DAWN PUTNAM SESSION

Putnam 1994 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1994 — Paper A **5**
1994 A1 5
1994 A2 7
1994 A3 9
1994 A4 12
1994 A5 14
1994 A6 17

Putnam 1994 — Paper B **21**
1994 B1 21
1994 B2 23
1994 B3 26
1994 B4 28
1994 B5 30
1994 B6 33

Putnam 1994 — Paper A

1994 A1

1994 A1. Suppose that a sequence $a_1, a_2, a_3, \ldots$ satisfies $0 < a_n \leq a_{2n} + a_{2n+1}$ for all $n \geq 1$. Prove that the series $\sum_{n=1}^{\infty} a_n$ diverges.

Fields: real analysis, series, inequalities on index trees **Position:** A1 — the opener **Answer:** The series diverges: summing the hypothesis over $n = 1, \ldots, N$ gives $S_{2N+1} \geq S_N + a_1$ for the partial sums $S_N = \sum_{n \leq N} a_n$, so the partial sums along $N = 2^k - 1$ grow at least like $k\, a_1$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You have been trained on comparison tests, so the reflex is to look for a lower bound on the individual term a_n that is itself non-summable — something like $a_n \geq c/n$ — and then compare with the harmonic series.

So you try to push the inequality downwards. From $a_1 \leq a_2 + a_3$ you know the pair (a_2, a_3) has total at least a_1. From $a_2 \leq a_4 + a_5$ and $a_3 \leq a_6 + a_7$, the quadruple $a_4, \ldots, a_7$ has total at least a_1. It looks as if each term at "depth" k (indices $2^k \leq n < 2^{k+1}$) should be about $a_1/2^k$, i.e. about a_1/n, and you go looking for a proof of $a_n \geq c/n$.

This is the wall. The hypothesis says nothing about how a sum is split between its two parts. Take $a_1 = 1$, and at every step give almost all of the mass to the right child: $a_{2n+1} = a_n$ and $a_{2n} = \varepsilon_n$ for any positive ε_n you like. The inequality $a_n \leq a_{2n} + a_{2n+1}$ holds, yet a_{2n} can be 10^{-n}, far below $1/n$. There is no pointwise lower bound to be had. A student who commits to termwise comparison either burns the time trying to prove a false lemma or, worse, writes down a "by induction $a_n \geq a_1/n$" that is wrong.

The model chain. Let $S_N = \sum_{n=1}^{N} a_n$. Add the hypothesis for $n = 1, 2, \ldots, N$:

$$S_N = \sum_{n=1}^{N} a_n \;\leq\; \sum_{n=1}^{N} (a_{2n} + a_{2n+1}) \;=\; \sum_{m=2}^{2N+1} a_m \;=\; S_{2N+1} - a_1.$$

The middle equality holds because, as n runs over $1, \ldots, N$, the indices $2n$ and $2n+1$ run over every integer from 2 to $2N+1$ exactly once. So

$$S_{2N+1} \;\geq\; S_N + a_1 \qquad (N \geq 1). \tag{$*$}$$

Put $N_k = 2^k - 1$, so that $2N_k + 1 = N_{k+1}$. Then $(*)$ reads $S_{N_{k+1}} \geq S_{N_k} + a_1$, and since $S_{N_1} = S_1 = a_1$, induction gives $S_{2^k - 1} \geq k\, a_1$ for all $k \geq 1$. As $a_1 > 0$, the partial sums are unbounded; since all terms are positive the partial sums are increasing, so $S_N \to \infty$ and the series diverges.

(Equivalently, by contradiction: if the series converged to S, letting $N \to \infty$ in $S_N \leq S_{2N+1} - a_1$ gives $S \leq S - a_1$, impossible because $a_1 > 0$.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same observation — the children of n are $2n$ and $2n+1$, and each generation carries at least the mass of the previous one. They split at the

question "what am I trying to bound?" The natural attempt asks for a bound on the *term* a_n; the model asks for a bound on the *sum*, and gets it by adding the given inequality over a whole range of n at once.

> *The question is about $\sum a_n$, and the hypothesis tells me nothing about individual terms, only about sums of pairs. So I should not try to control a_n by itself — I should add the inequality up over $n = 1, \dots, N$ and see what the right-hand side becomes as a sum.*

> *The right-hand side, summed, is almost the same series again, missing only a_1. A series that is at least itself plus a fixed positive amount cannot be finite.*

Layer 3 THE WHY

Why is summing the inequality the natural move and not a lucky one? Look at what the hypothesis is. The map $n \mapsto \{2n, 2n+1\}$ sends each positive integer to two children, and every integer $m \geq 2$ is the child of exactly one parent, namely $\lfloor m/2 \rfloor$. So the positive integers form a binary tree rooted at 1, and the condition says that mass does not decrease when you pass from a node to its children. That is a statement about how mass is *distributed*, not about any single node. A conservation-type (or growth-type) law about distribution can only be exploited by totalling: you add up over a generation, or over the first N nodes, and the law turns into an inequality between totals. Termwise bounds are the wrong currency, which is exactly why the counterexample with $a_{2n} = \varepsilon_n$ exists.

Once you sum, the reason the answer drops out is the bijection $\{1, \dots, N\} \times \{0, 1\} \to \{2, \dots, 2N+1\}$, $(n, j) \mapsto 2n + j$. The children of the first N nodes are precisely the nodes $2, \dots, 2N+1$ — an initial run of the tree with only the root left out. The root is the only node that is nobody's child; that is why a_1, and only a_1, appears as the gap. So the whole proof is: "each generation weighs at least as much as the root, and there are infinitely many generations." Writing it in the tree picture: if $B_k = \sum_{2^k \leq n < 2^{k+1}} a_n$ is the mass of generation k, then $B_{k+1} = \sum_{2^k \leq n < 2^{k+1}} (a_{2n} + a_{2n+1}) \geq B_k$, so every $B_k \geq B_0 = a_1$ and $\sum a_n = \sum_k B_k = \infty$. This is Cauchy condensation's grouping $2^k \leq n < 2^{k+1}$, forced on you by the indices $2n, 2n+1$.

The positivity $a_n > 0$ is used twice and should be seen to be used: once to guarantee $a_1 > 0$ (so the gap is genuine), and once so that the partial sums are monotone and grouping is harmless. Without positivity the statement is false ($a_n = 0$ for all n satisfies the non-strict inequality).

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the index pattern a_n versus $a_{2n} + a_{2n+1}$. On seeing it, write down: "the children of $1, \dots, N$ are exactly $2, \dots, 2N+1$," or equivalently draw the first three levels of the binary tree.

2. It suggests summing the inequality over a range of n (or over a generation $2^k \leq n < 2^{k+1}$) and comparing partial sums, $S_N \leq S_{2N+1} - a_1$.

3. Hesitation happens when you try to prove $a_n \geq c/n$ and cannot. The reinforcement: build the lopsided example $a_{2n+1} = a_n$, a_{2n} tiny, in ten seconds, see that no termwise bound exists, and switch to sums.

4. The trap is the comparison test reflex — an attempted induction "$a_n \geq a_1/n$" that is false, and that a grader will catch at once.

5. Any hypothesis relating a_n to a_{2n}, a_{2n+1} (or $a_{kn}, \dots, a_{kn+k-1}$) is a tree hypothesis; the natural unit is the generation, and the natural operation is to total it.

6. Reusable rule: when a condition constrains how a quantity is shared among several pieces but not how it is split, bound totals, not pieces. Sum the condition over its whole range and see which terms are left over.

- ▷ **Cue:** an inequality links a_n to $a_{2n} + a_{2n+1}$, and you are asked about $\sum a_n$ $\longrightarrow$ **Move:** sum the inequality over $n = 1, \ldots, N$; the right side is the series from 2 to $2N+1$, so you get $S_{2N+1} \geq S_N + a_1$
- ▷ **Cue:** you are tempted to prove a termwise lower bound like $a_n \geq c/n$ from a hypothesis about sums of several terms $\longrightarrow$ **Move:** first test it against a lopsided example where one piece takes all the mass; if it fails, bound partial sums instead
- ▷ **Cue:** indices $2n, 2n+1$ or dyadic blocks $2^k \leq n < 2^{k+1}$ appear $\longrightarrow$ **Move:** group the series by generation (Cauchy condensation grouping) and compare consecutive block sums

1994 A2

1994 A2. Let A be the area of the region in the first quadrant bounded by the line $y = \frac{1}{2}x$, the x-axis, and the ellipse $\frac{1}{9}x^2 + y^2 = 1$. Find the positive number m such that A is equal to the area of the region in the first quadrant bounded by the line $y = mx$, the y-axis, and the ellipse $\frac{1}{9}x^2 + y^2 = 1$.

Fields: geometry, integration, linear maps and area **Position:** A2 — early, routine if you reroute **Answer:** $m = \frac{2}{9}$. Under $(x, y) \mapsto (x/3, y)$ the ellipse becomes the unit circle, both regions become circular sectors, and equal sectors need the lines $y = \frac{3}{2}u$ and $y = 3m\,u$ to make equal angles with the u-axis and the y-axis respectively, i.e. $3m = \frac{2}{3}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Areas bounded by curves: integrate. The line $y = \frac{1}{2}x$ meets the ellipse where $\frac{x^2}{9} + \frac{x^2}{4} = 1$, i.e. $x = 6/\sqrt{13}$, $y = 3/\sqrt{13}$. The first region is the triangle with vertices $(0,0)$, $(6/\sqrt{13}, 0)$, $(6/\sqrt{13}, 3/\sqrt{13})$ plus the piece under the ellipse from $x = 6/\sqrt{13}$ to $x = 3$:

$$A = \frac{9}{13} + \int_{6/\sqrt{13}}^{3} \sqrt{1 - \tfrac{x^2}{9}}\, dx.$$

With $x = 3\sin\theta$ the integral is $\frac{3}{2}\big[\theta + \sin\theta\cos\theta\big]$ between $\theta_0 = \arcsin(2/\sqrt{13})$ and $\pi/2$. Since $\sin\theta_0 \cos\theta_0 = \frac{2}{\sqrt{13}} \cdot \frac{3}{\sqrt{13}} = \frac{6}{13}$, this is $\frac{3}{2}\big(\frac{\pi}{2} - \theta_0\big) - \frac{9}{13}$, and the $\frac{9}{13}$ cancels the triangle: $A = \frac{3}{2}\big(\frac{\pi}{2} - \arcsin\frac{2}{\sqrt{13}}\big)$.

Now do the same for the second region with a general m: the intersection is at $x_m = 3/\sqrt{1+9m^2}$, the region is "under the ellipse, above the line" from 0 to x_m, you get another arcsin with argument depending on m, and another triangle term that you hope cancels. Then you set the two expressions equal and solve a transcendental-looking equation for m, relying on arcsin/arctan identities to make it algebraic. All of this works. It is also a full page of calculus in which a single sign slip or a lost factor of 3 destroys the answer, and the cancellation of the triangle terms — which is the whole reason the answer is clean — arrives as a surprise rather than as something you expected.

The model chain. Let $T(x, y) = (x/3,\, y)$, and write (u, y) for the new coordinates. T is linear with determinant $\frac{1}{3}$, so it multiplies every area by $\frac{1}{3}$, and it maps lines through the origin to lines through the origin, axes to axes, and the ellipse $\frac{1}{9}x^2 + y^2 = 1$ to the unit circle $u^2 + y^2 = 1$.

The first region becomes the part of the unit disc in the first quadrant between the u-axis and the line $y = \frac{3}{2}u$ (since $y = \frac{1}{2}x = \frac{3}{2}u$). That is a circular sector of angle $\alpha = \arctan\frac{3}{2}$, with area $\frac{1}{2}\alpha$. So $A = 3 \cdot \frac{1}{2}\alpha = \frac{3}{2}\arctan\frac{3}{2}$.

The second region becomes the part of the disc between the line $y = 3mu$ and the y-axis: a sector of angle $\frac{\pi}{2} - \beta$ with $\beta = \arctan(3m)$, area $\frac{1}{2}(\frac{\pi}{2} - \beta)$, so original area $\frac{3}{2}(\frac{\pi}{2} - \beta)$.

The areas are equal iff $\alpha = \frac{\pi}{2} - \beta$, i.e. iff the line $y = 3mu$ makes the same angle with the y-axis as $y = \frac{3}{2}u$ makes with the u-axis. Reflecting in $y = u$ swaps these, so $y = 3mu$ must be the reflection of $y = \frac{3}{2}u$, namely $y = \frac{2}{3}u$. Hence $3m = \frac{2}{3}$ and

$$m = \frac{2}{9}.$$

(Check: $\arctan\frac{3}{2} + \arctan\frac{2}{3} = \frac{\pi}{2}$ because the two arguments are reciprocals of each other and both positive.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree the answer depends on the ellipse. They split at the first line: the natural attempt sets up an integral in x and y; the model changes coordinates *before* computing anything, so that the curve becomes a circle and "area bounded by two rays and the curve" becomes "sector area", which is proportional to angle.

> *Both regions are bounded by two lines through the origin and the ellipse — that is a sector shape. For a circle the area of a sector is half the angle, with no integration. An ellipse is a circle stretched along one axis, and a stretch scales every area by the same factor, so I should squash x by 3 and compare angles instead of computing integrals.*

> *After the squash, one sector hangs off the horizontal axis and the other off the vertical axis. Equal angles from the two axes means one line is the mirror image of the other in the diagonal, so the slopes are reciprocals.*

Layer 3 THE WHY

(No official solution is in the archive's sources for 1994; the chain above is constructed and checked here, including numerically against the direct integral.) Every short solution to this problem begins with the substitution $x \mapsto x/3$. Here is where that comes from. A linear map L of the plane multiplies every area by $|\det L|$; this is the determinant's definition as a volume scale factor, and for the diagonal map $\mathrm{diag}(\frac{1}{3}, 1)$ it is visible directly: every rectangle with sides parallel to the axes shrinks horizontally by 3, and any region's area is a limit of sums of such rectangles. So the *ratio* or *equality* of two areas is unchanged by L. Since the problem asks only for an equality of two areas, you are free to apply any linear map that makes the areas easier to compute, and the ellipse $\frac{1}{9}x^2 + y^2 = 1$ is, by its equation, the image of the unit circle under $(u, y) \mapsto (3u, y)$ — the equation is literally $u^2 + y^2 = 1$ with $u = x/3$ written in. The coefficient $\frac{1}{9}$ is the problem telling you which stretch to undo.

