DAWN PUTNAM SESSION

Putnam 2018 — Papers A and B

Putnam 120 Archive

Contents

Putnam 2018 — Paper A **5**
2018 A1 5
2018 A2 7
2018 A3 10
2018 A4 13
2018 A5 16
2018 A6 19

Putnam 2018 — Paper B **23**
2018 B1 23
2018 B2 25
2018 B3 28
2018 B4 30
2018 B5 34
2018 B6 37

Putnam 2018 — Paper A

2018 A1

2018 A1. Find all ordered pairs (a, b) of positive integers for which

$$\frac{1}{a} + \frac{1}{b} = \frac{3}{2018}.$$

Fields: Number theory (Diophantine equations, factoring by completing the rectangle, divisor counting) **Position:** A1 – the opener **Answer:** Six ordered pairs: $(673, 1358114)$, $(674, 340033)$, $(1009, 2018)$, $(2018, 1009)$, $(340033, 674)$, $(1358114, 673)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The equation has one obvious first move: clear denominators. Multiplying through by $2018ab$ gives

$$2018(a + b) = 3ab,$$

and from here the natural instinct is to isolate one variable. Solving for b,

$$b = \frac{2018a}{3a - 2018},$$

and now the problem is: for which positive integers a is this a positive integer? You can get a bound on a immediately: since $1/a < 3/2018$ (because $1/b > 0$), $a > 2018/3 = 672.67\ldots$, so $a \geq 673$; and by symmetry, taking $a \leq b$ WLOG, $2/a \geq 3/2018$, so $a \leq 1345.3\ldots$. That pins a to the range 673 to 1345 – about 673 candidate values. For each you would need to check whether $3a - 2018$ divides $2018a$, which is a real divisibility test, not something you can eyeball. This is not an unreasonable plan; it terminates, and if you had a calculator you would find the answer. But under exam conditions, checking a range of nearly 700 integers by hand for a divisibility condition is not a method, it is a threat – the problem has been reduced to a search rather than solved.

The model chain. Return to $2018(a + b) = 3ab$, i.e. $3ab - 2018a - 2018b = 0$. Multiply by 3 to make the coefficient on ab a perfect square coefficient:

$$9ab - 3 \cdot 2018\, a - 3 \cdot 2018\, b = 0.$$

This is exactly the cross-term expansion of $(3a - 2018)(3b - 2018)$, missing only the constant: $(3a - 2018)(3b - 2018) = 9ab - 3 \cdot 2018a - 3 \cdot 2018b + 2018^2$. So adding 2018^2 to both sides of the equation above,

$$(3a - 2018)(3b - 2018) = 2018^2.$$

The search over a has vanished. What remains is a finite, structured problem: factor 2018^2 and pair its divisors. Since $2018 = 2 \times 1009$ and 1009 is prime, $2018^2 = 2^2 \cdot 1009^2$, which has $(2 + 1)(2 + 1) = 9$ positive divisors and 9 corresponding negative ones.

Now note a constraint that costs nothing extra: for any integer a, $3a - 2018 \equiv -2018 \pmod 3$, and since $2018 = 3 \cdot 672 + 2$, $-2018 \equiv -2 \equiv 1 \pmod 3$. So *every* value $3a - 2018$ that can ever arise is automatically congruent to 1 mod 3 – this is not a condition to impose, it is a filter that the divisor $d = 3a - 2018$ must already pass before $a = (d + 2018)/3$ can even be an integer. Writing $2018^2 = 2^i \cdot 1009^j$ with $i \in \{0,1,2\}$, $j \in \{0,1,2\}$, and using $1009 \equiv 1 \pmod 3$ and $2 \equiv 2, 4 \equiv 1 \pmod 3$, a divisor is $\equiv 1 \pmod 3$ exactly when $i \in \{0,2\}$. That is 6 positive divisors: $1, 4, 1009, 4 \cdot 1009, 1009^2, 4 \cdot 1009^2$. Each gives a pair $(d, 2018^2/d)$, hence $a = (d + 2018)/3$, $b = (2018^2/d + 2018)/3$, both automatically positive integers. The three negative divisors $\equiv 1 \pmod 3$ (namely $-2, -2 \cdot 1009, -2 \cdot 1009^2$) each force one of a, b negative, since a negative d close to 0 makes $2018^2/d$ a large-magnitude negative number, and $(2018^2/d + 2018)/3 < 0$. These are discarded. The six surviving divisors give the six ordered pairs listed above.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same cleared-denominator equation $2018(a + b) = 3ab$. They split at the very next step: the natural attempt isolates b as a function of a and turns the problem into a bounded search over a; the model instead multiplies by 3 and completes $9ab - 3 \cdot 2018a - 3 \cdot 2018b$ into the product $(3a - 2018)(3b - 2018)$ by adding 2018^2 to both sides, converting the search into a factorization.

> *The equation has a single ab term together with linear terms in a and b – that is exactly the shape of a product of two linear expressions minus a constant, so instead of solving for one variable in terms of the other, I should complete the rectangle: multiply through until the ab-coefficient is a perfect square, add back whatever constant makes the left side factor as (linear in a)(linear in b), and turn a search over one variable into a finite pairing of divisors of a fixed number.*

> *Before I even factor* 2018^2*, I should ask what values* $3a - 2018$ *can possibly take as a ranges over the integers – it is always* $\equiv 1 \pmod 3$ *regardless of which a I pick, so this is a constraint on which divisors of* 2018^2 *are even eligible to appear, not something I need to check case by case after the fact.*

- ▷ **Cue:** a Diophantine equation has exactly one bilinear term ab together with linear terms in a and b (no a^2 or b^2) $\longrightarrow$ **Move:** multiply through so the ab-coefficient becomes a perfect square, then add the matching constant to complete $(\alpha a + \beta)(\alpha b + \gamma)$, turning the equation into a product equal to a fixed number
- ▷ **Cue:** a factored equation $(\text{expr in } a)(\text{expr in } b) = N$ has been reached $\longrightarrow$ **Move:** stop thinking in terms of a; switch to thinking in terms of the divisors of N, since every solution corresponds to exactly one ordered factor pair of N

Layer 3 THE WHY

Why is completing the rectangle the forced move, not a lucky guess? Because of what the equation $3ab - 2018a - 2018b = 0$ *is*: a bilinear form in a and b plus linear corrections, and the algebraic identity $(\alpha a + \beta)(\alpha b + \gamma) = \alpha^2 ab + \alpha\gamma a + \alpha\beta b + \beta\gamma$ is the unique way to write a bilinear-plus-linear expression as a product of two linear factors. Any equation of the shape $pab + qa + rb = s$ can be forced into this mould by multiplying through by p and choosing $\beta = r, \gamma = q$ appropriately; this is the same completing-the-square identity one level up, applied to a cross term instead of a squared term. Once the equation reads $(\text{linear})(\text{linear}) = N$ with N a fixed integer, the problem stops being about the real-number function $b(a)$ and becomes about the finite, discrete set of factorizations of N – and a finite set is precisely what "find all ordered pairs" is asking to enumerate. Searching over the bounded range of a works too, in the sense that it eventually terminates, but it treats a as a continuous parameter to be swept, when the equation's true content is arithmetic (about the divisors of 2018^2), not analytic (about the shape of a curve $b(a)$).

The mod-3 observation deserves the same scrutiny: it looks like an extra filter applied after factoring, but it is in fact structural. $3a - 2018$ is built by adding a fixed residue $-2018 \bmod 3$ to a multiple of 3; no choice of integer a can change that residue. Recognising this in advance means the divisor list to check is 6 candidates, not 18 – the congruence is discovered by the form of the substitution, not verified afterwards by trial.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is seeing a single ab term with linear terms in a and b and no squared terms – this should immediately be flagged, on first read, as "completable into a product of two linear factors," before any attempt to solve for one variable.
2. That cue suggests multiplying through by the coefficient needed to make the ab-term a perfect square, then adding the constant that completes the product – converting the equation into $(\text{linear})(\text{linear}) = N$.
3. Hesitation occurs right after clearing denominators, when isolating $b = b(a)$ feels like the obvious next step because it is the standard technique for "one equation, two unknowns." It is reinforced away by asking what kind of object the equation actually describes: with a bilinear term present, it is not a curve to be swept over a, it is a factorization waiting to be exposed.
4. The trap is the false sense of progress from bounding a (a genuine, correct bound) and mistaking "the search is finite" for "the search is tractable" – a bound of 673 candidate values is finite but not something a human enumerates by hand under time pressure.
5. Recognition would be faster next time by treating "$1/a + 1/b = k/N$" and its relatives as a named family: reciprocal Diophantine equations always clear to a bilinear-plus-linear form, and that form always completes into $(pa - N)(pb - N) = N^2$ (or the analogous constant) for some integer p built from the coefficients.
6. The reusable decision rule: whenever a Diophantine equation contains exactly one product term xy together with terms linear in x and y separately, complete it algebraically into a product of two linear expressions equal to a constant, and solve by pairing divisors of that constant – never isolate one variable and search a numeric range when a factoring identity is available.

2018 A2

2018 A2. Let $S_1, S_2, \ldots, S_{2^n-1}$ be the nonempty subsets of $\{1, 2, \ldots, n\}$ in some order, and let M be the $(2^n - 1) \times (2^n - 1)$ matrix whose (i, j) entry is

$$m_{ij} = \begin{cases} 0 & \text{if } S_i \cap S_j = \emptyset; \\ 1 & \text{otherwise.} \end{cases}$$

Calculate the determinant of M.

Fields: Linear algebra (determinants, block matrices, row/column reduction), combinatorics (subset structure) **Position:** A2 – second problem, harder than the position suggests **Answer:** $\det M = 1$ if $n = 1$, and $\det M = -1$ for every $n > 1$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The determinant is well-defined and, in principle, computable by cofactor expansion for any fixed n, so the natural first move is to just build M for the smallest cases and see what comes out. For $n = 1$: one subset, $\{1\}$, and $M = (1)$, so $\det M_1 = 1$. For $n = 2$: subsets

$\{1\}, \{2\}, \{1,2\}$, giving

$$M_2 = \begin{pmatrix} 1 & 0 & 1 \\ 0 & 1 & 1 \\ 1 & 1 & 1 \end{pmatrix}, \qquad \det M_2 = -1.$$

Already a pattern suggests itself: $1, -1, \ldots$ – perhaps $\det M_n = -1$ for all $n > 1$. But a guessed pattern is not a proof, and the honest next step is to check $n = 3$: a 7×7 matrix built from the 7 nonempty subsets of $\{1,2,3\}$. Writing it out entry by entry (36 off-diagonal pairs to classify as intersecting or not, plus the diagonal of all 1s) and expanding a 7×7 determinant by hand is a very large amount of error-prone bookkeeping – cofactor expansion of a 7×7 matrix already costs 7! terms in the naive count, and even with expansion tricks it is not something to attempt under exam time pressure, let alone for the general n the problem actually asks about. Pattern-spotting from $n = 1, 2$ (and maybe a painfully-earned $n = 3$) gives a conjecture, but conjecture is not proof, and this route gives no mechanism for verifying $n = 4, 5, \ldots$ short of repeating the ordeal.

The model chain. The index set is not an arbitrary list – it is *all* nonempty subsets of $\{1, \ldots, n\}$, and a set of subsets of an n-element set has one canonical recursive structure: split on whether the distinguished element n belongs to the subset or not. Order the $2^n - 1$ nonempty subsets of $\{1, \ldots, n\}$ as follows: first the $2^{n-1} - 1$ nonempty subsets $S'_1, \ldots, S'_{2^{n-1}-1}$ of $\{1, \ldots, n-1\}$ (unchanged, call this list block **I**); then the same $2^{n-1} - 1$ sets each with n adjoined, $S'_i \cup \{n\}$ (block **II**); then finally the singleton $\{n\}$ on its own (block **III**).

Now read off the intersection structure in blocks. Within block **I**, intersections are exactly the M_{n-1} problem, unchanged. Between two sets of block **II**, both contain n, so they always intersect – every entry there is 1. Between block **I** and block **II**: $S'_i \cap (S'_j \cup \{n\}) = S'_i \cap S'_j$ (since S'_i never contains n), so this off-diagonal block is again a copy of M_{n-1}'s entries. Block **III**, the row/column for $\{n\}$: it intersects S'_i never (block **I** entries are 0) and it intersects every $S'_j \cup \{n\}$ always (block **II** entries are 1), and the diagonal entry for $\{n\}$ itself is 1. This gives exactly the block form

$$M_n = \begin{pmatrix} M_{n-1} & M_{n-1} & 0 \\ M_{n-1} & J & \mathbf{1} \\ 0 & \mathbf{1}^T & 1 \end{pmatrix},$$

where J is the all-ones matrix of size $2^{n-1} - 1$ and $\mathbf{1}$ is the all-ones column vector.

This block form is now amenable to elementary row and column operations, which do not change the determinant. Subtract the last row (the $\{n\}$ row) from every row of block **II**: since that row is $(0, \ldots, 0, 1, \ldots, 1, 1)$, this turns the J block into the zero matrix and leaves the $\mathbf{1}$ column of block **II** as zero (each row loses its trailing 1). Then do the same for columns: subtract the last column from every column of block **II**. The result is block-diagonal apart from the isolated 1 in the bottom-right corner and the unchanged M_{n-1} blocks on the top-left and (now) below it:

$$\begin{pmatrix} M_{n-1} & 0 & 0 \\ M_{n-1} & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

Removing the last row/column (contributing a factor of 1), swap the first $2^{n-1} - 1$ rows with the next $2^{n-1} - 1$ rows – a permutation of $2^{n-1} - 1$ blocks of size 1 against $2^{n-1} - 1$ others, an odd permutation when $2^{n-1} - 1$ is odd (which it is, since 2^{n-1} is even for $n \geq 2$), contributing a factor of -1. What remains is the block-triangular $\begin{pmatrix} M_{n-1} & 0 \\ M_{n-1} & M_{n-1} \end{pmatrix}$, whose determinant is $(\det M_{n-1})^2$. Hence

$$\det M_n = -(\det M_{n-1})^2 \qquad (n > 1),$$

and since $\det M_1 = 1$, induction gives $\det M_n = -1$ for every $n > 1$ (once negative, squaring and negating keeps it at -1 forever: $-((-1)^2) = -1$).

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that small cases should be computed – $n = 1, 2$ are cheap and instructive, and the model chain uses exactly the same $\det M_1 = 1, \det M_2 = -1$ as data points. The paths split at what to do once the pattern is suspected: the natural attempt tries to confirm it by brute-force computing $n = 3$ (a 7×7 determinant) and has no route to the general n; the model instead asks what structural feature of the index set – "all nonempty subsets of $\{1, \ldots, n\}$" – could turn the single instance n into a relationship with the instance $n - 1$, and finds it by splitting every subset according to whether it contains the element n.

> *The rows and columns of M are indexed by nonempty subsets of $\{1, \ldots, n\}$, and every such subset either contains n or doesn't – that single binary split partitions the whole index set into two blocks of size $2^{n-1} - 1$ (plus the lone singleton $\{n\}$), and since n is exactly the "new" element going from $\{1, \ldots, n-1\}$ to $\{1, \ldots, n\}$, I should expect the blocks to relate to M_{n-1} rather than trying to compute M_n from scratch.*

> *Once I have the matrix in 3×3 block form with a repeated M_{n-1} block and an all-ones block J, I don't need to expand the determinant directly – row and column operations preserve the determinant, and subtracting the all-ones row/column from the block that contains it is exactly the move that kills J and decouples the blocks, because every entry of J and the adjacent $\mathbf{1}$-vector was produced by that same row/column being present.*

- ▷ **Cue:** a matrix (or array) is indexed by all subsets of an n-element set $\longrightarrow$ **Move:** split the index set by whether the distinguished element n belongs to each subset, expecting the resulting blocks to relate recursively to the same structure built on $\{1, \ldots, n-1\}$
- ▷ **Cue:** a block of a matrix consists entirely of 1s (or another constant), produced because every pair of indices in that block shares some fixed feature (e.g. both contain n) $\longrightarrow$ **Move:** use the row (or column) that generated that constant block – typically the row/column for the smallest or most extreme index in the block – to subtract it out via elementary row/column operations, collapsing the constant block to zero before computing the determinant

Layer 3 THE WHY

Why is splitting on membership of n the forced move, rather than one recursive idea among many that happened to work? Because of what the index set $2^{\{1,\ldots,n\}} \setminus \{\emptyset\}$ *is*: subsets of $\{1, \ldots, n\}$ are built, canonically and without any arbitrary choice, from subsets of $\{1, \ldots, n-1\}$ by deciding, for each one, whether to adjoin n. This is the same decomposition that proves $|2^{\{1,\ldots,n\}}| = 2 \cdot |2^{\{1,\ldots,n-1\}}|$ – it is not a technique borrowed from elsewhere, it is the definition of how the power set of an n-set is built from the power set of an $(n-1)$-set. Any object indexed by $P(\{1, \ldots, n\})$ – a sum, a matrix, a generating function – inherits this decomposition automatically, and a recursion in n is the natural shape to look for as a result. The matrix entries respect the split cleanly because intersection with a set containing n behaves predictably: $S \cap (T \cup \{n\})$ depends on $n \in S$, giving exactly the three regimes (both without n: recurse; both with n: automatically nonempty; one with, one without: reduces to intersection ignoring n) that produce the block form.

The row/column-elimination step deserves its own justification. Subtracting one row from another is legal precisely because it changes nothing about which linear combinations of rows can produce the zero vector – the determinant, as a measure of that independence, is invariant under it. Here it is not a generic simplification trick: the row being subtracted (the $\{n\}$ row) is literally the row responsible for every entry in the constant block J and the adjacent $\mathbf{1}$ vector, since those entries are 1 exactly because the other set in the pair contains n, which is exactly the condition recorded by the $\{n\}$ row. Removing that row's contribution from every row that

shares its cause is what collapses the block to zero – the algebra mirrors the combinatorial reason the block was constant in the first place.

