identify the candidate limit, then prove convergence to it through an inequality on the error

Layer 3 THE WHY

Why is the error the right object? Because of what convergence is. The definition of $x_n \to L$ is a statement about $|x_n - L|$ and nothing else. Once the limiting equation has told you what L must be, the problem is literally "show $|x_n - L|$ is eventually small", so the variable you should be writing recurrences for is $x_n - L$. Tracking x_n directly carries the constant L around as dead weight and hides the structure. This is the same move as centring a circle problem on the centre: the definition names the object of interest, so you put it at the origin.

Why does the error recurrence come out so clean? Because the relation is linear and L is chosen as a fixed point. Subtracting $3L = L + 2L$ from $y_n = x_{n-1} + 2x_n$ removes the constant exactly and leaves the homogeneous part $e_{n-1} + 2e_n$ driven by $d_n = y_n - y$. Linear relations always split this way: a particular solution (the constant L) plus the solution of the

homogeneous equation plus a forced part. The homogeneous equation $e_{n-1} + 2e_n = 0$ has solutions $e_n = C(-\frac{1}{2})^n$, and whether x_n converges is entirely the question of whether $|-\frac{1}{2}| < 1$. The model's inequality (‡) is that question asked in a form that survives the forcing term.

Why must the lim sup step be preceded by boundedness? The inequality $M \le \frac{1}{2}M$ forces $M = 0$ only when M is a finite number; with $M = \infty$ it reads $\infty \le \infty$ and says nothing. That is where the geometric factor is used a second time: (‡) keeps $|e_n|$ below $\max(|e_1|, 2D)$ because a factor of $\frac{1}{2}$ cannot let a bounded push accumulate. Skipping this is the single most common gap in written solutions to this problem.

How sharp is the hypothesis? Replace 2 by any non-zero c and consider $y_n = x_{n-1} + c\,x_n$. The same proof, with contraction factor $1/|c|$, works whenever $|c| > 1$. When $0 < |c| \le 1$ with $c \neq -1$, the sequence $x_n = (-1/c)^n$ has $y_n = (-1/c)^{n-1}(1-1) = 0$ yet does not converge; when $c = -1$, $x_n = n$ gives $y_n = -1$. So the statement holds for non-zero c exactly when $|c| > 1$, and the number 2 in the problem is there to make the characteristic root $-\frac{1}{2}$ lie inside the unit disc. Seeing this tells you which fact the proof must spend.

Layer $+\alpha$ BACKTESTING

1. The first cue is that y_n is a fixed linear combination of consecutive x's and only y_n is controlled. On seeing it, write down the only possible limit, $L = \frac{1}{3}\lim y_n$, and immediately define $e_n = x_n - L$.

2. A contraction estimate: rewrite the relation as $e_n = -\frac{1}{2}e_{n-1} + \frac{1}{2}d_n$ with $d_n \to 0$, bound the sequence, and close with lim sup (or an explicit ε-iteration).

3. Hesitation comes after finding L: the student feels they have "shown" the limit is $y/3$ and cannot see what is left to prove. It is reinforced away by stating out loud that the limiting equation only identifies the candidate and says nothing about existence; the proof is a statement about $|x_n - L|$.

4. Two traps. First, taking limits in $y_n = x_{n-1} + 2x_n$ before convergence of x_n is known – circular. Second, applying lim sup to (‡) without first proving boundedness, so that $M = \infty$ is not excluded. A third, quieter trap is reaching for monotonicity, which fails on the first numerical example.

5. Any time a problem controls a combination of neighbouring terms and asks about the terms themselves, ask what the homogeneous version does to a small perturbation. Swapping the coefficients and testing $x_n = r^n$ takes ten seconds and shows whether the result is even true.

6. Decision rule: to prove convergence of a sequence defined implicitly by a relation, first read off the only possible limit from the relation, then pass to the error and show the relation contracts it (factor strictly less than 1 in absolute value) up to a vanishing forcing term; prove boundedness first, then conclude with lim sup $M \le kM$, $k < 1$. If the factor is ≥ 1, look for a counterexample of the form r^n instead.

Putnam 1969 — Paper B

1969 B1

1969 B1. Let n be a positive integer such that $n+1$ is divisible by 24. Prove that the sum of all the divisors of n is divisible by 24.

Caveat. The MAA 1965–1984 OCR renders the hypothesis as "$n+t$ is divisible by 24" – a single-character substitution of the digit `1` for the letter `t`, a common scan artefact in this volume. The book's own solution line settles the reading beyond doubt: it opens "The condition $24|n+1$ is equivalent to $n \equiv -1 \pmod 3$ and $n \equiv -1 \pmod 8$", which only parses if the hypothesis is $24 \mid n+1$. No other content in the statement is affected, and the solution was independently re-derived in full below (Layer 1) to confirm the reading.

Fields: Number theory: congruences, divisor sums, the Chinese Remainder Theorem **Position:** B1 (opener) **Answer:** Divisors of n pair off as $\{d, n/d\}$, and the hypothesis forces every such pair to sum to a multiple of 24, so the whole sum is.

Layer 1 THE WHAT — surface analysis

The natural attempt. "Sum of divisors" summons σ, the multiplicative arithmetic function, on sight: write $n = p_1^{a_1} \cdots p_k^{a_k}$ and use $\sigma(n) = \prod_i \sigma(p_i^{a_i}) = \prod_i (1 + p_i + \cdots + p_i^{a_i})$. Since $24 \mid n+1$, certainly $\gcd(n, 24) = 1$, so every prime factor p_i is odd and not 3, hence p_i lies in one of the eight residue classes mod 24 coprime to 24: $\{1, 5, 7, 11, 13, 17, 19, 23\}$. The natural next move is to try to show that, for a prime power p^a with p in one of these classes, the block $1 + p + \cdots + p^a$ behaves predictably mod 24, and that these blocks multiply together to force the total to be -1 times something... but this immediately fans out. A single prime with $n = p$ works in one line ($\sigma(p) = p + 1 \equiv 0$). But n could be p^2, or pq for two primes in any of the eight classes, or p^2q^3, and so on – there is no bound on how many prime-power blocks appear or which residues they occupy, only the constraint that their *product* is $\equiv -1 \pmod{24}$. Multiplicativity turns a statement about n's residue into a statement about a product of blocks, but the blocks $1 + p + \cdots + p^a$ do not carry p's residue mod 24 in any simple way (a geometric series of length $a+1$ mixes residues), and checking every combination of residues and exponents that could multiply to -1 mod 24 is a case explosion with no natural stopping point. It can be pushed through for small, fixed factorisation shapes, but as a general argument it does not close.

The model chain. Work with n mod 24 directly, never with its factorisation. Since $24 \mid n+1$,

$$n \equiv -1 \equiv 2 \pmod 3, \qquad n \equiv -1 \equiv 7 \pmod 8.$$

First, n is not a perfect square: squares are $\equiv 0$ or $1 \pmod 3$ (as $0^2 \equiv 0$, $1^2 \equiv 1$, $2^2 \equiv 1$), and $n \equiv 2 \pmod 3$ is neither. So every divisor d of n is paired with a *different* divisor n/d; the divisors of n split into complementary pairs $\{d, n/d\}$ with $d \neq n/d$, and

$$\sigma(n) = \sum_{d \mid n} d = \sum_{\text{pairs}} \left(d + \tfrac{n}{d}\right).$$

It remains to show every pair-sum $d + n/d$ is a multiple of 24; by the Chinese Remainder Theorem (3 and 8 coprime, $3 \cdot 8 = 24$) it suffices to check this mod 3 and mod 8 separately.

Mod 3: since $n \not\equiv 0$, d is a unit mod 3, so $d \in \{1, 2\}$, and every unit mod 3 satisfies $d^2 \equiv 1 \pmod 3$ (check both: $1^2 = 1$, $2^2 = 4 \equiv 1$), i.e. $d^{-1} \equiv d \pmod 3$. From $d \cdot (n/d) \equiv n \equiv 2 \pmod 3$, $n/d \equiv 2d^{-1} \equiv 2d \pmod 3$. If $d \equiv 1$: $n/d \equiv 2$, sum $\equiv 1 + 2 = 3 \equiv 0$. If $d \equiv 2$: $n/d \equiv 4 \equiv 1$, sum $\equiv 2 + 1 = 3 \equiv 0$. Either way $d + n/d \equiv 0 \pmod 3$.

