DAWN PUTNAM SESSION

Putnam 2016 — Papers A and B

Putnam 120 Archive

Contents

Putnam 2016 — Paper A **5**
2016 A1 5
2016 A2 7
2016 A3 9
2016 A4 12
2016 A5 16
2016 A6 18

Putnam 2016 — Paper B **23**
2016 B1 23
2016 B2 25
2016 B3 28
2016 B4 31
2016 B5 34
2016 B6 37

Putnam 2016 — Paper A

2016 A1

2016 A1. Find the smallest positive integer j such that for every polynomial $p(x)$ with integer coefficients and for every integer k, the integer $p^{(j)}(k) = \frac{d^j}{dx^j}p(x)\Big|_{x=k}$ (the j-th derivative of $p(x)$ at k) is divisible by 2016.

Fields: number theory (falling factorials, binomial coefficients, Legendre's formula) **Position:** A1 – opener, but the two-directional argument (necessity and sufficiency) is a real idea, not arithmetic
Answer: $j = 8$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The quantifiers – "for every polynomial" and "for every integer k" – are intimidating, and a natural instinct is to fight them directly rather than reduce them. A student differentiates $p(x) = \sum a_m x^m$ term by term, gets $p^{(j)}(k) = \sum_{m \geq j} m(m-1)\cdots(m-j+1)a_m k^{m-j}$, and now has a genuine object in hand: a sum of terms, each a product of j consecutive integers times an arbitrary integer coefficient times a power of k. The temptation is to try to show $2016 \mid p^{(j)}(k)$ by casework on k modulo small numbers – k even or odd, $k \equiv 0, 1, 2 \pmod 3$, and so on – treating $m(m-1)\cdots(m-j+1)$ and k^{m-j} as interacting factors whose product must be tracked together. This can be made to work in principle, but it multiplies the casework by every residue class of every prime power dividing 2016, and it never isolates which part of the expression is actually doing the divisibility work: is it the falling factorial, or is it k? Worse, it gives no way to find the *smallest* j, since there is no natural stopping point in a congruence chase – the student can prove $j = 10$ works this way and have no idea whether $j = 9$ or $j = 8$ also work.

The model chain. Split the problem in two: find a necessary condition on j by testing the simplest possible polynomial, then show that condition is also sufficient.

Necessity. Take $p(x) = x^j$. Then $p^{(j)}(x) = j!$, a constant, so $p^{(j)}(k) = j!$ for every k. The hypothesis forces $2016 \mid j!$.

Sufficiency. For general $p(x) = \sum_{m=0}^{n} a_m x^m$ with integer coefficients,

$$p^{(j)}(k) = \sum_{m=j}^{n} m(m-1)\cdots(m-j+1)\, a_m\, k^{m-j} = j! \sum_{m=j}^{n} \binom{m}{j} a_m k^{m-j}.$$

Since $\binom{m}{j}$ is an integer for $m \geq j \geq 0$, and a_m, k^{m-j} are integers, every term of the sum is an integer multiple of $j!$. So $p^{(j)}(k)$ is always divisible by $j!$, regardless of p or k. Combined with necessity, the condition "$2016 \mid p^{(j)}(k)$ for every p, k" is equivalent to $2016 \mid j!$.

It remains to find the smallest such j. Factor $2016 = 2^5 \cdot 3^2 \cdot 7$, and use Legendre's formula $v_2(j!) = \sum_{i \geq 1} \lfloor j/2^i \rfloor$. Check $j = 7$: $v_2(7!) = \lfloor 7/2 \rfloor + \lfloor 7/4 \rfloor + \lfloor 7/8 \rfloor = 3 + 1 + 0 = 4 < 5$, so $2^5 \nmid 7!$ and $j = 7$ fails. Check $j = 8$: $v_2(8!) = 4 + 2 + 1 = 7 \geq 5$; also $v_3(8!) \geq v_3(6!) = 2$ (already met at $j = 6$) and $v_7(8!) \geq 1$ (met at $j = 7$). So $2016 \mid 8! = 40320 = 20 \cdot 2016$, and $j = 8$ works. Since $j = 7$ already fails on the power of 2 alone, $j = 8$ is minimal.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths reach the same differentiated sum $\sum m(m-1)\cdots(m-j+1)a_m k^{m-j}$. They split immediately after that: the natural attempt treats it as a mixed object to be attacked by congruences on k; the model chain recognises $j!\binom{m}{j}$ inside it and stops looking at k entirely.