Why is that the right thing to do and not a trick? Because the regions are not arbitrary: each is bounded by two rays from the centre of the ellipse and an arc of the ellipse. Linear maps preserve "ray from the origin" and "line through the origin". So after the map, each region is still bounded by two rays from the centre and an arc — now of a circle — which is the definition of a sector. The area of a circular sector is $\frac{1}{2}r^2\theta$, and that formula is precisely what the integral in the natural attempt was reconstructing by hand (the θ in $\frac{3}{2}[\theta + \sin\theta\cos\theta]$ is the sector, and the $\sin\theta\cos\theta$ is the triangle that cancelled). The cancellation that looked like good luck is the fact that a sector is the natural shape.

What a linear map does *not* preserve is angle. So you must translate the lines honestly: $y = \frac{1}{2}x$ becomes $y = \frac{3}{2}u$, not $y = \frac{1}{2}u$. This is the single place the model chain can go wrong, and it is worth writing the substitution out rather than doing it in your head.

Finally, the symmetry step: the sector from the horizontal axis up to angle α and the sector from angle $\pi/2 - \alpha$ up to the vertical axis are mirror images in the diagonal $y = u$, and that reflection sends slope s to slope $1/s$. So equal areas force $3m = \left(\frac{3}{2}\right)^{-1}$.

Layer $+\alpha$ BACKTESTING

1. The cue is an ellipse with axes on the coordinate axes, and regions bounded by lines through its centre. Write down $u = x/3$ and the transformed equations $u^2 + y^2 = 1$, $y = \frac{3}{2}u$, $y = 3mu$ immediately.

2. The method is an area-scaling linear change of variables, followed by the sector formula $\frac{1}{2}\theta$ for the unit circle.

3. Hesitation: "am I allowed to change coordinates when the question is about area?" Reinforce with the determinant fact: every area scales by the same constant, so equalities of areas survive. Second hesitation: transforming the slopes — write the substitution out.

4. The trap is keeping slope $\frac{1}{2}$ after squashing x, which gives $m = 2$ or $m = \frac{1}{2}$ depending on how you then argue; both are wrong. A secondary trap is the full integration, which is correct but costs ten minutes and invites arithmetic error.

5. Whenever a conic or quadric appears only with its axes as coordinate axes, and the question is about area, volume, or ratios along lines through the centre, the problem is a circle problem in disguise.

6. Reusable rule: if a question asks only for equalities or ratios of areas (or volumes), first apply the linear map that makes the figure symmetric; compute in the symmetric picture, then translate back only the objects the map distorts (angles, slopes), never the ratios.

▷ **Cue:** an ellipse $\frac{x^2}{a^2} + \frac{y^2}{b^2} = 1$ and areas cut off by lines through its centre $\longrightarrow$ **Move:** rescale $x \mapsto x/a$, $y \mapsto y/b$ to get the unit circle; areas scale by $1/(ab)$ and become sector areas $\frac{1}{2}\theta$

▷ **Cue:** a linear change of variables has been made in a problem involving lines $\longrightarrow$ **Move:** rewrite each line's equation explicitly in the new coordinates; slopes and angles are not preserved even though area ratios are

▷ **Cue:** two sectors, one measured from the horizontal axis and one from the vertical axis, must be equal $\longrightarrow$ **Move:** reflect in the diagonal: the bounding lines have reciprocal slopes

1994 A3

1994 A3. Show that if the points of an isosceles right triangle of side length 1 are each colored with one of four colors, then there must be two points of the same color which are at least a distance $2 - \sqrt{2}$ apart.

Fields: combinatorial geometry, graph colouring, pigeonhole **Position:** A3 — mid-paper, short once the right seven points are found **Answer:** True. Seven explicit points (the three vertices, two points on the legs at distance $\sqrt{2} - 1$ from the right angle, two points on the hypotenuse at distance $2 - \sqrt{2}$ from the acute vertices) cannot be four-coloured so that every pair at distance $\geq 2 - \sqrt{2}$ gets different colours. ("Side length 1" is read as legs of length 1, hypotenuse $\sqrt{2}$; the source statement has the typo "whch", corrected here.)

Layer 1 THE WHAT — surface analysis

The natural attempt. Four colours, "two points of the same colour": pigeonhole. The reflex is to find five points that are pairwise at distance at least $d = 2 - \sqrt{2} \approx 0.586$; then two of the five share a colour and you are done.

So you hunt. Put the right angle at $C = (0,0)$ and the other vertices at $A = (1,0)$, $B = (0,1)$. The three vertices are pairwise $1, 1, \sqrt{2}$ apart — three points. The midpoint $M = (\frac{1}{2}, \frac{1}{2})$ of the hypotenuse is at distance $\frac{\sqrt{2}}{2} \approx 0.707$ from each vertex — four points. Now a fifth. It must be at least 0.586 from A, from B, from C and from M. Try points on the legs: $(t, 0)$ is within 0.586 of C if $t < 0.586$ and within 0.586 of A if $t > 0.414$, so every point of the leg CA is too close to C or A, except those in $[0.414, 0.586]$, and those are within 0.51 of M. You can try moving M along the hypotenuse, you can try other quadruples, and nothing works.

The wall appears to be real, not a failure of imagination: a numerical search (thousands of random starts, constrained optimisation) finds that the largest possible minimum distance among five points of the triangle is about 0.536, apparently $4 - 2\sqrt{3}$, which is below $d \approx 0.586$. Pure pigeonhole on a set of mutually far points is too crude. You can burn the whole time on it.

Caveat. The claim that no five points of the triangle are pairwise at distance $\geq 2 - \sqrt{2}$ is supported by numerical optimisation only, not by a proof. The model chain below does not depend on it; it only explains why the pigeonhole route stalls.

The model chain. Let $d = 2 - \sqrt{2}$ and $s = \sqrt{2} - 1$, so $s + d = 1$, $s\sqrt{2} = d$, and $0 < s < d < 1$ (since $2\sqrt{2} < 3$). Place $C = (0,0)$, $A = (1,0)$, $B = (0,1)$ and consider the seven points

$$C,\ A,\ B, \qquad E = (s,0),\ F = (0,s), \qquad G = (d,s),\ H = (s,d).$$

$E \in CA$ and $F \in CB$; G and H lie on the hypotenuse $x + y = 1$. Call two points *far* if their distance is at least d. Suppose, for contradiction, that no two far points have the same colour. The following pairs are far:

- A, B, C pairwise: distances $1, 1, \sqrt{2}$.
- G with each vertex: $GA = \sqrt{(1-d)^2 + s^2} = s\sqrt{2} = d$; $GB = \sqrt{d^2 + (1-s)^2} = d\sqrt{2}$; $GC = \sqrt{d^2 + s^2} > d$. By the symmetry $(x, y) \mapsto (y, x)$, which swaps $A \leftrightarrow B$, $E \leftrightarrow F$, $G \leftrightarrow H$ and fixes C, also H is far from A, B, C.
- E with A: $EA = 1 - s = d$. E with B: $\sqrt{s^2 + 1} > 1$. E with H: $H - E = (0, d)$. By the same symmetry, F is far from B, A and G.
- E with F: $EF = s\sqrt{2} = d$.

Now colour. A, B, C are pairwise far, so they use three different colours; call them $1, 2, 3$ respectively. G is far from all three, so G has colour 4; likewise H has colour 4. E is far from A (colour 1), B (colour 2) and H (colour 4), so E has colour 3. F is far from B, A and G, so F also has colour 3. But $EF = d$, so E and F are far and share a colour — contradiction. Hence some two points at distance at least $2 - \sqrt{2}$ have the same colour.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the first move: the three vertices are mutually far, so they take three different colours. They split immediately after. The natural attempt insists on finding points far from *everything* already chosen (a five-point clique). The model accepts points that are far from only *some* of the chosen points, and uses them to force colours: a point far from three differently coloured points has its colour determined.

> *I cannot find five mutually far points, and I should stop trying, because the contradiction cannot come from pigeonhole alone. I need a chain of forced colours: once A, B, C have three colours, any point far from all three is forced to the fourth colour, and any point far from two vertices and one fourth-colour point is forced to the third vertex's colour. I should look for points that get forced, and then two forced-equal points that are far apart.*

> *The number $2 - \sqrt{2}$ is $\sqrt{2}(\sqrt{2} - 1)$, and $\sqrt{2} - 1 = 1 - (2 - \sqrt{2})$. So the points at distance $2 - \sqrt{2}$ from a vertex, measured along an edge, are exactly the points where these two numbers meet. The constant is telling me where to put the points.*

Layer 3 THE WHY

(No official solution is in the archive's sources for 1994; the seven-point chain above is constructed here, every distance has been checked exactly, and an exhaustive computer check confirms the seven points admit no valid four-colouring.)

Why must you give up the clique? Rephrase the problem as a graph. Join two points of the triangle when they are far. A colouring with no two far points sharing a colour is a proper colouring of this graph, and the problem says the graph has chromatic number at least 5. There are two ways a graph can need five colours: it contains five mutually joined vertices (a K_5), or it does not, and colours are forced along a chain until something collides. The first is pigeonhole; the second is the ordinary argument you would use to show, say, that an odd cycle needs three colours — no triangle, yet two colours fail. When the clique search stalls, that is the signal to switch to the forcing argument, not to search harder.

Why these seven points? Work backwards from what the forcing needs. After A, B, C take colours $1, 2, 3$, the useful points are of two kinds. A point far from all three vertices is forced to colour 4. On the hypotenuse, a point is always far from C (the distance from C to the hypotenuse is $\frac{\sqrt{2}}{2} > d$), so it is forced to colour 4 as soon as it is at least d from both A and B, i.e. anywhere in the middle stretch of length $\sqrt{2} - 2d$. The extreme such points are G and H, at distance exactly d from A and B respectively. Those extremes are what the problem's own constant prescribes, and extremes are where you get the most room for the next step.

A point on a leg near the right angle is too close to C to be forced by the vertices alone, so it needs a colour-4 point to push it: E on CA must be far from A (so $EA \geq d$, putting E at $x \leq s$) and far from H. The point directly below H on the leg is $(s, 0)$, and it sits at distance exactly d from H and exactly d from A — both constraints are tight simultaneously. That is not luck; it is because $s + d = 1$. Then E and its mirror image F are both forced to C's colour, and they are at distance $s\sqrt{2} = d$ apart. Every binding constraint (EA, GA, HB, HE, FG, EF) is an equality at the chosen points, which is the sign that $2 - \sqrt{2}$ is the constant this configuration was built to realise.

Note that "at least" in the statement is essential: EA, GA, HE, EF are all *exactly* d. With "more than d" this configuration proves nothing. (Whether $2 - \sqrt{2}$ is the best possible constant is not asked, and is not claimed here.)

Layer $+\alpha$ BACKTESTING

1. The first cue is "four colours" together with "two points of the same colour at distance $\geq d$". Write down the graph reformulation: join points that are far; the claim is that this graph is not four-colourable. Then write the constant as $d = 2 - \sqrt{2} = \sqrt{2}(\sqrt{2} - 1)$, $1 - d = \sqrt{2} - 1$.

2. It suggests a finite witness: fix the colours of a mutually far triple (the vertices), then look for points whose colours are forced, and hunt for two forced-equal points that are far apart.

3. Hesitation occurs when the five-point clique cannot be found. Reinforce by recalling the odd-cycle example: chromatic number can exceed the largest clique, and the proof then is a forcing chain.

4. The trap is the pure pigeonhole search, which stalls because (numerically) the clique does not exist; and, secondarily, reading "at least" as "more than" and discarding points at distance exactly d, which are the only ones that work.

5. A constant like $2 - \sqrt{2}$ inside a right isosceles triangle is built from 1 and $\sqrt{2}$; decompose it into the leg and hypotenuse lengths and place points at those distances from the vertices.

6. Reusable rule: to show a "distance graph" needs more than k colours, first try a $(k+1)$-clique; if none exists, fix colours on a maximal clique and chain forced colours, choosing each new point at an extreme position allowed by the constraints, until two forced-equal points are joined.

▷ **Cue:** "colour with k colours; show two points of the same colour are at distance at least d" ⟶ **Move:** recast as: the graph joining points at distance $\geq d$ is not k-colourable; find a

finite subgraph that is not

- ▷ **Cue:** no $k+1$ mutually far points can be found $\longrightarrow$ **Move:** fix colours on a large mutually far set and force further colours one point at a time until two forced-equal points are far apart
- ▷ **Cue:** a strange constant such as $2-\sqrt{2}$ in a figure with sides 1 and $\sqrt{2}$ $\longrightarrow$ **Move:** rewrite it as a combination of the side lengths $(\sqrt{2}(\sqrt{2}-1), 1-(\sqrt{2}-1))$ and place points at exactly those distances from vertices
- ▷ **Cue:** the statement says "at least" a distance $\longrightarrow$ **Move:** expect the extremal configuration to use points at exactly that distance; do not discard equality cases

1994 A4

1994 A4. Let A and B be 2×2 matrices with integer entries such that A, $A+B$, $A+2B$, $A+3B$, and $A+4B$ are all invertible matrices whose inverses have integer entries. Show that $A+5B$ is invertible and that its inverse has integer entries.

Fields: linear algebra (determinants); algebra (polynomials, root counting); pigeonhole **Position:** A4 — mid-paper, and one of the gentler A4s once the determinant is written as a polynomial in the multiplier of B. **Answer:** $\det(A+tB)$ is a polynomial of degree at most 2 in t that equals ± 1 at $t=0,1,2,3,4$; three of those values coincide, so the polynomial is constant, hence $\det(A+5B)=\pm 1$ and $(A+5B)^{-1}$ has integer entries.

Layer 1 THE WHAT — surface analysis

The natural attempt. You write $A=\begin{pmatrix} a & b \\ c & d \end{pmatrix}$, $B=\begin{pmatrix} p & q \\ r & s \end{pmatrix}$ and start computing inverses. The inverse of $A+kB$ is

$$\frac{1}{(a+kp)(d+ks)-(b+kq)(c+kr)}\begin{pmatrix} d+ks & -(b+kq) \\ -(c+kr) & a+kp \end{pmatrix},$$

and "integer entries" becomes a divisibility condition: the determinant divides each of the four entries of the adjugate, for $k=0,\dots,4$. That is twenty divisibility conditions in eight unknowns. You try special cases — B diagonal, $B=I$ — and they do not generalise. You try to show $A+5B$ is a combination of the earlier matrices, e.g. $A+5B=5(A+B)-4A$, but a linear combination of invertible matrices tells you nothing about invertibility, let alone about the inverse.