Mod 8: d is a unit mod 8, so $d \in \{1, 3, 5, 7\}$, and every unit mod 8 squares to 1 ($1^2 = 1$, $3^2 = 9 \equiv 1$, $5^2 = 25 \equiv 1$, $7^2 = 49 \equiv 1$), so again $d^{-1} \equiv d \pmod 8$. From $d \cdot (n/d) \equiv n \equiv 7 \pmod 8$, $n/d \equiv 7d^{-1} \equiv 7d \pmod 8$. Tabulate: $d \equiv 1 \Rightarrow n/d \equiv 7$ (sum $\equiv 0$); $d \equiv 3 \Rightarrow n/d \equiv 21 \equiv 5$ (sum $\equiv 0$); $d \equiv 5 \Rightarrow n/d \equiv 35 \equiv 3$ (sum $\equiv 0$); $d \equiv 7 \Rightarrow n/d \equiv 49 \equiv 1$ (sum $\equiv 0$). In every case $d + n/d \equiv 0 \pmod 8$.

So $d + n/d \equiv 0 \pmod 3$ and $\pmod 8$ for every pair, hence $d + n/d \equiv 0 \pmod{24}$ by CRT. Summing over all pairs, $\sigma(n)$ is a sum of multiples of 24, so $24 \mid \sigma(n)$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains use only that $24 \mid n + 1$. They split on what to build the proof *out of*. The natural attempt sums over the prime-power blocks $\Pi_i \sigma(p_i^{a_i})$, hoping the mod-24 arithmetic will factor along with n's own factorisation. The model chain instead sums over the complementary divisor pairs $\{d, n/d\}$, controlling $d + n/d$ directly from the single identity $d \cdot (n/d) = n$ taken mod 3 and mod 8 – never asking what d actually is.

> *The hypothesis only pins down n* mod 24 *– nothing about how n factors – so I should build the whole proof out of statements about residues mod* 24*, and never introduce n's actual prime factors. Divisor pairing, $d \leftrightarrow n/d$, is the one algebraic move that turns "n* mod 24 *is known" directly into "$d + n/d$* mod 24 *is known," with no need to know d itself.*

Layer 3 THE WHY

Start from what "multiplicative" means: $\sigma(mn) = \sigma(m)\sigma(n)$ for $\gcd(m, n) = 1$ is a statement about how σ interacts with a *given* factorisation. It is the correct tool when you know the factorisation and want $\sigma(n)$ exactly. Here the hypothesis is not a factorisation, it is a single congruence, $n \equiv -1 \pmod{24}$, and infinitely many different factorisations produce the same residue – so any argument that routes through $n = p_1^{a_1} \cdots p_k^{a_k}$ has to re-derive the global residue fact from arbitrarily many different starting configurations, which is exactly backwards: it is manufacturing structure the hypothesis never supplied.

The pairing $d \leftrightarrow n/d$, by contrast, is forced by the definition of σ itself: $\sigma(n) = \sum_{d \mid n} d$, and divisors of n come with a built-in involution $d \mapsto n/d$ satisfying the *identity* $d \cdot (n/d) = n$ – not an approximation, an equality, true for every divisor regardless of what n factors as. Once that identity is in hand, and the target modulus is $24 = 3 \times 8$ with $\gcd(3, 8) = 1$, the Chinese Remainder Theorem says exactly that checking mod 3 and mod 8 separately is equivalent to checking mod 24 – this is not a computational convenience, it is the content of $\mathbb{Z}/24\mathbb{Z} \cong \mathbb{Z}/3\mathbb{Z} \times \mathbb{Z}/8\mathbb{Z}$. And the reason the residue tables collapse to a one-line rule in each case is Lagrange's theorem applied to the tiny groups $(\mathbb{Z}/3)^\times$ (order 2) and $(\mathbb{Z}/8)^\times$ (order 4, exponent 2): every element of a group of exponent 2 is its own inverse, so $d^{-1} \equiv d$ in both moduli, and $n/d \equiv (n \bmod m) \cdot d$ can be read off in one step rather than solved for. Finally, ruling out $d = n/d$ reduces to n not being a perfect square, which is forced by $n \equiv 2 \pmod 3$ since squares mod 3 can only be 0 or 1 – again a one-line consequence of what a quadratic residue is, not an extra hypothesis.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the hypothesis itself: $n \equiv -1 \pmod{24}$. On seeing a congruence hypothesis about n (not a factorisation) paired with a question about $\sigma(n)$, write down immediately: "$24 = 3 \times 8$; work mod 3 and mod 8 via CRT."

2. That cue suggests controlling the divisors of n by pairing d with n/d, rather than factoring n, because the given information is a residue, and $d \cdot (n/d) = n$ is the one identity that turns a residue on n into information about d and n/d jointly.

3. Hesitation strikes right after reaching for σ's multiplicativity – the reflexive first move for any $\sigma(n)$ problem. It is reinforced away by noticing there is no bound on how many primes n could have or which residues they occupy, so a factorisation-based argument has no natural place to stop, whereas checking $d + n/d$ only ever needs a finite residue table mod 3 and mod 8.

4. The trap is treating "sum of divisors" as automatically a multiplicative-function problem. Multiplicativity is the right tool for *computing* $\sigma(n)$ from a known factorisation; it is the wrong tool for proving a congruence that is supposed to hold for every n in a residue class, regardless of factorisation.

5. To recognise this faster: whenever a divisor-sum congruence is to be proved from a congruence on n itself (not from n's factorisation), suspect divisor pairing $d \leftrightarrow n/d$ first – it is the one identity that converts "$n \equiv r \pmod m$" into information about complementary divisors without ever needing to know what those divisors are.

6. Decision rule: to prove $\sigma(n) \equiv c \pmod m$ from a hypothesis $n \equiv r \pmod m$, pair divisors as $\{d, n/d\}$ and derive $d + n/d \pmod m$ directly from $d \cdot (n/d) \equiv r \pmod m$, splitting m via CRT into coprime prime-power factors where each unit group has small, known exponent; separately confirm n is not a perfect square (using the same congruence) so the pairing never collapses to a fixed point.

▷ **Cue:** a congruence is asked about $\sigma(n)$ (or any function built from the divisors of n), given only a congruence on n itself $\longrightarrow$ **Move:** pair divisors as d and n/d, and derive $d + n/d \pmod m$ directly from the identity $d \cdot (n/d) = n$, splitting m via the Chinese Remainder Theorem into coprime factors

▷ **Cue:** you need $d^{-1} \bmod m$ for a small modulus m $\longrightarrow$ **Move:** check whether every unit mod m squares to 1 (exponent-2 unit group); if so $d^{-1} \equiv d$, and residues can be read off directly instead of solved for

▷ **Cue:** ruling out a self-paired divisor $d = n/d$ $\longrightarrow$ **Move:** show n is not a perfect square by checking n's residue mod a small prime against that prime's possible quadratic residues

1969 B2

1969 B2. Show that a finite group cannot be the union of two of its proper subgroups. Does the statement remain true if "two" is replaced by "three"?

Fields: Group theory: Lagrange's theorem, inclusion–exclusion, elementary abelian 2-groups **Position:** B2 **Answer:** No group is the union of two proper subgroups; with three, the statement fails – the Klein four-group is the union of its three subgroups of order 2.

Layer 1 THE WHAT — surface analysis

The natural attempt. Suppose $G = H \cup K$ with H, K proper subgroups. A first instinct is to chase elements directly: pick $h \in H \setminus K$ and $k \in K \setminus H$ (both exist, else one subgroup contains the other and $G = H \cup K$ collapses to the bigger one, which is proper – contradiction already). Where does hk live? If $hk \in H$, then since $h \in H$, $h^{-1}(hk) = k \in H$, contradicting $k \notin H$; so $hk \in K$. Symmetrically $kh \in H$. This is correct and can be pushed further (consider hkh, khk, ...) to squeeze out a contradiction, but it is a slow route: each new product has to be relocated into H or K by the same case-check, and it is not obvious in advance how many products must be chased before the contradiction surfaces, nor does this approach extend cleanly to the three-subgroup question, where three-way products give an even larger branching tree with no natural place to stop.

Counting is cheaper. $H \cap K$ contains at least the identity, so

$$|H \cup K| = |H| + |K| - |H \cap K| \leq |H| + |K| - 1.$$

By Lagrange's theorem, $|H|$ and $|K|$ divide $|G|$, and since H, K are proper, each has index at least 2, so $|H|, |K| \leq |G|/2$. Hence $|H \cup K| \leq |G| - 1 < |G|$, and $H \cup K \neq G$: this already finishes the two-subgroup case in three lines, no element-chasing needed. Now push the same idea to three subgroups H, K, L: by inclusion–exclusion,

$$|H \cup K \cup L| \leq |H| + |K| + |L| - 2,$$

using only that each of the three pairwise intersections contains the identity (this crude bound ignores the triple intersection term, which only helps). With $|H|, |K|, |L| \leq |G|/2$,

$$|H \cup K \cup L| \leq \tfrac{3}{2}|G| - 2.$$

For this to force $|H \cup K \cup L| < |G|$, one would need $\frac{3}{2}|G| - 2 < |G|$, i.e. $|G| < 4$ – but $|G| \geq 4$ is exactly the interesting range. The natural attempt, having worked perfectly for two subgroups, produces no contradiction at all for three once $|G| \geq 4$. It is tempting to hunt for a sharper inequality (bound the pairwise intersections more carefully, or invoke more group structure) to try to rescue the impossibility claim – but that hunt is chasing the wrong conclusion.