> *The product $m(m-1)\cdots(m-j+1)$ of j consecutive integers is always divisible by $j!$ – that is exactly what makes $\binom{m}{j} = m(m-1)\cdots(m-j+1)/j!$ an integer, a count of subsets. So the divisibility of every term by* 2016*, if it happens at all, must come entirely from $j!$ dividing* 2016*; the value of k contributes nothing, and I don't need to do casework on it.*

> *To find the smallest j that could possibly work, I don't need a general polynomial at all – plug in the one polynomial where the falling-factorial coefficient is the* entire *answer, $p(x) = x^j$, which kills the k^{m-j} factor by making $m - j = 0$. That converts "for all p, for all k" into a single clean question: is* $2016 \mid j!$*?*

Layer 3 THE WHY

The move that resolves this problem is recognising that $m(m-1)\cdots(m-j+1)/j! = \binom{m}{j}$, and a binomial coefficient is an integer for the same reason it exists at all: it counts the number of j-element subsets of an m-element set, and a count of anything is a nonnegative integer by definition. So "$j!$ divides a product of j consecutive integers" is not a number-theoretic fact requiring proof from scratch each time – it is a restatement, in different notation, of what a binomial coefficient *is*. Once that identification is made, the general-p case adds nothing beyond a sum of integer multiples of $j!$, which is why testing the single monomial $p(x) = x^j$ is not a lucky special case but the extremal case: it is the one monomial for which the "extra" factor k^{m-j} collapses to $k^0 = 1$, so it isolates exactly the quantity $(j!)$ that every other term is merely a multiple of. This is why necessity and sufficiency meet at the same number $j!$ rather than at two different bounds that need reconciling.

The remaining step – finding the least j with $2016 \mid j!$ – is where the factorisation $2016 = 2^5 \cdot 3^2 \cdot 7$ matters structurally, not decoratively: because $j!$'s prime factorisation grows in each prime independently (Legendre's formula treats each prime separately), the three divisibility requirements $2^5 \mid j!$, $3^2 \mid j!$, $7 \mid j!$ can be checked one at a time, and the answer is the maximum of the three thresholds – here 2^5 is the binding constraint because powers of 2 accumulate the slowest relative to 2016's demand on them among the three prime powers present.

Caveat. I re-derived both directions independently: necessity via $p(x) = x^j \Rightarrow p^{(j)}(k) = j!$, and sufficiency via the binomial-coefficient identity, then computed $v_2(7!) = 4$, $v_2(8!) = 7$ by hand and confirmed $8! = 40320 = 20 \times 2016$. This matches Kedlaya and Ng's stated answer of $j = 8$ exactly, so the chain is fully verified, not merely paraphrased.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "for every polynomial p and every integer k" attached to a fixed target divisor – this is a signal to find the extremal/simplest case that forces a necessary condition, rather than trying to reason about the general case first.
2. That cue suggests testing the single monomial $p(x) = x^j$, since differentiating a pure power j times leaves only a constant, collapsing "for every k" to a non-issue.
3. Hesitation occurs right after writing out $p^{(j)}(k) = \sum m(m-1)\cdots(m-j+1)a_m k^{m-j}$, because the expression looks like it needs casework on k; the reinforcement is to notice the falling factorial and divide by $j!$ on sight, converting it to $\binom{m}{j}$.

4. The trap is chasing congruences on k across all the prime-power factors of 2016 – this can be forced to work but never produces the minimal j, because there is no stopping criterion inside a case-by-case congruence argument.

5. Recognisable faster next time by treating any "j-th derivative of an integer polynomial at an integer" as $j!\sum\binom{m}{j}a_m k^{m-j}$ on sight – the factor of $j!$ should be written down the moment a j-th derivative of a monomial appears.

6. Reusable rule: when a problem demands "N divides [some derivative-or-difference expression] for every choice of integer inputs," first test the single simplest input that collapses all the free variables but one, to get a necessary condition; then prove sufficiency separately using an integrality identity (here, the binomial coefficient) rather than direct congruence casework.