Where this stalls: the condition "the inverse has integer entries" has been left as a statement about entries, and there are far too many of them. The attempt never compresses the hypothesis to a single number.

The model chain. *Step 1: integer inverse means determinant* ± 1. Let M be an integer matrix. If M^{-1} exists and has integer entries, then $\det M$ and $\det M^{-1}$ are integers (determinants of integer matrices are integers, being sums of products of entries) with $\det M\cdot\det M^{-1}=\det I=1$, so $\det M=\pm 1$. Conversely, if $\det M=\pm 1$ then $M^{-1}=\frac{1}{\det M}\operatorname{adj} M=\pm\operatorname{adj} M$, and the adjugate of an integer matrix is an integer matrix. So for integer 2×2 matrices,

$$M^{-1}\text{ exists and is integral}\iff \det M\in\{1,-1\}.$$

Step 2: the determinant along the line is a quadratic. With the entries above, define

$$f(t)=\det(A+tB)=(a+tp)(d+ts)-(b+tq)(c+tr)=\det A+t\,(as+dp-br-cq)+t^2\det B.$$

This is a polynomial in t of degree at most 2 with integer coefficients.

Step 3: pigeonhole. By hypothesis and Step 1, $f(0),f(1),f(2),f(3),f(4)\in\{1,-1\}$. Five values in a two-element set: some value $\varepsilon\in\{\pm 1\}$ is taken at least $\lceil 5/2\rceil=3$ times.

Step 4: root counting. The polynomial $f(t) - \varepsilon$ has degree at most 2 and at least three distinct roots. A nonzero polynomial of degree at most 2 has at most 2 roots, so $f(t) - \varepsilon$ is the zero polynomial: $f \equiv \varepsilon$.

Step 5: conclude. $\det(A + 5B) = f(5) = \varepsilon = \pm 1$, so by Step 1, $A + 5B$ is invertible with integer inverse. In fact $A + tB$ has an integral inverse for every integer t. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the phrase "invertible matrices whose inverses have integer entries". The natural attempt translates it entry by entry into divisibility conditions. The model translates it into one number: $\det = \pm 1$. Everything after that is a statement about a single polynomial, and the problem is over in three lines.

> *Integer matrix, integer inverse — the determinant of each is an integer and they multiply to* 1. *So the whole hypothesis is "*$\det = \pm 1$*". I should be tracking one number per matrix, not four entries.*

> *The matrices are* $A + kB$ *for* $k = 0, 1, 2, 3, 4$ *— equally spaced points on a line in matrix space. The determinant along that line is a polynomial in* k*, and for* 2×2 *matrices its degree is at most* 2. *A quadratic that hits only* ± 1 *at five points has nowhere to go.*

Layer 3 THE WHY

Why is the determinant the right compression, and why is "polynomial in t" the right frame? Go to definitions.

An integer matrix with integer inverse is a unit in the ring of integer matrices. The determinant is multiplicative, $\det(MN) = \det M \det N$, and sends integer matrices to integers. A multiplicative map sends units to units, and the only units in $\mathbb{Z}$ are ± 1. That is the entire content of Step 1: it is the statement "units go to units", and the reverse direction is the adjugate formula. So whenever a problem says "invertible over the integers", the first thing to write down is $\det = \pm 1$. It converts a condition on the whole matrix (the inverse, four entries, each a rational function of eight unknowns) into a single integer.

Once there is one integer per matrix, the question becomes: how does that integer depend on k? The problem has handed you an arithmetic progression of matrices $A, A + B, \ldots, A + 4B$. Anything polynomial in the entries of a matrix — and the determinant is, by definition, a sum of signed products of entries — becomes a polynomial in k along that progression. Its degree is at most the size of the matrix, because each product in the determinant has exactly n factors, each linear in k. For 2×2 that is degree 2.

Now the count 5 explains itself. A degree-2 polynomial is pinned down by 3 values. Five values in $\{\pm 1\}$ guarantee three equal ones by pigeonhole, and three equal values of a quadratic force it to be constant. The numbers in the problem are not arbitrary: $5 = 2 \cdot 2 + 1$ is exactly the least number of points at which a two-valued function must repeat a value $3 = 2 + 1$ times.

Is 5 really needed? Yes, and checking this is how you confirm you understand the mechanism. With four matrices the pigeonhole only guarantees two equal values, and a quadratic can alternate. Take $A = \begin{pmatrix} 0 & -1 \\ 1 & -3 \end{pmatrix}$, $B = I$. Then

$$\det(A + tB) = t(t - 3) + 1 = t^2 - 3t + 1,$$

which takes the values $1, -1, -1, 1$ at $t = 0, 1, 2, 3$ and the value 5 at $t = 4$. So $A, \ldots, A + 3B$ all have integer inverses, but $A + 4B$ does not. The hypothesis on five matrices is sharp, and the proof uses all five.

There is no official solution in the archive's sources for 1994. The solution above is constructed and verified here; every step is elementary and has been checked, including the sharpness example.

The rigour points are small but real. First, Step 1 needs both directions: the hypothesis uses "integral inverse $\Rightarrow$ det $= \pm 1$", the conclusion uses "det $= \pm 1 \Rightarrow$ integral inverse". A script that proves only one direction has not finished. Second, "f has degree at most 2" must be justified by writing f out; the t^2 coefficient is $\det B$, which may be zero, so f may be linear or constant — the argument covers those cases because it only uses "at most 2".

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "invertible matrices whose inverses have integer entries". On seeing it, write $\det(A + kB) = \pm 1$ for $k = 0, \ldots, 4$.
2. Treat the determinant as a function of the parameter k: $f(t) = \det(A + tB)$ is a polynomial of degree at most 2. Then pigeonhole plus root counting.
3. Hesitation comes after writing det $= \pm 1$: the sign is unknown, so you cannot immediately say "$f - 1$ has five roots". Reinforce it away with pigeonhole: five values in two classes put three in one class, and three is enough for a quadratic.
4. The trap is working with entries of the inverse — twenty divisibility conditions — instead of one determinant per matrix. A secondary trap is trying to write $A + 5B$ as a combination of the given matrices; invertibility is not preserved by linear combination.
5. Slogan: *"invertible over $\mathbb{Z}$" means "determinant is a unit"*. And: *a family of matrices linear in a parameter has a determinant polynomial in that parameter, of degree at most the size of the matrix.*
6. Decision rule: when a quantity is known to take few values at many points and depends polynomially on a parameter, count: if the number of points exceeds (number of possible values)$\times$(degree), some value is taken more than degree times and the polynomial is constant.

▷ **Cue:** an integer matrix whose inverse has integer entries $\longrightarrow$ **Move:** write $\det M = \pm 1$ (units map to units under the multiplicative determinant)
▷ **Cue:** a list of matrices $A, A + B, A + 2B, \ldots$ in arithmetic progression $\longrightarrow$ **Move:** study $f(t) = \det(A + tB)$, a polynomial in t of degree at most n
▷ **Cue:** a polynomial of degree at most d that takes one of v values at more than vd points $\longrightarrow$ **Move:** pigeonhole gives $d + 1$ equal values; the polynomial is constant
▷ **Cue:** a hypothesis with a specific count (here five matrices) $\longrightarrow$ **Move:** check the count is sharp by constructing a counterexample with one fewer

1994 A5

1994 A5. Let $(r_n)_{n \geq 0}$ be a sequence of positive real numbers such that $\lim_{n \to \infty} r_n = 0$. Let S be the set of numbers representable as a sum

$$r_{i_1} + r_{i_2} + \cdots + r_{i_{1994}},$$

with $i_1 < i_2 < \cdots < i_{1994}$. Show that every nonempty interval (a, b) contains a nonempty subinterval (c, d) that does not intersect S.

Fields: real analysis (closure, nowhere dense sets); countability; sequential compactness **Position:** A5 — hard end of the paper. The statement is point-set topology in disguise; once that is seen the proof is short, but seeing it is the problem. **Answer:** The closure of S is contained in the countable set of all sums of at most 1994 terms (the empty sum 0 included); a closed countable set contains no interval, so every (a, b) has a point outside $\overline{S}$, and a small interval around it misses S.

Layer 1 THE WHAT — surface analysis

The natural attempt. You read "every interval contains a subinterval missing S" and try to build the gap directly. Take (a, b). The terms r_n tend to 0, so all but finitely many are below any

$\delta > 0$. You try to argue that sums near a point x must use "mostly small" terms and then show there is some room between them. Try $k = 1$ first: $S_1 = \{r_n\}$ is a sequence tending to 0; its only accumulation point is 0, so away from 0 it is a discrete set and gaps are easy. For $k = 2$: $S_2 = \{r_i + r_j : i < j\}$ accumulates at each r_i (take $j \to \infty$) and at 0. So near r_i you need a gap that avoids both r_i and the sums $r_i + r_j$ converging to it, and the sums $r_{i'} + r_j$ for other i'. It gets messy fast: the accumulation points of S_2 are themselves accumulating at 0, and you have to track which ones lie in (a, b).

For $k = 1994$ this direct construction is a nested mess of accumulation points of accumulation points. You can see it should work, but writing it for general k by explicit ε-bookkeeping is a long, error-prone induction.

The model chain. Throughout, write S_k for the set of sums $r_{i_1} + \cdots + r_{i_k}$ with $i_1 < \cdots < i_k$, for $k \geq 1$, and $S_0 = \{0\}$. So $S = S_{1994}$. Let $\overline{X}$ denote the closure of $X \subseteq \mathbb{R}$: the set of limits of convergent sequences from X.

Step 1: each S_k is countable. S_k is the image of the set of k-element subsets of $\{0, 1, 2, \dots\}$, which is countable (it injects into $\mathbb{N}^k$). So $T = S_0 \cup S_1 \cup \cdots \cup S_{1994}$ is countable.

Step 2: $\overline{S_k} \subseteq S_0 \cup \cdots \cup S_k$. Let $x = \lim_{j\to\infty} x_j$ with

$$x_j = r_{i_1^{(j)}} + \cdots + r_{i_k^{(j)}}, \qquad i_1^{(j)} < \cdots < i_k^{(j)}.$$

Pass to subsequences one index position at a time. Look at the integer sequence $(i_1^{(j)})_j$. If it is bounded, it takes some value infinitely often; pass to a subsequence on which it is constant. If it is unbounded, pass to a subsequence on which it tends to ∞. Then do the same for $i_2^{(j)}$ on the new subsequence, and so on up to $i_k^{(j)}$. Subsequences of convergent sequences converge to the same limit, so still $x_j \to x$.

After this, let m be the number of positions p at which $i_p^{(j)}$ is constant, and suppose position p is not constant, so $i_p^{(j)} \to \infty$. For every $p' > p$, $i_{p'}^{(j)} > i_p^{(j)} \to \infty$, so position p' also tends to infinity. Hence the constant positions are exactly $p = 1, \dots, m$, with values $c_1 < \cdots < c_m$, and for $p > m$ we have $i_p^{(j)} \to \infty$, so $r_{i_p^{(j)}} \to 0$ because $r_n \to 0$. Taking the limit of the (fixed-length) sum,

$$x = \lim_j x_j = r_{c_1} + \cdots + r_{c_m} + \sum_{p=m+1}^{k} \lim_j r_{i_p^{(j)}} = r_{c_1} + \cdots + r_{c_m} \in S_m.$$

So $x \in S_0 \cup \cdots \cup S_k$.

Step 3: a closed countable set contains no interval. By Step 2, $\overline{S} \subseteq T$, which is countable by Step 1. The interval (a, b) is uncountable, so there is $y \in (a, b) \setminus \overline{S}$.

Step 4: the gap. The complement of a closed set is open, so there is $\delta > 0$ with $(y - \delta, y + \delta) \cap \overline{S} = \varnothing$. Shrink δ so that also $(y - \delta, y + \delta) \subseteq (a, b)$. Set $c = y - \delta$, $d = y + \delta$. Then (c, d) is a nonempty subinterval of (a, b) that misses $\overline{S} \supseteq S$. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains know that S piles up at the limits of sums in which some indices run off to infinity. The natural attempt tries to build the gap by hand, around each accumulation point. The model stops at the observation "S is countable", notices that this is not enough (the rationals are countable and dense), and asks the next question: is the *closure* of S countable? Once it is, the gap comes for free.

> *"Every interval contains a subinterval missing S" — that is exactly "S is nowhere dense", which is the same as "$\overline{S}$ has no interior". I don't want to build gaps; I want to know what $\overline{S}$ looks like.*

> *S is countable, but so is $\mathbb{Q}$, and $\mathbb{Q}$ is everywhere. What separates S from $\mathbb{Q}$? If I take limits of sums in S, some indices stay put and the rest run off, and the ones that run off contribute 0. So a limit point is a shorter sum. The closure is still countable.*

Layer 3 THE WHY

Why reach for closure and countability instead of building gaps? Start from what the conclusion means. "Every nonempty open interval contains a nonempty open subinterval disjoint from S" is word for word the definition of a *nowhere dense* set. The standard equivalent is that $\overline{S}$ has empty interior. You need not know the name; you need to recognise that the conclusion is a statement about how S sits in the line, and that the natural object for that is $\overline{S}$ — because the property "misses an open set" is insensitive to replacing S by $\overline{S}$. An open interval misses S if and only if it misses $\overline{S}$ (any point of $\overline{S}$ inside an open interval has points of S arbitrarily close, hence inside the interval). So the problem is about $\overline{S}$ from the start, and the problem statement is only phrased in terms of S.

Now the question is what $\overline{S}$ can contain, and the hypothesis $r_n \to 0$ answers it. A sum of 1994 terms is determined by 1994 indices. In a convergent sequence of such sums, each index either settles down or escapes to infinity (after passing to a subsequence — this is the discrete version of Bolzano–Weierstrass: a sequence of integers is either bounded, hence has a constant subsequence, or has a subsequence tending to infinity). The escaping indices contribute terms that tend to 0, which is precisely what $r_n \to 0$ means. So limits of k-term sums are m-term sums with $m \leq k$. That is the entire content of the hypothesis; nothing else about (r_n) is used. Positivity is not even needed.