The model chain. Part 1 (two subgroups). As above: H, K proper subgroups of a finite group G have $|H|, |K| \leq |G|/2$ by Lagrange's theorem (a proper subgroup has index ≥ 2). Since $e \in H \cap K$,

$$|H \cup K| = |H| + |K| - |H \cap K| \leq \tfrac{|G|}{2} + \tfrac{|G|}{2} - 1 = |G| - 1 < |G|,$$

so $H \cup K$ omits at least one element of G; it cannot equal G.

Part 2 (three subgroups – read the failed bound, don't fight it). The natural attempt's bound $|H \cup K \cup L| \leq \frac{3}{2}|G| - 2$ stops being less than $|G|$ once $|G| \geq 4$. Rather than search for a better inequality, ask what it would take for the bound itself to be *tight*, i.e. for equality to hold: every subgroup used must be as large as Lagrange allows (index exactly 2, so $|H| = |K| = |L| = |G|/2$), and the pairwise intersections must be as small as possible (exactly $\{e\}$, no larger overlap). Under these two conditions,

$$|H \cup K \cup L| = 3\Big(\frac{|G|}{2} - 1\Big) + 1 = \frac{3|G|}{2} - 2,$$

and setting this equal to $|G|$ gives $|G| = 4$.

So look for G of order 4 with three subgroups of order 2 meeting pairwise only in $\{e\}$. If every one of G's three non-identity elements lies in one of these three order-2 subgroups (which is what "pairwise intersecting only in e, covering G" forces), then every non-identity element of G, being the sole non-identity element of a subgroup of order 2, has order exactly 2. A finite group in which every non-identity element has order 2 is automatically abelian (for any x, y: $(xy)^2 = e$ forces $xy = (xy)^{-1} = y^{-1}x^{-1} = yx$) and, having order 4 with no element of order 4, cannot be cyclic. The unique such group is the Klein four-group $V_4 = \{e, x, y, xy\}$ with $x^2 = y^2 = e$, $xy = yx$. Its three subgroups of order 2,

$$H = \{e, x\}, \qquad K = \{e, y\}, \qquad L = \{e, xy\},$$

pairwise intersect only in $\{e\}$, and together they cover all four elements of V_4: $H \cup K \cup L = V_4$. So the statement is false once "two" becomes "three": V_4 is a finite group that *is* the union of three of its proper subgroups.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both attempts compute the same inclusion–exclusion bound, $|H \cup K \cup L| \leq \frac{3}{2}|G| - 2$ (from Lagrange plus "every pairwise intersection contains e"), applied to three

subgroups. They split at the very next move. The natural attempt treats the fact that this bound no longer beats $|G|$ as an obstacle – something to patch by finding a sharper inequality that still forces a contradiction. The model chain treats it as data: it asks what configuration makes the bound an *equality*, and builds exactly that configuration, arriving at $|G| = 4$ and the Klein four-group.

> *My counting bound doesn't rule out three subgroups the way it ruled out two – instead of hunting for a sharper inequality that the numbers no longer support, I should ask what equality in my* own *bound would actually require, and go build that case.*

Layer 3 THE WHY

Part 1's proof is not a trick; it is Lagrange's theorem and the subgroup axioms read literally. "H is a subgroup" means G is partitioned into cosets of H, each of size $|H|$, so $|H|$ divides $|G|$ – that is what Lagrange's theorem *is*. "H is proper" means there is more than one coset, so the number of cosets is at least 2, forcing $|H| \leq |G|/2$. "Every subgroup contains the identity" is a subgroup axiom, giving $e \in H \cap K$ for free. Inclusion–exclusion, $|H \cup K| = |H| + |K| - |H \cap K|$, is just careful counting, not group theory at all. Every ingredient of the two-subgroup proof is forced by definitions; nothing is chosen.

The same ingredients, run for three subgroups, give a bound that is no longer strictly below $|G|$ once $|G| > 4$ – and this is the moment the axioms stop settling the question by themselves. A bound that fails to give a contradiction is not evidence of a stronger bound waiting to be found; it is evidence that the question's answer may genuinely change, and the bound's own equality case tells you exactly where to look. Equality in $|H| \leq |G|/2$ demands index exactly 2 for every subgroup used (the largest a proper subgroup can be); equality in "pairwise intersections are size 1" demands no accidental overlap beyond the identity, which the definitions do not forbid. Combining these two *equality* conditions – not the inequality, its boundary case – pins $|G| = 4$ exactly, because $3(|G|/2 - 1) + 1 = |G|$ has the unique positive solution $|G| = 4$.

From there, the group structure is forced, not guessed: three subgroups of order 2 covering $G \setminus \{e\}$ with no overlap means every non-identity element of G belongs to exactly one order-2 subgroup, hence has order 2. An element of order 2 in every non-identity slot forces $(xy)^2 = e$ for all x, y (expand and cancel), which is exactly the statement that G is abelian – again forced, not assumed. A group of order 4 that is not cyclic and has every non-identity element of order 2 has no other possible multiplication table than V_4's: this is why the counterexample is not a lucky find but the unique object satisfying the equality case of the very inequality that proved impossibility for two subgroups. (One more axiom-level fact worth naming: a *cyclic* group can never be covered by proper subgroups at all, for any number of them, because its generator lies in none of them by definition of "proper". This is why the search for a counterexample must leave cyclic groups behind entirely, and it is consistent with V_4, which is not cyclic.)

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "union of proper subgroups" together with "finite group." On seeing it, write down Lagrange's theorem immediately: any proper subgroup H satisfies $|H| \leq |G|/2$.

2. That cue suggests inclusion–exclusion counting: bound $|H \cup K|$ (or $|H \cup K \cup L|$) using the subgroup order bounds together with the fact that every subgroup contains the identity, forcing overlap.

3. For the three-subgroup question, hesitation occurs exactly where the counting bound, extended from two terms to three, stops being strictly below $|G|$. It is reinforced away by treating that failure as information – where does the bound become an equality? – rather than as an obstacle to be pushed through with a sharper estimate.

4. The trap is assuming the two-subgroup proof method must settle the three-subgroup question the same way (impossibility). A counting method that proves impossibility at one value

of a parameter can become neutral – or even point straight at a construction – at the next value, and the arithmetic itself has to be read to find out which.

5. To recognise this faster: whenever an inequality proof stops being strict as some parameter increases, immediately check whether equality is attainable; if it is, solve the equality case explicitly to build the extremal example, rather than searching for a stronger inequality that no longer exists.

6. Decision rule: for a "union of k proper subgroups" question, bound $|H_1 \cup \cdots \cup H_k|$ by $\sum |H_i| -$ (overlaps forced by the shared identity) via Lagrange and inclusion–exclusion; if this beats $|G|$, impossibility is proved; if it does not, solve the equality case of that same bound (largest allowed subgroup sizes, smallest allowed overlaps) to construct the counterexample instead of chasing a tighter inequality.

▷ **Cue:** a finite group is claimed to be (or not be) the union of k proper subgroups $\longrightarrow$ **Move:** bound each proper subgroup's order by $|G|/2$ via Lagrange (index ≥ 2), then apply inclusion–exclusion using that every subgroup shares the identity, and compare the resulting bound to $|G|$

▷ **Cue:** your own counting bound stops being strictly below the target as a parameter (here, the number of subgroups) increases $\longrightarrow$ **Move:** solve the equality case of the bound to construct the extremal example, rather than hunting for a tighter inequality that no longer exists

▷ **Cue:** every non-identity element of a group is forced to lie in some subgroup of order 2 $\longrightarrow$ **Move:** conclude every non-identity element has order 2, hence the group is abelian and elementary; for order 4 and non-cyclic, this pins the group down as the Klein four-group

1969 B3

1969 B3. The terms of a sequence $T_1, T_2, T_3, \ldots$ satisfy

$$T_n T_{n+1} = n \qquad (n = 1, 2, 3, \ldots),$$

and $\lim_{n\to\infty} \dfrac{T_n}{T_{n+1}} = 1$. Show that $\pi T_1^2 = 2$.

Caveat. This is the most heavily damaged statement of the three problems in this batch. The MAA 1965–1984 OCR reduces it to a fragment with the defining relation itself missing entirely: "The terms of a sequence T_n satisfy …Show that $\pi T_1^2 = 2$… $(n = 1, 2, 3, \ldots)$ and $\lim T_n / T_{n+1} = 1$" – the two-column layout has been read straight across, scrambling the order of the pieces and dropping the displayed recurrence formula altogether (a common failure mode for this scan whenever the relation is a displayed equation rather than inline text).