▷ **Cue:** j-th derivative of an integer-coefficient polynomial evaluated at an integer $\longrightarrow$ **Move:** rewrite as $j!\sum\binom{m}{j}a_m k^{m-j}$ and read off divisibility by $j!$ immediately, ignoring k

▷ **Cue:** "smallest j such that N divides [expression] for every choice of polynomial/input" $\longrightarrow$ **Move:** test the simplest extremal case (e.g. a pure monomial) to get a necessary condition on j, then verify sufficiency via an integrality/combinatorial identity

2016 A2

2016 A2. Given a positive integer n, let $M(n)$ be the largest integer m such that $\binom{m}{n-1} > \binom{m-1}{n}$. Evaluate $\lim_{n\to\infty}\frac{M(n)}{n}$.

Fields: combinatorics / algebra (binomial coefficient ratios, quadratic optimisation, limits) **Position:** A2 – second problem, algebra-heavy rather than conceptually deep **Answer:** $\dfrac{3+\sqrt{5}}{2}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The definition of $M(n)$ invites direct computation: work out $M(1), M(2), M(3), \ldots$ by hand, testing small m against the inequality for each n, and try to spot a pattern in $M(n)/n$ numerically – perhaps guessing it tends to e, or 2, or some other familiar constant, and then trying to justify the guess after the fact. This has two costs. First, computing $\binom{m}{n-1}$ and $\binom{m-1}{n}$ for even moderate n by hand is slow and error-prone, and the trend only becomes visible after enough terms to make the arithmetic painful. Second, and more seriously, guessing the constant from numerics gives no route to proving it is exactly right – $(3+\sqrt{5})/2$ is not a value anyone would recognise from a handful of decimal approximations of $M(n)/n$, so a numerics-first approach risks converging on the wrong guess (something like 2.6 could easily be mistaken for e − something or a rational approximation) and then spending the rest of the problem trying to prove a false claim. A related but more promising natural move is to try to bound $M(n)$ using Stirling's approximation on the binomial coefficients directly; this can work, but turns a clean algebraic identity into an asymptotic argument with error terms that have to be controlled, which is far more machinery than the problem needs.

The model chain. Form the ratio, not the difference, of the two binomial coefficients – consecutive binomial coefficients are related by a clean rational function, so dividing collapses the inequality to algebra. For $m > n+1$ (so both coefficients are positive),

$$\frac{\binom{m}{n-1}}{\binom{m-1}{n}} = \frac{m!\,n!\,(m-n-1)!}{(m-1)!\,(n-1)!\,(m-n+1)!} = \frac{mn}{(m-n+1)(m-n)}.$$

So the condition $\binom{m}{n-1} > \binom{m-1}{n}$ is equivalent to

$$(m-n+1)(m-n) - mn < 0.$$

Fix n and treat this as a quadratic in m. Expanding: $(m-n+1)(m-n) - mn = m^2 - (2n-1)m + n(n-1) - mn = m^2 - (3n-1)m + n(n-1)$. Its roots are

$$\alpha(n), \beta(n) = \frac{(3n-1) \pm \sqrt{(3n-1)^2 - 4n(n-1)}}{2} = \frac{3n-1 \pm \sqrt{5n^2-2n+1}}{2},$$

both real since $5n^2 - 2n + 1 = 4n^2 + (n-1)^2 > 0$. The quadratic is negative exactly for $\beta(n) < m < \alpha(n)$, so the inequality holds exactly on that open interval. Since $\alpha(n) - \beta(n) = \sqrt{5n^2-2n+1} > 1$ for $n \geq 1$, there is always at least one integer strictly inside it.

$M(n)$ is therefore the largest integer strictly less than $\alpha(n)$, so $\alpha(n) - 1 \leq M(n) < \alpha(n)$. Divide by n and let $n \to \infty$:

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

The same limit holds for $(\alpha(n) - 1)/n$, so by the sandwich theorem $\lim_{n\to\infty} M(n)/n = \dfrac{3+\sqrt{5}}{2}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths recognise the problem asks for the boundary of an inequality between binomial coefficients that grows with n. They split at how that boundary is located: the natural attempt tries to locate it empirically (compute $M(n)$ for small n, or estimate the coefficients' sizes via Stirling), while the model chain forms the exact ratio $\binom{m}{n-1}/\binom{m-1}{n} = \dfrac{mn}{(m-n+1)(m-n)}$ and solves the resulting quadratic in m exactly.