Layer $+\alpha$ BACKTESTING

1. The first useful cue, seen at the moment the problem states M is indexed by *all* nonempty subsets of $\{1,\dots,n\}$, is that this index set has a canonical recursive construction from $\{1,\dots,n-1\}$ – this should be written down before attempting to compute any entries.
2. That cue suggests partitioning the rows/columns by membership of the element n, expecting a block matrix whose blocks relate to M_{n-1}, and looking for a recursion $\det M_n = f(\det M_{n-1})$ rather than a closed-form computed from scratch.
3. Hesitation occurs after computing $\det M_1$ and $\det M_2$ and seeing the tempting pattern $1, -1$ – the instinct is to verify by brute-force computing $n = 3$. It is reinforced away by recognising that no amount of small-case checking proves a formula for all n; only an inductive/recursive argument does, so the effort belongs in finding the recursion, not in extending the case-check.
4. The trap is treating the problem as "compute a determinant" in the sense of direct expansion, when the size $2^n - 1$ makes direct expansion infeasible for anything beyond $n = 2$ – the problem is really testing recognition of subset-indexed recursive structure, disguised as a computation.
5. Recognition would be faster next time by treating "object indexed by $P(\{1,\dots,n\})$" as an automatic trigger for the split-on-the-largest-element decomposition, the same reflex used to prove $|P(\{1,\dots,n\})| = 2^n$ by induction.
6. The reusable decision rule: whenever a matrix, sum, or other object is indexed by all subsets of an n-element set, decompose by membership of one distinguished element to relate the n-instance to the $(n-1)$-instance, then use elementary row/column operations – guided by which row or column is responsible for any constant sub-block – to expose block-diagonal or block-triangular structure before computing the determinant.

2018 A3

2018 A3. Determine the greatest possible value of $\sum_{i=1}^{10} \cos(3x_i)$ for real numbers $x_1, x_2, \dots, x_{10}$ satisfying $\sum_{i=1}^{10} \cos(x_i) = 0$.

Fields: Calculus/algebra (triple-angle identity, constrained optimization over a compact domain, smoothing arguments) **Position:** A3 – third problem, a genuine optimization **Answer:** $\frac{480}{49}$, achieved with three of the $y_i = \cos x_i$ equal to 1 and seven equal to $-3/7$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The triple-angle identity $\cos(3\theta) = 4\cos^3\theta - 3\cos\theta$ is the obvious way in: writing $y_i = \cos x_i \in [-1,1]$, the constraint is $\sum y_i = 0$ and the objective is

$$\sum_{i=1}^{10} \cos(3x_i) = \sum_{i=1}^{10} (4y_i^3 - 3y_i) = 4\sum y_i^3 - 3\sum y_i = 4\sum y_i^3,$$

using the constraint to kill the linear term. So the problem reduces to: maximize $\sum y_i^3$ subject to $\sum y_i = 0$, $y_i \in [-1,1]$. This is now a textbook-shaped constrained optimization, and the textbook tool is Lagrange multipliers: set $L = \sum y_i^3 - \lambda \sum y_i$ and solve $\partial L/\partial y_i = 3y_i^2 - \lambda = 0$ for all i, giving $y_i = \pm\sqrt{\lambda/3}$ for every i. Since $\sum y_i = 0$, if p of the y_i take the value $+v = \sqrt{\lambda/3}$ and q take $-v$, the constraint forces $pv = qv$, i.e. $p = q$ (assuming $v \neq 0$); with 10 variables this means $p = q = 5$. But then $\sum y_i^3 = 5v^3 - 5v^3 = 0$, giving objective 0. That is a valid critical point, but it is obviously not the maximum – pushing even one variable to $y = 1$ and rebalancing

the rest gives a strictly positive sum, as a moment's experimentation shows. The Lagrange calculation has found a stationary point of the unconstrained-in-$\mathbb{R}^{10}$ problem, but it has not found the maximum over the actual feasible region, which is the box $[-1,1]^{10}$ intersected with a hyperplane – and the interior stationary point it hands back is not even a local maximum there.

The model chain. The reason Lagrange's interior critical point fails is visible in the second derivative: $\frac{d^2}{dy^2}(y^3) = 6y$, which is positive for $y > 0$ and negative for $y < 0$. The cubic is convex on the positive half and concave on the negative half – it is not globally concave, so there is no guarantee that an interior stationary point of a linear-constrained sum of cubics is a maximum; on a non-concave objective over a compact convex domain, maxima are generally found at the boundary, not in the interior. This licenses a smoothing/exchange argument instead: fix any two indices $i < j$ and their sum $s = y_i + y_j$, and study how $y_i^3 + y_j^3 = y_i^3 + (s - y_i)^3$ varies as y_i ranges over $[\max(-1, s-1), \min(1, s+1)]$ with s fixed. Its second derivative in y_i is $6y_i - 6(s - y_i) = 6(2y_i - s)$... more directly, differentiating $g(y) = y^3 + (s-y)^3$ twice gives $g''(y) = 6y - 6(s-y) = 6(2y - s)$, but the cleaner route the official solution takes is to note g has *constant* second derivative when expressed the right way: actually $g''(y) = 6y + 6(s-y)\cdot(-1)\cdot(-1)$ – recomputing carefully, $g(y) = y^3 + (s-y)^3$, $g'(y) = 3y^2 - 3(s-y)^2$, $g''(y) = 6y + 6(s-y) = 6s$, a constant depending only on s, not on y. So g is entirely convex (if $s > 0$) or entirely concave (if $s < 0$) as a function of y on the whole segment – meaning its maximum over an interval occurs at an endpoint of that interval whenever $s > 0$ (convexity), or, if $s < 0$, at a genuine interior critical point (concavity) where $g'(y) = 0$, i.e. $y = -(s-y)$, i.e. $y_i = -y_j$... but that requires $s = 0$.

Putting this together: at a maximizing configuration, for every pair i, j, either one of y_i, y_j is extremal ($= \pm 1$), or $y_i = y_j$ with $s < 0$ (the concave, interior-critical case, which by $g'(y) = 0$ symmetric in y_i, y_j forces them equal), or $y_i = -y_j$ ($s = 0$). In the last case the pair contributes $y_i^3 + y_j^3 = 0$ and can simply be deleted, reducing to a smaller instance of the same problem – so assume this never happens. Then every non-extremal y_i shares one common negative value $y < 0$ (from the $y_i = y_j$, $s < 0$ case pairing up all non-extremal indices), and since $\sum y_i = 0$ cannot be satisfied using only $y = -1$ and $y < 0$ values (their sum is always negative), some $y_i = 1$ must appear, and $y = -1$ cannot appear alongside $y = 1$ without contradicting y being the unique non-extremal negative value unless $y = -1$ itself is exactly that value – ruled out by needing $y > -1$ for a genuine non-extremal solution. So the extremal configuration has a copies of 1 and b copies of a single value $y = -a/b \in (-1, 0)$ with $a + b = 10$, $a < b$. The objective becomes

$$4\sum y_i^3 = 4\left(a - b\cdot\frac{a^3}{b^3}\right) = 4a\left(1 - \frac{a^2}{b^2}\right),$$

increasing in b for fixed a, so the best choice takes $a + b = 10$ exactly (not smaller). Checking the four integer splits with $a < b$, $a + b = 10$: $(1,9) \to 320/81 \approx 3.95$, $(2,8) \to 15/2 = 7.5$, $(3,7) \to 480/49 \approx 9.80$, $(4,6) \to 80/9 \approx 8.89$. The maximum is $\boxed{480/49}$, at $a = 3, b = 7$: three x_i with $\cos x_i = 1$ (i.e. $x_i = 0$) and seven with $\cos x_i = -3/7$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the reduction via the triple-angle identity to maximizing $\sum y_i^3$ subject to $\sum y_i = 0$ – that step is forced by the problem's own phrasing and nobody skips it. The split occurs immediately after: the natural attempt applies Lagrange multipliers as if the objective were globally well-behaved (concave, or at least such that interior stationary points are maxima), and gets an interior critical point that turns out to be a saddle, not a maximum. The model chain instead checks the second derivative of the cubic, sees it changes sign, and concludes the maximum cannot in general be interior – forcing the analysis onto the boundary of the box via a pairwise smoothing argument.

> *Before trusting the Lagrange critical point as the answer, I should check whether the objective is actually concave on the feasible region – y^3 has second derivative $6y$, which changes sign at $y = 0$, so the function is not globally concave, and an unconstrained interior stationary point of a non-concave function under a linear constraint is not guaranteed to be a maximum; I need to examine the boundary of the box $[-1, 1]^{10}$ instead.*

> *Rather than optimizing all ten variables at once, I should look at what happens to just two of them, y_i and y_j, with their sum $s = y_i + y_j$ held fixed – this reduces a 10-variable problem to a 1-variable function $g(y) = y^3 + (s - y)^3$ whose curvature ($g'' = 6s$, a constant sign depending only on s) tells me immediately whether the pair's optimum sits at an endpoint of its range or at the point where $y_i = y_j$.*

- ▷ **Cue:** an optimization objective built from a function that is convex on part of its domain and concave on another part (e.g. y^3, sign-changing second derivative) $\longrightarrow$ **Move:** do not trust an interior Lagrange stationary point as the maximum; instead examine the boundary of the feasible box directly, typically via a pairwise smoothing/exchange argument that fixes the sum of two variables and studies the resulting 1-variable function's curvature
- ▷ **Cue:** a symmetric objective over many variables with one linear constraint, where a two-variable exchange leaves the constraint unchanged $\longrightarrow$ **Move:** fix the sum of a pair $y_i + y_j = s$ and analyse $g(y) = y_i^3 + (s - y_i)^3$ (or the analogous function for the true objective) as a function of one variable; its sign-constant second derivative g'' tells you whether the pair's contribution is maximized at an endpoint (convex, push to extreme values) or where the pair is equal (concave, keep them tied)

Layer 3 THE WHY

Why must the pairwise smoothing argument replace Lagrange multipliers here, rather than being merely a faster alternative? Because Lagrange's theorem only certifies stationary points as candidates for extrema of a function on a manifold defined by equality constraints – it says nothing about whether a given stationary point is a maximum, a minimum, or a saddle, and that classification depends on the objective's curvature, which for $\sum y_i^3$ is not fixed in sign. A maximum of a genuinely concave function under linear constraints is always interior and always found by Lagrange; a maximum of a function that is convex somewhere on the domain is, by the defining property of convexity, never interior to the region where it is convex (a convex function's supremum on an interval is at an endpoint, not the middle) – so as soon as $6y$ can be positive, the search for a maximum must include the boundary $y = \pm 1$, and the Lagrange interior point can only ever be relevant on the concave, $y < 0$ portion.

The pairwise reduction is forced by the same logic at a smaller scale: fixing $y_i + y_j = s$ turns the two-variable slice of the objective into the single-variable function $g(y) = y^3 + (s - y)^3$, and its second derivative $g''(y) = 6s$ is *constant* in y – meaning g is entirely convex or entirely concave along that whole slice, with no sign change to worry about within the pair. This is exactly why the case analysis collapses to three clean possibilities (extremal, tied-and-negative, or cancelling-and-removable): a function that is convex throughout an interval has its maximum at an endpoint, full stop, and a function that is concave throughout has it where the derivative vanishes, full stop. The messiness of a sign-changing second derivative, which is what defeated the naive global Lagrange approach, is precisely what the two-at-a-time restriction eliminates: at fixed s, the pair's own curvature no longer changes sign, so the one-variable calculus that failed globally succeeds locally.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the triple-angle identity $\cos 3\theta = 4\cos^3\theta - 3\cos\theta$ appearing alongside a constraint on $\sum \cos x_i$ – this should immediately be written down as a substitution $y_i = \cos x_i$ that turns a trigonometric optimization into a pure algebraic one in $y_i \in [-1, 1]$.

2. That cue suggests treating the problem as constrained optimization of $\sum y_i^3$ subject to $\sum y_i = 0$ – but the presence of a cubic, rather than a quadratic or linear objective, should immediately raise the question of concavity before reaching for Lagrange multipliers.

3. Hesitation occurs at the moment Lagrange multipliers hand back a critical point with objective value 0, which is obviously not the answer (the objective can clearly be made positive). It is reinforced away by checking the second derivative of the objective's building block, y^3, and recognising that its sign change means interior stationary points cannot be trusted as global maxima.

4. The trap is applying Lagrange multipliers mechanically to a non-concave objective and accepting the resulting critical point as the answer without checking whether it is even a local maximum, let alone the global one.

5. Recognition would be faster next time by treating "maximize a sum of a function $f(y_i)$ subject to a linear constraint, where f'' changes sign on the domain" as an automatic trigger for boundary/smoothing analysis rather than Lagrange multipliers – concavity (or its absence) should be checked before, not after, setting up the Lagrangian.

6. The reusable decision rule: for a symmetric optimization $\max \sum f(y_i)$ under a linear constraint and box bounds, first check whether f is globally concave. If it is, Lagrange's interior critical point is the answer. If f'' changes sign, abandon the global Lagrangian approach and use pairwise smoothing – fix two variables' sum and exploit that the resulting one-variable function has constant-sign curvature, forcing each pair to either sit at the box boundary or be tied together, and reduce the problem to choosing how many variables sit at each extreme.

2018 A4

2018 A4. Let m and n be positive integers with $\gcd(m,n) = 1$, and let

$$a_k = \left\lfloor \frac{mk}{n} \right\rfloor - \left\lfloor \frac{m(k-1)}{n} \right\rfloor$$

for $k = 1, 2, \ldots, n$. Suppose that g and h are elements in a group G and that

$$gh^{a_1} gh^{a_2} \cdots gh^{a_n} = e,$$

where e is the identity element. Show that $gh = hg$. (As usual, $\lfloor x \rfloor$ denotes the greatest integer less than or equal to x.)

Fields: group theory (words in two generators, change of generators); number theory (the Euclidean algorithm, floor functions); combinatorics (lattice paths / Christoffel words) **Position:** A4 — the first genuinely structural problem on the paper. Nothing is computed; everything depends on seeing which operation shrinks the problem. **Answer:** Proof: the substitutions $(g,h) \mapsto (gh,h)$ (when $m > n$) and $(g,h) \mapsto (g,gh)$ (when $m < n$) run the subtractive Euclidean algorithm on (m,n) down to $(1,1)$, where the relation is $gh = e$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You read "show $gh = hg$" and try to manipulate the relation. Conjugate it, cyclically rotate it (if $xy = e$ then $yx = e$ in a group), multiply on the left and right, and hope a commutator falls out. For general m, n that is hopeless — you have not even written the word down — so you do small cases.

$m = 1$: here $a_k = 0$ for $k < n$ and $a_n = 1$, so the relation is $g^n h = e$. Then $h = g^{-n}$, which commutes with g. Done, and it felt easy.

$m = 2$, $n = 3$: $a_1 = \lfloor 2/3 \rfloor = 0$, $a_2 = \lfloor 4/3 \rfloor - 0 = 1$, $a_3 = 2 - 1 = 1$, so the relation is $g\,gh\,gh = e$. You stare at it, rotate it to $gh\,gh\,g = e$, try to isolate h, and get a tangle. With luck you notice the repeated block gh and write $x = gh$: the relation is $gx^2 = e$, so $g = x^{-2}$ commutes with x, and therefore with $h = g^{-1}x$.

$m = 3, n = 5$: $a = (0, 1, 0, 1, 1)$, relation $g\,gh\,g\,gh\,gh = e$. With $x = gh$ it becomes $g\,x\,g\,x\,x = e$. There are still isolated g's, it does not collapse, and you are stuck again, because you are solving each case by inspection.

That is the wall. Each case yields to a local trick, but you have no mechanism that turns the $(3, 5)$ case into a case you have already done. The natural attempt treats the relation as an equation to be solved; it needs to be treated as one member of a family, to be reduced to a smaller member.

The model chain. Notation. For positive integers m, n let $W_{m,n}(g, h) = gh^{a_1}gh^{a_2}\cdots gh^{a_n}$ with a_k as in the statement. By telescoping, $a_1 + \cdots + a_k = \lfloor mk/n \rfloor$, so the word has n letters g and m letters h, and each a_k is $\lfloor m/n \rfloor$ or $\lfloor m/n \rfloor + 1$.

Position lemma. The j-th letter h of $W_{m,n}$ ($1 \le j \le m$) lies in the block of the k-th g exactly when $\lfloor m(k-1)/n \rfloor < j \le \lfloor mk/n \rfloor$; since j is an integer this is $m(k-1)/n < j \le mk/n$, i.e. $k = \lceil nj/m \rceil$. It is preceded by k letters g and $j - 1$ letters h, so its position in the word is

$$q_j = j + \lceil nj/m \rceil.$$

A word in two letters of known length is determined by the positions of one letter, so this formula pins $W_{m,n}$ down.

Claim. If $\gcd(m, n) = 1$ and $W_{m,n}(g, h) = e$ in a group, then $gh = hg$. Strong induction on $m + n$.

Base, $m = n = 1$. The relation is $gh = e$, so $h = g^{-1}$ and $hg = e = gh$. Since $\gcd(m, n) = 1$, $m = n$ forces $m = n = 1$; every other case has $m \ne n$.

Case $m > n$. Put $m' = m - n \ge 1$, so $\gcd(m', n) = 1$. Then

$$a_k = \left\lfloor \tfrac{m'k}{n} + k \right\rfloor - \left\lfloor \tfrac{m'(k-1)}{n} + k - 1 \right\rfloor = b_k + 1, \qquad b_k = \left\lfloor \tfrac{m'k}{n} \right\rfloor - \left\lfloor \tfrac{m'(k-1)}{n} \right\rfloor,$$

so $gh^{a_k} = (gh)h^{b_k}$ and $W_{m,n}(g, h) = W_{m',n}(gh, h)$. By induction gh commutes with h: $ghh = hgh$, and cancelling h on the right gives $gh = hg$.