Then countability finishes it, because intervals are uncountable. This is the one fact from beyond A-level that the proof uses: $\mathbb{R}$ (and any interval) is uncountable, by Cantor's diagonal argument. If you prefer to avoid it, Step 3 can be replaced by induction on k using only "finitely many accumulation layers", but that is exactly the messy path the natural attempt was on. Countability is the cheap route.

Why is the number 1994 there? Only to mark the year. The proof works for every fixed k. What would break is letting k vary: the set of all finite sums $\bigcup_k S_k$ can be dense (take $r_n = 2^{-n}$; finite sums of distinct terms give every dyadic rational in $(0, 2)$). The escaping-index argument needs a fixed number of indices, so that the limit is a sum of *finitely many* fixed terms plus *finitely many* vanishing terms. Seeing why k must be fixed is how you see that the proof is right.

There is no official solution in the archive's sources for 1994. The argument above is constructed and checked here; the subsequence extraction in Step 2 is the only place requiring care, and it is written out in full. Note in particular the monotonicity used: because $i_1 < i_2 < \cdots$, once one position escapes, all later positions escape too, so the fixed positions form an initial block.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the conclusion: "every interval contains a subinterval missing S". On seeing it, write: *it suffices to show $\overline{S}$ contains no interval*.

2. Identify $\overline{S}$ by taking limits of sequences from S and extracting subsequences index by index; then use countability of $\overline{S}$ against uncountability of intervals.

3. Hesitation occurs at "S is countable, so done?" — which is false, since $\mathbb{Q}$ is countable and dense. Reinforce it away by remembering that what must be countable (or otherwise small) is the closure, because the conclusion is about open sets.

4. The trap is building gaps by hand around nested layers of accumulation points. A second trap is thinking 1994 matters; a third is forgetting that the index positions must be handled in order, so that escaping indices form a tail.

5. Slogan: *"every interval contains a subinterval missing S" means "$\overline{S}$ has empty interior"; a closed countable set always qualifies.* And: *limits of fixed-length sums from a null sequence are shorter sums.*

6. Decision rule: when a conclusion is about open sets avoiding a set S, replace S by $\overline{S}$ and ask what the limit points of S are. If limits are "degenerate versions" of the elements (fewer terms, lower rank, smaller support), the closure is a finite union of such families and is often countable or otherwise small.

- ▷ **Cue:** "every nonempty interval contains a nonempty subinterval disjoint from S" $\longrightarrow$ **Move:** prove $\overline{S}$ has empty interior; a closed countable set suffices
- ▷ **Cue:** a sequence of sums with a fixed number of indices, where $r_n \to 0$ $\longrightarrow$ **Move:** pass to subsequences so each index is constant or tends to infinity; escaping indices contribute 0 in the limit
- ▷ **Cue:** a countable set that you want to be "sparse" $\longrightarrow$ **Move:** countability alone is not enough ($\mathbb{Q}$); show the closure is countable
- ▷ **Cue:** a sequence of integers $\longrightarrow$ **Move:** either bounded (constant subsequence) or unbounded (subsequence tending to infinity)

1994 A6

1994 A6. Let $f_1, \ldots, f_{10}$ be bijections of the set of integers such that for each integer n, there is some composition $f_{i_1} \circ f_{i_2} \circ \cdots \circ f_{i_m}$ of these functions (allowing repetitions) which maps 0 to n. Consider the set of 1024 functions

$$\mathcal{F} = \{f_1^{e_1} \circ f_2^{e_2} \circ \cdots \circ f_{10}^{e_{10}}\},$$

$e_i = 0$ or 1 for $1 \le i \le 10$. (f_i^0 is the identity function and $f_i^1 = f_i$.) Show that if A is any nonempty finite set of integers, then at most 512 of the functions in $\mathcal{F}$ map A to itself.

Fields: combinatorics (pairing / involution counting); group actions (orbits, invariant sets) **Position:** A6 — last problem of the morning paper. Hard to find, short to write: the whole proof is one pairing, and the difficulty is choosing which exponent to pair on. **Answer:** Some f_i does not map A onto itself; take the largest such i. Pair each exponent vector with the one differing only in e_i. In each pair at most one function maps A to itself, so at most $1024/2 = 512$ do.

Layer 1 THE WHAT — surface analysis

The natural attempt. You notice $512 = 1024/2$ and guess the answer is "half". That suggests pairing, or showing that some single exponent e_i is "bad". The first thing you try is: some f_i must fail to preserve A (otherwise everything generated by the f_i preserves A, and the orbit of 0 could not be all of $\mathbb{Z}$). So fix such an i and pair $g = u \circ v$ with $g' = u \circ f_i \circ v$, where $u = f_1^{e_1} \circ \cdots \circ f_{i-1}^{e_{i-1}}$ and $v = f_{i+1}^{e_{i+1}} \circ \cdots \circ f_{10}^{e_{10}}$. If both map A to itself, then

$$u(v(A)) = A = u(f_i(v(A))),$$

so, u being injective, $f_i(v(A)) = v(A)$. You wanted a contradiction with "f_i does not preserve A", but what you have is that f_i preserves $v(A)$, which is a different set. For an arbitrary choice of i, v can move A to a set that f_i happens to preserve, and the argument stalls.

Many students stop here and try something heavier: think of $\mathcal{F}$ as a cube $\{0,1\}^{10}$, look for an independent set, or try induction on the number of functions. None of these touch the real issue, which is what v does to A.

The model chain. Say that a function g *preserves* A if $g(A) \subseteq A$. For an injective g and finite A, $g(A) \subseteq A$ forces $g(A) = A$, since $|g(A)| = |A|$. Compositions are applied right to left, so in $f_1^{e_1} \circ \cdots \circ f_{10}^{e_{10}}$ the function $f_{10}^{e_{10}}$ acts first.

Step 1: some f_i does not preserve A. Suppose every f_i satisfies $f_i(A) = A$. Then every composition h of the f_i satisfies $h(A) = A$; since h is a bijection of $\mathbb{Z}$, it also maps $\mathbb{Z} \setminus A$ onto $\mathbb{Z} \setminus A$. Now split on whether $0 \in A$.

- If $0 \in A$: for each $n \in \mathbb{Z}$ choose h with $h(0) = n$; then $n \in h(A) = A$. So $A = \mathbb{Z}$, contradicting finiteness.
- If $0 \notin A$: for every composition h, $h(0) \in \mathbb{Z} \setminus A$. Since A is nonempty, pick $n \in A$; no composition maps 0 to n, contradicting the hypothesis.

So some f_i has $f_i(A) \neq A$.

Step 2: choose the last one. Let i be the *largest* index with $f_i(A) \neq A$. Then $f_j(A) = A$ for every $j > i$.

Step 3: the pairing. Group the 1024 exponent vectors $(e_1, \ldots, e_{10})$ into 512 pairs, pairing each vector with the one that differs from it only in the i-th coordinate. A pair gives functions

$$g_0 = u \circ v, \qquad g_1 = u \circ f_i \circ v,$$

where $u = f_1^{e_1} \circ \cdots \circ f_{i-1}^{e_{i-1}}$ and $v = f_{i+1}^{e_{i+1}} \circ \cdots \circ f_{10}^{e_{10}}$ are the same for both.

Step 4: at most one of each pair preserves A. Each factor of v is either the identity or some f_j with $j > i$, and each of these maps A onto A by Step 2, so $v(A) = A$. Suppose both g_0 and g_1 preserve A. Then $g_0(A) = A$ and $g_1(A) = A$, that is,

$$u(A) = u(v(A)) = A = u(f_i(v(A))) = u(f_i(A)).$$

Since u is injective, $u(A) = u(f_i(A))$ gives $A = f_i(A)$, contradicting the choice of i.

Step 5: count. Each of the 512 pairs contributes at most one exponent vector whose function preserves A. So at most 512 exponent vectors, and hence at most 512 functions of $\mathcal{F}$ (fewer if distinct vectors give the same function), map A to itself. □

The bound is attained. Take $A = \{0\}$, $f_1(n) = n + 1$, $f_2(n) = n - 1$, $f_3 = \cdots = f_{10} = \mathrm{id}$; compositions of f_1 alone reach every positive integer and of f_2 alone every negative one, and $f_1 \circ f_2$ gives 0, so the orbit hypothesis holds. The functions with $e_1 = e_2$ are the identity and preserve A; those with $e_1 \neq e_2$ are shifts by ± 1 and do not. That is exactly 512 exponent vectors.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains find an f_i that moves A, and both write the same pair $u \circ v$, $u \circ f_i \circ v$. Both reach

$$f_i(v(A)) = v(A).$$

At that line the natural attempt is stuck, because $v(A)$ is not A. The model has already arranged, before pairing, that $v(A) = A$: it chose i as the *largest* index that moves A, so everything acting before f_i fixes A.

> *Pairing on e_i leaves me with "f_i preserves $v(A)$", and I need $v(A)$ to be A. The functions in v act first — they are $f_{i+1}, \ldots, f_{10}$. If all of those preserve A, I'm done. So I shouldn't pick any old f_i that moves A; I should pick the last one.*

> *Why must some f_i move A? Because A is finite and the f_i can reach everything from 0. If all of them fixed A setwise, the orbit of 0 would be stuck inside A or stuck outside it.*

Layer 3 THE WHY

Where does "take the largest index" come from? Drive it down to what the composition is.

A function in $\mathcal{F}$ is a word $f_1^{e_1} \circ \cdots \circ f_{10}^{e_{10}}$, and the word is read right to left: the input meets $f_{10}^{e_{10}}$ first. To ask whether g maps A to itself, follow A through the word from the right. Toggling e_i inserts or deletes one letter in the middle. The letters to the left, u, are injective and are applied to both versions, so they cannot tell the two versions apart: $u(X) = u(Y)$ implies $X = Y$. So the

comparison boils down to whether f_i changes the set that arrives at position i, and the set that arrives there is $v(A)$, not A. The whole difficulty is that you do not control $v(A)$.

There are two ways to control it. You could quantify over all possible $v(A)$ — hopeless, since v ranges over 2^{10-i} functions. Or you could make $v(A) = A$ for every v, and that only requires each letter to the right of position i to preserve A. Among the letters that move A, choose the one furthest right; then everything to its right preserves A by construction. This is the extremal principle applied to position in a word: the extreme element has the property that everything beyond it is "good". It is the same choice as picking the first time a process goes wrong: before that moment everything is under control.

Why is the answer a half, with pairing as the mechanism? A cube $\{0,1\}^{10}$ splits into 512 edges in any one fixed direction. "At most half" says: there is a direction i such that no edge in direction i has both endpoints good. That is exactly what Steps 3 and 4 prove. When a problem's bound is 2^{n-1} out of 2^n subsets or sign vectors, look for a coordinate to toggle so that toggling destroys the property.

Why is the orbit hypothesis there? Only to guarantee Step 1. Without it, all ten functions could fix A and all 1024 would preserve it. The hypothesis is stated in a heavy form (every n is reachable from 0), but the proof uses a very small part of it: some point of A and some point outside A are connected by the f_i. That is the meaning of "no finite nonempty set is invariant under all the generators". Note also that the proof does not use inverses of the f_i: only the semigroup orbit of 0 is given, and Step 1 is written to use only that.

There is no official solution in the archive's sources for 1994. The proof above is constructed and verified here, step by step. The one subtlety worth checking twice is Step 4's use of $v(A) = A$: it needs $f_j(A) = A$, not merely $f_j(A) \subseteq A$, and for bijections of $\mathbb{Z}$ and finite A these are the same, by counting.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $512 = 1024/2$ together with the product structure $\mathcal{F} = \{f_1^{e_1} \circ \cdots \circ f_{10}^{e_{10}}\}$. On seeing it, write: *find one exponent e_i such that toggling it always destroys the property.*

2. Pairing along one coordinate of the cube $\{0,1\}^{10}$; the choice of coordinate is made by the extremal principle (the largest i with $f_i(A) \neq A$).

3. Hesitation occurs at $f_i(v(A)) = v(A)$. Reinforce it away by asking what v is made of and choosing i so that v is built from A-preserving letters only.

4. The trap is to pick any f_i that moves A — the pairing then fails at $v(A)$. A second trap is to argue Step 1 with inverses of the f_i, which the hypothesis does not supply; a case split on whether $0 \in A$ avoids them.

5. Slogan: *a "half" bound on 2^n options is a pairing along one coordinate.* And: *in a composition, choose the extremal bad letter so that everything that acts before it is good.*

6. Decision rule: when a pairing argument leaves a residue that depends on the rest of the configuration, go back and choose the paired element extremally (largest, last, first to fail), so that the rest of the configuration is forced into its "good" state.

▷ **Cue:** bound of 2^{n-1} out of 2^n choices of n binary exponents $\longrightarrow$ **Move:** pair along one coordinate so that at most one element of each pair qualifies

▷ **Cue:** a composition $f_1^{e_1} \circ \cdots \circ f_n^{e_n}$ with one letter toggled $\longrightarrow$ **Move:** cancel the injective left part; what remains is the toggled letter acting on the image of the right part

▷ **Cue:** the right part of a word must preserve a set $\longrightarrow$ **Move:** choose the toggled letter to be the last one that fails to preserve it

▷ **Cue:** every integer reachable from 0 by the generators, and a finite nonempty set A $\longrightarrow$ **Move:** some generator does not map A onto itself (split on $0 \in A$ or not)

▷ **Cue:** injective map sending a finite set into itself $\longrightarrow$ **Move:** it maps the set onto itself

Putnam 1994 — Paper B

1994 B1

1994 B1. Find all positive integers n that are within 250 of exactly 15 perfect squares.

Fields: number theory (elementary counting with squares), inequalities **Position:** B1 – opening problem of the second session **Answer:** Exactly the integers with $315 \le n \le 325$ or $332 \le n \le 350$ (reading "within 250" as $|n - k^2| \le 250$ and counting $0 = 0^2$ as a perfect square).