The statement above was reconstructed from two independent sources, cross-checked against each other and against the book's own solution:

(1) The book's solution to B-3 (transcribed faithfully below, itself part of the OCR but readable) begins "The first relation implies that $T_n = \dfrac{(n-1)(n-3)\cdots 3\cdot 1}{(n-2)(n-4)\cdots 2}\, T_1$ if n is even, $T_n = \dfrac{(n-1)(n-3)\cdots 2}{(n-2)(n-4)\cdots 1}\, T_1$ if n is odd," and later computes T_n/T_{n+1} as T_1^2 times a partial Wallis product, concluding $\pi T_1^2 = 2$. A short induction (carried out independently in Layer 1 below) shows that these closed forms are generated *exactly* by the pair of relations $T_1 T_2 = 1$, $T_2 T_3 = 2$, $T_3 T_4 = 3, \ldots$, i.e. $T_n T_{n+1} = n$ for every $n \geq 1$ – and by no simpler relation: attempting the alternative reading $T_1 = T_2$ together with the two-step

recurrence $nT_{n+2} = (n+1)T_n$ produces a sequence with *no* free scaling parameter (the ratio T_n/T_{n+1} becomes independent of T_1 and converges to $2/\pi \neq 1$, which would make the stated hypothesis false), so that reading was ruled out by direct calculation, not merely disfavoured.

(2) An independently worded but substance-preserving paraphrase of this exact problem (Putnam 1969 B3, in a public archive of reworded historical Putnam problems maintained to avoid reproducing MAA copyrighted text verbatim) states: "The sequence $a_1, a_2, a_3, \ldots$ satisfies $a_1a_2 = 1$, $a_2a_3 = 2$, $a_3a_4 = 3$, $a_4a_5 = 4, \ldots$. Also $\lim_{n\to\infty} a_n/a_{n+1} = 1$. Prove that $a_1 = \sqrt{2/\pi}$." This confirms the relation $T_nT_{n+1} = n$ exactly, and $a_1 = \sqrt{2/\pi}$ is precisely equivalent to $\pi T_1^2 = 2$.

Both independent checks agree, and the reconstructed relation reproduces the book's own displayed formulas and final line without alteration (verified in full in Layer 1). No mathematical content beyond what these two sources jointly force has been invented.

Fields: Real analysis: recurrence sequences, double factorials, the Wallis product for π **Position:** B3 (lower middle) **Answer:** $\pi T_1^2 = 2$, i.e. $T_1^2 = 2/\pi$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The relation $T_nT_{n+1} = n$ makes T_n computable in closed form by unwinding two interleaved chains – odd-indexed terms in terms of T_1, even-indexed terms in terms of $T_2 = 1/T_1$ – and each chain is a ratio of double factorials, the exact object that shows up in the Gamma function via $\Gamma\left(k+\frac{1}{2}\right) = \dfrac{(2k-1)!!}{2^k}\sqrt{\pi}$ and $(2k)!! = 2^k k!$. A natural, entirely correct route is therefore to write $T_{2k+1} = T_1 \cdot \dfrac{(2k)!!}{(2k-1)!!} = T_1 \cdot \dfrac{2^{2k}(k!)^2}{(2k)!}$, recognise this as (a multiple of) a central binomial coefficient, and bring in Stirling's approximation $n! \sim \sqrt{2\pi n}\,(n/e)^n$ to extract its growth rate: $\binom{2k}{k} \sim \dfrac{4^k}{\sqrt{\pi k}}$, hence $\dfrac{2^{2k}(k!)^2}{(2k)!} = \dfrac{4^k}{\binom{2k}{k}} \sim \sqrt{\pi k}$. Chasing the constants through correctly eventually reproduces $T_1 = \sqrt{2/\pi}$ – but only after introducing the Gamma function, Stirling's formula for factorials (itself a nontrivial asymptotic theorem needing care with its own error terms), and a fair amount of bookkeeping to keep the 2's and π's straight through several substitutions. It is correct, but it imports far more machinery than the problem's own recurrence actually needs, and a single slip in the Stirling constant is easy to make and hard to catch.

A second, subtler trap: from $T_nT_{n+1} = n$ one can derive, cheaply, that $\dfrac{T_n}{T_{n+1}} = \dfrac{T_n^2}{T_nT_{n+1}} = \dfrac{T_n^2}{n}$, so the hypothesis $\lim T_n/T_{n+1} = 1$ says exactly $T_n^2/n \to 1$, i.e. $T_n \sim \sqrt{n}$. It is tempting to treat this as already the answer – "the sequence grows like $\sqrt{n}$, so of course T_1 is whatever makes that consistent" – and to hunt for T_1 by further purely algebraic manipulation of $T_nT_{n+1} = n$ and $T_n \sim \sqrt{n}$ alone. This hunt cannot succeed: every relation obtainable by finite algebraic manipulation of a recurrence with rational coefficients, applied to n, is itself rational in n and in the initial value; no finite chain of additions, multiplications and divisions can manufacture the transcendental number π. The asymptotic $T_n \sim \sqrt{n}$ is a true and useful fact, but by itself it is silent about T_1's precise value, because it holds with the *same* leading constant 1 no matter what T_1 actually is – the recurrence, run forward, automatically re-normalises any starting value into agreement with $T_n^2/n \to 1$; that is not evidence about T_1, it is a tautology of the recurrence. Something outside the algebra has to be brought in.

The model chain. *Step 1 (unwind the recurrence).* From $T_nT_{n+1} = n$ and $T_{n+1}T_{n+2} = n+1$, divide: $T_{n+2}/T_n = (n+1)/n$ for all $n \geq 1$. This is a single two-step recurrence; solve the odd-indexed

chain by induction. Claim: for $k \geq 0$,

$$T_{2k+1} = T_1 \cdot \frac{(2k)!!}{(2k-1)!!}.$$

Base case $k = 0$: both sides are T_1 (empty products equal 1). Inductive step: if $T_{2k+1} = T_1(2k)!!/(2k-1)!!$, then using $T_{2k+3} = \dfrac{2k+2}{2k+1} T_{2k+1}$ and $(2k+2)!! = (2k+2)(2k)!!$, $(2k+1)!! = (2k+1)(2k-1)!!$,

$$T_{2k+3} = \frac{2k+2}{2k+1} \cdot T_1 \frac{(2k)!!}{(2k-1)!!} = T_1 \cdot \frac{(2k+2)(2k)!!}{(2k+1)(2k-1)!!} = T_1 \cdot \frac{(2k+2)!!}{(2k+1)!!},$$

completing the induction.

Step 2 (extract two independent asymptotics for the same quantity). On one hand, the given hypothesis, via $T_n/T_{n+1} = T_n^2/n$ (immediate from $T_n T_{n+1} = n$), says

$$\lim_{n\to\infty} \frac{T_n^2}{n} = 1,$$

and this holds along every subsequence, in particular $\lim_{k\to\infty} \dfrac{T_{2k+1}^2}{2k+1} = 1$.

On the other hand, Step 1 gives $T_{2k+1}^2 = T_1^2 \left(\dfrac{(2k)!!}{(2k-1)!!} \right)^2$ exactly, so

$$\frac{T_{2k+1}^2}{2k+1} = T_1^2 \cdot \frac{\left[(2k)!!\right]^2}{\left[(2k-1)!!\right]^2 (2k+1)}.$$

The fraction on the right is exactly the k-th partial product of Wallis' formula for $\pi/2$: writing $\pi/2 = \prod_{j=1}^{\infty} \dfrac{2j}{2j-1} \cdot \dfrac{2j}{2j+1}$ (itself proved by comparing the integrals $I_n = \int_0^{\pi/2} \sin^n x\, dx$ via the inequality $\sin^{n+1} x < \sin^n x < \sin^{n-1} x$ on $(0, \pi/2)$, which forces $I_{2n+1} < I_{2n} < I_{2n-1}$ and squeezes the ratio to $\pi/2$), the partial product up to $j = k$ is precisely $\left[(2k)!!\right]^2 / \left\{ \left[(2k-1)!!\right]^2 (2k+1) \right\}$, and this converges to $\pi/2$ as $k \to \infty$.

Step 3 (equate the two limits). Combining,

$$1 = \lim_{k\to\infty} \frac{T_{2k+1}^2}{2k+1} = T_1^2 \cdot \lim_{k\to\infty} \frac{\left[(2k)!!\right]^2}{\left[(2k-1)!!\right]^2 (2k+1)} = T_1^2 \cdot \frac{\pi}{2}.$$

Hence $\pi T_1^2 = 2$, as required. (The even-indexed chain, run through the symmetric computation, gives the same equation; the problem's own hypothesis – a single limit over the whole sequence – already guarantees the two subsequential limits agree, so checking one suffices.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same cheap identity, $T_n/T_{n+1} = T_n^2/n$, and both correctly conclude $T_n \sim \sqrt{n}$. They split at the very next question: *where does π come from?* The natural attempt keeps working algebraically – trying to extract T_1 by manipulating the recurrence further, or by reaching for the heaviest available machinery (Gamma functions, Stirling's formula) that happens to also produce the right asymptotic. The model chain stops manipulating and instead asks which classical, self-contained theorem the closed-form quantity $\left[(2k)!!\right]^2 / \left\{ \left[(2k-1)!!\right]^2 (2k+1) \right\}$ already matches – recognises it as Wallis' own partial product for $\pi/2$ – and imports that theorem wholesale rather than re-deriving the same convergence from scratch through the Gamma function.