Case $m < n$. Put $n' = n - m \ge 1$, so $\gcd(m, n') = 1$. Now every $a_k \in \{0, 1\}$, so every h in $W_{m,n}$ is immediately preceded by a g. Replace each adjacent pair gh by a single letter x. I claim the result is exactly $W_{m,n'}(g, x)$, that is,

$$W_{m,n}(g, h) = W_{m,n'}(g, gh). \tag{$*$}$$

Both sides, read as words in g and x, have length n and m letters x. Before the j-th h of $W_{m,n}$ stand $\lceil nj/m \rceil$ letters g and $j - 1$ letters h; after merging, the $j - 1$ earlier h's have each absorbed one earlier g, and the last of those g's is absorbed into the j-th x itself. So the number of letters before the j-th x is $(\lceil nj/m \rceil + j - 1) - (j - 1) - 1 = \lceil nj/m \rceil - 1$, and the j-th x sits at position $\lceil nj/m \rceil$. In $W_{m,n'}(g, x)$ the position lemma puts the j-th x at $j + \lceil n'j/m \rceil = j + \lceil nj/m \rceil - j = \lceil nj/m \rceil$. Same positions, so the words coincide. (Checked also by machine for all $1 \le m < n < 40$.) By induction g commutes with $x = gh$: $g\,gh = gh\,g$, and cancelling g on the left gives $gh = hg$.

Each non-base case strictly decreases $m + n$ and preserves coprimality, so the induction closes. □

Comparison with the official solution. Kedlaya–Ng handle $m < n$ differently: they invert the relation and rename $\tilde{g} = h^{-1}$, $\tilde{h} = g^{-1}$, which turns $W_{m,n}$ into $W_{n,m}$ (a lattice-path reflection). That swap is correct (I verified it by the same position count), but it does not decrease $m + n$, so it must be followed by a subtraction inside the same step. The substitution $h \mapsto gh$ in $(*)$ is a subtraction in its own right and needs no swap.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree up to the $(2,3)$ case, where the block gh appears and $x = gh$ works. At that line they split. The natural attempt records "$x = gh$ worked here" as a trick for this case and moves on to $(3,5)$ by inspection. The model asks what the substitution did to the *parameters*: $g\,gh\,gh = g\,x^2 = W_{2,1}(g,x)$, so $(2,3)$ became $(2,1)$ — one step of the subtractive Euclidean algorithm. And $(3,5)$ with $x = gh$ is $g\,x\,g\,x\,x = W_{3,2}(g,x)$: not stuck at all, it is $(3,2)$, which the other substitution sends to $(1,2)$, then $(1,1)$.

> *Replacing h by gh, or g by gh, does not change whether the two elements commute — g commutes with gh exactly when it commutes with h. So I am free to change generators, and I should pick the change that makes the word shorter.*

> *$g\,x\,g\,x\,x$ is not a failure. It has 5 letters, 3 of them x and 2 of them g — I should ask whether it is the same kind of word with smaller numbers, not whether it collapses.*

> *$\gcd(m,n) = 1$ is not being used for any divisibility. It is there so that something terminates. The thing that terminates because of coprimality is Euclid, and it ends at $(1,1)$ — where the relation is $gh = e$.*

Layer 3 THE WHY

Why must the change-of-generators thought be thought? Start from the conclusion, not the hypothesis. "g and h commute" is a property of the *pair*, and it is invariant under a family of replacements: $(g,h) \to (gh,h)$, $(g,h) \to (g,gh)$, $(g,h) \to (h,g)$, $(g,h) \to (g^{-1},h)$. Each is invertible and preserves commutation in both directions — if g and h commute so do gh and h, and conversely $g = (gh)h^{-1}$. When the target is invariant under a set of operations, those operations are free moves: you can apply them as often as you like without losing anything. So the only question is which free move simplifies the hypothesis, and the hypothesis is a word whose shape is controlled by two integers. A free move that shrinks those integers is exactly what you want.

Now the phrasing. The problem hands you $\gcd(m,n) = 1$ and floor functions of mk/n. Coprimality is not a hypothesis you need for a computation; it is a hypothesis you need for a *termination*. The one process in elementary mathematics that terminates at $(1,1)$ precisely because of coprimality is the Euclidean algorithm. So the problem is announcing that the proof descends through pairs (m,n) and lands on the trivial relation $gh = e$. Once you expect Euclid, you look for two moves, $(m,n) \to (m-n,n)$ and $(m,n) \to (m,n-m)$, and the two free substitutions supply them.

Why do the substitutions have that effect? Draw the word. Walk from $(0,0)$ to (n,m), each g a step right and each h a step up; $W_{m,n}$ is the lattice path that hugs the line $y = mx/n$ from below (after the k-th right step you have climbed exactly $\lfloor mk/n \rfloor$). Merging "gh" into a single diagonal step is the shear $(x,y) \mapsto (x-y,y)$, which sends the line of slope m/n to the line of slope $m/(n-m)$ and the hugging path to the hugging path. Merging "gh" into the g when $m > n$ is the shear $(x,y) \mapsto (x,y-x)$. Shears are the two generators of $SL_2(\mathbb{Z})$, and the Euclidean algorithm is exactly the process of shearing a primitive vector (n,m) back to $(1,1)$. So the proof is not a trick: floor functions of mk/n *are* a line of rational slope, and the natural symmetries of lattice lines are shears. Kedlaya's first solution uses the other natural symmetry (reflection in $y = x$), and Sarkar's third solution says the same thing topologically: an element of $GL_2(\mathbb{Z})$ carries the segment to (n,m) onto the segment to $(1,0)$.

Why must $\gcd(m,n) = 1$ really be there? Because without it the descent stops at (d,d), and (d,d) is not trivial. At $(2,2)$ we have $a_1 = a_2 = 1$ and the relation is $(gh)^2 = e$, which says

only that gh has order dividing 2. In S_3 take $g = (1\,2\,3)$ and $h = g^{-1}(1\,2)$. Then $gh = (1\,2)$, so $(gh)^2 = e$, while $hg = g^{-1}(1\,2)\,g$ is the conjugate of $(1\,2)$ by an element that does not commute with it, hence a different transposition (checked by composing the permutations). So $gh \neq hg$. The hypothesis is exactly the condition that Euclid lands on a relation of length two.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the pair "$\gcd(m,n) = 1$" and "$\lfloor mk/n \rfloor$". On seeing it, write down $a_1 + \cdots + a_k = \lfloor mk/n \rfloor$ (telescoping), $a_k \in \{\lfloor m/n \rfloor, \lfloor m/n \rfloor + 1\}$, and the words "Euclid: which substitution sends (m,n) to $(m-n,n)$, and which to $(m,n-m)$?"

2. Strong induction on $m+n$ following the subtractive Euclidean algorithm, each step being a change of generators that preserves the conclusion.

3. Hesitation occurs at $g\,x\,g\,x\,x$, where the substitution seems to have failed. What reinforces it away: count letters. Five letters, three x, two g — it is $W_{3,2}(g,x)$. A second hesitation occurs at $(*)$, where the pattern is convincing but unproved; the reinforcement is that a two-letter word is determined by the positions of one letter, so compute positions on both sides.

4. Two traps. (i) Solving each small case by an ad hoc manipulation and never finding a reduction — zombie maths in its group-theoretic form. (ii) If you take the official swap route, writing "induction on $m+n$" and then applying the swap, which does *not* decrease $m+n$; the swap must be followed by a subtraction within the same step.

5. Faster next time via the slogan: *a coprimality hypothesis on two parameters is a promise that some process reduces them by Euclid to* $(1,1)$. Look for it in lattice problems, tilings of rectangles by squares, continued fractions, and any statement indexed by a slope m/n. When a substitution turns a case into a word you do not recognise, count its letters before abandoning it.

6. The reusable decision rule: **when the conclusion is invariant under a family of transformations, treat those transformations as free moves and use them to shrink the hypothesis until it is trivial.** Here "g,h commute" is invariant under change of generators; elsewhere it might be "the determinant is ± 1" (invariant under row operations) or "the gcd is 1" (invariant under subtracting one entry from another).

▷ **Cue:** $\gcd(m,n) = 1$ appears in a statement where no divisibility is being computed $\longrightarrow$ **Move:** expect a Euclidean-algorithm induction: find a move $(m,n) \to (m-n,n)$ and a move $(m,n) \to (m,n-m)$

▷ **Cue:** you must prove two group elements commute $\longrightarrow$ **Move:** change generators freely: $(g,h) \to (gh,h)$, $(g,h) \to (g,gh)$, $(g,h) \to (h^{-1},g^{-1})$ all preserve commutation in both directions

▷ **Cue:** exponents given by differences $\lfloor mk/n \rfloor - \lfloor m(k-1)/n \rfloor$ $\longrightarrow$ **Move:** draw the lattice path from $(0,0)$ to (n,m) hugging $y = mx/n$ from below; shears and the reflection $x \leftrightarrow y$ act on the word

▷ **Cue:** two words in two letters must be shown equal $\longrightarrow$ **Move:** compare the positions of one of the letters; that determines the word

▷ **Cue:** a substitution turns a small case into an unfamiliar word $\longrightarrow$ **Move:** count its letters and ask whether it is the same family with smaller parameters before discarding it

2018 A5

2018 A5. Let $f : \mathbb{R} \to \mathbb{R}$ be an infinitely differentiable function satisfying $f(0) = 0$, $f(1) = 1$, and $f(x) \geq 0$ for all $x \in \mathbb{R}$. Show that there exist a positive integer n and a real number x such that $f^{(n)}(x) < 0$.

Fields: real analysis (Taylor's theorem with integral remainder, smooth but non-analytic functions, absolutely monotone functions) **Position:** A5 — hard, but short once the right object is in view. The

difficulty is entirely in deciding what to prove, not in the proof. **Answer:** Proof by contradiction: if every derivative were nonnegative, then f vanishes on $(-\infty, 0]$, which forces $f^{(n)}(1) \geq n!$ for every n, and then Taylor at 1 gives $f(2) \geq N + 1$ for every N.

Layer 1 THE WHAT — surface analysis

The natural attempt. The statement is negative ("there exist n and x with $f^{(n)}(x) < 0$"), so you assume the opposite: $f^{(n)}(x) \geq 0$ for all $n \geq 0$ and all x. Now you want a contradiction, and you do not know where to look, so you hunt for a counterexample to see what breaks.

Polynomials: $f(x) = x^2$ has $f \geq 0$, $f(0) = 0$, $f(1) = 1$, but $f'(x) = 2x < 0$ for $x < 0$. Every nonconstant polynomial fails the same way, because $f' \geq 0$ everywhere makes f nondecreasing, so $f \leq f(0) = 0$ on $(-\infty, 0]$, so with $f \geq 0$ you get $f \equiv 0$ on $(-\infty, 0]$. A polynomial vanishing on a half-line is zero.

So you need a function that is identically zero on the left and rises on the right: the standard smooth-but-not-analytic example, $f(x) = e^{-1/x}$ for $x > 0$ and 0 for $x \leq 0$, rescaled so $f(1) = 1$. Compute: $f'(x) = x^{-2}e^{-1/x} \geq 0$. Good. $f''(x) = (x^{-4} - 2x^{-3})e^{-1/x} = x^{-4}(1 - 2x)e^{-1/x}$, negative for $x > 1/2$. Fails. You try e^{-1/x^2}, $e^{-1/x}$ composed with e^x, sums, products. Every candidate has some derivative that dips negative somewhere. You now believe the statement, but you have a pile of failed examples and no proof. You try to show directly that some $f^{(n)}$ must be negative, perhaps by an intermediate value argument on f'', and get nowhere: for any fixed n there *are* functions with $f, f', \ldots, f^{(n)} \geq 0$ satisfying the conditions. (Take a smooth nondecreasing step s with $s = 0$ on $(-\infty, 0]$ and $s > 0$ on $(0, \infty)$, let f be its n-fold iterated integral from 0, and rescale so $f(1) = 1$; then $f^{(k)} \geq 0$ for $k \leq n + 1$.) So no finite-order argument can work.

The model chain. Suppose, for contradiction, that $f^{(n)}(x) \geq 0$ for all $n \geq 0$ and all real x.

Step 1: f is flat at 0. Since $f' \geq 0$, f is nondecreasing, so for $x \leq 0$ we have $0 \leq f(x) \leq f(0) = 0$. Thus $f \equiv 0$ on $(-\infty, 0]$, so every derivative vanishes on $(-\infty, 0)$, and by continuity of $f^{(n)}$,

$$f^{(n)}(0) = 0 \qquad (n \geq 0).$$

Step 2: each derivative is nondecreasing. Since $f^{(n+1)} \geq 0$, $f^{(n)}$ is nondecreasing for every n.

Step 3: $f^{(n)}(1) \geq n!$. Taylor's theorem about 0 with integral remainder, using Step 1 to kill the polynomial part, gives for $n \geq 1$

$$1 = f(1) = \int_0^1 \frac{(1-t)^{n-1}}{(n-1)!} f^{(n)}(t)\, dt \;\leq\; f^{(n)}(1) \int_0^1 \frac{(1-t)^{n-1}}{(n-1)!}\, dt = \frac{f^{(n)}(1)}{n!},$$

where the inequality uses Step 2 ($f^{(n)}(t) \leq f^{(n)}(1)$ on $[0, 1]$) and the nonnegativity of the kernel. So $f^{(n)}(1)/n! \geq 1$ for all $n \geq 1$, and also for $n = 0$ since $f(1) = 1$.

Step 4: the contradiction at $x = 2$. Taylor's theorem about 1 with integral remainder gives, for any N,

$$f(2) = \sum_{k=0}^{N} \frac{f^{(k)}(1)}{k!} + \int_1^2 \frac{(2-t)^N}{N!} f^{(N+1)}(t)\, dt \;\geq\; \sum_{k=0}^{N} \frac{f^{(k)}(1)}{k!} \;\geq\; N + 1,$$

because the remainder integrand is nonnegative and each term of the sum is at least 1 by Step 3. Since $f(2)$ is a fixed real number and N is arbitrary, this is impossible. Hence some $f^{(n)}(x) < 0$. □

(This is the official first solution compressed: Kedlaya–Ng prove the stronger $f(x) \leq f^{(n)}(1)x^n/n!$ on $[0, 1]$ by induction and conclude $f \equiv 0$; I only need the case $x = 1$, which the integral remainder gives in one line.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree up to the contradiction hypothesis and the observation that $f \equiv 0$ on $(-\infty, 0]$ — the natural attempt even finds this, in the course of killing polynomials.

At that line they split. The natural attempt reads "f vanishes on the left" as a *design constraint* for a counterexample and goes shopping for flat functions. The model reads it as *information*: all derivatives vanish at 0, so Taylor's formula about 0 has no polynomial part and f equals its remainder — and the remainder has a sign.

> *All derivatives are nonnegative. That is precisely the hypothesis that makes every Taylor remainder nonnegative. So every Taylor polynomial is a lower bound for f to the right of its centre. I should write down Taylor's theorem with the integral remainder and read off inequalities, not hunt for examples.*

> *f is flat at 0, so going from 0 to 1 the function has to climb from nothing to 1 using only its n-th derivative, for every n. That forces $f^{(n)}(1)$ to be at least $n!$. Big derivatives at 1 mean the Taylor series at 1 grows — and its partial sums are lower bounds for $f(2)$.*

> *No argument using only finitely many derivatives can work, because for each fixed n there are functions with the first n derivatives nonnegative. The contradiction must come from letting $n \to \infty$, so I want an inequality with n in it.*

Layer 3 THE WHY

Why must the Taylor-remainder thought be thought? Go to what the hypothesis *is*. "$f^{(n)} \geq 0$ for all n" is a statement about infinitely many derivatives simultaneously, and the only standard tool in which infinitely many derivatives appear together is Taylor's theorem. Moreover Taylor's theorem with integral remainder,

$$f(x) = \sum_{k<n} \frac{f^{(k)}(a)}{k!}(x-a)^k + \int_a^x \frac{(x-t)^{n-1}}{(n-1)!} f^{(n)}(t)\,dt,$$

has exactly one place where the sign of $f^{(n)}$ enters: the remainder. For $x > a$ the kernel $(x-t)^{n-1}$ is nonnegative, so a sign condition on $f^{(n)}$ is *literally* a sign condition on the remainder. The hypothesis is tailor-made for this formula — it is not a lucky choice, it is the only formula the hypothesis talks to.

Why does the contradiction have to be "growth" against "flatness"? Because that is the real content of the problem, and it is worth knowing. A function all of whose derivatives are nonnegative on an interval is called *absolutely monotone*, and Bernstein's theorem says such a function is real-analytic there — it equals its Taylor series. The official remark (Elkies, attributing the argument to Shalen) proves exactly this from the same two Taylor inequalities used above. Analytic functions cannot vanish on a half-line without vanishing everywhere. That is why every counterexample hunt dies: the flat-on-the-left functions like $e^{-1/x}$ are the textbook *non-analytic* smooth functions, and the hypothesis forbids non-analyticity. The problem is "absolutely monotone implies analytic" in disguise, and the disguise is that you are only asked for the contradiction.

Seen that way, the two steps of the model chain are the two halves of the analyticity argument. Step 3 says that flatness at 0 forces the Taylor coefficients $f^{(n)}(1)/n!$ at 1 to stay at least 1 — so the Taylor series at 1 cannot converge at distance 1. Step 4 says that nonnegative remainders make the partial sums at 2 bounded by $f(2)$ — so the Taylor series at 1 *must* converge at distance 1. Divergent and convergent at once: contradiction. Kedlaya's solution introduces the name "ultraconvex" and the induction $f(x) \leq f^{(n)}(1)x^n/n!$ without saying where it comes from. Here is where: it is the statement that the remainder at order n from 0 is dominated by its value using the largest derivative on the interval, which is the largest derivative at the right endpoint because each derivative is nondecreasing.