Layer 1 THE WHAT — surface analysis

The natural attempt. You translate the condition into a count. The perfect squares within 250 of n are the k^2 (with $k \ge 0$) lying in the closed window $[n-250,\, n+250]$, so you want

$$\#\{k \ge 0 : n - 250 \le k^2 \le n + 250\} = 15.$$

The instinct from calculus is to approximate: the number of squares in $[a, b]$ is about $\sqrt{b} - \sqrt{a}$, so you set $\sqrt{n+250} - \sqrt{n-250} \approx 15$. Rationalising, $\sqrt{n+250} - \sqrt{n-250} = 500/(\sqrt{n+250} + \sqrt{n-250})$, so $\sqrt{n+250} + \sqrt{n-250} \approx 100/3$, giving $\sqrt{n} \approx 16.7$ and $n \approx 278$. Now you are stuck at the wall this method always hits: the true count is

$$\left\lfloor \sqrt{n+250} \right\rfloor - \left\lceil \sqrt{n-250} \right\rceil + 1,$$

and floors and ceilings can shift the approximation by one in either direction. The estimate $n \approx 278$ is not even inside the answer set (at $n = 278$ the window $[28, 528]$ contains $6^2, \ldots, 22^2$, which is 17 squares). So you begin tabulating $n = 270, 280, 290, \ldots$, computing two square roots each time, hunting for where the count equals 15 and – the part that is easy to get wrong – where it stops equalling 15. You will probably find a block near 320 and a block near 340, and have no proof that there is nothing else, nor that the gap $326 \le n \le 331$ is real.

The model chain. Stop counting as a function of n; count as a function of the squares. If exactly 15 squares lie in the window, they are consecutive squares, since the window is an interval: call them $m^2, (m+1)^2, \ldots, (m+14)^2$, with $m \ge 0$ the smallest root. Exactly-15 is then four inequalities:

(i) $m^2 \ge n - 250$, (ii) $(m+14)^2 \le n + 250$, (iii) $(m-1)^2 < n - 250$ (if $m \ge 1$), (iv) $(m+15)^2 > n + 250$.

Each pair of opposite inequalities eliminates n and bounds m.

From (i) and (ii): $(m+14)^2 - m^2 \le 500$, i.e. $28m + 196 \le 500$, so $m \le 10$.

The case $m = 0$ is impossible: then the window contains $0, 1, 4, \ldots, 196$ and (iv) says $225 > n + 250$, contradicting $n \ge 1$. So $m \ge 1$ and (iii) is in force. From (iii) and (iv): $(m+15)^2 - (m-1)^2 > 500$, i.e. $32m + 224 > 500$, so $m \ge 9$.

Hence $m \in \{9, 10\}$, and for each the four inequalities cut out an interval of n:

$$m = 9: \ \max(23^2 - 250,\ 8^2 + 251) \le n \le \min(9^2 + 250,\ 24^2 - 251), \text{ i.e. } 315 \le n \le 325;$$

$$m = 10: \ \max(24^2 - 250,\ 9^2 + 251) \le n \le \min(10^2 + 250,\ 25^2 - 251), \text{ i.e. } 332 \le n \le 350.$$

(Here $23^2 - 250 = 279$, $8^2 + 251 = 315$, $9^2 + 250 = 331$, $24^2 - 251 = 325$; and $24^2 - 250 = 326$, $9^2 + 251 = 332$, $10^2 + 250 = 350$, $25^2 - 251 = 374$.) Conversely every n in these ranges satisfies (i)–(iv) for the stated m, so has exactly 15 squares in its window. The answer is $315 \le n \le 325$ or $332 \le n \le 350$: thirty integers.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree on the translation "count the k with $n - 250 \le k^2 \le n + 250$". They split at the next line. The natural attempt treats n as the unknown and the count as a function of n, which forces floors of square roots. The model names the fifteen squares – $m^2, \ldots, (m+14)^2$ – and treats m as the unknown, so the condition becomes four polynomial inequalities with no floor anywhere.

> *Counting squares in a window as a function of n means floors of square roots, and I will be off by one somewhere. The count is fixed at 15, so I should name the fifteen squares instead: call the smallest one m^2, and then "exactly 15" is "the first and last are in, their neighbours are out".*

> *The window has width 500 whatever n is. So fifteen consecutive squares must fit in width 500, and seventeen consecutive squares (the fifteen plus both neighbours) must not. That is a two-sided bound on m before I have looked at n at all.*

Layer 3 THE WHY

Why is naming the extreme square the natural move rather than a trick? Go back to what "exactly 15" means as a logical statement. "At least 15" is an existence claim; "at most 15" is a non-existence claim; and a non-existence claim about a run of consecutive objects is always checked at the boundary. If the squares inside an interval are consecutive – and they must be, because an interval contains everything between its members and the squares are increasing in k – then "exactly 15" is equivalent to: the 15 are inside, and the one before and the one after are outside. The whole condition lives on four specific squares, $(m-1)^2, m^2, (m+14)^2, (m+15)^2$. Once you see that, introducing m is not a clever parametrisation; it is writing the definition of "exactly" in its only checkable form.

The second reason is structural. The window $[n - 250, n + 250]$ moves with n but never changes width. A width-500 ruler slides along the squares, whose gaps $2k + 1$ grow steadily. A fixed ruler against a growing spacing means the count decreases (roughly) as the ruler moves right – so there is a narrow band of positions where it reads exactly 15. Eliminating n between (i) and (ii), and between (iii) and (iv), is the algebraic form of "the ruler has width 500". Both eliminations are subtraction of two squares, $(m+a)^2 - (m+b)^2$, which is linear in m because the m^2 cancels. That cancellation is why the method is cheap: difference-of-squares turns every constraint into a linear inequality in m.

The calculus approximation is not wrong; it answers a different question ("roughly where?"). The problem asks "exactly which?", and floor functions are exactly where approximations fail. When a problem says "exactly N" about a discrete count, the approximation can locate the region, but the proof has to be discrete.

One interpretive point deserves stating. "Within 250" is read as $|n - k^2| \le 250$ and 0 is counted as a perfect square (it is 0^2); this is the standard reading and gives the answer above, which I have also checked by direct enumeration over $1 \le n \le 5000$. Were 0 excluded, the case $m = 1$ would reopen: the window for $1 \le n \le 5$ contains exactly $1, 4, \ldots, 225$, fifteen squares, and $1 \le n \le 5$ would join the answer. The model chain handles this in the $m = 0$ line; the reader should know that line is where the convention bites.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the word "exactly" attached to a count of consecutive objects in an interval. On seeing it, write down: "smallest one is m^2, largest is $(m+14)^2$, neighbours $(m-1)^2$ and $(m+15)^2$ are outside."

2. It suggests boundary parametrisation: replace "count as a function of n" by "the extreme members are in, their neighbours are out", and eliminate n between opposite inequalities.

3. Hesitation occurs after the estimate $n \approx 278$, when the count keeps coming out as 16 or 17 and you start doubting the arithmetic. What reinforces it away is the rule that floors cannot be approximated away in an "exactly" question; switch to discrete bookkeeping immediately.

4. The trap is the continuous estimate: it points at $n \approx 278$, which is not an answer, and tempts you into tabulation with no proof of completeness. A second, quieter trap is the gap $326 \le n \le 331$: at $n = 326$ the window is $[76, 576]$, containing 9^2 through 24^2 – sixteen squares – so the answer really is two blocks, not one.

5. Whenever a fixed-width window slides over a sequence with growing gaps and a count is fixed, the problem is about the extreme members of the window. Name them before touching the window's position.

6. The reusable rule: to impose "exactly N consecutive terms lie in an interval", impose "first and last in, both neighbours out"; subtract opposite inequalities to kill the interval's position, and then read off the position range for each surviving case.

▷ **Cue:** a count of squares (or any monotone sequence) inside an interval is fixed at an exact value $\longrightarrow$ **Move:** name the extreme members m^2 and $(m+N-1)^2$ and the two neighbours; write four inequalities

▷ **Cue:** an interval $[n-r, n+r]$ of fixed width slides with the parameter n $\longrightarrow$ **Move:** subtract upper and lower constraints to eliminate n, leaving a constraint on the sequence alone

▷ **Cue:** a difference $(m+a)^2 - (m+b)^2$ appears in a constraint $\longrightarrow$ **Move:** expand: it is linear in m, so the constraint is a linear inequality

▷ **Cue:** a continuous estimate lands on a value that fails the exact condition $\longrightarrow$ **Move:** abandon approximation; the answer set is cut out by discrete inequalities, not by an equation

1994 B2

1994 B2. For which real numbers c is there a straight line that intersects the curve

$$x^4 + 9x^3 + cx^2 + 9x + 4$$

in four distinct points?

Fields: algebra, calculus (Rolle's theorem, convexity, polynomials) **Position:** B2 – second problem of the second session **Answer:** Exactly for $c < \frac{243}{8}$.

Layer 1 THE WHAT — surface analysis

Throughout, write $P(x) = x^4 + 9x^3 + cx^2 + 9x + 4$; the curve is $y = P(x)$.

The natural attempt. A vertical line meets the graph once, so the line is $y = mx + b$, and the intersections are the real roots of

$$Q(x) = P(x) - mx - b = x^4 + 9x^3 + cx^2 + (9-m)x + (4-b).$$

You want: for which c do there exist m, b making Q have four distinct real roots? The first things you reach for are the tools you know for "how many real roots". The discriminant of a quartic is a degree-6 polynomial in its coefficients, with m and b still free, and even when computed it only says "positive means four real or zero real roots". So you try Vieta: roots $r_1, \ldots, r_4$ with

$\sum r_i = -9$ and $\sum_{i<j} r_i r_j = c$. Then $\sum r_i^2 = 81 - 2c$, and real distinct roots make $\sum r_i^2$ positive, giving only $c < 81/2$. That is a necessary condition, not the answer, and there is no clear way to sharpen it. Alternatively, the coefficients $1, 9, c, 9, 4$ look almost palindromic, so you try dividing by x^2 and substituting $u = x + 1/x$ – but $4 \neq 1$, the symmetry is false, and the line terms $mx + b$ break it anyway. You are three techniques in and have a bound, not an answer.

The model chain. Two facts about $Q = P - (mx + b)$ do all the work.

(a) The line is invisible to the second derivative. $Q''(x) = P''(x) = 12x^2 + 54x + 2c$, whatever m and b are.

Necessity. Suppose Q has four distinct real roots $r_1 < r_2 < r_3 < r_4$. By Rolle's theorem Q' vanishes at some $s_1 \in (r_1, r_2)$, $s_2 \in (r_2, r_3)$, $s_3 \in (r_3, r_4)$, which are distinct. Applying Rolle to Q', Q'' vanishes at some $t_1 \in (s_1, s_2)$ and $t_2 \in (s_2, s_3)$, again distinct. So the quadratic $12x^2 + 54x + 2c$ has two distinct real roots: its discriminant is positive,

$$54^2 - 4 \cdot 12 \cdot 2c > 0 \iff 2916 > 96c \iff c < \tfrac{243}{8}.$$

(b) Sufficiency: remove the cubic term. Substitute $x = t - \frac{9}{4}$. Direct expansion gives

$$P\left(t - \tfrac{9}{4}\right) = t^4 + \left(c - \tfrac{243}{8}\right) t^2 + \left(\tfrac{801}{8} - \tfrac{9c}{2}\right) t + \left(\tfrac{81c}{16} - \tfrac{23843}{256}\right).$$

(The t^2 coefficient is $6 \cdot \frac{81}{16} - 27 \cdot \frac{9}{4} + c = c - \frac{243}{8}$.) Write $p = c - \frac{243}{8} < 0$, and let $\alpha t + \beta$ be the linear and constant part. Since t is x plus a constant, $\alpha t + \beta - \frac{p^2}{8}$ is a linear function of x, so it is a line $y = mx + b$. It meets the curve where

$$t^4 + pt^2 = -\tfrac{p^2}{8}.$$

Put $u = t^2$: $u^2 + pu + \frac{p^2}{8} = 0$, so $u = \frac{-p \pm |p|/\sqrt{2}}{2}$. Because $p < 0$, $-p = |p|$ and both values $u = \frac{|p|}{2}\left(1 \pm \frac{1}{\sqrt{2}}\right)$ are distinct and positive. Each gives two values $t = \pm\sqrt{u}$, so there are four distinct real t, hence four distinct x, hence four distinct intersection points.

Both directions together: a line meets the curve in four distinct points if and only if $c < \frac{243}{8}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that the question is "when can $Q = P - mx - b$ have four distinct real roots". They split at which property of Q to interrogate. The natural attempt interrogates the roots directly (discriminant, Vieta, symmetry), where m and b remain tangled in every formula. The model interrogates Q'', the one derivative in which m and b have vanished, and uses Rolle to convert "four roots of Q" into "two roots of Q''".

> *I get to choose the line, so the answer should depend only on the shape of the curve, not on m and b. Which feature of the curve is unchanged when I subtract a line? The second derivative. So the answer is a condition on P''.*

> *Four roots, differentiate twice, Rolle each time: four roots of Q give three of Q' give two of Q''. Q'' is a quadratic I can write down, so its discriminant is the necessary condition.*

> *For the converse I need to exhibit a line. The cubic term is what spoils symmetry, so I translate it away; then $t^4 + pt^2$ is even, and with $p < 0$ it is a W shape. A horizontal line slightly below the central bump cuts a W four times.*

Layer 3 THE WHY

Why is "look at the second derivative" natural rather than inspired? Start from what a line is: the set of functions whose second derivative is zero. Subtracting a line is therefore exactly

the operation that leaves P'' unchanged. The problem quantifies over all lines – "is there a straight line" – and the invariant of that quantifier is P''. Any answer that depends only on c must be readable from what survives every choice of line, and P'' is that survivor. This is the "of course" level the Blueprint asks for: the problem is about lines, lines are the kernel of d^2/dx^2, so differentiate twice.

Geometrically the same fact says: a line can cut a graph four times only if the graph bends one way, then the other, then back – two inflection points, i.e. P'' changes sign twice. When $c \geq \frac{243}{8}$, $P'' \geq 0$ everywhere, the graph is convex, and a convex graph lies above each of its chords between their endpoints, so a line meets it at most twice. Rolle's theorem is the rigorous form of that picture, and it handles the borderline $c = \frac{243}{8}$ cleanly: there P'' has one (double) root, not two, so four roots are impossible.