> *Every identity I can derive from $T_nT_{n+1} = n$ alone is built entirely out of the integer n and rational operations on it and on T_1 – it will hold no matter what real number I initially pick for T_1, so by itself it can never force a particular value, let alone one involving π, because π is transcendental and no finite chain of additions, multiplications and divisions applied to integers can produce it. I need an outside, analytic fact about how $(2k)!!/(2k-1)!!$ actually grows, and the fact I need already has a name: Wallis' product.*

Layer 3 THE WHY

Every step before Step 2 is algebra: the recurrence $T_nT_{n+1} = n$ is a rational relation between integers and the T_n's, and unwinding it by induction (Step 1) produces T_{2k+1} as T_1 times a ratio of two double factorials – itself nothing more than repeated multiplication. No matter how cleverly this algebra is rearranged, it cannot, on its own, produce π: π is transcendental, meaning it is not the root of any polynomial with rational coefficients, and a fortiori it cannot arise from finitely many field operations applied to integers. This is not a heuristic; it is the reason the natural attempt's second trap (trying to pin down T_1 by pure algebra on $T_nT_{n+1} = n$ and $T_n \sim \sqrt{n}$) is structurally doomed, not merely unlucky. Whatever route is taken, an analytic limit theorem — an actual argument about convergence, not finitely many arithmetic operations — has to be imported from outside the recurrence.

The question is which limit theorem, and the answer is forced by the *shape* of what Step 1 already produced. Double factorials are, by definition, the closed form of the Wallis sine-power integrals $I_n = \int_0^{\pi/2} \sin^n x\,dx$: integration by parts gives $I_n = \frac{n-1}{n} I_{n-2}$, the same two-step recurrence structure as $T_{n+2} = \frac{n+1}{n} T_n$, and unwinding it produces exactly the double-factorial ratios $(2k-1)!!/(2k)!! \cdot \frac{\pi}{2}$ and $(2k)!!/(2k+1)!!$. Wallis' theorem is the statement that the elementary two-sided squeeze $I_{2n+1} < I_{2n} < I_{2n-1}$ – itself nothing more than $\sin^{2n+1} x < \sin^{2n} x < \sin^{2n-1} x$ on $(0, \pi/2)$, a one-line inequality – forces the ratio $I_{2n}/I_{2n+1} \to 1$, and translating that ratio back through the closed form yields $\left[(2k)!!\right]^2 / \left\{\left[(2k-1)!!\right]^2 (2k+1)\right\} \to \pi/2$. So the presence of π in the final answer is not an accident of this particular problem: *any* recurrence whose closed-form solution is a ratio of double factorials, subjected to a ratio-tends-to-one hypothesis on consecutive terms, is structurally required to reduce to Wallis' theorem, because that theorem is exactly the classical statement of how such ratios converge, and there is no cheaper route to π available once the problem has been reduced to this shape. Reaching for Stirling's approximation to the Gamma function instead is not wrong – it proves the same asymptotic – but it is strictly more expensive: it needs the Gamma function's definition, its half-integer values, and Stirling's own (harder) asymptotic proof, to establish a fact that the elementary sine-integral squeeze gives directly and with exact control of the limiting constant at every finite stage.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the closed-form solution itself: once unwinding the recurrence produces a ratio of double factorials, $(2k)!!/(2k-1)!!$ or its reciprocal, write down immediately: "Wallis' product; this converges like $\sqrt{\pi k}$."
2. That cue suggests: stop manipulating the recurrence algebraically and instead import Wallis' theorem as an external, already-proved analytic fact, matching its known limiting constant against whatever the problem's own hypothesis independently forces on the same quantity.
3. Hesitation strikes right after deriving $T_n/T_{n+1} = T_n^2/n$ and noticing it says nothing about π. It is reinforced away by remembering that π is transcendental, so it cannot be produced by iterating a rational recurrence, however cleverly rearranged – an analytic input is not optional, so further algebraic hunting is a dead end and should be abandoned immediately, not eventually.
4. The trap is believing that, because both the hypothesis ($\lim T_n/T_{n+1} = 1$) and the recurrence look like "just algebra and limits," some clever telescoping or symmetry argument will

locate T_1 without ever explicitly invoking π – a plausible-looking dead end, since π can only enter from an outside theorem, never from inside the recurrence's own arithmetic.

5. To recognise this faster: whenever a recurrence's closed form is a ratio of double factorials (equivalently, of central-binomial-coefficient type) and a limiting condition is imposed on consecutive terms, expect Wallis' product (or, more expensively, Stirling's approximation to the Gamma function) to be the required external theorem – and reach for the cheaper of the two, proved by nothing more than the elementary sine-power squeeze, rather than the heavier one built on Stirling's asymptotic.

6. Decision rule: to pin down an unknown scaling constant in a sequence whose closed form is a ratio of double factorials, (i) unwind the recurrence by induction to get the closed form exactly, (ii) separately extract what the given asymptotic hypothesis forces on the same quantity (here, $T_n^2/n \to 1$), and (iii) equate the two independently-obtained limits using Wallis' product, choosing it over Stirling whenever the closed form already matches Wallis' own shape.

▷ **Cue:** a two-step recurrence, $n\,T_{n+2} = (n+1)\,T_n$, unwinds to a ratio of consecutive odd/even integer products (a quotient of double factorials) $\longrightarrow$ **Move:** recognise this as the Wallis-product shape and expect its limit to involve π; stop manipulating the recurrence algebraically, since no rational chain of operations can ever produce a transcendental constant

▷ **Cue:** a hypothesis gives $\lim T_n/T_{n+1} = 1$ for a sequence satisfying $T_nT_{n+1} = n$ $\longrightarrow$ **Move:** rewrite as $T_n^2/n \to 1$ (i.e. $T_n \sim \sqrt{n}$) and treat this as an independent asymptotic, to be matched against the closed form's own asymptotic rather than mistaken for the final answer

▷ **Cue:** you need the precise growth rate of $(2k)!!/(2k-1)!!$ $\longrightarrow$ **Move:** cite Wallis' product, $\big[(2k)!!\big]^2/\big\{[(2k-1)!!]^2(2k+1)\big\} \to \pi/2$, proved via the elementary squeeze $I_{2n+1} < I_{2n} < I_{2n-1}$ on the sine-power integrals $\int_0^{\pi/2} \sin^n x\,dx$, rather than via Stirling's approximation to the Gamma function

1969 B4

> **1969 B4.** Show that any curve of unit length can be covered by a closed rectangle of area $1/4$.

Fields: Plane geometry, covering problems, arc length, inequalities (AM–GM, triangle inequality) **Position:** B4 (upper middle) **Answer:** Rotate the curve so that the chord joining its endpoints is horizontal; the axis-parallel bounding box then has width a and height b with $a^2 + 4b^2 \le 1$, so $ab \le \frac{1}{4}$, and the box can be enlarged to a closed rectangle of area exactly $\frac{1}{4}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. A curve of length 1 cannot wander far, so draw its bounding box and bound the box's area. Put the curve anywhere in the plane, with any orientation, and let the smallest axis-parallel rectangle containing it have width a and height b. The curve touches the left side and the right side, so it contains two points whose x-coordinates differ by a; the arc between them has length at least the straight-line distance, so $a \le 1$. Likewise $b \le 1$. That gives area ≤ 1, four times too much.

Sharpen it. Let L and R be points of the curve on the left and right sides, T and B points on the top and bottom. Take them in the order they occur along the curve and join them by the polygonal path through them; its length is at most 1. Its horizontal movements add up to at least a and its vertical movements to at least b, so by the triangle inequality its length is at least $\sqrt{a^2+b^2}$. Hence $a^2 + b^2 \le 1$, and

$$ab \le \tfrac{1}{2}(a^2+b^2) \le \tfrac{1}{2}.$$

Area $\frac{1}{2}$: better, and still twice the target. A straight segment of length 1 placed along the diagonal of a square of side $1/\sqrt{2}$ shows that nothing more can be squeezed out of this set-up: the bounding box in that position really has area $\frac{1}{2}$. The wall is not in the estimates. The orientation was chosen badly, and no estimate can repair that.