Why was $f \geq 0$ on all of $\mathbb{R}$ given, rather than on $[0,1]$? Because it is what forces flatness at 0. Without the left half-line, $f(x) = x$ on $\mathbb{R}$ has $f' = 1$, $f^{(n)} = 0$ for $n \geq 2$, satisfies $f(0) = 0, f(1) = 1$, and fails only because $f(x) < 0$ for $x < 0$. The hypothesis $f \geq 0$ on the left is what converts "nondecreasing" into "zero on the left" — that is the one line where the whole real line is used.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "$f^{(n)}(x) < 0$ for some n" — a statement about all derivatives at once. On seeing it, write the contrapositive hypothesis "$f^{(n)} \geq 0$ for all n" and directly under it Taylor's theorem with integral remainder, noting that the remainder is nonnegative to the right of the centre.

2. Taylor's theorem with integral remainder, used twice: once about the flat point 0 to get a lower bound on the derivatives at 1, once about 1 to turn those derivatives into a lower bound on $f(2)$.

3. Hesitation occurs at "$f \equiv 0$ on $(-\infty, 0]$". It feels like a dead end ("so the function must be of $e^{-1/x}$ type") and pushes you toward constructing examples. What reinforces it away: flatness means *all derivatives at* 0 *are zero*, and a Taylor expansion with zero polynomial part is pure remainder — the most information-rich situation Taylor's theorem allows.

4. The trap is the counterexample hunt with bump-type functions. Each attempt is plausible, each fails at some high derivative, and the failures teach nothing transferable. A second trap: trying a fixed-order argument (for example, only f, f', f''), which cannot succeed because the statement is false if only finitely many derivatives are constrained.

5. Faster next time: *a sign condition on every derivative is a sign condition on every Taylor remainder.* Whenever a hypothesis quantifies over all derivatives, expand in Taylor series with integral remainder before doing anything else. And whenever a smooth function is forced to vanish on an interval, ask whether the hypotheses force analyticity — if so, the contradiction is automatic.

6. The reusable decision rule: **when a statement is about all orders n at once, the contradiction must come from $n \to \infty$; find an inequality with n in it that grows on one side and is bounded on the other.** Here: $f(2) \geq N + 1$ for all N. Elsewhere: $a_n \geq c^n$ against $a_n \leq$ polynomial, or a series forced both to diverge and to converge.

▷ **Cue:** a hypothesis or conclusion quantifies over all derivatives $f^{(n)}$ $\longrightarrow$ **Move:** write Taylor's theorem with integral remainder; a sign on $f^{(n)}$ is a sign on the remainder, so Taylor polynomials become one-sided bounds

▷ **Cue:** a smooth function is forced to vanish identically on a half-line $\longrightarrow$ **Move:** all derivatives vanish at the endpoint; expand about that endpoint and f equals its remainder

▷ **Cue:** $f^{(n+1)} \geq 0$ on an interval $\longrightarrow$ **Move:** $f^{(n)}$ is nondecreasing there, so bound $f^{(n)}(t)$ by its value at the right endpoint inside an integral

▷ **Cue:** the statement fails if only finitely many conditions are imposed $\longrightarrow$ **Move:** the proof must let $n \to \infty$; aim for an inequality depending on n that becomes impossible in the limit

▷ **Cue:** a counterexample search keeps failing at higher and higher order $\longrightarrow$ **Move:** stop searching; the failures are evidence of a rigidity theorem (here: absolutely monotone implies analytic)

2018 A6

2018 A6. Suppose that A, B, C, and D are distinct points, no three of which lie on a line, in the Euclidean plane. Show that if the squares of the lengths of the line segments AB, AC,

AD, BC, BD, and CD are rational numbers, then the quotient

$$\frac{\text{area}(\triangle ABC)}{\text{area}(\triangle ABD)}$$

is a rational number.

Fields: geometry (vectors, areas as determinants); linear algebra (Gram determinants, $\det(M^{\mathsf{T}}N) = \det M \det N$); algebra (rationality arguments) **Position:** A6 — the last problem of the morning, but one of the more approachable A6s. The whole difficulty is one reframing; after it the proof is four lines. **Answer:** Proof: with $\mathbf{b} = \vec{AB}$, $\mathbf{c} = \vec{AC}$, $\mathbf{d} = \vec{AD}$, all six dot products are rational, and the ratio equals $\det(\mathbf{b},\mathbf{c})\det(\mathbf{b},\mathbf{d}) / \det(\mathbf{b},\mathbf{d})^2$, a quotient of two Gram determinants.

Layer 1 THE WHAT — surface analysis

The natural attempt. Areas from squared side lengths: that is Heron's formula, and you know its squared-length form. For a triangle with squared sides p, q, r,

$$16\,\text{area}^2 = 4pq - (p+q-r)^2.$$

Apply it to $\triangle ABC$ with $p = AB^2$, $q = AC^2$, $r = BC^2$, all rational: $\text{area}(\triangle ABC)^2 \in \mathbb{Q}$. Same for $\triangle ABD$. So the square of the quotient is rational.

And now the wall. $\sqrt{2}$ has rational square. You need the quotient itself rational, i.e. you need

$$\text{area}(\triangle ABC) \cdot \text{area}(\triangle ABD) \in \mathbb{Q},$$

because then the quotient is (product)/(area of ABD)^2, a rational over a nonzero rational. Heron gives you each factor's square, never the product. Notice also that you have not used CD at all — a sure sign the argument is incomplete, since CD is the only datum that links the two triangles.

The next instinct is coordinates. Put A at the origin and B on the x-axis: $B = (\sqrt{a}, 0)$ with $a = AB^2$. Write $C = (x_1, y_1)$, $D = (x_2, y_2)$. The areas are $\frac{1}{2}\sqrt{a}\,|y_1|$ and $\frac{1}{2}\sqrt{a}\,|y_2|$, so the quotient is $|y_1/y_2|$. Now grind: $AC^2 = x_1^2 + y_1^2$, $BC^2 = (x_1 - \sqrt{a})^2 + y_1^2$, so subtracting gives $2x_1\sqrt{a} \in \mathbb{Q}$, hence $x_1^2 \in \mathbb{Q}$, hence $y_1^2 \in \mathbb{Q}$; similarly for D; then CD^2 gives $x_1x_2 + y_1y_2 \in \mathbb{Q}$, and $x_1x_2 = (x_1\sqrt{a})(x_2\sqrt{a})/a \in \mathbb{Q}$, so $y_1y_2 \in \mathbb{Q}$, so $y_1/y_2 = y_1y_2/y_2^2 \in \mathbb{Q}$. This works — it is the official first solution — and a strong student who pushes through will get there. But notice what made it work: at the end you needed y_1y_2, a *product*, exactly the thing Heron could not provide. The coordinates found it by accident of bookkeeping.

The model chain. Let $\mathbf{b} = \vec{AB}$, $\mathbf{c} = \vec{AC}$, $\mathbf{d} = \vec{AD}$, as column vectors in $\mathbb{R}^2$.

Step 1: every dot product is rational. $\mathbf{b}\cdot\mathbf{b} = AB^2$, $\mathbf{c}\cdot\mathbf{c} = AC^2$, $\mathbf{d}\cdot\mathbf{d} = AD^2$ are rational by hypothesis. By polarisation,

$$BC^2 = |\mathbf{c}-\mathbf{b}|^2 = \mathbf{b}\cdot\mathbf{b} + \mathbf{c}\cdot\mathbf{c} - 2\,\mathbf{b}\cdot\mathbf{c},$$

so $\mathbf{b}\cdot\mathbf{c} = \frac{1}{2}(AB^2 + AC^2 - BC^2) \in \mathbb{Q}$. Likewise $\mathbf{b}\cdot\mathbf{d} = \frac{1}{2}(AB^2 + AD^2 - BD^2)$ and $\mathbf{c}\cdot\mathbf{d} = \frac{1}{2}(AC^2 + AD^2 - CD^2)$ are rational.

Step 2: areas are determinants. Writing $\det(\mathbf{u},\mathbf{v})$ for the determinant of the 2×2 matrix with columns $\mathbf{u}, \mathbf{v}$,

$$\text{area}(\triangle ABC) = \tfrac{1}{2}|\det(\mathbf{b},\mathbf{c})|, \qquad \text{area}(\triangle ABD) = \tfrac{1}{2}|\det(\mathbf{b},\mathbf{d})|,$$

and $\det(\mathbf{b},\mathbf{d}) \neq 0$ because A, B, D are not collinear.

Step 3: a product of two determinants is a determinant of dot products. With $M = (\mathbf{b}\ \mathbf{c})$ and $N = (\mathbf{b}\ \mathbf{d})$,

$$\det(\mathbf{b},\mathbf{c})\det(\mathbf{b},\mathbf{d}) = \det(M^{\mathsf{T}}N) = \det\begin{pmatrix} \mathbf{b}\cdot\mathbf{b} & \mathbf{b}\cdot\mathbf{d} \\ \mathbf{c}\cdot\mathbf{b} & \mathbf{c}\cdot\mathbf{d} \end{pmatrix} = (\mathbf{b}\cdot\mathbf{b})(\mathbf{c}\cdot\mathbf{d}) - (\mathbf{b}\cdot\mathbf{d})(\mathbf{b}\cdot\mathbf{c}) \in \mathbb{Q},$$

using $\det M^{\mathsf{T}} = \det M$ and multiplicativity. The same identity with $M = N$ gives $\det(\mathbf{b},\mathbf{d})^2 = |\mathbf{b}|^2|\mathbf{d}|^2 - (\mathbf{b}\cdot\mathbf{d})^2 \in \mathbb{Q}$, nonzero by Step 2.

Step 4: conclude.

$$\frac{\text{area}(\triangle ABC)}{\text{area}(\triangle ABD)} = \left|\frac{\det(\mathbf{b},\mathbf{c})}{\det(\mathbf{b},\mathbf{d})}\right| = \left|\frac{\det(\mathbf{b},\mathbf{c})\det(\mathbf{b},\mathbf{d})}{\det(\mathbf{b},\mathbf{d})^2}\right| \in \mathbb{Q}. \qquad \square$$

(This is Bhargava's second official solution, written with the multiplication step made explicit. Only $\mathbf{b}\cdot\mathbf{c}$ enters via BC, only $\mathbf{c}\cdot\mathbf{d}$ via CD: every one of the six hypotheses is used.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree up to "each area has rational square; I need the product of the two areas to be rational". At that line they split. The natural attempt, having got the squares from Heron, goes back to coordinates and grinds. The model asks for a formula in which the product of two areas appears as a polynomial in lengths — and the formula exists because area is a determinant, and determinants multiply.

> *Rational squares are not enough — $\sqrt{2}$ has a rational square. I need the product* $\text{area}(ABC)\cdot\text{area}(ABD)$ *to be rational. So I want a formula for the product of two areas that only uses lengths.*

> *Area is a determinant, and* $\det M \det N = \det(M^{\mathsf{T}}N)$. *The entries of $M^{\mathsf{T}}N$ are dot products, and dot products are what squared lengths give me by polarisation. That is the whole proof.*

> *I have not used CD yet. Whatever I do next must involve CD^2, and the only place it can enter a formula in $\mathbf{b},\mathbf{c},\mathbf{d}$ is through $\mathbf{c}\cdot\mathbf{d}$.*

Layer 3 THE WHY

Why must the product thought be thought? Because of what rationality of a quotient means. The problem asks for $X/Y \in \mathbb{Q}$, and every tool that turns lengths into areas (Heron, Pythagoras, the cross product) produces the *square* of an area. From squares alone you can reach X^2/Y^2 and no further. The only algebraic route from squares to the quotient itself is $X/Y = XY/Y^2$. So the product XY is not an inspired intermediate; it is *forced* by the arithmetic of square roots. Once you write $X/Y = XY/Y^2$, the question becomes "is there a length-only formula for a product of two areas?", and the answer has to be sought.

Why is the determinant the right home for that question? Go to the definition of area in the plane. The signed area of the parallelogram on $\mathbf{u},\mathbf{v}$ is $\det(\mathbf{u},\mathbf{v})$ — that is what a 2×2 determinant is, geometrically. Lengths and angles, meanwhile, are all encoded in dot products. The bridge between the two is the one identity linking determinants to dot products: $\det(M)\det(N) = \det(M^{\mathsf{T}}N)$, whose matrix is the table of dot products between the columns of M and those of N. When $M = N$ this is the familiar $\det(\mathbf{u},\mathbf{v})^2 = |\mathbf{u}|^2|\mathbf{v}|^2 - (\mathbf{u}\cdot\mathbf{v})^2$, which is Heron in disguise. The mixed case $M \neq N$ is the product formula you were looking for. So the natural attempt was one generalisation away: Heron is the diagonal case of a Gram determinant, and the problem needs the off-diagonal case.

Why is *the plane* essential? Because in the plane both parallelograms have normal vectors along the same axis, so the product of their areas is the product of two signed numbers. In space it is not. Take $A=(0,0,0)$, $B=(1,0,0)$, $C=(0,1,0)$, $D=(0,1,1)$: the six squared distances are $AB^2=1$, $AC^2=1$, $AD^2=2$, $BC^2=2$, $BD^2=3$, $CD^2=1$, all rational, no three points collinear. But $\text{area}(\triangle ABC) = \frac{1}{2}$ and $\text{area}(\triangle ABD) = \frac{1}{2}|(1,0,0)\times(0,1,1)| = \frac{1}{2}\sqrt{2}$, so the quotient is $1/\sqrt{2}$. In $\mathbb{R}^3$ the identity reads $(\mathbf{b}\times\mathbf{c})\cdot(\mathbf{b}\times\mathbf{d}) = (\mathbf{b}\cdot\mathbf{b})(\mathbf{c}\cdot\mathbf{d}) - (\mathbf{b}\cdot\mathbf{d})(\mathbf{b}\cdot\mathbf{c})$, still rational, but it equals the product of the areas (times 4) only when the two normals are parallel

— i.e. when the four points are coplanar. That is the hypothesis "in the Euclidean plane" doing its work, and it is exactly the volume term V^2 in Rusin's third official solution. Kedlaya's first solution places B on the x-axis and grinds; here is where the product y_1y_2 it eventually needs comes from: with $\mathbf{b} = (\sqrt{a}, 0)$, $\det(\mathbf{b}, \mathbf{c}) \det(\mathbf{b}, \mathbf{d}) = a\, y_1 y_2$, so the coordinate computation is the Gram determinant written out in one particular frame.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "squares of the lengths ... are rational" together with "quotient of areas is rational". On seeing it, write $\dfrac{X}{Y} = \dfrac{XY}{Y^2}$ and underneath: "I get Y^2 from Heron; I need XY."

2. Vectors from a common vertex, dot products by polarisation, areas as determinants, and the product rule $\det M \det N = \det(M^{\mathsf{T}} N)$ (the Gram determinant).

3. Hesitation occurs immediately after Heron, when you have $X^2, Y^2 \in \mathbb{Q}$ and feel you are nearly done. What reinforces it away: test against $\sqrt{2}$ — squares being rational says nothing about the quotient — and check which hypotheses you have used. CD is unused, so you are not nearly done.

4. The trap is Heron's formula itself: it answers a question adjacent to the one asked (each area squared) and feels like progress. A second trap is forgetting the absolute value: areas are $|\det|/2$, and the proof should say that the quotient of absolute values of a rational ratio is rational.

5. Faster next time via the slogan: *to show X/Y is rational when only X^2, Y^2 are visibly rational, prove XY is rational.* And: *in the plane, a product of two areas is a polynomial in dot products.*

6. The reusable decision rule: **when a quantity only comes to you squared, look for the bilinear version of the formula that squared it.** $|\mathbf{u}|^2$ generalises to $\mathbf{u} \cdot \mathbf{v}$; $\det(\mathbf{u}, \mathbf{v})^2$ generalises to $\det(\mathbf{u}, \mathbf{v}) \det(\mathbf{x}, \mathbf{y})$; $\|f\|^2$ generalises to $\langle f, g \rangle$. The mixed version is almost always what the problem needs.

▷ **Cue:** you know X^2 and Y^2 are rational but need X/Y rational $\longrightarrow$ **Move:** write $X/Y = XY/Y^2$ and prove XY rational

▷ **Cue:** squared distances between points are given $\longrightarrow$ **Move:** pass to vectors from one point; polarisation turns squared distances into all pairwise dot products

▷ **Cue:** an area in the plane appears $\longrightarrow$ **Move:** write it as $\frac{1}{2}|\det(\mathbf{u}, \mathbf{v})|$; products of areas then become $\det(M^{\mathsf{T}} N)$, a determinant of dot products

▷ **Cue:** a hypothesis (here CD^2) has not been used $\longrightarrow$ **Move:** you are not done; identify the only quantity it can control (here $\mathbf{c} \cdot \mathbf{d}$) and find the formula that needs it

▷ **Cue:** a planar statement about lengths and areas $\longrightarrow$ **Move:** test it in $\mathbb{R}^3$ to see which step uses planarity; that step is the heart of the proof

Putnam 2018 — Paper B

2018 B1

2018 B1. Let $\mathcal{P}$ be the set of vectors defined by

$$\mathcal{P} = \left\{ \begin{pmatrix} a \\ b \end{pmatrix} \middle| 0 \le a \le 2, 0 \le b \le 100, \text{ and } a, b \in \mathbb{Z} \right\}.$$

Find all $\mathbf{v} \in \mathcal{P}$ such that the set $\mathcal{P} \setminus \{\mathbf{v}\}$ obtained by omitting vector $\mathbf{v}$ from $\mathcal{P}$ can be partitioned into two sets of equal size and equal sum.

Fields: Combinatorics, invariants **Position:** B1 – opener of the second half; disguises a one-line necessary condition behind a search space of 303 candidate vectors, and then demands you actually build the partition, not just name it **Answer:** $\mathbf{v} = (1, b)$ for any even b with $0 \le b \le 100$ (fifty-one vectors).