> *This is a strict inequality between two binomial coefficients with adjacent, closely related parameters – I should form the* ratio $\binom{m}{n-1}/\binom{m-1}{n}$*, not try to compare or subtract them directly, because a ratio of nearby binomial coefficients collapses to a rational function of m, n in one line, while a difference does not simplify at all.*

> *After clearing denominators, the condition is a genuine quadratic in m for fixed n – I should stop thinking of m as ranging over integers found by search and instead solve the quadratic exactly with the quadratic formula, treating m as continuous; the true boundary $\alpha(n)$ is an exact algebraic number, and $M(n)$ is then pinned between $\alpha(n) - 1$ and $\alpha(n)$ automatically, with no need to locate $M(n)$ itself precisely.*

Layer 3 THE WHY

The reason forming the ratio is the right first move, rather than an arbitrary trick, is the defining recursive structure of binomial coefficients: adjacent entries of Pascal's triangle are related multiplicatively, $\binom{m}{k+1}/\binom{m}{k} = (m-k)/(k+1)$, not additively – the factorial definition $\binom{m}{k} = m!/(k!(m-k)!)$ makes ratios of binomial coefficients telescope into products of linear factors, while differences of binomial coefficients have no such closed form. So whenever a problem compares two binomial coefficients with an inequality (as opposed to asking for a sum or an exact difference), forming the ratio is not a guess to try among several – it is the operation for which the factorial structure was built to simplify.

Once the ratio is a rational function $\dfrac{mn}{(m-n+1)(m-n)}$, comparing it to 1 and clearing denominators is legitimate specifically because the denominator $(m-n+1)(m-n)$ is a product of two positive integers once $m > n+1$ (which must be checked, and is the reason the problem's regime of interest is m large relative to n) – clearing denominators when the sign of what you're multiplying by is known is not a step that needs re-justifying each time; it is exactly what "the

inequality is equivalent to" means. The resulting object being a quadratic in m is then forced by the algebra, not chosen: a quadratic's sign is negative exactly between its two roots, by the very definition of what a parabola opening upward does, so $M(n)$ being "the largest integer below the larger root" is not something to search for – it falls straight out of locating $\alpha(n)$.

Finally, the limit computation $\alpha(n)/n \to (3+\sqrt{5})/2$ is the standard device for a $\sqrt{\text{quadratic in } n}/n$ limit: divide every term inside the square root by n^2 before taking the limit, so that $\sqrt{5n^2-2n+1}/n \to \sqrt{5}$ rather than being left as an indeterminate ∞/∞ expression. This is worth stating explicitly because it is the one place a numerics-only approach could not have found the answer: $(3+\sqrt{5})/2 \approx 2.618$ is recognisable algebraically (it is φ^2 for the golden ratio φ) but not something a handful of computed values of $M(n)/n$ would visibly converge to without already suspecting the closed form.

Caveat. I re-derived the ratio identity $\binom{m}{n-1}/\binom{m-1}{n} = mn/[(m-n+1)(m-n)]$, the quadratic $m^2-(3n-1)m+n(n-1)$, its roots, and the limit $(3+\sqrt{5})/2$ independently by direct algebra, and each step matches Kedlaya and Ng's stated solution exactly, so this chain is fully verified.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is an inequality (not an equation) between two binomial coefficients with related, nearby parameters; the move to write down immediately is the ratio $\binom{m}{n-1}/\binom{m-1}{n}$, since ratios of binomial coefficients are always rational functions of the indices.

2. That ratio's rational-function form suggests clearing denominators (once their sign is known) to convert the comparison into a polynomial inequality, here a quadratic in m.

3. Hesitation occurs when deciding whether to solve the quadratic exactly or estimate $M(n)$ by search; the reinforcement is to notice that "largest m satisfying a quadratic inequality" is always "the integer part of the larger root," so the quadratic formula gives the answer directly, with no search needed.

4. The trap is trying to intuit the limiting constant from a handful of computed values of $M(n)/n$; irrational algebraic constants like $(3+\sqrt{5})/2$ are not recognisable from a short numerical list, and time spent guessing is time not spent deriving the exact quadratic.

5. Recognisable faster next time by treating "compare two binomial coefficients with an inequality" as an automatic cue to form the ratio and simplify via the factorial definition, rather than reasoning about the coefficients' relative sizes qualitatively.