Why is shifting by $\frac{9}{4}$ natural for the converse? Because $-\frac{9}{4}$ is the midpoint of the two roots of P'' (their sum is $-\frac{54}{12} = -\frac{9}{2}$), i.e. the centre of the two inflection points. Depressing a quartic – substituting $x = t - \frac{a_3}{4}$ to kill the cubic coefficient – is the standard move, and the reason it helps here is sharp: after the shift, everything of degree at most one is "line" and can be absorbed into the choice of m, b. What is left, $t^4 + pt^2$, is a single even function with one parameter, and its shape is readable at sight: $t^2(t^2 + p)$ is 0 at $t = 0$ and dips to $-\frac{p^2}{4}$ at $t = \pm\sqrt{-p/2}$. Any horizontal level strictly between $-\frac{p^2}{4}$ and 0 is crossed four times; $-\frac{p^2}{8}$ is the midpoint, chosen so that the roots come out explicitly.

The Vieta bound $c < \frac{81}{2}$ is worth understanding as a near-miss: it uses only the first two symmetric functions of the roots, which know nothing about whether the roots are real *and* distinct in the right configuration. The second-derivative route uses the real-variable structure (Rolle) that the algebraic route cannot see.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "is there a straight line": the problem quantifies over lines. On seeing it, write down $Q = P - (mx + b)$ and $Q'' = P''$.

2. It suggests counting real roots via Rolle's theorem applied twice, which transfers the condition to the fixed quadratic P''.

3. Hesitation occurs after the necessary condition is found: "is $c < \frac{243}{8}$ also sufficient, and how do I pick the line?" What reinforces it away is the habit of depressing the polynomial: after $x = t - \frac{9}{4}$, the free linear part is absorbed by the line, leaving an even quartic whose graph you can draw.

4. Two traps. The near-palindromic coefficients $1, 9, c, 9, 4$ invite the reciprocal substitution $u = x + 1/x$, which fails because $4 \neq 1$ and because the line's terms break any symmetry. And the quartic discriminant is a trap of cost: huge, with two free parameters, and not even decisive between four and zero real roots.

5. Any time a problem lets you add an arbitrary polynomial of degree $< k$ (a line, a constant, a quadratic), look at the k-th derivative: it is the invariant of the freedom.

6. The reusable rule: n distinct real roots of f force $n - j$ distinct real roots of $f^{(j)}$; use this to turn a question about roots of a family into a question about one fixed derivative, and prove the converse by removing the degree-$(n-1)$ term and choosing the free low-degree part explicitly.

▷ **Cue:** the problem asks whether some line (any $mx + b$) meets a graph in k points $\longrightarrow$ **Move:** subtract the line; the second derivative is unchanged, so the condition is on P'' alone

▷ **Cue:** a polynomial must have n distinct real roots $\longrightarrow$ **Move:** apply Rolle j times: its j-th derivative must have $n - j$ distinct real roots

▷ **Cue:** a quartic with a cubic term and a free linear part $\longrightarrow$ **Move:** substitute $x = t - a_3/4$ to depress it; absorb all degree ≤ 1 terms into the free part

▷ **Cue:** an even quartic $t^4 + pt^2$ with $p < 0$ $\longrightarrow$ **Move:** it is a W with bump 0 and minima $-p^2/4$; a horizontal level strictly between meets it four times

▷ **Cue:** coefficients look palindromic but the end coefficients differ ⟶ **Move:** do not use the reciprocal substitution; the symmetry is false

1994 B3

1994 B3. Find the set of all real numbers k with the following property: For any positive, differentiable function f that satisfies $f'(x) > f(x)$ for all x, there is some number N such that $f(x) > e^{kx}$ for all $x > N$.

Fields: analysis (differential inequalities, integrating factors) **Position:** B3 – third problem of the second session **Answer:** The set is $\{k : k < 1\} = (-\infty, 1)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You test the hypothesis on the functions you know satisfy it. For $f(x) = e^{ax}$, $f' = af > f$ exactly when $a > 1$. Such an f eventually exceeds e^{kx} if and only if $k < a$. Since a can be any number above 1, a k works for all these examples if and only if $k < a$ for every $a > 1$, i.e. $k \leq 1$. So you write down the answer $k \leq 1$.

Then you try to prove it, starting from $f' > f$. You might divide by f (allowed, since $f > 0$) to get $(\ln f)' > 1$, integrate from 0 to x, and obtain $\ln f(x) - \ln f(0) > x$ for $x > 0$, so $f(x) > f(0)e^x$. For $k < 1$ this finishes: $f(0)e^x > e^{kx}$ once $e^{(1-k)x} > 1/f(0)$, i.e. for $x > \ln(1/f(0))/(1-k)$. But at $k = 1$ the inequality $f(x) > f(0)e^x$ is useless when $f(0) < 1$, and you cannot close the argument. The exponential test family has misled you: it never probes $k = 1$, because every member grows strictly faster than e^x.

The model chain. Rewrite the hypothesis exactly. Put $g(x) = e^{-x}f(x)$, which is positive and differentiable. Then

$$g'(x) = e^{-x}\big(f'(x) - f(x)\big),$$

so $f' > f$ everywhere $\iff$ $g' > 0$ everywhere. Thus the functions allowed by the problem are precisely $f = e^x g$ with g positive, differentiable and with $g' > 0$ (in particular strictly increasing).

Every $k < 1$ has the property. For $x > 0$, $g(x) > g(0) > 0$, so $f(x) > g(0)e^x$. Then $f(x) > e^{kx}$ holds whenever $g(0)e^{(1-k)x} > 1$, i.e. for all $x > N := \max\big(0, \frac{-\ln g(0)}{1-k}\big)$.

No $k \geq 1$ has the property. Choose g increasing but bounded by 1: $g(x) = e^{-e^{-x}}$, so

$$f(x) = e^{x - e^{-x}}.$$

This f is positive and differentiable, and $f'(x) = (1 + e^{-x})f(x) > f(x)$ for all x. But $f(x) < e^x \leq e^{kx}$ for every $x \geq 0$ and every $k \geq 1$, so no N can work.

Hence the set is exactly $k < 1$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach $f(x) > f(0)e^x$ (the natural attempt via $(\ln f)' > 1$, the model via g increasing). They split at $k = 1$. The natural attempt, carrying the conjecture $k \leq 1$ from its test family, tries to push the inequality to prove $k = 1$ works. The model asks what the hypothesis permits in full – all $e^x g$ with g increasing – and sees at once that g may be bounded, which kills $k = 1$.

> *$f' > f$ compares f with the solutions of $f' = f$, which are Ce^x. So I should write $f = e^x g$ and see what the inequality says about g. It says exactly $g' > 0$ – nothing more.*

> *My test functions e^{ax} all have $a > 1$ strictly, so none of them says anything about the boundary $k = 1$. The borderline question is: can f grow no faster than e^x? That means g increasing but bounded – and increasing bounded positive functions exist.*

Layer 3 THE WHY

Why is $g = e^{-x}f$ the natural substitution rather than a trick? Because $f' - f$ is a linear differential operator, and the one thing you always know about a linear operator is its kernel: $f' - f = 0$ has solutions Ce^x. The inequality $f' - f > 0$ is a statement that f fails to be in the kernel in a particular direction. The standard way to measure a function against the kernel is variation of parameters: replace the constant C by a function $g(x)$. Then $f' - f = e^x g'$ – the operator acting on $e^x g$ only sees g'. This is the same algebra as the integrating factor for $y' - y = h$ at A-level; the only difference is that h is an unknown positive function rather than a given one. Once you write it, the hypothesis is rewritten as its own simplest equivalent: g is increasing. Nothing is lost and nothing is added, and that is exactly what a clean characterisation should do.

That characterisation also explains the answer's shape. The question "does $e^x g(x)$ eventually beat e^{kx}?" is "does $g(x)$ eventually beat $e^{(k-1)x}$?". If $k < 1$, the right side decays to 0 while g is bounded below by $g(0) > 0$ on $x > 0$: yes. If $k \geq 1$, the right side is at least 1 for $x \geq 0$, so the question is whether an increasing positive g must eventually exceed 1 – and it need not, because increasing functions can be bounded. The counterexample is not inspired; it is the first increasing bounded positive function you can write whose derivative is strictly positive, with "bounded by 1" chosen to match $e^{(k-1)x} \geq 1$. Any of $e^{-e^{-x}}$, $\frac{1}{2} + \frac{1}{\pi}\arctan x$, or $1/(1+e^{-x})$ would serve; I used the first because $f'/f = 1 + e^{-x}$ is then visibly greater than 1.

The natural attempt's error is methodological. A test family can suggest an answer, but at a boundary it has to be designed to probe that boundary. Every e^{ax} with $a > 1$ is strictly above the threshold, so the family cannot distinguish "$k \leq 1$" from "$k < 1$". Whenever a problem's answer has an endpoint, the endpoint must be decided by its own argument – a proof or a counterexample aimed precisely at it.

The problem is not hard once g is written; its difficulty is entirely that the boundary is open, and the obvious examples all say closed.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $f'(x) > f(x)$ – a first-order linear differential inequality. On seeing it, write down $g = e^{-x}f$ and $g' = e^{-x}(f' - f) > 0$.

2. It suggests the integrating factor and a complete characterisation of the admissible f as $e^x \times$ (increasing positive function), after which both directions become questions about increasing functions.

3. Hesitation occurs at $k = 1$: the lower bound $f > f(0)e^x$ nearly works, and the test family says $k = 1$ should be fine. What reinforces it away is asking of every endpoint "which example lives exactly at the boundary?"; here that is g bounded.

4. The trap is the exponential test family e^{ax}, $a > 1$: it gives the closed answer $k \leq 1$, which is wrong, because it never touches the boundary growth rate e^x.

5. A strict inequality in a hypothesis ($f' > f$, not $f' \geq f + \varepsilon$) is a signal that the boundary may be approached but not attained. Look for extremal functions that satisfy it with arbitrarily small margin, such as $f'/f = 1 + e^{-x} \to 1$.

6. The reusable rule: for a linear differential inequality $Lf > 0$, factor out a kernel solution and rewrite it as monotonicity of the quotient; then decide each endpoint of the answer by an example built to sit exactly on it.

▷ **Cue:** a differential inequality of the form $f' > af$ or $f' - af > 0$ $\longrightarrow$ **Move:** set $g = e^{-ax}f$; the hypothesis becomes $g' > 0$

- ▷ **Cue:** a family of test examples all lying strictly on one side of a threshold $\longrightarrow$ **Move:** do not trust it at the endpoint; construct an example at the threshold itself
- ▷ **Cue:** asked whether an increasing positive function must exceed a fixed level $\longrightarrow$ **Move:** no: increasing functions can be bounded, e.g. $e^{-e^{-x}}$ or $1/(1+e^{-x})$
- ▷ **Cue:** "for all $x > N$, $f(x) > e^{kx}$" with $f = e^x g$ $\longrightarrow$ **Move:** divide through: compare g with $e^{(k-1)x}$, which decays iff $k < 1$

1994 B4

1994 B4. For $n \geq 1$, let d_n be the greatest common divisor of the entries of $A^n - I$, where

$$A = \begin{pmatrix} 3 & 2 \\ 4 & 3 \end{pmatrix} \quad \text{and} \quad I = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix}.$$

Show that $\lim_{n\to\infty} d_n = \infty$.

Fields: number theory; linear algebra; Pell equations **Position:** B4 — middle of the paper; a two-line computation once you look at the determinant, and a long detour if you do not **Answer:** Proof: $A^n = \begin{pmatrix} x_n & y_n \\ 2y_n & x_n \end{pmatrix}$ with $x_n^2 - 2y_n^2 = \det A^n = 1$, so $(x_n - 1)(x_n + 1) = 2y_n^2$; splitting this into coprime pieces gives $d_n \geq \sqrt{x_n - 1} \geq \sqrt{3^n - 1}$. In fact d_n is exactly $2p_n$ (n odd) or $2q_n$ (n even), where $(1+\sqrt{2})^n = p_n + q_n\sqrt{2}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Compute. $A - I = \begin{pmatrix} 2 & 2 \\ 4 & 2 \end{pmatrix}$, so $d_1 = 2$. $A^2 = \begin{pmatrix} 17 & 12 \\ 24 & 17 \end{pmatrix}$, so $d_2 = \gcd(16, 12, 24) = 4$. $A^3 = \begin{pmatrix} 99 & 70 \\ 140 & 99 \end{pmatrix}$, so $d_3 = \gcd(98, 70, 140) = 14$. $A^4 = \begin{pmatrix} 577 & 408 \\ 816 & 577 \end{pmatrix}$, so $d_4 = \gcd(576, 408) = 24$. The sequence $2, 4, 14, 24, 82, 140, \ldots$ is growing, but not monotonically in any obvious pattern.

The strong student now reaches for a divisibility structure. As polynomials, $t^n - 1$ divides $t^{mn} - 1$, so $A^{mn} - I = (A^n - I)(A^{n(m-1)} + \cdots + I)$, and every entry of a product of integer matrices is an integer combination of entries of the left factor. Hence $d_n \mid d_{mn}$. Good: along multiples, d never drops.

Then the wall. For n prime, d_n is only guaranteed to be a multiple of $d_1 = 2$. Divisibility along the lattice of divisors says nothing about d_{101} versus d_{100}. A second attempt diagonalises: the eigenvalues are $3 \pm 2\sqrt{2}$, and $A^n - I$ has entries like $\frac{1}{2}\left((3+2\sqrt{2})^n + (3-2\sqrt{2})^n\right) - 1$. Now one needs the gcd of two such irrational-looking expressions, and gcds of closed forms are exactly what closed forms are bad at. The student has size information and no arithmetic information.

The model chain. **Step 1: shape of A^n.** Claim: $A^n = \begin{pmatrix} x_n & y_n \\ 2y_n & x_n \end{pmatrix}$ with x_n, y_n positive integers. True for $n = 1$ ($x_1 = 3, y_1 = 2$). If true for n, then

$$A^{n+1} = A^n A = \begin{pmatrix} 3x_n + 4y_n & 2x_n + 3y_n \\ 6y_n + 4x_n & 4y_n + 3x_n \end{pmatrix},$$

which has the same shape with $x_{n+1} = 3x_n + 4y_n$, $y_{n+1} = 2x_n + 3y_n$. In particular $x_{n+1} \geq 3x_n$, so $x_n \geq 3^n$.