Layer 1 THE WHAT — surface analysis

The natural attempt. $\mathcal{P}$ has $3 \times 101 = 303$ points, arranged in three columns $a = 0, 1, 2$ against the rows $b = 0, \ldots, 100$. The natural first move is to look for a symmetry to exploit and land, correctly, on the reflection

$$\varphi(a, b) = (2 - a,\ 100 - b),$$

which maps $\mathcal{P}$ to itself and pairs every point with another point of the same "type," since $a \mapsto 2 - a$ swaps columns $0, 2$ and fixes column 1, while $b \mapsto 100 - b$ reflects each row. A student who spots this pairing is tempted to try to build the partition directly, for a general candidate $\mathbf{v}$, by throwing half the φ-pairs into S_1 and half into S_2 and hoping it works out. This is a reasonable instinct and it is most of the correct construction – but attempted blind, before asking what $\mathbf{v}$ is even allowed to be, it runs into two separate obstructions that have to be discovered by trial: first, φ has a single fixed point, $(1, 50)$, which cannot be one of a pair and so must be handled specially; second, removing an arbitrary $\mathbf{v}$ breaks whichever pair $\mathbf{v}$ sat in, leaving an odd point behind that also needs rehoming. Chasing this construction for an unrestricted $\mathbf{v}$, without first knowing which $\mathbf{v}$ can possibly work, means testing the construction against candidates that are doomed before the arithmetic even starts – real work spent on a search that a two-line argument would have pruned first.

The model chain. Before building anything, ask what any valid $\mathbf{v}$ is *forced* to satisfy. Compute the sum of every vector in $\mathcal{P}$. The a-coordinate is $0, 1, 2$ each appearing 101 times, and the b-coordinate is $0, \ldots, 100$ each appearing 3 times, so

$$\Sigma(\mathcal{P}) = \left(101(0 + 1 + 2),\ 3\textstyle\sum_{b=0}^{100} b\right) = (303,\ 15150).$$

If $\mathcal{P} \setminus \{\mathbf{v}\}$ splits into S_1, S_2 with $|S_1| = |S_2|$ and $\Sigma(S_1) = \Sigma(S_2)$, then

$$\Sigma(\mathcal{P}) = \mathbf{v} + \Sigma(S_1) + \Sigma(S_2) = \mathbf{v} + 2\Sigma(S_1),$$

so $\mathbf{v} \equiv \Sigma(\mathcal{P}) = (303, 15150) \pmod 2$ in each coordinate: $\mathbf{v}$'s first entry must be odd and its second must be even. Since $a \in \{0,1,2\}$, the only odd choice is $a = 1$. So any working $\mathbf{v}$ has the form $(1, b)$ with b even – this is necessary, and it collapses the search from 303 candidates to 51.

Now build the partition for such a $\mathbf{v} = (1, b)$. The involution $\varphi(a, b) = (2 - a, 100 - b)$ has the unique fixed point $(1, 50)$, so it pairs the other 302 points of $\mathcal{P}$ into 151 pairs, each summing to $(2, 100)$.

Case $b \neq 50$**.** The pair containing $\mathbf{v} = (1, b)$ is $\{(1, b), (1, 100 - b)\}$. Set aside this pair together with two more pairs from the same involution, $\{(2, b), (0, 100 - b)\}$ and $\{(2, 25 + \frac{b}{2}), (0, 75 - \frac{b}{2})\}$ (both are genuine φ-pairs, and both lie in $\mathcal{P}$ since b is even and $0 \leq b \leq 100$). That leaves 148 pairs; split 74 into S_1 and 74 into S_2 – equal size, equal sum $74(2, 100)$ on each side. Of the six points in the three set-aside pairs, discard $\mathbf{v} = (1, b)$ (it is being removed from $\mathcal{P}$), add back the fixed point $(1, 50)$ (which *is* present in $\mathcal{P} \setminus \{\mathbf{v}\}$ when $b \neq 50$), and assign

$$S_1 \mathrel{+}= \{(0, 100 - b),\ (2, 25 + \tfrac{b}{2}),\ (1, 50)\}, \qquad S_2 \mathrel{+}= \{(1, 100 - b),\ (2, b),\ (0, 75 - \tfrac{b}{2})\}.$$

Both additions sum to $(3,\ 175 - \frac{b}{2})$: check S_1's three points, $(0 + 2 + 1,\ 100 - b + 25 + \frac{b}{2} + 50) = (3, 175 - \frac{b}{2})$, and S_2's, $(1 + 2 + 0,\ 100 - b + b + 75 - \frac{b}{2}) = (3, 175 - \frac{b}{2})$ – identical. So S_1, S_2 each have $148 + 3 = 151$ points and equal sums, exactly the required partition of the 302 points of $\mathcal{P} \setminus \{\mathbf{v}\}$.

Case $b = 50$**.** Now $\mathbf{v} = (1, 50)$ is the fixed point of φ itself, already excluded from the 151 pairs, so $\mathcal{P} \setminus \{\mathbf{v}\}$ *is* the 302-point paired set directly. Set aside the three pairs $\{(2, 50), (0, 50)\}$, $\{(1, 51), (1, 49)\}$, $\{(0, 49), (2, 51)\}$; split the remaining 148 into 74-74 as before; assign $S_1 \mathrel{+}= \{(2, 50), (1, 51), (0, 49)\}$ (sum $(3, 150)$) and $S_2 \mathrel{+}= \{(0, 50), (1, 49), (2, 51)\}$ (sum $(3, 150)$) – equal again. Both cases give $|S_1| = |S_2| = 151$ and $\Sigma(S_1) = \Sigma(S_2)$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths notice the same reflection $\varphi(a, b) = (2 - a, 100 - b)$ early. The natural attempt reaches for φ and tries to *use* it immediately, on an as-yet-unrestricted $\mathbf{v}$. The model reaches for φ's cousin – not the pairing, but the total sum – first, and only turns to construction once the candidate set has been cut from 303 points to 51.

> *Before I try to build a partition, I should ask what invariant is forced by the very shape of the problem – equal sums on two halves mean the total sum minus* $\mathbf{v}$ *must be even in every coordinate, so let me just compute the total sum of* $\mathcal{P}$ *and read off the parity of* $\mathbf{v}$ *directly, instead of testing constructions against candidates I haven't yet shown can work.*

> *The vector* $(1, 50)$ *is exactly the fixed point of my pairing involution* φ*, so I cannot simply delete an arbitrary* $\mathbf{v}$ *from an already-built pairing of* $\mathcal{P} \setminus \{(1, 50)\}$ *– when* $\mathbf{v} \neq (1, 50)$*, the point* $(1, 50)$ *itself is still in play and needs a seat, and when* $\mathbf{v} = (1, 50)$*, there is no partner to strand. These are genuinely different bookkeeping problems and have to be built, and checked, separately.*

Layer 3 THE WHY

The parity argument is not a trick specific to this grid – it is the generic fact that whenever a problem asks "remove one element so the rest splits into two parts of equal size and equal sum," the equation $\Sigma(\text{full set}) = \mathbf{v} + 2\Sigma(S_1)$ is forced by the definitions alone, with no reference to what $\mathcal{P}$ actually looks like. The moment "equal sum, two parts" appears, this identity is available for free, and it constrains $\mathbf{v}$ modulo 2 in every coordinate before a single partition is attempted. Computing $\Sigma(\mathcal{P})$ itself is cheap precisely because $\mathcal{P}$ is a product of two ranges, so its coordinate sums split into (number of points) $\times$ (average value) – the sum of a rectangle of lattice points is never a reason to enumerate.

Why does the reflection $\varphi(a,b) = (2-a, 100-b)$ appear at all, and why is it the right tool for sufficiency rather than just a heuristic? Because a ranges over $[0,2]$ and b over $[0,100]$, both intervals symmetric about their own midpoints (1 and 50), reflection through the midpoint of a symmetric index range is the canonical involution with no moving parts: it is an automorphism of $\mathcal{P}$ as a set, and it forces every pair to carry the identical sum $(2,100)$, which is exactly the raw material a constant-sum bulk assignment needs. Once the necessary condition has trimmed $\mathbf{v}$ down to $(1,b)$, b even, the only remaining work is bookkeeping around the one point that has no partner under φ – and the case split on whether $\mathbf{v}$ *is* that fixed point is not an accident of casework but the direct consequence of φ having exactly one fixed point on an odd-sized set.

Caveat. None required: the necessary condition, both explicit constructions (for $b \neq 50$ and $b = 50$), and every stated sum were independently re-derived and checked coordinate by coordinate against the official solution, and all match.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "remove one element, then split the rest into two parts of equal size and equal sum" – on seeing it, immediately write $\Sigma(\text{full set}) = \mathbf{v} + 2\Sigma(S_1)$.

2. That identity suggests computing the total sum of the full set and reading off the removed element's residue class mod 2 in each coordinate, before attempting any construction.

3. Hesitation occurs at the construction step, specifically around the vector $(1,50)$: it is the unique fixed point of the natural pairing involution, so it needs separate treatment whether or not it is the vector being removed. The reinforcement is to explicitly identify the fixed point of any pairing symmetry before building on top of it.

4. The trap is treating the parity condition as sufficient on its own – "find all $\mathbf{v}$" demands an actual working partition for every one of the 51 surviving candidates, not just a proof that no other candidate can work.

5. Recognisable faster next time by naming the pattern: "delete-one, equal-size-equal-sum partition" is always an invariant problem first and a construction problem second, in that order.

6. Reusable rule: in any "remove one element, then equal-sum-equal-size partition" problem, compute the sum of the full set first; the removed element is forced into that sum's residue class mod 2 (coordinatewise, or mod whatever the partition structure forces), and only then attempt a construction, ideally via a natural symmetry of the index set (reflection through the midpoint of a symmetric range) that pairs up the remainder with a constant sum.

▷ **Cue:** remove one element from a set, then split the rest into two parts of equal size and equal sum $\longrightarrow$ **Move:** compute the full set's sum first; the removed element is forced into that sum's residue class mod 2 coordinatewise, cutting the candidate list before any construction is attempted

▷ **Cue:** an index set is a product of ranges each symmetric about its own midpoint $\longrightarrow$ **Move:** pair points via reflection through the midpoints; this involution has a constant pair-sum and at most one fixed point, which must be handled separately

2018 B2

2018 B2. Let n be a positive integer, and let $f_n(z) = n + (n-1)z + (n-2)z^2 + \cdots + z^{n-1}$. Prove that f_n has no roots in the closed unit disk $\{z \in \mathbb{C} : |z| \leq 1\}$.

Fields: Complex analysis, polynomials **Position:** B2 – a short statement whose difficulty is entirely in choosing the right algebraic manipulation before any inequality is applied **Answer:** No roots in $|z| \leq 1$; proved by multiplying by $(z-1)$ to telescope f_n into $z + z^2 + \cdots + z^n - n$, bounding the sum by the triangle inequality, and isolating $z = 1$ as the sole equality case, which is checked directly not to be a root.

Layer 1 THE WHAT — surface analysis

The natural attempt. The coefficients of f_n decrease from n down to 1, so the instinct is to isolate the leading term and bound everything else against it: for $|z| \le 1$,

$$|f_n(z)| \ge n - \left|(n-1)z + (n-2)z^2 + \cdots + z^{n-1}\right| \ge n - [(n-1) + (n-2) + \cdots + 1] = n - \frac{n(n-1)}{2}.$$

This is a completely legitimate use of the reverse triangle inequality, and for $n = 1, 2, 3$ it even gives a positive (or non-negative) lower bound. But the tail $(n-1) + (n-2) + \cdots + 1 = \binom{n}{2}$ grows quadratically while the leading term n grows linearly, so for every $n \ge 4$ the bound is negative and says nothing. The student has correctly identified that the coefficients need to be compared to each other, but peeling off one term at a time and bounding the rest crudely is the wrong shape of comparison: it discards the fact that the coefficients decrease by exactly 1 each step, treating $n, n-1, \ldots, 1$ as though they were an arbitrary bounded sequence instead of an arithmetic progression with real structure to exploit.

The model chain. The coefficients $n, n-1, \ldots, 1$ have constant difference 1, and a sum with linearly-varying coefficients is precisely what multiplying by $(z-1)$ is built to simplify: writing $f_n(z) = \sum_{k=0}^{n-1}(n-k)z^k$,

$$(z-1)f_n(z) = \sum_{k=0}^{n-1}(n-k)z^{k+1} - \sum_{k=0}^{n-1}(n-k)z^k.$$

Reindexing the first sum by $j = k+1$ (so j runs 1 to n) and separating the $k = 0$ term from the second, the coefficient of z^j for $1 \le j \le n-1$ is $(n-j+1) - (n-j) = 1$, the coefficient of z^n is 1 (from the $j = n$ term of the first sum alone), and there is a leftover constant $-n$ from the $k = 0$ term of the second sum. So every linear coefficient collapses to 1 and

$$(z-1)f_n(z) = z + z^2 + \cdots + z^n - n.$$

(Check at $n = 2$: $f_2(z) = 2 + z$, and $(z-1)(2+z) = z^2 + z - 2$, matching $z + z^2 - 2$.)

For $|z| \le 1$, the triangle inequality gives

$$|z + z^2 + \cdots + z^n| \le |z| + |z|^2 + \cdots + |z|^n \le n,$$

so $|(z-1)f_n(z)| = |z + \cdots + z^n - n| \ge n - |z + \cdots + z^n| \ge 0$, with equality in the middle step only when the displayed triangle inequality is itself an equality. Equality in $|z + \cdots + z^n| \le n$ requires *both* $|z| = 1$ *and* that $z, z^2, \ldots, z^n$ all point in the same direction in $\mathbb{C}$. Writing $z = e^{i\theta}$, this second condition needs $\theta \equiv 2\theta \pmod{2\pi}$ (comparing the arguments of z and z^2 alone already forces this), i.e. $\theta \equiv 0$, so $z = 1$. Thus for every $z \ne 1$ with $|z| \le 1$, the bound is strict: $|z + \cdots + z^n| < n$, hence $z + \cdots + z^n - n \ne 0$, hence $(z-1)f_n(z) \ne 0$; since also $z - 1 \ne 0$, this forces $f_n(z) \ne 0$.

The one case not covered, $z = 1$, is checked directly: $f_n(1) = n + (n-1) + \cdots + 1 = \frac{n(n+1)}{2} > 0$, so $z = 1$ is not a root either. Every point of the closed unit disk has now been accounted for, so f_n has no roots there.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same observation – the coefficients $n, n-1, \ldots, 1$ are an arithmetic progression – but they use that fact differently. At the line "$n - \frac{n(n-1)}{2}$" the natural attempt has already spent the observation: it used "the coefficients are bounded by n" and stopped there, discarding the fact that they decrease by exactly 1 at every step. The model attempt instead uses the *constant difference* itself as the lever, multiplying by $(z-1)$ so that the difference structure of the coefficients converts the whole weighted sum into an unweighted one before any inequality is applied at all.

> *Bounding the tail $(n-1)+(n-2)+\cdots+1=\binom{n}{2}$ against the leading term n is doomed before I even compute it – the tail has quadratically more weight than the head, so no direct peeling is going to work. What I actually need is an operation that turns this linearly-weighted sum into an unweighted one, and multiplying by $(z-1)$ is exactly the standard device for a coefficient sequence with constant difference: it is the discrete analogue of differentiating away a linear term.*

> *Equality in $|z+z^2+\cdots+z^n|\le n$ happens only when every term already has modulus 1 and they all point the same way in the complex plane – and forcing z and z^2 alone to share an argument pins down $z=1$ immediately, so the one case my inequality can't rule out is exactly the single point I can dispose of by direct computation, not by more inequality work.*

Layer 3 THE WHY

Why does multiplying by $(z-1)$ work, and why was it the right thing to reach for? A coefficient sequence with constant difference is, by definition, a discrete arithmetic progression, and $(z-1)$ is the operator whose effect on $\sum c_k z^k$ is to replace each coefficient with the *difference* of consecutive coefficients (up to boundary terms). Since the differences here are all 1 by construction of f_n, the differenced sum collapses to the simplest possible object, a pure sum of powers of z with all coefficients equal to 1 – which is exactly the shape the triangle inequality handles cleanly, because $|z^k| = |z|^k$ needs no further estimate. This is the same move as summation by parts (Abel summation): whenever the coefficients themselves carry structure – here, linear structure – multiplying by $(z-1)$ (or summing by parts against the sequence $1,1,\ldots,1$) trades that structure for a shorter, flatter sum, and a flat sum is always easier to bound tightly than a weighted one.

Why is the equality case in the triangle inequality worth working out explicitly, rather than just asserting "the inequality is usually strict"? Because $|f_n(z)|\ge 0$ alone is not enough – the problem asks for $f_n(z)\neq 0$, and the bound $|z+\cdots+z^n|\le n$ only forces $(z-1)f_n(z)\neq 0$ where the inequality is *strict*. The single point where it can fail to be strict, geometrically, is where the n unit vectors $z, z^2, \ldots, z^n$ are forced into perfect alignment, and on the unit circle the map $\theta\mapsto k\theta$ can only align two of its outputs, θ and 2θ, when $\theta=0$. That is not a coincidence about this polynomial; it is a fact about how rotation by a fixed angle spreads points around a circle unless the angle is zero. Once that single case is isolated, it is cheap to check directly, which is exactly why the archive favours funnelling an inequality argument down to one residual case and disposing of it by hand rather than trying to extend the inequality to cover everything.