The model chain. Step 1 (orientation). Let the curve have endpoints P_0 and P_5 (if the curve is closed, $P_0 = P_5$). Rotate and translate so that P_0 and P_5 both lie on the x-axis. Let the smallest closed axis-parallel rectangle containing the curve be $[x_-, x_+] \times [y_-, y_+]$, with width $a = x_+ - x_-$ and height $b = y_+ - y_-$. By compactness of the curve, each side of the rectangle contains a point of the curve; call these points P_1, P_2, P_3, P_4, relabelled in the order they occur along the curve (two may coincide if the curve passes through a corner).

Step 2 (a polygon no longer than the curve). The broken line $P_0P_1P_2P_3P_4P_5$ has vertices taken in order along the curve. Each segment P_iP_{i+1} is no longer than the arc of the curve from P_i to P_{i+1}, so the broken line has total length $\ell \le 1$.

Step 3 (horizontal budget). Write $P_i = (x_i, y_i)$. Among the x_i are the values x_- and x_+, so by the triangle inequality on the real line,

$$H := \sum_{i=0}^{4} |x_{i+1} - x_i| \;\ge\; x_+ - x_- = a.$$

Step 4 (vertical budget — the doubling). Both endpoints have the same height $y_0 = y_5 = 0$, which lies in $[y_-, y_+]$. The path visits height y_+ at some vertex and height y_- at another; call the earlier of these two heights u and the later w, so $\{u, w\} = \{y_+, y_-\}$. Splitting the path at those two vertices and applying the triangle inequality to each piece,

$$V := \sum_{i=0}^{4} |y_{i+1} - y_i| \;\ge\; |u - 0| + |w - u| + |0 - w| = b + \big(|y_+| + |y_-|\big) = b + b = 2b,$$

because $y_- \le 0 \le y_+$ gives $|y_+| + |y_-| = y_+ - y_- = b$.

Step 5 (combine). Replace each segment vector $(x_{i+1} - x_i,\, y_{i+1} - y_i)$ by $(|x_{i+1} - x_i|,\, |y_{i+1} - y_i|)$, which has the same length. The triangle inequality for vectors gives

$$1 \;\ge\; \ell = \sum_i \big|(|\Delta x_i|, |\Delta y_i|)\big| \;\ge\; \Big|\big(H,\, V\big)\Big| = \sqrt{H^2 + V^2} \;\ge\; \sqrt{a^2 + 4b^2}.$$

Step 6 (AM–GM). $a^2 + 4b^2 = a^2 + (2b)^2 \ge 2 \cdot a \cdot 2b = 4ab$, so $4ab \le 1$ and $ab \le \frac{1}{4}$. The rectangle has area at most $\frac{1}{4}$; enlarge it, keeping it around the curve, until the area is exactly $\frac{1}{4}$ (if $a = 0$, first widen it to a positive width, then lengthen the vertical side). This closed rectangle covers the curve.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains draw the bounding box and inscribe a polygon through the touching points. They split at the one decision taken before any estimate: where the curve sits. The natural attempt leaves the orientation arbitrary and so gets one traversal in each direction, $a^2 + b^2 \le 1$. The model chain fixes the orientation by the only canonical line the curve owns — the chord through its endpoints — and in that orientation the vertical direction is traversed twice, $a^2 + (2b)^2 \le 1$.

> *I'm getting $\frac{1}{2}$ and I need $\frac{1}{4}$, so somewhere I'm counting a length once that I ought to count twice. A path that starts and ends at the same height and has to touch both the top and the bottom must go up and come back; if I put both endpoints on the x-axis, the vertical extent gets paid for twice.*

> *The rectangle is mine to place. Before I estimate anything, I should ask which orientation makes the estimate strongest, and the endpoints are the only thing about an arbitrary curve that gives me a preferred direction.*

Layer 3 THE WHY

Start with what "length 1" means. Arc length is *defined* as the supremum of the lengths of inscribed polygons: points taken in order along the curve, joined by straight segments. So the only fact the hypothesis hands you is that every inscribed polygon has length at most 1. Every step of any proof has to go through an inscribed polygon. This is why both chains build one, and it is not a trick: the definition leaves you no other route.

Next, what can an inscribed polygon tell you about a box? A box has two numbers, a and b, and the target is a bound on the product ab. A product bound from a sum-of-squares bound is AM–GM, and AM–GM is exact: $a^2 + b^2 \le 1$ gives exactly $ab \le \frac{1}{2}$, while $a^2 + (\lambda b)^2 \le 1$ gives $ab \le 1/(2\lambda)$. Now read the target backwards. The answer $\frac{1}{4}$ is the problem telling you that $\lambda = 2$ — that one of the two directions must be covered twice. Working backwards from the stated constant to the inequality that would produce it is the move here, and it is available because AM–GM converts coefficients to constants with no slack.

Why would a direction be covered twice? A path covers a direction twice when it has to go out and come back, and that happens when it starts and ends at the same coordinate in that direction. An arbitrary curve has no such direction, until you make one: any two points lie on a common horizontal line once you rotate, and the curve's two endpoints are the natural choice, being the only two points singled out by the curve's own description. (For a closed curve the endpoints coincide and every direction is covered twice, which is stronger still.)

The segment-on-a-diagonal example in Layer 1 shows that the doubling has to come from orientation and not from better estimation: in the diagonal position the box really does have area $\frac{1}{2}$. The problem says "can be covered by a rectangle", not "by its bounding box in any position", so you are allowed to choose, and you must. Whenever a covering problem lets you choose the cover, the first question is which placement makes the object's own structure (here, the chord) line up with the cover's structure (here, the sides).

One more axiom-level point, about Step 5. Replacing each $(\Delta x_i, \Delta y_i)$ by $(|\Delta x_i|, |\Delta y_i|)$ does not change its length, and after the replacement all the vectors point into the first quadrant, so adding them loses no horizontal or vertical progress. That is why the total length is at least the length of the single vector (H, V). It is the triangle inequality, applied after the signs have been removed so that nothing cancels.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the number $\frac{1}{4}$ next to "unit length". On seeing it, write down: "the crude bound will give $\frac{1}{2}$ (from $a^2 + b^2 \le 1$); I need a factor 2 on one side — $a^2 + 4b^2 \le 1$."

2. Inscribe a polygon through the extreme points of the bounding box (forced by the definition of arc length), bound its horizontal and vertical variation separately, and combine with the triangle inequality and AM–GM.

3. Hesitation comes after the $\frac{1}{2}$ bound, when it seems the estimate itself must be improved. It goes away once you see that the diagonal segment really does need area $\frac{1}{2}$ in that position, so the fix has to be in the placement. After that, the chord is the only candidate.

4. The trap is treating "the rectangle" as the bounding box of the curve in whatever position it was drawn. The problem lets you choose the rectangle; a proof that never uses that freedom cannot reach $\frac{1}{4}$.

5. When a stated constant is off from your bound by an exact factor like 2, reverse-engineer which term needs doubling, and ask what in the configuration could be traversed or counted twice.

6. Decision rule: in a covering or bounding problem where you may place the cover, align the cover with the object's own canonical feature (endpoints, chord, axis of symmetry) before estimating; a path forced to return to its starting coordinate pays for that direction twice.

▷ **Cue:** a curve of given length must fit in a box or region of stated area $\longrightarrow$ **Move:** inscribe the polygon through the extreme points (arc length is the sup of inscribed polygons) and bound its horizontal and vertical variation separately

▷ **Cue:** your bound is exactly twice (or k times) the target constant $\longrightarrow$ **Move:** work backwards through AM–GM to see which coefficient must change, then find the geometric reason that direction is counted twice

▷ **Cue:** you are free to choose the position of the cover ("can be covered by a rectangle") $\longrightarrow$ **Move:** fix the orientation by the object's canonical line — here, put both endpoints on the x-axis so the path must go out and return vertically

1969 B5

1969 B5. Let $a_1 < a_2 < a_3 < \cdots$ be an increasing sequence of positive integers. Let the series

$$\sum_{n=1}^{\infty} 1/a_n$$

be convergent. For any number x, let $k(x)$ be the number of the a_n's which do not exceed x. Show that $\lim_{x\to\infty} k(x)/x = 0$.

Fields: Real analysis: infinite series (convergence), limits, counting functions **Position:** B5 (hard half of the paper, though short once seen) **Answer:** For any $\varepsilon > 0$ choose N with $\sum_{n>N} 1/a_n < \varepsilon$; each $a_n \le x$ with $n > N$ contributes at least $1/x$ to that tail, so $(k(x) - N)/x < \varepsilon$ and $\limsup k(x)/x \le \varepsilon$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The statement is a limit, so try contradiction. Suppose $k(x)/x \not\to 0$. Then there are $\varepsilon > 0$ and $x_1 < x_2 < \cdots \to \infty$ with $k(x_j) \ge \varepsilon x_j$. At each x_j there are at least εx_j terms $a_n \le x_j$, each with $1/a_n \ge 1/x_j$, so

$$\sum_{a_n \le x_j} \frac{1}{a_n} \;\ge\; \frac{k(x_j)}{x_j} \;\ge\; \varepsilon.$$

That is not a contradiction: $\sum 1/a_n$ is a finite number and can certainly be at least ε. The natural next idea is to add up many such blocks, one per x_j, to force divergence. But the blocks overlap: the terms counted at x_2 include those counted at x_1, so adding the lower bounds counts the same terms repeatedly. You can try to repair this by thinning the sequence so that x_{j+1} is huge compared with x_j, so that most of the terms below x_{j+1} lie above x_j. That can be done, but it becomes an argument with nested $\varepsilon/2$'s, and at this point in the exam the solver is manufacturing divergence by hand.