Step 2: the determinant. $\det A = 9 - 8 = 1$, so $\det A^n = 1$:

$$x_n^2 - 2y_n^2 = 1, \qquad \text{i.e.} \qquad (x_n - 1)(x_n + 1) = 2y_n^2. \tag{$*$}$$

The entries of $A^n - I$ are $x_n - 1$, y_n, $2y_n$, $x_n - 1$, so $d_n = \gcd(x_n - 1, y_n)$.
Step 3: split $(*)$ into coprime pieces. Write $x = x_n$, $y = y_n$. From $(*)$, x^2 is odd, so x is odd and $m = \frac{x-1}{2}$, $m+1 = \frac{x+1}{2}$ are consecutive positive integers (positive because $x \geq 3$), hence coprime. Dividing $(*)$ by 4: $m(m+1) = y^2/2$. So y^2 is even, $y = 2z$, and

$$m(m+1) = 2z^2.$$

For an odd prime r, $v_r(2z^2)$ is even and r divides at most one of $m, m+1$, so v_r of each factor is even. The prime 2 appears to the odd power $1 + 2v_2(z)$ in the product and divides only one factor. Therefore one of $m, m+1$ is a perfect square and the other is twice a perfect square: say $\{m, m+1\} = \{u^2, 2v^2\}$ with $u, v \geq 1$, and then $z = uv$, $y = 2uv$.
Step 4: the bound.

- If $m = u^2$: $x - 1 = 2u^2$ and $y = 2uv$, so $2u \mid d_n$, and $2u = \sqrt{2(x-1)}$.
- If $m = 2v^2$: $x - 1 = 4v^2$ and $y = 2uv$, so $2v \mid d_n$, and $2v = \sqrt{x-1}$.

Either way $d_n \geq \sqrt{x_n - 1} \geq \sqrt{3^n - 1} \to \infty$. □

Exact value (for the record). Since $\gcd(u, v) = 1$ and u is odd in the second case (as $u^2 = m + 1 = 2v^2 + 1$), the gcds are exactly $2u$ and $2v$. Matching with $A = B^2$, $B = \begin{pmatrix} 1 & 1 \\ 2 & 1 \end{pmatrix}$, $B^n = \begin{pmatrix} p_n & q_n \\ 2q_n & p_n \end{pmatrix}$, one finds $d_n = 2p_n$ for n odd and $2q_n$ for n even: $2, 4, 14, 24, 82, 140, \ldots$, which matches the table. I checked this identity by machine for $n \leq 39$; the proof above does not depend on it.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains compute A^n and see $d_n = \gcd(x_n - 1, y_n)$. They split at the next line. (N) asks how d_n relates to d_m — structure *across* n. (M) asks what x_n and y_n satisfy *at a single n* — and the answer is $\det A^n = 1$. From there the gcd question is about $x^2 - 1 = 2y^2$, a single Diophantine identity, and one identity with a square on one side factors.

> *I want a common factor of $x - 1$ and y. What equation ties x to y for a fixed n? The determinant: $x^2 - 2y^2 = 1$. So $x^2 - 1 = 2y^2$, and $x^2 - 1 = (x-1)(x+1)$. The factor $x - 1$ I care about is sitting inside a product that equals a square. That is a gcd machine.*

> *Consecutive integers are coprime, and their product is twice a square, so each is a square or twice a square. Then $x - 1$ is about u^2 and y is about $u \cdot v$ — they share u, which is about $\sqrt{x}$.*

▷ **Cue:** a gcd of entries of a matrix power, with $\det = 1$ ⟶ **Move:** write the determinant identity for A^n; it is a Diophantine equation linking exactly the entries you need
▷ **Cue:** $x^2 - 1 = c\,y^2$ (or any $(x-1)(x+1) =$ square-ish) ⟶ **Move:** halve to consecutive integers, use coprimality, conclude each factor is a square up to the constant
▷ **Cue:** a problem asking that a gcd tends to infinity ⟶ **Move:** produce an explicit common factor whose size you can bound below; do not try to compute the gcd exactly

Layer 3 THE WHY

Why is the determinant the natural object and not a lucky one? Because of what gcd is. A gcd of two numbers is invisible to their size; it is visible only to multiplicative relations between them. The natural attempt had plenty of size information (eigenvalues, growth) and one multiplicative relation across different n ($d_n \mid d_{mn}$). What it lacked was a multiplicative relation between $x_n - 1$ and y_n *at the same n*. For a 2×2 matrix power there are exactly two invariants that behave well under powers: the trace (additive, via Cayley–Hamilton) and the determinant

(multiplicative: $\det A^n = (\det A)^n$). Only one of them is an equation in the entries that holds at every n with a fixed right-hand side. Given $\det A = 1$, that equation is $x_n^2 - 2y_n^2 = 1$ for all n. So once you phrase the question as "what relation do the entries satisfy", the determinant is the only candidate.

Why does the factorisation then close the problem? Because $x^2 - 1 = 2y^2$ says a *product* is (twice) a square, and the fundamental theorem of arithmetic distributes a square across coprime factors. The coprimality is free, since $\frac{x-1}{2}$ and $\frac{x+1}{2}$ differ by one. The same move proves that every Pythagorean triple has the form $k(a^2 - b^2, 2ab, a^2 + b^2)$: a product of nearly coprime things is a square, therefore each thing is. So the step is a standard piece of machinery, not a trick.

Why does the bound come out as $\sqrt{x_n}$? $x - 1 \approx 2u^2$ and $y = 2uv$ share the factor u, and u is roughly the square root of x because $x - 1$ is a constant times u^2. Any factorisation of the form "$x - 1$ is a square times a constant, and y is a product of the two square roots" yields a common factor of size about $\sqrt{x}$. That is the whole mechanism.

There is a second viewpoint worth owning. $3 + 2\sqrt{2} = (1 + \sqrt{2})^2$, and correspondingly $A = B^2$ with $B = \begin{pmatrix} 1 & 1 \\ 2 & 1 \end{pmatrix}$, $\det B = -1$. Then $A^n - I = B^n (B^n - B^{-n})$, and $B^n - B^{-n}$ is $2q_n \begin{pmatrix} 0 & 1 \\ 2 & 0 \end{pmatrix}$ for n even and $2p_n I$ for n odd. This is the matrix form of $\lambda^{2n} - 1 = \lambda^n(\lambda^n - \lambda^{-n})$: when A has a square root, $A^n - I$ factors as a difference of two powers. It is the "u, v" of Step 3 seen from above — the u and v are p_n and q_n. You do not need it, but if you recognise $3 + 2\sqrt{2}$ as a unit square in $\mathbb{Z}[\sqrt{2}]$ it hands you the exact answer.

Layer $+\alpha$ BACKTESTING

1. The first cue is "greatest common divisor of the entries" of a power of a matrix with $\det A = 9 - 8 = 1$. On seeing it, write $\det A^n = 1$ as an equation in the entries of A^n.

2. Diophantine factorisation: $x^2 - 1 = 2y^2$, split $(x - 1)(x + 1)$ into coprime halves, conclude square / twice-square.

3. Hesitation comes after the table $2, 4, 14, 24$: the student sees growth but no pattern and starts hunting for a recursion for d_n. Reinforce it away with the question "what equation holds at a single n?"

4. The trap is working across n (divisibility $d_n \mid d_{mn}$, recursions), which is true and useless for prime n. A second trap is trying to compute d_n exactly: the problem asks only for a lower bound tending to infinity.

5. When a problem asks about the arithmetic (gcd, divisibility) of the entries of A^n, the determinant is an arithmetic identity in those entries for free. Look for it before looking for recursions.

6. Reusable decision rule: *to show a gcd of two quantities is large, find one polynomial identity linking them; if it says a product is a square (up to a constant), split into coprime factors and read the common factor off the square roots.*

1994 B5

1994 B5. For any real number α, define the function $f_\alpha(x) = \lfloor \alpha x \rfloor$. Let n be a positive integer. Show that there exists an α such that for $1 \le k \le n$,

$$f_\alpha^k(n^2) = n^2 - k = f_{\alpha^k}(n^2).$$

Fields: analysis; inequalities; floor functions **Position:** B5 — late on the paper but not deep; the difficulty is choosing to commit to one explicit α instead of intersecting n intervals **Answer:** $\alpha = 1 - \frac{1}{n^2}$

works. The iteration drops by exactly 1 each step because $0 < m/n^2 \le 1$ for $m \le n^2$; and $n^2 - k \le n^2(1-1/n^2)^k \le n^2 - k + \dfrac{k(k-1)}{2n^2} < n^2 - k + 1$ by first- and second-order Bernoulli.

Layer 1 THE WHAT — surface analysis

The natural attempt. Read the two conditions separately and turn each into a condition on α.

Iteration. $f_\alpha^k(n^2) = n^2 - k$ for all $k \le n$ means: starting from $m = n^2$, each application of f_α lowers m by exactly one, for $m = n^2, n^2-1, \dots, n^2-n+1$. So $m - 1 \le \alpha m < m$, i.e.

$$1 - \tfrac{1}{m} \le \alpha < 1 \quad \text{for } n^2 - n + 1 \le m \le n^2.$$

The strongest of these is $m = n^2$: $\alpha \in \left[1 - \frac{1}{n^2}, 1\right)$.

Powers. $\lfloor \alpha^k n^2 \rfloor = n^2 - k$ means $n^2 - k \le \alpha^k n^2 < n^2 - k + 1$, i.e.

$$\left(1 - \tfrac{k}{n^2}\right)^{1/k} \le \alpha < \left(1 - \tfrac{k-1}{n^2}\right)^{1/k}.$$

So one needs the intersection of $n+1$ intervals with k-th-root endpoints to be nonempty. The student now tries to order the endpoints. Are the lower endpoints increasing or decreasing in k? Compare $(1-k/N)^{1/k}$ with $(1-(k+1)/N)^{1/(k+1)}$ by raising to the power $k(k+1)\dots$ this is a real inequality between two unrelated-looking expressions, and it has to be done for the upper endpoints too, and then each lower endpoint must be compared with each upper endpoint. That is $O(n^2)$ comparisons of k-th roots. The approach is correct in principle and a swamp in practice.

The model chain. Let $N = n^2$ and $\varepsilon = \frac{1}{N}$, and set $\alpha = 1 - \varepsilon = 1 - \frac{1}{n^2}$. (For $n = 1$ this is $\alpha = 0$, and both sides equal 0 at $k = 1$; the argument below covers it.)

Step 1: the iteration. For an integer m with $1 \le m \le N$,

$$\alpha m = m - \frac{m}{N}, \qquad 0 < \frac{m}{N} \le 1,$$

so $m - 1 \le \alpha m < m$ and $f_\alpha(m) = m - 1$. Starting at N and applying this $k \le n$ times (every value passed through lies in $[N-n+1, N] \subseteq [1, N]$) gives $f_\alpha^k(N) = N - k$.

Step 2: two Bernoulli bounds. For $0 \le \varepsilon \le 1$ and every integer $k \ge 1$,

$$1 - k\varepsilon \;\le\; (1-\varepsilon)^k \;\le\; 1 - k\varepsilon + \binom{k}{2}\varepsilon^2.$$

Both by induction on k, with equality at $k = 1$. Multiply the claim for k by $1 - \varepsilon \ge 0$. Lower: $(1-\varepsilon)^{k+1} \ge (1-k\varepsilon)(1-\varepsilon) = 1 - (k+1)\varepsilon + k\varepsilon^2 \ge 1 - (k+1)\varepsilon$. Upper:

$$(1-\varepsilon)^{k+1} \le \left(1 - k\varepsilon + \tbinom{k}{2}\varepsilon^2\right)(1-\varepsilon) = 1 - (k+1)\varepsilon + \left(\tbinom{k}{2} + k\right)\varepsilon^2 - \tbinom{k}{2}\varepsilon^3 \le 1 - (k+1)\varepsilon + \tbinom{k+1}{2}\varepsilon^2.$$

Step 3: the powers. Multiply Step 2 by N, with $\varepsilon N = 1$:

$$N - k \;\le\; \alpha^k N \;\le\; N - k + \frac{k(k-1)}{2N}.$$

For $1 \le k \le n$, $k(k-1) \le n(n-1) < 2n^2 = 2N$, so the right side is strictly less than $N - k + 1$. Hence $\lfloor \alpha^k N \rfloor = N - k$, i.e. $f_{\alpha^k}(n^2) = n^2 - k$.

Steps 1 and 3 together give $f_\alpha^k(n^2) = n^2 - k = f_{\alpha^k}(n^2)$ for $1 \le k \le n$. □

I checked this α with exact rational arithmetic for every $n \le 59$ and every $k \le n$. With this α the power condition first fails near $k \approx \sqrt{2}\,n$ (at $k = 8$ for $n = 5$, $k = 15$ for $n = 10$, $k = 44$ for $n = 30$), which is exactly where $\frac{k(k-1)}{2n^2}$ reaches 1: the statement's range $k \le n$ sits comfortably inside.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains translate the two conditions into constraints on α and both find, from $k = 1$, that $\alpha \in [1 - \frac{1}{n^2}, 1)$. They split there. (N) keeps every constraint as an interval and tries to intersect them all. (M) stops and *picks* the left endpoint $\alpha = 1 - \frac{1}{n^2}$, then checks the remaining constraints for that single number. Checking a number is an inequality; intersecting intervals is a family of inequalities between inequalities.

> *The $k = 1$ case already pins α to $[1 - 1/n^2, 1)$, and the problem only asks that* some *α exists. So I'll stop solving for α and test the most natural candidate, the endpoint $1 - 1/n^2$. Existence proofs want a witness, not a region.*

> *$\alpha^k n^2 = n^2(1 - 1/n^2)^k$ should be about $n^2 - k$ — that is Bernoulli. The question is only whether the error stays under 1. The error is the second-order term, $\binom{k}{2}/n^2$, and $k \le n$ makes that less than $\frac{1}{2}$.*

- ▷ **Cue:** "show there exists α" with many constraints, one of which already pins α to a short interval $\longrightarrow$ **Move:** commit to an endpoint or natural point of that interval and verify the others; do not intersect the constraints
- ▷ **Cue:** $(1 - \varepsilon)^k$ must be located within an interval of width ε around $1 - k\varepsilon$ $\longrightarrow$ **Move:** use Bernoulli both ways: $1 - k\varepsilon \le (1 - \varepsilon)^k \le 1 - k\varepsilon + \binom{k}{2}\varepsilon^2$
- ▷ **Cue:** a floor iteration $\lfloor \alpha m \rfloor$ that must decrease by exactly one $\longrightarrow$ **Move:** write $\alpha = 1 - \varepsilon$; the condition is $0 < \varepsilon m \le 1$

Layer 3 THE WHY

Why commit to $\alpha = 1 - 1/n^2$? Two reasons, one logical and one structural.