Caveat. None required for this problem: the telescoping identity was verified both algebraically in general and numerically at $n=2$, the equality-case argument for the triangle inequality was independently re-derived from the argument condition $\theta\equiv 2\theta \pmod{2\pi}$, and $f_n(1)=n(n+1)/2>0$ was checked directly. All match the official solution's first method.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is coefficients forming an arithmetic progression – here $n, n-1, \ldots, 1$, decreasing by exactly 1 each step. On seeing it, write down "constant difference" and think "multiply by $(z-1)$" before attempting any direct bound.
2. That cue suggests the telescoping method: multiply by $(z-1)$ (or sum by parts) to collapse the linearly-weighted sum into a flat sum of pure powers, which the triangle inequality can then bound tightly.
3. Hesitation occurs right after reaching $|z+\cdots+z^n-n|\ge n-|z+\cdots+z^n|\ge 0$: it is tempting to declare victory, but a bound of ≥ 0 does not prove $\neq 0$. The reinforcement is to explicitly

work out the equality condition of the triangle inequality just used, rather than assuming the inequality is generically strict.

4. The trap is bounding the leading coefficient against the sum of the rest via the reverse triangle inequality without checking how the tail's total weight grows relative to the head – here quadratically versus linearly, so the bound fails for all but the smallest n.

5. Recognisable faster next time by naming the pattern: any polynomial whose coefficients form (or nearly form) an arithmetic progression is an instant cue for "multiply by $(z-r)$, where r matches the progression's structure, to telescope before bounding."

6. Reusable rule: to bound a polynomial with arithmetic-progression coefficients away from zero on a disk, multiply by $(z-1)$ to telescope it into a pure sum of powers, bound that sum by the triangle inequality, and isolate the single equality case (all terms of modulus 1 and equal argument) for direct, separate verification rather than trying to strengthen the inequality itself.

▷ **Cue:** polynomial coefficients form an arithmetic progression (constant difference between consecutive coefficients) $\longrightarrow$ **Move:** multiply by $(z-1)$ (difference the coefficient sequence) to telescope the sum into a flat sum of pure powers of z, then bound that with the triangle inequality

▷ **Cue:** a triangle-inequality bound on $|z_1 + \cdots + z_n| \leq n$ over $|z_i| \leq 1$ needs to be shown strict except at one point $\longrightarrow$ **Move:** isolate the equality condition – all terms of modulus 1 and identical argument – reduce it to a small system (comparing just the first two terms is often enough), and dispose of the single surviving case by direct computation

2018 B3

2018 B3. Find all positive integers $n < 10^{100}$ for which simultaneously n divides 2^n, $n-1$ divides $2^n - 1$, and $n-2$ divides $2^n - 2$.

Fields: Number theory (orders, Mersenne-type divisibility) **Position:** B3 – three divisibility conditions stacked on top of one another; the whole problem is finding one reusable lemma and applying it twice rather than inventing new machinery for each condition **Answer:** $n = 2^{2^\ell}$ for $\ell = 1, 2, 4, 8$, i.e. $n \in \{4,\ 16,\ 65536,\ 2^{256}\}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The first condition, $n \mid 2^n$, is easy: 2^n is a pure power of 2, so n can have no prime factor other than 2, forcing $n = 2^m$ for some $m \geq 0$ (and $m \geq 1$, since $n = 1$ makes $n - 1 = 0$, which cannot divide anything). A student then turns to $n - 1 \mid 2^n - 1$, i.e. $2^m - 1 \mid 2^{2^m} - 1$, and is tempted to attack it head-on: factor $2^m - 1$ for small m, compute $2^n \bmod (2^m - 1)$ by brute force, or invoke Euler's theorem and start worrying about $\varphi(2^m - 1)$ and the order of 2 modulo $2^m - 1$ from scratch, as though this were a fresh problem about an arbitrary modulus. This is real number theory and it can be pushed through case by case for small m, but it treats $2^m - 1$ and $2^n - 1$ as unrelated numbers needing separate factoring, rather than recognising that both are the *same shape* of number – a power of 2 minus 1 – linked by a single, general fact about such numbers. Worse, the third condition, $n - 2 \mid 2^n - 2$, does not even present itself as a difference of two "$2^{\text{something}} - 1$" numbers at first glance, and grinding on it directly (factoring $2^m - 2$, checking cases by hand) multiplies the casework instead of reusing what the second condition already taught.

The model chain. The reusable fact is: for positive integers a, b, $2^a - 1 \mid 2^b - 1$ if and only if $a \mid b$. (The order of 2 modulo $2^a - 1$ is exactly a, since $2^a \equiv 1$ and no smaller positive exponent $c < a$ can satisfy $2^c \equiv 1 \pmod{2^a - 1}$, because $0 < 2^c - 1 < 2^a - 1$; then $2^b \equiv 1 \pmod{2^a - 1}$ iff the order a divides b.)

First condition. $n \mid 2^n$ forces $n = 2^m$, $m \geq 1$ as above (excluding $n = 1$, where $n - 1 = 0$).

Second condition. $n-1 = 2^m - 1$ divides $2^n - 1 = 2^{2^m} - 1$ iff $m \mid 2^m$ (applying the lemma with $a = m, b = 2^m$). Now $m \mid 2^m$ iff m has no odd prime factor – if $p \mid m$ for an odd prime p, then $p \mid m \mid 2^m$ is impossible since 2^m is a pure power of 2 – so $m = 2^\ell$ for some $\ell \geq 0$ (and $\ell \geq 1$, since $\ell = 0$ gives $m = 1$, $n = 2$, where $n - 2 = 0$).

Third condition. $n - 2 = 2^m - 2 = 2(2^{m-1} - 1)$ and $2^n - 2 = 2(2^{n-1} - 1)$ share the extracted factor of 2, so $n - 2 \mid 2^n - 2$ iff $2^{m-1} - 1 \mid 2^{n-1} - 1$ iff $(m-1) \mid (n-1)$ (same lemma, $a = m-1$, $b = n-1$). Since $n - 1 = 2^m - 1$ and $m - 1 = 2^\ell - 1$, this says $(2^\ell - 1) \mid (2^m - 1)$, and since $m = 2^\ell$, the lemma applies once more ($a = \ell$, $b = m = 2^\ell$) to reduce this to $\ell \mid 2^\ell$ – exactly the same condition as in the second step, now on ℓ instead of m: by the identical argument, ℓ must itself be a power of 2.

Assembling the bound. So $\ell = 2^j$ for some $j \geq 0$, and $\ell \geq 1$ forces $j \geq 0$ genuinely (any $j \geq 0$ gives $\ell \geq 1$). The constraint $n < 10^{100}$ bounds everything from the top: $n = 2^m < 10^{100} < 2^{340}$ gives $m \leq 340$ (using $10^3 < 2^{10}$, so $10^{100} < 2^{340}$); and $m = 2^\ell \leq 340 < 512 = 2^9$ gives $\ell \leq 8$. Among the powers of two, $\ell \in \{1, 2, 4, 8\}$ are the ones satisfying $1 \leq \ell \leq 8$ (note $\ell = 16$ already exceeds 8). For these, $m = 2^\ell \leq 2^8 = 256$, so $n = 2^m \leq 2^{256} < 10^{100}$ (since $2^{256} < (2^3)^{86} = 8^{86} < 10^{86}$), confirming every one of the four candidates genuinely satisfies the size bound. Hence

$$n = 2^{2^\ell}, \qquad \ell \in \{1, 2, 4, 8\}, \quad \text{i.e. } n \in \{4,\ 16,\ 65536,\ 2^{256}\}.$$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree, at the very first step, that $n = 2^m$. The split happens at the second condition, $n - 1 \mid 2^n - 1$: the natural attempt reads $2^m - 1 \mid 2^{2^m} - 1$ as a fresh divisibility problem about a specific-looking number $2^m - 1$ and starts factoring it case by case; the model attempt recognises $2^m - 1$ and $2^{2^m} - 1$ as two instances of the same family, "$2^{\text{exponent}} - 1$," and reaches for the single lemma $2^a - 1 \mid 2^b - 1 \iff a \mid b$ instead of any fresh number theory.

> *The expression $n - 1 = 2^m - 1$ is exactly the shape $2^a - 1$, so instead of treating $n - 1 \mid 2^n - 1$ as a new number-theory problem about the specific integer $2^m - 1$, I should recognise it as another instance of the one fact I already know about differences of powers of 2: $2^a - 1$ divides $2^b - 1$ exactly when a divides b. That converts a divisibility statement about enormous numbers into a divisibility statement about the small exponents m and 2^m.*

> *The third condition, $n - 2 \mid 2^n - 2$, doesn't look like the lemma's shape at all because of the trailing -2 instead of -1 – but $n - 2 = 2(2^{m-1} - 1)$ and $2^n - 2 = 2(2^{n-1} - 1)$ share that factor of 2 exactly, so dividing it out of both sides turns the third condition into the identical Mersenne-divisibility lemma again, just between the shifted exponents $m - 1$ and $n - 1$. I don't need a new idea for the third condition; I need to apply the one I already have to a shifted pair.*

Layer 3 THE WHY

Why does the lemma $2^a - 1 \mid 2^b - 1 \iff a \mid b$ do all the work here, and why is reaching for it (instead of factoring specific numbers) the right instinct? Because the multiplicative order of 2 modulo $2^a - 1$ is forced to be exactly a: $2^a \equiv 1 \pmod{2^a - 1}$ by definition, and no smaller positive exponent can work, since $2^c - 1$ for $c < a$ is strictly between 0 and $2^a - 1$ and so cannot be a multiple of it. Once the order is pinned down as exactly a, "$2^a - 1$ divides $2^b - 1$" is by definition "$2^b \equiv 1$ modulo $2^a - 1$," which for an element of known order a happens exactly when $a \mid b$ – a fact about cyclic groups, not about the specific size of $2^a - 1$. This is why the lemma is reusable across both the second and third conditions even though the numbers involved

($2^m - 1$, then $2^\ell - 1$) are wildly different sizes: the lemma never depended on the size, only on the shape "power of 2, minus 1."

The recursive appearance of the very same sub-problem – "m divides a power of 2" reappearing as "ℓ divides a power of 2" – is not a coincidence either. Both come from the identical fact that a positive integer k divides 2^k if and only if k itself is a pure power of 2: any odd prime factor of k would have to divide 2^k, which is impossible, since 2^k's prime factorisation contains only 2's. So each condition in this problem, in turn, strips off one layer of "power of 2" structure and hands back an identical-shaped sub-problem one level down – which is exactly why the same argument gets reused twice rather than needing to be reinvented.

Caveat. None required for this problem: the order-of-2 lemma, its two applications (to $(m, 2^m)$ and to (ℓ, m)), the boundary exclusions at $n = 1$ and $n = 2$, and the numeric bound $2^{256} < 10^{100}$ were all independently re-derived and checked, and match the official solution and its final answer exactly.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "n divides 2^n" – on seeing it, write down immediately that n can have no prime factor other than 2 (since 2^n is a pure power of 2), forcing n to be a power of 2.

2. That cue suggests looking for the same "power of 2" structure in every subsequent condition, and reaching for the single lemma $2^a - 1 \mid 2^b - 1 \iff a \mid b$ wherever a modulus of the shape "power of 2 minus 1" appears.

3. Hesitation occurs at the third condition, $n - 2 \mid 2^n - 2$, because it isn't literally of the form $2^a - 1 \mid 2^b - 1$ – the reinforcement is to notice the shared factor of 2 in both $n - 2 = 2(2^{m-1} - 1)$ and $2^n - 2 = 2(2^{n-1} - 1)$, divide it out, and land back on the same lemma shape.

4. The trap is forgetting the boundary cases where a modulus is 0: $n = 1$ makes $n - 1 = 0$, and $n = 2$ makes $n - 2 = 0$, and "0 divides x" is only true for $x = 0$ – these must be excluded by hand before the general lemma is applied, not swept into the same case as everything else.

5. Recognisable faster next time by naming the shape: whenever a divisor is "(power of 2) $\pm$ small constant" and the dividend is "$2^{\text{something}} \pm$ the same small constant," look immediately for a way to reduce to $2^a - 1 \mid 2^b - 1 \iff a \mid b$, factoring out any shared constant first.

6. Reusable rule: a chain of conditions "(power of 2, shifted by a constant) divides (2^N, shifted by the same constant)" reduces, via the single lemma $2^a - 1 \mid 2^b - 1 \iff a \mid b$ (applied after factoring out any shared power of 2), to pure divisibility among the exponents – and when the resulting divisor condition is again "some integer divides a power of 2," it forces that integer to itself be a power of 2, which regenerates the identical sub-problem one level down.

▷ **Cue:** a modulus is exactly $2^a - 1$ for some exponent a, and the dividend is $2^b - 1$ $\longrightarrow$ **Move:** apply $2^a - 1 \mid 2^b - 1 \iff a \mid b$ directly – the order of 2 mod $2^a - 1$ is exactly a

▷ **Cue:** a modulus is of the form (power of 2) $\pm$ small constant c, but not literally $2^a - 1$ $\longrightarrow$ **Move:** factor out any shared power of 2 from both the modulus and the dividend to expose a genuine $2^a - 1 \mid 2^b - 1$ shape, then apply the lemma

▷ **Cue:** a positive integer k is required to divide 2^k (or another pure power of 2) $\longrightarrow$ **Move:** conclude k itself must be a power of 2, since any odd prime factor of k cannot divide a number whose only prime factor is 2

2018 B4

2018 B4. Given a real number a, we define a sequence by $x_0 = 1$, $x_1 = x_2 = a$, and $x_{n+1} = 2x_n x_{n-1} - x_{n-2}$ for $n \geq 2$. Prove that if $x_n = 0$ for some n, then the sequence is

periodic.

Fields: algebra (nonlinear recurrences); trigonometric substitution; elementary number theory (Fibonacci numbers modulo d) **Position:** B4 — the hinge of the B paper: routine once the substitution is seen, a swamp without it **Answer:** With $F_0 = 0$, $F_1 = 1$, $F_{k+1} = F_k + F_{k-1}$: if $|a| \le 1$ and $a = \cos b$ then $x_n = \cos(F_n b)$; if $|a| > 1$ no term vanishes. A zero forces b to be a rational multiple of π, and the periodicity of F_n modulo $4F_n$ makes x_n periodic.

Layer 1 THE WHAT — surface analysis

The natural attempt. Compute, which is right. From $x_0 = 1$, $x_1 = x_2 = a$:

$$x_3 = 2a^2 - 1, \qquad x_4 = 4a^3 - 3a, \qquad x_5 = 16a^5 - 20a^3 + 5a, \qquad x_6 = 128a^8 - 256a^6 + 160a^4 - 32a^2 + 1.$$

The degrees run $0, 1, 1, 2, 3, 5, 8$. That is the Fibonacci sequence, and it is the first real piece of information: the degree of x_{n+1} is $\deg x_n + \deg x_{n-1}$ because the product term dominates.

Now use the hypothesis the direct way. Suppose $x_N = 0$. The recurrence at the zero gives

$$x_{N+1} = 2x_N x_{N-1} - x_{N-2} = -x_{N-2}, \qquad x_{N+2} = 2x_{N+1}x_N - x_{N-1} = -x_{N-1},$$

$$x_{N+3} = 2x_{N+2}x_{N+1} - x_N = 2x_{N-1}x_{N-2} = x_N + x_{N-3} = x_{N-3}.$$

This looks like a reflection about the zero, and a strong student now hopes for $x_{N+k} = \pm x_{N-k}$ or similar, then a second mirror at the start (indeed, running the recurrence backwards gives $x_{-1} = x_{-2} = a$, so the sequence looks symmetric about 0 too), and two reflections compose to a translation. Push one more step:

$$x_{N+4} = 2x_{N+3}x_{N+2} - x_{N+1} = -2x_{N-3}x_{N-1} + x_{N-2},$$

and that is not $\pm x_{N-4}$. A numerical check with $a = \cos(\pi/10)$, where $x_5 = 0$, gives $x_9 \approx -0.309$ and $x_1 \approx 0.951$. The reflection breaks after three steps. That is the wall: the recurrence has a local symmetry at a zero, but the global structure is hidden, and manipulating the polynomials in a will not reveal it. Degree-F_n polynomials are not something you can see periodicity in.

The model chain. Let $F_0 = 0$, $F_1 = 1$, $F_{k+1} = F_k + F_{k-1}$, so $F_2 = 1$, $F_3 = 2$, $F_4 = 3$, and extend backwards by the same rule if needed. The key identity is

$$2\cos\theta\cos\varphi = \cos(\theta - \varphi) + \cos(\theta + \varphi). \tag{$*$}$$

Case $|a| \le 1$. Write $a = \cos b$ with $b \in [0, \pi]$. Claim: $x_n = \cos(F_n b)$ for all $n \ge 0$. It holds for $n = 0, 1, 2$ since $F_0 = 0$, $F_1 = F_2 = 1$. If it holds up to $n \ge 2$, then by $(*)$ with $F_n - F_{n-1} = F_{n-2}$ and $F_n + F_{n-1} = F_{n+1}$,

$$x_{n+1} = 2\cos(F_n b)\cos(F_{n-1}b) - \cos(F_{n-2}b) = \cos(F_{n+1}b).$$

Suppose $x_N = 0$. Then $N \ge 1$ (as $x_0 = 1$), so $F_N \ge 1$, and $\cos(F_N b) = 0$ gives $F_N b = \frac{k\pi}{2}$ with k odd. Put $d = 4F_N$; then $b = \frac{2\pi k}{d}$, and

$$x_n = \cos\left(\frac{2\pi k F_n}{d}\right)$$

depends only on F_n mod d. The map $T(u, v) = (v, u + v)$ on $(\mathbb{Z}/d\mathbb{Z})^2$ is a bijection (inverse $(u, v) \mapsto (v - u, u)$), and $(F_{n+1}, F_{n+2}) = T(F_n, F_{n+1})$. A bijection of a finite set has every orbit a cycle, so there is $L \ge 1$ with $T^L(0, 1) = (0, 1)$, whence $F_{n+L} \equiv F_n \pmod d$ for every $n \ge 0$. Therefore $x_{n+L} = x_n$ for all $n \ge 0$: the sequence is periodic.