A second natural attempt uses a remembered slogan: "if $\sum b_n$ converges then $b_n = o(1/n)$", so $1/a_n = o(1/n)$, i.e. $n/a_n \to 0$, and then (since $a_{k(x)} \le x$) $k(x)/x \le k(x)/a_{k(x)} \to 0$. The final step is right. The slogan is false in general: take $b_n = 1/n$ when n is a perfect square and $b_n = 1/n^2$ otherwise. Then $\sum b_n$ converges, but $nb_n = 1$ infinitely often. A solver who writes the slogan down has written a wrong proof and may not realise it.

The model chain. *Step 0 (finiteness).* Since the a_n are strictly increasing positive integers, $a_n \ge n$, so $a_n \to \infty$ and only finitely many a_n are $\le x$; $k(x)$ is a finite, non-decreasing function of x, and $k(x) = 0$ for $x < a_1$. We may take $x > 0$.

Step 1 (use convergence as "tails are small"). Fix $\varepsilon > 0$. Because $\sum 1/a_n$ converges, its tails tend to zero, so there is N with

$$\sum_{n=N+1}^{\infty} \frac{1}{a_n} < \varepsilon.$$

Step 2 (a count is a sum of 1's; compare term by term). Let x be large enough that $k(x) > N$. The terms counted by $k(x)$ are $a_1, \ldots, a_{k(x)}$, because the sequence is increasing (so the $a_n \le x$ form an initial segment). For each n with $N < n \le k(x)$ we have $a_n \le x$, hence $1/x \le 1/a_n$. Summing over those $k(x) - N$ indices,

$$\frac{k(x) - N}{x} \le \sum_{n=N+1}^{k(x)} \frac{1}{a_n} \le \sum_{n=N+1}^{\infty} \frac{1}{a_n} < \varepsilon.$$

If instead $k(x) \le N$, then $(k(x) - N)/x \le 0 < \varepsilon$ as well.

Step 3 (the head is negligible). So for every $x > 0$,

$$0 \le \frac{k(x)}{x} < \varepsilon + \frac{N}{x}.$$

With N fixed, $N/x \to 0$, so $\limsup_{x\to\infty} k(x)/x \le \varepsilon$. As $\varepsilon > 0$ was arbitrary, $\lim_{x\to\infty} k(x)/x = 0$.

Remark. Integrality was used only to see $a_n \to \infty$; for positive reals, $a_n \to \infty$ follows from $1/a_n \to 0$, which convergence gives. So the result holds for any increasing sequence of positive reals, as the official solution notes. Monotonicity is what matters. It is exactly what makes the counted terms an initial segment and each of them at most x.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same inequality: each $a_n \le x$ gives $1/a_n \ge 1/x$. They split at what it is compared with. The natural attempt compares $k(x)/x$ with the *whole* sum $\sum_{a_n \le x} 1/a_n$ and gets "$\ge \varepsilon$", which is not a contradiction. The model compares $(k(x) - N)/x$ with a *tail* $\sum_{n>N} 1/a_n$, which convergence makes as small as we like.

> *Convergence doesn't mean the sum is small. It means the tails are small. So I should throw away the first N terms, which only cost me N/x and that goes to zero, and compare the rest with the tail.*

> *Where does "increasing" come in? It makes the $a_n \le x$ the first $k(x)$ terms, and every one of them is at most x. That gives me a lower bound of $1/x$ on each term of a block of consecutive terms, which is what I need.*

Layer 3 THE WHY

Why the tail and not the sum? Go back to the definition. "$\sum 1/a_n$ converges" means the partial sums have a limit, and by the Cauchy criterion this is equivalent to: for every ε there is an N beyond which every block of consecutive terms sums to less than ε. The whole sum being a finite number is a weak fact; it cannot beat a lower bound of a fixed ε, which is why the contradiction attempt stalls. The small-tails fact is strong, and it is the content of convergence that lets you prove something tends to zero. So whenever the goal is "$\to 0$" and the hypothesis is "converges", the tail is what you use.

Why is throwing away the first N terms free? The quantity is $k(x)/x$, a count divided by x. Removing N counted items changes it by N/x, and N is fixed before x is sent to infinity. This order of quantifiers (choose N from ε, then let $x \to \infty$) is the structure of every "head plus tail" argument, and it is why the head never has to be controlled: the x in the denominator does that.

Why does a count compare with a sum of reciprocals? A count is a sum of 1's: $k(x) - N = \sum_{N<n\le k(x)} 1$. Dividing by x gives $\sum_{N<n\le k(x)} 1/x$, and now the comparison with $\sum 1/a_n$ is term

by term, using only $a_n \le x$. The step that looks inventive is writing $k(x)/x$ in the same form as the hypothesis: both are then sums over the same indices, and one sum is termwise smaller.

Why does the slogan fail while this proof succeeds? The false slogan "$\sum b_n < \infty \Rightarrow nb_n \to 0$" needs one more hypothesis: b_n non-increasing. With it, $(n-N)b_n \le b_{N+1} + \cdots + b_n < \varepsilon$, which is the model chain's inequality with $x = a_n$. (This corrected statement is sometimes called Olivier's theorem.) The problem's word "increasing" is that hypothesis in disguise, since a_n increasing means $1/a_n$ decreasing. So the second natural attempt had the right destination and dropped the one hypothesis that makes it true. When a hypothesis in the statement seems to do nothing, find the step that needs it; here, without monotonicity, terms $a_n \le x$ with late indices could be scattered among huge early ones and the block argument would fail.

Finally, the converse is false, which tells you the result is one-directional: the primes have $k(x)/x \to 0$ (by Chebyshev-type estimates, $k(x)$ is of order $x/\log x$), yet $\sum 1/p$ diverges. Density zero is necessary for convergence, not sufficient.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the combination "$\sum 1/a_n$ convergent" and "show $\ldots \to 0$". On seeing it, write down: "fix ε, choose N with $\sum_{n>N} 1/a_n < \varepsilon$."

2. Head plus tail: split off the first N indices (cost $N/x \to 0$) and compare the rest term by term with the tail, using $a_n \le x \Rightarrow 1/a_n \ge 1/x$.

3. Hesitation comes at the contradiction set-up, when "$\sum_{a_n \le x_j} 1/a_n \ge \varepsilon$" fails to contradict anything. It goes away once you remember that the contradiction has to be with the tail, so start the sum at $N+1$.

4. The trap is the false slogan that convergence of $\sum b_n$ forces $b_n = o(1/n)$. It needs monotonicity; a counterexample is $b_n = 1/n$ on perfect squares and $1/n^2$ elsewhere.

5. Whenever a goal is "$\to 0$" and the hypothesis is "converges", your first line should use small tails, not the finite total.

6. Decision rule: to show a normalised count tends to zero, write the count as a sum of 1's, divide by the normaliser to get a termwise comparison with the convergent series, and discard a fixed head whose contribution the normaliser kills.

▷ **Cue:** hypothesis "$\sum b_n$ converges", goal "something $\to 0$" $\longrightarrow$ **Move:** fix ε, choose N so the tail beyond N is $< \varepsilon$, and aim the whole argument at that tail

▷ **Cue:** a counting function $k(x)$ divided by x $\longrightarrow$ **Move:** write $k(x)/x$ as $\sum 1/x$ over the counted items and compare termwise with a known series

▷ **Cue:** a monotonicity hypothesis ("increasing") whose use is not obvious $\longrightarrow$ **Move:** find the step that needs it; here it makes the counted terms an initial segment, each bounded by x — the Olivier-type $nb_n \to 0$ step

1969 B6

1969 B6. Let A and B be matrices of size 3×2 and 2×3 respectively. Suppose that their product in the order AB is given by

$$AB = \begin{pmatrix} 8 & 2 & -2 \\ 2 & 5 & 4 \\ -2 & 4 & 5 \end{pmatrix}.$$

Show that the product BA is given by

$$BA = \begin{pmatrix} 9 & 0 \\ 0 & 9 \end{pmatrix}.$$

Caveat. The source OCR drops the $(1,1)$ entry of AB (it reads "[~ -2 / 2 -2] 5 4 / 4 5"). The matrix above is the standard statement and was checked against the book's own solution, which uses $(AB)^2 = 9AB$ and rank $AB = 2$: with 8 in the corner, $(AB)^2 = 9AB$ holds entry by entry, $\det AB = 8 \cdot 9 - 2 \cdot 18 - 2 \cdot 18 = 0$, and the first two rows are independent, so the rank is 2. No other entry makes both facts true together with the trace $18 = \operatorname{tr}(9I_2)$.