6. Reusable rule: when a problem defines an extremal integer m (largest/smallest) via an inequality between binomial coefficients (or other ratio-friendly combinatorial quantities) depending on a parameter n, form the exact ratio, clear denominators once their sign is controlled, solve the resulting polynomial inequality exactly, and read off the growth rate as $n \to \infty$ by dividing inside any square root by the highest power of n before taking the limit.

▷ **Cue:** a strict inequality between two binomial coefficients with adjacent or closely related parameters $\longrightarrow$ **Move:** form the exact ratio (not the difference) using the factorial definition, since ratios of binomial coefficients simplify to rational functions

▷ **Cue:** a limit of (root of a quadratic in n)$/n$ as $n \to \infty$ $\longrightarrow$ **Move:** divide every term inside any square root by the highest power of n before taking the limit, to avoid an indeterminate form

2016 A3

2016 A3. Suppose that f is a function from $\mathbb{R}$ to $\mathbb{R}$ such that $f(x)+f\left(1-\frac{1}{x}\right)=\arctan x$ for all real $x \neq 0$. (As usual, $y=\arctan x$ means $-\pi/2 < y < \pi/2$ and $\tan y = x$.) Find

$\int_0^1 f(x)\,dx$.

Fields: functional equations / analysis (Möbius substitutions of finite order, arctan identities) **Position:** A3 – requires spotting a 3-cycle and combining it with a second, domain-preserving symmetry; a genuine two-idea problem **Answer:** $\frac{3\pi}{8}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The target is a definite integral, and the functional equation links $f(x)$ to f at a shifted argument, $\tau(x) := 1 - \frac{1}{x}$. The standard trick for "integrate a functional equation" problems is to integrate both sides over the domain of interest and substitute to relate $\int f(\tau(x))\,dx$ back to $\int f$. So a strong student writes

$$\int_0^1 f(x)\,dx + \int_0^1 f\left(1 - \frac{1}{x}\right) dx = \int_0^1 \arctan x\,dx,$$

and tries to handle the second integral by the substitution $u = 1 - 1/x$. This is where the wall appears: as x ranges over $(0,1)$, $u = 1 - 1/x$ ranges over $(-\infty, 0)$, not over $(0,1)$. The substitution is perfectly legal as calculus – $du = dx/x^2$ – but it produces $\int_{-\infty}^0 f(u) \cdot (\text{something in } u)\,du$, an integral of f over a completely different, unbounded interval, which is a new unknown quantity, not a multiple of $\int_0^1 f(x)\,dx$. The single functional equation, integrated as given, does not close the loop: it relates the target integral to an integral of f somewhere else entirely. A student can spend real time trying to make this substitution work – attempting a second substitution to fold $(-\infty, 0)$ back onto $(0,1)$, or trying to bound the unwanted integral – before recognising that integrating the equation as literally given is the wrong operation to perform first.

The model chain. Solve for $f(x)$ pointwise before integrating anything. The substitution $\tau(x) = 1 - \frac{1}{x} = \frac{x-1}{x}$ has order 3:

$$\tau(x) = \frac{x-1}{x}, \qquad \tau(\tau(x)) = \frac{1}{1-x}, \qquad \tau(\tau(\tau(x))) = x.$$

So writing the original equation at x, at $\tau(x)$, and at $\tau^2(x)$ gives three linear equations in the three unknowns $f(x), f(\tau(x)), f(\tau^2(x))$:

$$\begin{aligned} f(x) + f\left(\tfrac{x-1}{x}\right) &= \arctan x, \\ f\left(\tfrac{x-1}{x}\right) + f\left(\tfrac{1}{1-x}\right) &= \arctan\left(\tfrac{x-1}{x}\right), \\ f\left(\tfrac{1}{1-x}\right) + f(x) &= \arctan\left(\tfrac{1}{1-x}\right). \end{aligned}$$

Adding the first and third, subtracting the second, cancels the two unwanted values and doubles $f(x)$:

$$2f(x) = \arctan x + \arctan\left(\frac{1}{1-x}\right) - \arctan\left(\frac{x-1}{x}\right).$$

Now, for $x \in (0,1)$, form $2f(x) + 2f(1-x)$ by writing the same formula with x replaced by $1-x$ and adding. The six arctan terms pair off using $\arctan(t) + \arctan(1/t) = \frac{\pi}{2}$ for $t > 0$:

$$\left[\arctan x + \arctan \tfrac{1}{x}\right] + \left[\arctan(1-x) + \arctan \tfrac{1}{1-x}\right] + \left[\arctan \tfrac{1-x}{x} + \arctan \tfrac{x}{1-x}\right] = \frac{3\pi}{2},$$