Case $a > 1$. Write $a = \cosh t$, $t > 0$. The identity $2\cosh\theta\cosh\varphi = \cosh(\theta-\varphi)+\cosh(\theta+\varphi)$ and the same induction give $x_n = \cosh(F_n t) \geq 1$, so no term vanishes.

Case $a < -1$. Write $a = -\cosh t$, $t > 0$, and claim $x_n = (-1)^{F_n}\cosh(F_n t)$. True for $n = 0, 1, 2$. For the step,

$$2x_n x_{n-1} = (-1)^{F_n+F_{n-1}}\big(\cosh(F_{n+1}t)+\cosh(F_{n-2}t)\big) = (-1)^{F_{n+1}}\cosh(F_{n+1}t)+(-1)^{F_{n-2}}\cosh(F_{n-2}t),$$

using $F_{n+1} - F_{n-2} = 2F_{n-1}$, which is even, so the two signs agree. Subtracting x_{n-2} leaves $x_{n+1} = (-1)^{F_{n+1}}\cosh(F_{n+1}t)$. Again $|x_n| \geq 1$ and no term vanishes.

So a zero can only occur when $|a| \leq 1$, and there the sequence is periodic. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree through the computation of x_3, x_4, x_5. They split at the line written immediately after $x_3 = 2a^2 - 1$ and $x_4 = 4a^3 - 3a$ appear on the page. The natural attempt reads them as "polynomials in a of Fibonacci degree" and keeps computing. The model reads them as $\cos 2b$ and $\cos 3b$ in disguise and substitutes $a = \cos b$ before computing anything else.

> *$2a^2 - 1$ and $4a^3 - 3a$ — I have seen both of those. They are $\cos 2\theta$ and $\cos 3\theta$ written in terms of $\cos\theta$. So I should put $a = \cos b$ and see whether every term is a cosine of a multiple of b. The multiples go $0, 1, 1, 2, 3, 5$: Fibonacci. Now the recurrence $2x_n x_{n-1}$ is the product-to-sum formula, and it will turn into an addition of angles.*

> *A zero of a cosine pins the angle to $\pi/2$ modulo π. If $\cos(F_N b) = 0$ then b is a rational multiple of π, and then the sequence only depends on F_n modulo a fixed integer. Sequences defined by an invertible recurrence modulo d are periodic — that is the whole proof.*

> *The substitution only covers $|a| \leq 1$. What about $|a| > 1$? Cosine has a hyperbolic twin: $a = \cosh t$ satisfies the same product-to-sum identity, and $\cosh \geq 1$ never vanishes. For $a < -1$ I need a sign; try $(-1)^{F_n}$ and check the parity.*

- ▷ **Cue:** a nonlinear recurrence producing $2a^2 - 1$, $4a^3 - 3a$, or $x_{n+1} = 2x_n x_{n-1} - x_{n-2}$ ⟶ **Move:** substitute $a = \cos b$ (or $\cosh t$ outside $[-1, 1]$) and read the product term as the product-to-sum formula
- ▷ **Cue:** degrees of the iterates follow a recognisable integer sequence ⟶ **Move:** that sequence is the multiplier of the angle; guess $x_n = \cos(c_n b)$ with c_n that sequence and verify by induction
- ▷ **Cue:** $\cos(\text{integer} \cdot b) = 0$ at one index, periodicity required ⟶ **Move:** conclude $b \in \pi\mathbb{Q}$ and reduce to periodicity of the integer multipliers modulo the denominator
- ▷ **Cue:** an integer sequence given by an invertible linear recurrence, reduced modulo d ⟶ **Move:** it is purely periodic: the state map on $(\mathbb{Z}/d)^k$ is a bijection of a finite set

Layer 3 THE WHY

Why must the cosine substitution be thought, rather than stumbled upon?

Start from the recurrence itself: $x_{n+1} + x_{n-2} = 2x_n x_{n-1}$. A *product* of two terms equals a *sum* of two terms. The everyday function for which a product of two values is a sum of two values is cosine (with its hyperbolic twin cosh), via $(*)$. That identity is what turns multiplication into addition of angles — the same job logarithms do, but for sums of pairs instead of single values. So the form of the recurrence is a request: make x_n a cosine, and the product becomes addition in the argument. Once you ask "which angle?", the constraints $F_n \pm F_{n-1}$ must be the indices $n+1$ and $n-2$, and the Fibonacci numbers are forced: $F_n - F_{n-1} = F_{n-2}$ and $F_n + F_{n-1} = F_{n+1}$

are the defining relation read two ways. The initial data $1, a, a$ are $\cos 0, \cos b, \cos b$, matching $F_0, F_1, F_2 = 0, 1, 1$. Nothing is guessed; each piece is determined by the one before.

The computation route supplies the same cue from the other side. $2a^2 - 1 = T_2(a)$ and $4a^3 - 3a = T_3(a)$ are the Chebyshev polynomials, defined by $T_m(\cos b) = \cos(mb)$. An A-level student knows the double- and triple-angle formulas; the step is to read them backwards, from polynomial to angle.

Why does a zero give periodicity? Because the zero set of cosine is an arithmetic progression of *rational* multiples of π. Periodicity of $\cos(F_n b)$ in n requires the angles $F_n b$ to repeat modulo 2π, and for a generic real b they never do. A zero is precisely the hypothesis that makes $b/2\pi$ rational; once it is, only F_n modulo the denominator matters, and a deterministic invertible process on a finite set is periodic. The problem is a disguised statement that b is a rational angle.

Why the case split on $|a|$? Because cos maps $\mathbb{R}$ onto $[-1, 1]$ only. Kedlaya's solution handles $|a| > 1$ first by a monotonicity induction ($|x_{n+1}| \geq |x_n|$), then substitutes. That is correct, but the reason it works is the hyperbolic picture: outside $[-1, 1]$, $a = \pm\cosh t$, all terms have modulus $\cosh(F_n t) \geq 1$, and the monotonicity is the growth of cosh. Kedlaya's closing remark says you may take $b \in \mathbb{C}$ and skip the split; that is the same statement, since $\cos(\pi + it) = -\cosh t$ and $\cos(it) = \cosh t$, and the zeros of complex cosine are all real.

Finally, why did the natural reflection stall? The model explains it. With $F_N b \equiv \pi/2 \pmod{\pi}$, the neighbours satisfy $F_{N+k} = F_k F_{N+1} + F_{k-1} F_N$, so x_{N+k} depends on $F_{N+1} b$ as well as $F_N b$; the "mirror" at the zero is an artefact of the first few F_k being tiny. The honest symmetry is a translation in the angle, not a reflection in the index.

Caveat. Kedlaya's closing remark lists "$x_2 = 2a^2 - 1$, $x_3 = 4a^3 - 3a$, $x_4 = 16a^5 - 20a^3 + 5a$"; with the statement's indexing ($x_1 = x_2 = a$) these are x_3, x_4, x_5 — an index shift in the remark only, not in the solution. I verified symbolically that $x_0, \ldots, x_6$ equal $T_0, T_1, T_1, T_2, T_3, T_5, T_8$ evaluated at a; checked periodicity numerically (80-digit precision) for $a = \cos(k\pi/(2F_N))$ with $N = 3, 4, 5, 6$ and $k = 1, 3$, using the period L of F_n modulo $4F_N$; and checked the formula $x_n = (-1)^{F_n}\cosh(F_n t)$ for $a = -\cosh(0.7)$ up to $n = 14$.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $x_3 = 2a^2 - 1$, the double-angle formula. On seeing it, write "$a = \cos b$?" and compute x_4 to confirm $4a^3 - 3a = \cos 3b$.

2. The substitution $a = \cos b$, then the guess $x_n = \cos(c_n b)$ with c_n determined by forcing the recurrence through product-to-sum; c_n comes out as F_n.

3. Hesitation occurs twice. First at $|a| > 1$, where cosine does not reach; reinforce by remembering cosh satisfies the same identity (or use complex b). Second at "periodic": the student has $b \in \pi\mathbb{Q}$ and does not know how to conclude. Reinforce with the rule that an invertible recurrence modulo d is purely periodic — not merely eventually periodic — because its state map is a bijection.

4. The trap is the local reflection at a zero: three steps of $x_{N+k} = \pm x_{N-k}$ look like a proof starting, and it dies at the fourth. The second trap is proving only eventual periodicity (pigeonhole on pairs) without invertibility; the problem wants periodic from x_0.

5. Before manipulating a nonlinear recurrence, compute three terms and *name* each polynomial you get. Double- and triple-angle formulas, cosh identities, tan addition, and $\frac{x+1/x}{2}$ all appear as the iterates of recognisable recurrences; recognising one converts a nonlinear problem into a linear one on the angle.

6. Reusable decision rule: *a recurrence where a product of terms equals a sum of terms is a homomorphism in disguise — find the function* (cos, cosh, exp, tan) *that converts one operation into the other, and move to its argument, where the recurrence is linear.* Then periodicity, growth or vanishing questions become questions about a linear integer sequence modulo something.

2018 B5

2018 B5. Let $f = (f_1, f_2)$ be a function from $\mathbb{R}^2$ to $\mathbb{R}^2$ with continuous partial derivatives $\frac{\partial f_i}{\partial x_j}$ that are positive everywhere. Suppose that

$$\frac{\partial f_1}{\partial x_1}\frac{\partial f_2}{\partial x_2} - \frac{1}{4}\left(\frac{\partial f_1}{\partial x_2} + \frac{\partial f_2}{\partial x_1}\right)^2 > 0$$

everywhere. Prove that f is one-to-one.

Fields: multivariable calculus; quadratic forms (positive definiteness) **Position:** B5 — hard-looking and short: the difficulty is entirely in reading the hypothesis **Answer:** The hypothesis says the symmetric part of the Jacobian is positive definite. Along the segment from a to a', $g(t) = (a' - a) \cdot f(a + t(a' - a))$ has $g' > 0$, so $f(a) \neq f(a')$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Write $f_{ij} = \partial f_i / \partial x_j$ and $J = \begin{pmatrix} f_{11} & f_{12} \\ f_{21} & f_{22} \end{pmatrix}$. The reflex for "prove f is injective" in several variables is the Jacobian determinant. Since $\left(\frac{p+q}{2}\right)^2 \geq pq$ for all reals p, q,

$$\det J = f_{11}f_{22} - f_{12}f_{21} \;\geq\; f_{11}f_{22} - \tfrac{1}{4}(f_{12} + f_{21})^2 \;>\; 0.$$

So $\det J > 0$ everywhere, the inverse function theorem applies, and f is a local diffeomorphism: every point has a neighbourhood on which f is injective.

That is not the question. Local injectivity does not imply global injectivity: $E(x, y) = (e^x \cos y,\, e^x \sin y)$ has $\det DE = e^{2x} > 0$ everywhere and $E(0, 0) = E(0, 2\pi)$. The determinant alone cannot finish.

The next natural move is to use positivity of the partials. If $a \leq a'$ componentwise with $a \neq a'$, then f_1 increases strictly along the segment (its directional derivative is $f_{11}v_1 + f_{12}v_2 > 0$), so $f(a) \neq f(a')$. That handles comparable pairs. The remaining case is $a_1 < a_1'$, $a_2 > a_2'$ — one coordinate up, one down — and there the sign of $f_{11}v_1 + f_{12}v_2$ is undetermined. This is the wall most students hit.

It can be climbed, at a cost. Since $f_{12} > 0$, each vertical line meets the level set $\{f_1 = c\}$ at most once, and by the implicit function theorem that level set is the graph of a C^1 function $x_2 = \varphi(x_1)$ on an interval (the interval claim needs an intermediate-value argument using $f_{11} > 0$), with $\varphi' = -f_{11}/f_{12}$. Along the graph,

$$\frac{d}{dx_1} f_2(x_1, \varphi(x_1)) = f_{21} - \frac{f_{22}f_{11}}{f_{12}} = -\frac{\det J}{f_{12}} < 0,$$

so f_2 is strictly monotone on each level curve of f_1 and two distinct points with the same f_1 have different f_2. This route is correct; it uses $\det J > 0$, $f_{11} > 0$ and $f_{12} > 0$. But it costs an implicit-function argument, a connectedness argument for the domain of φ, and careful bookkeeping — and it never uses the specific form of the hypothesis, which is a sign that the setter meant something else.

The model chain. Let $a \neq a'$ in $\mathbb{R}^2$ and put $v = a' - a \neq 0$, $\gamma(t) = a + tv$ for $t \in [0, 1]$. Define

$$g(t) = v \cdot f(\gamma(t)) = v_1 f_1(\gamma(t)) + v_2 f_2(\gamma(t)).$$

Because the partials of f are continuous, f is differentiable and the chain rule gives $\frac{d}{dt} f(\gamma(t)) = J(\gamma(t))\, v$. Hence

$$g'(t) = v^{\mathsf{T}} J\, v = f_{11}v_1^2 + (f_{12} + f_{21})v_1 v_2 + f_{22}v_2^2,$$

evaluated at $\gamma(t)$. Complete the square (legitimate since $f_{11} > 0$):

$$g'(t) = f_{11}\Big(v_1 + \frac{f_{12}+f_{21}}{2f_{11}}v_2\Big)^2 + \frac{f_{11}f_{22} - \frac{1}{4}(f_{12}+f_{21})^2}{f_{11}}\,v_2^2.$$

Both coefficients are positive by hypothesis. If $v_2 \neq 0$ the second term is positive; if $v_2 = 0$ then $v_1 \neq 0$ and the first term is $f_{11}v_1^2 > 0$. So $g'(t) > 0$ on $[0,1]$, and by the mean value theorem $g(1) > g(0)$, that is,

$$v \cdot \big(f(a') - f(a)\big) > 0.$$

In particular $f(a') \neq f(a)$. So f is one-to-one. □

Nothing here used $f_{12} > 0$ or $f_{21} > 0$; the argument works verbatim in $\mathbb{R}^n$ whenever the symmetric part $\frac{1}{2}(J + J^\mathsf{T})$ is positive definite at every point.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths start by staring at the displayed inequality. They split at the first thing written about it. The natural attempt writes "this implies $\det J > 0$" — converting the hypothesis into the familiar injectivity currency and discarding its shape. The model writes "this is $ac - \frac{1}{4}b^2 > 0$ for the quadratic form $f_{11}v_1^2 + (f_{12}+f_{21})v_1v_2 + f_{22}v_2^2$" — keeping the shape and asking what object it describes.

> *Why $\frac{1}{4}(f_{12}+f_{21})^2$ and not $f_{12}f_{21}$? If the setter wanted $\det J > 0$ they would have written it. $AC - \frac{1}{4}B^2 > 0$ is the discriminant condition for $Ax^2 + Bxy + Cy^2$ to be positive definite. So the hypothesis is about the quadratic form $v^\mathsf{T}Jv$, where the off-diagonal entries only ever appear as the sum $f_{12} + f_{21}$.*

> *A quadratic form $v^\mathsf{T}Jv$ turns up when I differentiate f along a direction v and then dot with v. So take the segment from a to a', look at the component of f in the direction $v = a' - a$, and that one real function of t has derivative $v^\mathsf{T}Jv > 0$. A strictly increasing function of one variable cannot take the same value twice.*

> *Injectivity in one variable is monotonicity. In several variables I cannot ask for monotone in every direction, but I can ask: is f monotone along each segment, measured in the direction of that segment? That is exactly $(f(a') - f(a)) \cdot (a' - a) > 0$.*

- ▷ **Cue:** a hypothesis of the form $AC - \frac{1}{4}B^2 > 0$ or $AC > \frac{1}{4}(\cdot)^2$ ⟶ **Move:** read it as positive definiteness of $Ax^2 + Bxy + Cy^2$; look for where that quadratic form arises
- ▷ **Cue:** off-diagonal partials appear only as the sum $f_{12} + f_{21}$ ⟶ **Move:** the symmetric part of the Jacobian is in play, i.e. $v^\mathsf{T}Jv$; differentiate $v \cdot f(a + tv)$
- ▷ **Cue:** prove a map $\mathbb{R}^n \to \mathbb{R}^n$ is injective ⟶ **Move:** restrict to the segment between two points and dot with the segment's direction; reduce to strict monotonicity of a single real function
- ▷ **Cue:** Jacobian determinant positive everywhere ⟶ **Move:** this gives only local injectivity; do not stop there — check the exponential map $(e^x \cos y, e^x \sin y)$ as the standard counterexample

Layer 3 THE WHY

Why is reading the hypothesis as a quadratic form natural rather than lucky?

Go to the definition of injectivity. For $f : \mathbb{R} \to \mathbb{R}$ the working criterion is strict monotonicity, and the calculus test is $f' > 0$. What makes $f' > 0$ imply injectivity is the one-dimensional mean value theorem: $f(a') - f(a) = f'(\xi)(a' - a)$, and a positive number times a nonzero number

is nonzero. In two dimensions there is no mean value theorem for vector-valued f — the intermediate point depends on the component — so the only tool that survives is the scalar one. To use a scalar mean value theorem you must reduce to a scalar function of one real variable. The segment $a + tv$ gives the one variable. The scalar is a choice: which linear functional of f to track. Dotting with v is the canonical choice because it is the only direction available that is attached to the pair (a, a'); and then the one-dimensional statement "$(f(a') - f(a))(a' - a) > 0$" generalises verbatim to "$(f(a') - f(a)) \cdot (a' - a) > 0$". This is monotonicity of an operator, and its derivative test is $v^\mathsf{T} Jv > 0$.

So the chain of "of course" is: injective ⇒ monotone ⇒ mean value theorem ⇒ need a scalar ⇒ dot with v ⇒ derivative $v^\mathsf{T} Jv$ ⇒ need this positive ⇒ discriminant condition. The problem's hypothesis is the last link, printed in the statement. Reading it backwards gives the whole proof.