Fields: Linear algebra: matrices, rank, the AB/BA relationship **Position:** B6 (last of the paper; short once the right identity is found) **Answer:** $(AB)^2 = 9AB$ and rank $AB = 2$ force BA to be invertible, and $(BA)^3 = B(AB)^2A = 9(BA)^2$; cancelling $(BA)^2$ gives $BA = 9I$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You are given a product and asked about the reverse product, so try to find A and B. The columns of AB span its column space; the vector $w = (1,-2,2)^{\mathsf{T}}$ is orthogonal to all three rows, so AB has rank 2 with column space $w^{\perp}$. Pick a basis of $w^{\perp}$ with small integers, say $(2,1,0)^{\mathsf{T}}$ and $(-2,0,1)^{\mathsf{T}}$, and set

$$A_0 = \begin{pmatrix} 2 & -2 \\ 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Solve $A_0B_0 = AB$ row by row: row 2 of AB equals the first row of B_0, so $r_1 = (2,5,4)$; row 3 equals the second, $r_2 = (-2,4,5)$; row 1 checks, $2r_1 - 2r_2 = (8,2,-2)$. Then

$$B_0A_0 = \begin{pmatrix} 2 & 5 & 4 \\ -2 & 4 & 5 \end{pmatrix} \begin{pmatrix} 2 & -2 \\ 1 & 0 \\ 0 & 1 \end{pmatrix} = \begin{pmatrix} 9 & 0 \\ 0 & 9 \end{pmatrix}.$$

The answer is right, but nothing has been proved yet. The problem is about *every* pair (A,B) with this product, and there are infinitely many (A_0P and $P^{-1}B_0$ for any invertible 2×2 P, at least). To finish you need a lemma: every factorisation is of that form. (Proof: $\operatorname{col} A \supseteq \operatorname{col} AB$, which has dimension 2, so A has rank 2 and $\operatorname{col} A = \operatorname{col} A_0$, giving $A = A_0P$ with P invertible; then $A_0(PB - B_0) = 0$ and A_0 has independent columns, so $B_0 = PB$.) After that, $BA = P^{-1}B_0A_0P = P^{-1}(9I)P = 9I$. This route works, but it needs a guessed basis, a row-by-row solve, a uniqueness lemma and a conjugation.

The model chain. *Step 1 (the identity AB satisfies).* Direct multiplication gives

$$(AB)^2 = \begin{pmatrix} 72 & 18 & -18 \\ 18 & 45 & 36 \\ -18 & 36 & 45 \end{pmatrix} = 9\,AB.$$

(For instance, the $(1,1)$ entry is $64+4+4 = 72$ and the $(2,3)$ entry is $-4+20+20 = 36$.)

Step 2 (rank of AB). $\det AB = 8(25-16) - 2(10+8) - 2(8+10) = 0$, while the first two rows $(8,2,-2)$ and $(2,5,4)$ are not proportional. So rank $AB = 2$.

Step 3 (BA is invertible). For any matrices, $\operatorname{rank}(XYZ) \le \operatorname{rank} Y$. Hence

$$2 = \operatorname{rank}(9AB) = \operatorname{rank}\big((AB)^2\big) = \operatorname{rank}\big(A\,(BA)\,B\big) \le \operatorname{rank}(BA) \le 2.$$

So the 2×2 matrix BA has rank 2, i.e. it is invertible.

Step 4 (sandwich the identity). By associativity,

$$(BA)^3 = B\,(AB)(AB)\,A = B\,(AB)^2\,A = B\,(9AB)\,A = 9\,(BA)^2.$$

Step 5 (cancel). Multiply on the left by $\big((BA)^{-1}\big)^2$: $BA = 9I_2$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same observation: the only information is the single matrix AB. They split at what they do with it. The natural attempt tries to recover A and B, which the data cannot determine, and so has to prove afterwards that the choice did not matter. The model asks what *equations* AB satisfies ($(AB)^2 = 9AB$), and moves them across to BA using the rearrangement $B(AB)^kA = (BA)^{k+1}$.

> *I can't recover A and B from AB; they're only determined up to P and P^{-1}. So whatever I prove about BA has to come from some property of AB that I can carry across. What identity does AB satisfy? $(AB)^2 = 9AB$. Put B on the left and A on the right and it becomes $(BA)^3 = 9(BA)^2$.*

> *$(BA)^3 = 9(BA)^2$ isn't enough by itself, because BA could be singular. I need BA invertible to cancel. Rank: $A(BA)B = (AB)^2$ has rank 2, so the middle factor BA has rank at least 2.*

Layer 3 THE WHY

Kedlaya-style writeups (and the MAA book) begin with "observe that $ABAB = 9AB$". Here is where that comes from.

Start with the one structural fact that relates AB and BA at all: associativity. Every word in A and B can be bracketed however you like, so

$$B\,(AB)^k\,A = (BA)^{k+1} \qquad \text{for all } k \geq 0.$$

Because of this identity, any *polynomial* relation $p(AB) = 0$ gives $B\,p(AB)\,A = 0$, which is $x\,p(x)$ evaluated at BA: $(BA)\,p(BA) = 0$. So AB and BA satisfy each other's polynomial identities up to a factor of x. This is also why they have the same nonzero eigenvalues. The problem hands you a numerical AB and asks you to identify BA. The only channel between the two is polynomial identities, so the first job is to find one that AB satisfies. Computing $(AB)^2$ is the cheapest test there is, and it comes out as $9AB$. This is the step that looks like luck. It isn't: once you know identities are the only channel, squaring is the first thing to try, and the entries (a symmetric matrix with trace 18 that should become $9I_2$) point to eigenvalues $9, 9, 0$ and hence to $x^2 = 9x$.

Why is invertibility needed, and why is it the rank that supplies it? The transferred identity is $(BA)^2(BA - 9I) = 0$, and the extra factor of x from the sandwich is exactly what can fail to cancel. A 2×2 matrix with $M^3 = 9M^2$ could have $M = 0$, or $M = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}$, or $M = \mathrm{diag}(9, 0)$. What rules these out is dimension. AB has rank 2, and a rank-2 product $A \cdot (BA) \cdot B$ cannot pass through a middle factor of smaller rank, since rank can only drop under multiplication. So BA is a full-rank 2×2 matrix. The numbers 3 and 2 in the problem are doing this work: AB is 3×3 of rank 2, exactly the largest rank a product through $\mathbb{R}^2$ can have, so nothing is lost in passing to BA.

The natural attempt was not wrong, only expensive, and its expensive part (the uniqueness lemma) is where the same rank fact appears in another form: "A has rank 2, so A has independent columns" is the model's Step 3 in different words. Both routes rest on the same fact. The model avoids choosing coordinates.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "given AB, show $BA = \ldots$". On seeing it, write down: "$B(AB)^kA = (BA)^{k+1}$ — find a polynomial identity for AB." Then compute $(AB)^2$.

2. Transfer a polynomial identity from AB to BA by sandwiching, then cancel using invertibility, which comes from a rank count.

3. Hesitation comes at $(BA)^3 = 9(BA)^2$, where it is unclear whether dividing is legal. It goes away with the inequality $\operatorname{rank}(XYZ) \leq \operatorname{rank} Y$ applied to $A(BA)B = (AB)^2$.

4. The trap is exhibiting one factorisation A_0, B_0, computing $B_0A_0 = 9I$, and stopping. The problem quantifies over all A, B; one example proves nothing without the uniqueness-up-to-P lemma.

5. When a problem gives you a product of matrices but not the factors, look for identities the product satisfies (powers, minimal polynomial) rather than trying to recover the factors.

6. Decision rule: information about XY reaches YX only through $Y(XY)^kX = (YX)^{k+1}$; find the minimal-type polynomial of XY, transfer it, and use rank or dimension to remove the extra factor of x.

▷ **Cue:** given the numerical product AB of non-square matrices, asked for BA $\longrightarrow$ **Move:** compute $(AB)^2$ (and higher powers if needed) to find a polynomial identity, then sandwich: $B\,p(AB)\,A = (BA)\,p(BA)$

▷ **Cue:** an identity like $M^3 = 9M^2$ where you want $M = 9I$ $\longrightarrow$ **Move:** prove M invertible (here by $\operatorname{rank}(XYZ) \leq \operatorname{rank} Y$) and cancel

▷ **Cue:** a factorisation that is not unique (factors determined only up to P, P^{-1}) $\longrightarrow$ **Move:** do not choose factors; work only with quantities invariant under $A \mapsto AP$, $B \mapsto P^{-1}B$, such as AB and polynomial identities