The logical one: the goal is $\exists \alpha$. An existence statement is proved by a witness, and the natural attempt forgot that the problem never asked for the set of all good α. Intersecting all the constraints is what you do when asked to *characterise*; when asked to *exhibit*, you pick the point the most binding constraint suggests and see whether it survives. This is the same rule as "to show a system of inequalities is consistent, try the boundary case first". The endpoint is where the $k = 1$ condition holds with equality, so it is the only candidate that the problem itself names.

The structural one explains why this candidate works. Put $\alpha = 1 - \varepsilon$. The iteration condition says each step loses exactly one unit, which requires the loss εm to be in $(0, 1]$ for every m visited; the largest m visited is n^2, so $\varepsilon \le 1/n^2$ is forced. The power condition says $n^2(1 - \varepsilon)^k$ loses between $k - 1$ (exclusive) and k (inclusive) units. To first order $n^2(1 - \varepsilon)^k$ loses $kn^2\varepsilon$ units, so to first order we need $n^2\varepsilon = 1$ exactly — and that is the endpoint again. Both conditions point at the same number. Everything left is the second-order correction, and whether it is less than one unit.

That also explains the peculiar numbers in the statement. Why start at n^2 and only go to $k = n$? The second-order error is $\binom{k}{2}\varepsilon^2 \cdot n^2 = \binom{k}{2}/n^2$, which stays below 1 while $k \lesssim \sqrt{2}\,n$. The problem-setter chose the starting value to be the square of the number of steps so that the linear approximation to $(1 - \varepsilon)^k$ is accurate to within one unit over the whole range — with a margin, since the error at $k = n$ is under $\frac{1}{2}$. The computer check confirms the threshold: the first failure is at $k \approx \sqrt{2}n$. So "n^2" in the statement is a cue: it tells you that a quadratic error term is the thing being controlled.

Why is the two-sided Bernoulli the right tool rather than, say, ln and e^{-x} estimates? Because the target is an *integer* window of width one, and the lower bound must hold with the equality case $k = 1$ allowed (at $k = 1$, $\alpha n^2 = n^2 - 1$ exactly). Polynomial bounds in ε keep that equality exact; exponential estimates blur it into strict inequalities that you then have to repair.

Layer $+\alpha$ BACKTESTING

1. The first cue is the $k = 1$ instance: $\lfloor \alpha n^2 \rfloor = n^2 - 1$, which pins $\alpha \in [1 - \frac{1}{n^2}, 1)$. On seeing it, write $\alpha = 1 - \varepsilon$ with $\varepsilon \le 1/n^2$, and try $\varepsilon = 1/n^2$.
2. Exhibit-and-verify with Bernoulli's inequality in both directions.
3. Hesitation comes when the student has n interval constraints and starts comparing k-th roots. It is reinforced away by rereading the quantifier: "there exists an α".
4. The trap is characterising the set of admissible α — correct, and far more work than the problem requires. A smaller trap is a strict lower bound: at $k = 1$ the value αn^2 lands exactly on the integer $n^2 - 1$, so the lower estimate must be non-strict.
5. Watch for the shape "something at scale N, a number of steps at scale $\sqrt{N}$". It signals that a first-order approximation is exact to within one unit because the second-order error is k^2/N.
6. Reusable decision rule: *for an existence claim with many constraints, let the most binding constraint name the witness (usually its boundary), then verify the rest by a two-sided expansion whose error term the problem's scaling makes smaller than the tolerance.*

1994 B6

1994 B6. For any integer n, set

$$n_a = 101a - 100 \cdot 2^a.$$

Show that for $0 \le a, b, c, d \le 99$, $n_a + n_b \equiv n_c + n_d \pmod{10100}$ implies $\{a, b\} = \{c, d\}$.

Fields: number theory; Chinese remainder theorem; primitive roots; Vieta **Position:** B6 — the last problem, but short once the modulus is read correctly; the hard part is seeing that the two residues carry a sum and a product **Answer:** Proof: $n_a \equiv a \pmod{100}$ and $n_a \equiv 2^a \pmod{101}$. So $a + b \equiv c + d \pmod{100}$, and since 2 has order 100 modulo 101, $2^a 2^b \equiv 2^c 2^d$ as well as $2^a + 2^b \equiv 2^c + 2^d \pmod{101}$. Equal sum and product force $\{2^a, 2^b\} = \{2^c, 2^d\}$ in $\mathbb{F}_{101}$, hence $\{a, b\} = \{c, d\}$.

(The phrase "for any integer n" in the statement is as printed; the variable being defined is a, and n_a is a name, not n times anything.)

Layer 1 THE WHAT — surface analysis

The natural attempt. $10100 = 100 \cdot 101$ with $\gcd(100, 101) = 1$, so the strong student splits by the Chinese remainder theorem — that part is reflex. Modulo 100: $101a \equiv a$ and $100 \cdot 2^a \equiv 0$, so $n_a \equiv a$ and the hypothesis gives

$$a + b \equiv c + d \pmod{100}. \tag{1}$$

Modulo 101: $101a \equiv 0$ and $-100 \equiv 1$, so $n_a \equiv 2^a$ and

$$2^a + 2^b \equiv 2^c + 2^d \pmod{101}. \tag{2}$$

Now the student works on (2). Rearranged: $2^a - 2^c \equiv 2^d - 2^b$. Factor out the smaller powers: $2^{\min(a,c)}(2^{|a-c|} - 1) \equiv \pm 2^{\min(b,d)}(2^{|b-d|} - 1)$. This leads to a case analysis on which exponent is smallest, and on whether $2^j - 1$ can equal $2^i(2^l - 1)$ modulo 101 — and it can, for many triples, because modulo 101 nothing stops $2^j - 1$ from being any residue. Equation (2) on its own has plenty of solutions with $\{a, b\} \ne \{c, d\}$: it is one equation in a field with 101 elements and four unknowns. The student then tries to feed (1) in by substituting $d \equiv a + b - c$, getting

$$2^a + 2^b \equiv 2^c + 2^{a+b-c} \pmod{101},$$

and stares at it. That is the wall: the student has the right equation and does not see it as a factorisation.

The model chain. **Step 1: CRT split.** As above, $n_a \equiv a \pmod{100}$ and $n_a \equiv 2^a \pmod{101}$, and since $10100 = 100 \cdot 101$ with coprime factors, the hypothesis is equivalent to (1) and (2) together. **Step 2: 2 has order 100 modulo 101.** 101 is prime, so by Fermat $2^{100} \equiv 1$, and the order of 2 divides $100 = 2^2 \cdot 5^2$. If the order were a proper divisor, it would divide 50 or 20. Compute: $2^{10} = 1024 = 10 \cdot 101 + 14 \equiv 14$; $2^{20} \equiv 196 \equiv 95 \not\equiv 1$. And $2^{50} = 2^{20} \cdot 2^{20} \cdot 2^{10} \equiv 95 \cdot 95 \cdot 14$; $95 \equiv -6$, so this is $36 \cdot 14 = 504 = 4 \cdot 101 + 100 \equiv -1 \not\equiv 1$. So the order is exactly 100. (Alternatively: $101 \equiv 5 \pmod 8$, so 2 is a quadratic non-residue and $2^{50} \equiv -1$ by Euler's criterion.)

Two consequences. (i) $2^x \equiv 2^y \pmod{101}$ iff $x \equiv y \pmod{100}$; in particular $a \mapsto 2^a$ is injective on $\{0, \ldots, 99\}$. (ii) From (1), $2^{a+b} \equiv 2^{c+d}$, that is,

$$2^a \cdot 2^b \equiv 2^c \cdot 2^d \pmod{101}. \tag{3}$$

Step 3: Vieta in $\mathbb{F}_{101}$. Put $X = 2^a$, $Y = 2^b$, $U = 2^c$, $V = 2^d$, as elements of the field $\mathbb{F}_{101}$. By (2) and (3), $X + Y = U + V$ and $XY = UV$, so

$$(t - X)(t - Y) = t^2 - (X + Y)t + XY = t^2 - (U + V)t + UV = (t - U)(t - V)$$

as polynomials over $\mathbb{F}_{101}$. Substituting $t = U$: $(U - X)(U - Y) = 0$, and a field has no zero divisors, so $U = X$ or $U = Y$. If $U = X$, then $V = (U + V) - U = (X + Y) - X = Y$; if $U = Y$, then $V = X$. Either way $\{U, V\} = \{X, Y\}$ as multisets.
Step 4: back to exponents. By (i), $2^c = 2^a$ with $0 \le a, c \le 99$ forces $c = a$, and so on. Hence $\{c, d\} = \{a, b\}$. □

I also confirmed by machine that the 5050 values $n_a + n_b$ ($0 \le a \le b \le 99$) are pairwise distinct modulo 10100, and that $2^{20} \equiv 95$, $2^{50} \equiv 100 \pmod{101}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach (1) and (2). They split at what to do with (1). (N) uses (1) *additively* — to eliminate d — and is left with a single exponential equation. (M) *exponentiates* (1): because 2 has order 100 modulo 101, a congruence of exponents modulo 100 is a congruence of powers modulo 101. Then (1) becomes a statement about the product $2^a 2^b$, and (2) is a statement about the sum. Sum and product of two numbers determine the pair.

> *The modulus is* $100 \cdot 101$, *and* $100 = 101 - 1$. *That is not a coincidence:* 100 *is the exponent that Fermat attaches to* 101. *So the mod-*100 *information* $a + b \equiv c + d$ *is really information about exponents of* 2 *modulo* 101.

> *Exponentiate (1):* $2^a 2^b \equiv 2^c 2^d$. *And (2) says* $2^a + 2^b \equiv 2^c + 2^d$. *I know the sum and the product of* $\{2^a, 2^b\}$ *— they are the roots of one quadratic. So the pair is determined.*

- ▷ **Cue:** a modulus of the form $(p - 1) \cdot p$, or a number congruent to x mod $p - 1$ and to g^x mod p ⟶ **Move:** treat the mod-$(p - 1)$ part as an exponent: $x \equiv y \pmod{p - 1} \Rightarrow g^x \equiv g^y \pmod p$
- ▷ **Cue:** two unknowns whose sum and product are both known (in a field) ⟶ **Move:** they are the two roots of $t^2 - st + p$; the unordered pair is determined
- ▷ **Cue:** a formula $A \cdot x + B \cdot y$ with $A \equiv 1, B \equiv 0$ modulo one factor and $A \equiv 0, B \equiv 1$ modulo the other ⟶ **Move:** it is a CRT encoding of the pair (x, y); read off each component separately
- ▷ **Cue:** "2 has order 100 modulo 101" ⟶ **Move:** check $2^{20} \not\equiv 1$ and $2^{50} \not\equiv 1$ (the maximal proper divisors of 100), or use $p \equiv \pm 3 \pmod 8 \Rightarrow 2$ is a non-residue

Layer 3 THE WHY

Why is exponentiating (1) natural and not a trick? Look at how n_a is built. The coefficient 101 is $\equiv 1 \pmod{100}$ and $\equiv 0 \pmod{101}$; the coefficient -100 is $\equiv 0 \pmod{100}$ and $\equiv 1 \pmod{101}$. So $n_a = 101 \cdot a + (-100) \cdot 2^a$ is precisely the Chinese-remainder encoding of the pair $(a \bmod 100,\ 2^a \bmod 101)$. The problem-setter has glued an exponent to its power. A residue modulo $p-1$ sitting next to a residue modulo p is the discrete logarithm sitting next to the exponential. Once you read n_a that way, (1) and (2) are not two unrelated facts; they are the same pair of numbers $\{2^a, 2^b\}$ seen through two lenses — a multiplicative one (the exponents add, so the powers multiply) and an additive one.

Why does "sum and product" finish it? Because the unordered pair $\{X, Y\}$ is exactly the information in the polynomial $(t-X)(t-Y)$, and a monic quadratic is determined by its two coefficients, which are the sum and product. That is Vieta read backwards. The only structural fact needed is that $\mathbb{F}_{101}$ is a field (a product of nonzero elements is nonzero), which is why the problem uses the prime 101 and not a composite: with zero divisors, a quadratic can have more than two roots, and the argument would fail. So the prime modulus and the order-100 generator are exactly the two hypotheses the proof consumes.

Why is (2) alone hopeless? One congruence in a 101-element field cannot pin down a pair out of 5050; there are on average fifty pairs per residue. Unordered pairs need two independent symmetric functions, and the problem supplies the second one in disguise modulo 100. The natural attempt had both and used one of them to eliminate a variable instead of as a symmetric function. Elimination destroys symmetry; the quantity $\{a, b\}$ in the conclusion is symmetric; so the working should stay symmetric.

Finally, why the range $0 \le a, b, c, d \le 99$? Because 2^x modulo 101 has period 100, so exponents are defined only modulo 100, and the conclusion $\{a, b\} = \{c, d\}$ is about honest integers. The range is exactly one period. The statement is a disguised version of: *the map* $\{a, b\} \mapsto (a + b \bmod 100,\ 2^a + 2^b \bmod 101)$ *is injective on unordered pairs from* $\mathbb{Z}/100$ — a Sidon-type property built from the exponential map.

Layer $+\alpha$ BACKTESTING

1. The first cue is the modulus $10100 = 100 \cdot 101$ with $100 = 101 - 1$, together with a formula containing both a and 2^a. On seeing it, write $n_a \equiv a \pmod{100}$ and $n_a \equiv 2^a \pmod{101}$, and note that 100 is the Fermat exponent for 101.

2. CRT split, then convert the mod-100 exponent information into a product modulo 101 via the order of 2, then Vieta.

3. Hesitation comes at the pair (1), (2): the student uses (1) to substitute for d. Reinforce with the question "the conclusion is symmetric in a, b; which symmetric functions of $2^a, 2^b$ do I know?"

4. The trap is additive elimination, which yields a single exponential congruence with many spurious solutions. A second trap is assuming 2 is a primitive root modulo 101 without checking; the proof needs order exactly 100 for both directions of Step 2.

5. Recognise a CRT-encoded pair: coefficients that are 1 and 0 modulo complementary factors. Then ask what relation ties the two components — here, one is the logarithm of the other.

6. Reusable decision rule: *to show an unordered pair is determined by some data, extract two symmetric functions of the pair (typically sum and product) from the data and invoke Vieta in a field; if one of them is hidden in the exponent, a generator of known order converts it.*