Why does $v^\mathsf{T} Jv$ only see $f_{12} + f_{21}$? Because $v^\mathsf{T} Jv$ is a number equal to its own transpose, $v^\mathsf{T} J^\mathsf{T} v$, so it equals $v^\mathsf{T} \frac{1}{2}(J + J^\mathsf{T})v$; the antisymmetric part of J contributes nothing. The factor $\frac{1}{4}$ is $(\frac{1}{2})^2$ from averaging the two off-diagonal entries. The hypothesis is Sylvester's criterion for the symmetric matrix $\begin{pmatrix} f_{11} & \frac{1}{2}(f_{12} + f_{21}) \\ \frac{1}{2}(f_{12} + f_{21}) & f_{22} \end{pmatrix}$: positive top-left entry and positive determinant.

Kedlaya's solution opens with "Define a real-valued function g by $g(t) = (v_1, v_2) \cdot f(\gamma(t))$." Here is where that comes from: it is the one-variable monotonicity test, transplanted by the only reduction to one variable that the pair of points supplies. The positivity of the individual partials in the statement is a decoy for this route (Kedlaya's remark says as much), but it is exactly what the level-curve route of the natural attempt needs — which is why that route feels natural and costs so much.

Caveat. Kedlaya's remark says "the hypotheses $f_{11}, f_{12} > 0$ are also superfluous" and then argues via the sign of $f_{11} f_{22}$; from that argument, f_{11}, f_{22} is meant. I verified the model chain independently. The level-curve argument in the natural attempt is my own; I have checked its steps (uniqueness of the level point on each vertical line from $f_{12} > 0$; the domain of φ is an interval by the intermediate value theorem using $f_{11} > 0$; the derivative computation) but present it only as the costly alternative.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape $f_{11} f_{22} - \frac{1}{4}(f_{12} + f_{21})^2$ — a discriminant, with the two off-diagonal partials averaged. On seeing it, write the quadratic form $f_{11} v_1^2 + (f_{12} + f_{21}) v_1 v_2 + f_{22} v_2^2 > 0$ for $v \neq 0$.

2. It suggests differentiating along a line and dotting with the line's direction: $g(t) = v \cdot f(a + tv)$, whose derivative is that quadratic form.

3. Hesitation comes after $\det J > 0$, when the inverse function theorem seems to promise more than it does. Reinforce it away with the exponential-map counterexample, which kills the idea in one line and sends you back to the hypothesis.

4. The trap is the positivity of all four partials. It invites the componentwise-monotonicity argument, which handles comparable pairs and then stalls; and it invites converting the hypothesis to $\det J > 0$, which throws away the information that makes the short proof work.

5. When a hypothesis is stated in an unusual algebraic form, ask why the setter did not write the familiar one. The gap between $\frac{1}{4}(f_{12} + f_{21})^2$ and $f_{12} f_{21}$ is the message.

6. Reusable decision rule: *to prove a multivariable map is injective, reduce to one variable along the segment joining two points and track the component in the direction of the segment; injectivity then follows from positivity of $v^\mathsf{T} Jv$.* More generally, whenever a one-variable argument hinges on

the mean value theorem, the multivariable version is "restrict to a segment and project onto a well-chosen direction".

2018 B6

2018 B6. Let S be the set of sequences of length 2018 whose terms are in the set $\{1,2,3,4,5,6,10\}$ and sum to 3860. Prove that the cardinality of S is at most

$$2^{3860} \cdot \left(\frac{2018}{2048}\right)^{2018}.$$

Fields: combinatorics; generating functions; (implicitly) probability and the saddle-point bound **Position:** B6 — the last problem of the paper, and it looks it: a four-digit count, an odd step set, a bound with a strange base. The proof, once found, is four lines. Finding it is the hard part. **Answer:** $|S|$ is the coefficient of x^{3860} in $P(x)^{2018}$, where $P(x) = x + x^2 + \cdots + x^6 + x^{10}$. All coefficients are nonnegative, so for $x > 0$ that coefficient is at most $x^{-3860}P(x)^{2018}$. At $x = \frac{1}{2}$, $P(\frac{1}{2}) = \frac{1009}{1024} = \frac{2018}{2048}$, which is the stated bound.

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong student counts. A sequence in S is fixed, up to order, by how many times each value occurs. Write c_s for the number of terms equal to s, for $s \in \{1,2,3,4,5,6,10\}$. Then

$$\sum_s c_s = 2018, \qquad \sum_s s\,c_s = 3860,$$

and

$$|S| = \sum_{(c_s)} \binom{2018}{c_1, c_2, c_3, c_4, c_5, c_6, c_{10}},$$

the sum running over all nonnegative integer solutions. Subtracting the two equations gives $\sum_s (s-1)c_s = 1842$, so the "excess" over an all-ones sequence is 1842, spread in chunks of sizes $1,2,3,4,5,9$.

The first thing to try is a crude bound, to see how far off it is. Drop the restriction on the values: every sequence of 2018 positive integers summing to 3860 is counted by stars and bars, so

$$|S| \le \binom{3859}{2017} \approx 10^{1158.06}.$$

The target is $2^{3860}(2018/2048)^{2018} \approx 10^{1149.04}$. The crude bound is off by a factor of about 10^9. The restriction to $\{1,\dots,6,10\}$ has to be used, and used hard.

The second thing to try is "number of terms times largest term". The largest multinomial can be estimated with Stirling's formula (entropy), and the number of admissible (c_s) is the number of ways to write 1842 as a sum of parts from $\{1,2,3,4,5,9\}$ — I computed it: about 1.7×10^{11}. So this route loses a factor of up to 10^{11} before it starts. How much room is there to lose? Computing $|S|$ exactly (by dynamic programming on length and sum) gives $|S| \approx 8.048 \times 10^{1146}$, and the target is about 1.10×10^{1149}. The ratio is about 137. The whole budget for slack is a factor of 137 out of a number with 1149 digits. A method that wastes even 10^3 is dead.

That is the wall. Every standard counting estimate that goes through "how many types" or "how big is the biggest" leaks polynomial factors in 2018, and the problem has been set so that there is almost no room for leaks. What survives is a method that is exact up to one inequality, and that inequality has to be the right one.

The model chain. **Encode.** Let $P(x) = x + x^2 + x^3 + x^4 + x^5 + x^6 + x^{10}$. Expanding $P(x)^{2018}$ as a product of 2018 factors, each choice of one monomial x^{s_i} from the i-th factor contributes

$x^{s_1+\cdots+s_{2018}}$. So the coefficient of x^n in $P(x)^{2018}$ is the number of sequences $(s_1,\ldots,s_{2018})$ with terms in $\{1,\ldots,6,10\}$ and sum n. In particular

$$P(x)^{2018} = \sum_n a_n x^n, \qquad a_{3860} = |S|, \qquad a_n \geq 0 \text{ for all } n.$$

Drop every other term. For any real $x > 0$, every term $a_n x^n$ is nonnegative, so

$$|S|\, x^{3860} = a_{3860} x^{3860} \leq \sum_n a_n x^n = P(x)^{2018}, \qquad \text{i.e.} \qquad |S| \leq x^{-3860}\, P(x)^{2018}. \qquad (*)$$

Choose x. Take $x = \frac{1}{2}$. Then

$$P(\tfrac{1}{2}) = \tfrac{1}{2} + \tfrac{1}{4} + \tfrac{1}{8} + \tfrac{1}{16} + \tfrac{1}{32} + \tfrac{1}{64} + \tfrac{1}{1024} = \tfrac{63}{64} + \tfrac{1}{1024} = \tfrac{1008+1}{1024} = \tfrac{1009}{1024} = \tfrac{2018}{2048}.$$

Substituting into $(*)$,

$$|S| \leq 2^{3860} \left(\frac{2018}{2048}\right)^{2018}. \qquad \square$$

(The inequality is in fact strict, since $P(x)^{2018}$ has other positive terms, e.g. x^{2018}.)

The same proof without generating functions. Pick the terms $s_1,\ldots,s_{2018}$ independently at random, each equal to s with probability $q_s = \frac{2^{-s}}{P(1/2)} = \frac{1024}{1009}\, 2^{-s}$; these seven numbers are positive and sum to 1. A particular sequence $(s_1,\ldots,s_{2018}) \in S$ then occurs with probability

$$\prod_i q_{s_i} = \left(\tfrac{1024}{1009}\right)^{2018} 2^{-(s_1+\cdots+s_{2018})} = \left(\tfrac{1024}{1009}\right)^{2018} 2^{-3860},$$

the *same* for every member of S, because the sum is fixed. Distinct sequences are disjoint events, so $|S|$ times this probability is at most 1, which is the bound again.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree on what $|S|$ is. They split at the first thing done with the bound. The natural attempt treats the right-hand side as a number to be approached from below — "estimate $|S|$ and hope it comes in under 10^{1149}". The model reads the right-hand side as a formula and asks what it is a formula *for*: 2^{3860} is x^{-n} at $x = \frac{1}{2}$ with n the prescribed sum, and $(2018/2048)^{2018}$ is something raised to the power k, the prescribed length. Once written as $x^{-n}Q^k$, the question becomes "what function of x, raised to the length, equals $1009/1024$ at $x = \frac{1}{2}$?", and $P(\frac{1}{2})$ answers it on the first try.

> *The bound is $2^{sum} \cdot (something)^{length}$. Length as an exponent means a product of one factor per term; sum as an exponent means powers of a variable tracking the total. That is a generating function $P(x)^{2018}$ looked at through the coefficient of x^{3860}. And 2^{3860} is $(1/2)^{-3860}$, so I should put $x = \frac{1}{2}$ and see whether $P(\frac{1}{2})$ is $2018/2048$.*

> *Why would $2018/2048$ appear? $2048 = 2^{11}$ and $2018/2048 = 1009/1024$. Anything over $1024 = 2^{10}$ smells like a sum of powers of $\frac{1}{2}$ reaching down to 2^{-10} — and 10 is the odd one out in the step set. So $\frac{1}{2} + \cdots + \frac{1}{64} + \frac{1}{1024}$. It works.*

> *I only need an upper bound on one coefficient of a polynomial with nonnegative coefficients. Evaluating the whole polynomial at a positive point bounds every coefficient at once, because I am only throwing away nonnegative terms.*

- ▷ **Cue:** an upper bound of the shape $c^n \cdot d^k$ for the number of length-k sequences summing to n ⟶ **Move:** write the count as $[x^n]P(x)^k$ with P the step-set polynomial; bound by $x^{-n}P(x)^k$ at $x = 1/c$ and check $P(1/c) = d$
- ▷ **Cue:** a single coefficient of a power series with nonnegative coefficients must be bounded above ⟶ **Move:** evaluate the whole series at a positive point: $a_n \le x^{-n}F(x)$ for every $x > 0$; then choose x
- ▷ **Cue:** a denominator that is a power of 2 (or any base) in a stated bound ⟶ **Move:** try to write the fraction as a sum of negative powers of that base, one for each allowed step
- ▷ **Cue:** a counting bound with very little slack (computed or suspected) ⟶ **Move:** abandon "number of types times largest type"; use an exact identity plus one inequality, i.e. a generating-function or weighting argument
- ▷ **Cue:** counting objects whose total "weight" is fixed ⟶ **Move:** give each object probability proportional to x^{weight}; all objects in the set are then equally likely, so the count is at most one over that common probability

Layer 3 THE WHY

Why is the generating function the natural object and not a lucky one? Go back to what a sequence is. A sequence of length k with terms from a set A is an element of the Cartesian product A^k; counting elements of a product is multiplying; and a constraint on the *sum* of the coordinates is exactly what exponents do when you multiply: $x^{s_1}x^{s_2}\cdots x^{s_k} = x^{s_1+\cdots+s_k}$. So $P(x)^k = \left(\sum_{s\in A} x^s\right)^k$ is not a trick applied to the problem; it is the definition of "length-k sequence from A, recorded by its sum" written in algebra. Of course the count lives in $P(x)^{2018}$.

Why evaluate at a point? A coefficient is hard to compute; a value is easy. The one piece of structure that lets a value control a coefficient is positivity: all coefficients are counts, hence nonnegative, so for $x > 0$ the value $P(x)^k$ is a sum of nonnegative pieces, one of which is $|S|x^n$. You cannot do this with signed coefficients. This is the whole content of the proof, and it is the same move as Markov's inequality: a nonnegative quantity cannot exceed its total.

Why $x = \frac{1}{2}$? There are two answers, and a solver should hold both.

The first is reading the statement. The bound $x^{-n}P(x)^k$ holds for every $x > 0$, and the problem prints $2^{3860}\cdot(\,\cdot\,)^{2018}$. The setter has told you x.

The second answer says why the setter chose $\frac{1}{2}$ and why the bound is so good. As a function of $x > 0$, $\log\left(x^{-n}P(x)^k\right) = k\log P(x) - n\log x$ is minimised where

$$\frac{xP'(x)}{P(x)} = \frac{n}{k}.$$

Kedlaya reports Givental's remark that $n = 3860$ was chosen to make $x = \frac{1}{2}$ the minimiser, and I checked it with exact arithmetic: $\frac{1}{2}P'(\frac{1}{2})/P(\frac{1}{2}) = \frac{1930}{1009}$ and $2018\cdot\frac{1930}{1009} = 3860$. The left side has a meaning. Under the random choice in the model chain (s drawn with probability $\propto x^s$), $xP'(x)/P(x)$ is the expected value of one term. So the optimal x is the one that makes the *average* random sequence have exactly the prescribed sum: $3860/2018 \approx 1.913$ per term. With that x, the event "sum $= 3860$" is the most likely single value of the sum, not a rare tail. By the local central limit theorem its probability is about $1/\sqrt{2\pi k\sigma^2}$, where σ^2 is the variance of one term. Here $\sigma^2 \approx 1.483$, so

$$\frac{\text{bound}}{|S|} \approx \sqrt{2\pi\cdot 2018\cdot 1.483} \approx 137.1,$$

and the exact computation gives 137.14. That is why only an exact-up-to-one-inequality method could work, and why "types times largest type" could not: the true slack is a factor of 137, which is $\sqrt{k}$ size, and anything that loses more than $\sqrt{k}$ fails.

Kedlaya's official solution (due to Bhargava) proves $a(k,n) < 2^n(2018/2048)^k$ by double induction on $n+k$ and $n-k$. Here is where that comes from. The recursion

$$a(k,n) = a(k-1,n-1) + \cdots + a(k-1,n-6) + a(k-1,n-10)$$

is "remove the last term". Try a bound of the form $c^n d^k$: feeding it into the recursion gives $c^n d^{k-1}(c^{-1} + \cdots + c^{-6} + c^{-10})$, and the induction closes exactly when $d \geq P(1/c)$. So the induction is the generating-function argument unrolled one term at a time, and the constant $2018/2048$ is $P(\frac{1}{2})$ in both proofs. Soundararajan's remark in the official solutions makes the same identification.

Caveat. Kedlaya's remark says the stated bound is "about 10^{1149}" (correct: $10^{1149.04} \approx 1.10 \times 10^{1149}$) and that it "differs from the true value by a factor of about 150". I recomputed $|S|$ exactly by dynamic programming over length and sum in Python's big integers: $|S|$ has 1147 digits, $|S| \approx 8.04809 \times 10^{1146}$ (matching the leading digits of the Sage output quoted in the official solutions), and the ratio bound$/|S|$ is 137.14, not 150. The local-CLT prediction $\sqrt{2\pi k\sigma^2} \approx 137.13$ agrees. I use 137 throughout; "about 150" is a loose rounding in the source. I also checked exactly that $P(\frac{1}{2}) = 1009/1024$, that $2^{3860}P(\frac{1}{2})^{2018}$ equals the stated bound, that $\binom{3859}{2017} \approx 10^{1158.06}$, and that the number of admissible (c_s) is 169,016,587,187 $\approx 1.7 \times 10^{11}$. On the induction: at $k = n = 0$ one has $a(0,0) = 1 = 2^0(2018/2048)^0$, so the official solution's strict inequality needs its base at $k \geq 1$ as stated there; the generating-function proof gives $\leq$ with no base case at all.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the shape of the bound: a base raised to the *sum* (2^{3860}) times a base raised to the *length* ($(\cdot)^{2018}$). On seeing it, write $|S| = [x^{3860}]\,P(x)^{2018}$ with $P(x) = x + \cdots + x^6 + x^{10}$, and next to it $2018/2048 = 1009/1024$.

2. The cue suggests a coefficient bound by evaluation: for $x > 0$, $[x^n]F(x) \leq x^{-n}F(x)$ when F has nonnegative coefficients. Then evaluate at the x the statement dictates, $x = \frac{1}{2}$.

3. Hesitation comes at $2018/2048$, which looks like the year dropped in for decoration. Reinforce it away by reducing the fraction: $1009/1024$ has denominator 2^{10}, and the largest step is 10. Also hesitation before evaluating at all ("surely I have to compute the coefficient"): remind yourself the problem asks for an upper bound, and upper bounds may throw away nonnegative terms.

4. The trap is the size of the numbers and the odd step set, which pull you towards explicit counting: multinomials, Stirling, stars and bars. Each is either off by 10^9 or leaks polynomial factors, against a true slack of only 137. The other trap is 3860 itself, which looks as if it must be used arithmetically; it matters only because it makes $\frac{1}{2}$ the optimal evaluation point, and the proof never touches it beyond the exponent.

5. When a bound is a product of powers whose exponents are the parameters of the problem, read each exponent as a statement of what is being multiplied or tracked, and look for the object — generating function, product of probabilities, weight — whose evaluation produces exactly that product.

6. Reusable decision rule: *to bound the number of objects with a fixed total weight, weight each object by x^{weight} and sum over all objects; the fixed-weight ones contribute $(count) \cdot x^{target}$, and nonnegativity of the rest gives count $\leq x^{-target} \cdot (total)$. Choose x so that the average weight equals the target.* This is the same bound as Chernoff's inequality and the saddle-point method, and it is sharp up to a factor of order $\sqrt{\text{size}}$.