(each bracket is a reciprocal pair with both arguments positive on $(0,1)$), so $f(x) + f(1-x) = \frac{3\pi}{4}$ for all $x \in (0,1)$. Finally, using $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$ (substitute $u = 1-x$, which *does* map $(0,1)$ onto itself),

$$2\int_0^1 f(x)\,dx = \int_0^1 \big(f(x) + f(1-x)\big)dx = \frac{3\pi}{4}, \qquad \int_0^1 f(x)\,dx = \frac{3\pi}{8}.$$

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths start from the same equation and both eventually need a symmetry that lets $\int_0^1 f$ appear on both sides of an equation. They split at *which* substitution is used to try to close that loop. The natural attempt integrates the given equation directly and substitutes $u = 1 - 1/x$, the map the equation itself supplies – but that map sends $(0,1)$ to $(-\infty, 0)$. The model chain instead solves pointwise for $f(x)$ using the fact that τ has order 3, and only afterwards introduces $x \mapsto 1 - x$, a different map entirely, chosen specifically because it *does* send $(0,1)$ to itself.

> *$1 - 1/x$ does not stay inside $(0,1)$ for $x \in (0,1)$, so integrating the functional equation directly over $x \in [0,1]$ cannot relate $\int_0^1 f(x)\,dx$ to itself – I need a closed-form expression for $f(x)$ first, as a function of x alone, before integrating anything.*

> *$\tau(x) = 1 - 1/x$ applied three times returns to x – so instead of one equation in infinitely many unknown values of f, substituting x, $\tau(x)$, $\tau^2(x)$ gives exactly three equations in exactly three unknowns, $f(x)$, $f(\tau x)$, $f(\tau^2 x)$. That is a closed linear system to be solved by elimination, not a single relation to be integrated.*

> *Once I have $2f(x)$ as a sum of three arctan terms, the integration trick I actually need is $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$, which uses the reflection $x \mapsto 1 - x$ – a completely different map from τ, chosen because it is the one substitution that keeps everything inside $[0,1]$, unlike τ itself.*

Layer 3 THE WHY

The reason the pointwise approach succeeds where direct integration fails is that $\tau(x) = 1 - 1/x$ is a Möbius transformation, and composing it with itself is governed by matrix multiplication on its coefficients; the concrete fact that $\tau^3 = \mathrm{id}$ is what turns "f satisfies one equation relating two unknown values" into "f satisfies a closed system of three equations in three unknowns along a finite orbit." A functional equation of the form $f(x) + f(\tau x) = g(x)$ is solvable in closed form precisely when the orbit of x under τ is finite – if τ had infinite order, writing down the equation at $x, \tau(x), \tau^2(x), \ldots$ would never close, and no finite elimination would isolate $f(x)$. So checking $\tau^3 = \mathrm{id}$ before attempting to eliminate is not a computational nicety; it is the structural fact that makes elimination possible at all, and it is why "add equation 1, subtract equation 2, add equation 3" produces $2f(x)$ cleanly: it is Gaussian elimination on a system whose coefficient pattern (each unknown appears in exactly two of the three equations) is forced by the 3-cycle.

The identity $\arctan(t) + \arctan(1/t) = \pi/2$ for $t > 0$ is itself not an isolated trick to memorise for this problem; it is the statement that in a right triangle, the two non-right angles are complementary, and t, $1/t$ are the tangents of those two angles. Recognising the three reciprocal pairs buried inside the six-term sum $2f(x) + 2f(1-x)$ is recognising this geometric fact three times over, which is why the messy-looking sum of six arctan terms collapses so completely: nothing is being estimated or cancelled by luck, each bracket is exactly a pair of complementary angles.

Finally, $\int_0^1 f(x)\,dx = \int_0^1 f(1-x)\,dx$ is not a trick specific to this problem either – it is the definition of the interval $[0,1]$ being symmetric about its midpoint, so integration over it is invariant under the reflection $x \mapsto 1 - x$. This is precisely the substitution the natural attempt was missing: the functional equation's own substitution τ does not respect the domain $[0,1]$, but the domain's own reflection symmetry does, and the two maps are genuinely different objects serving different purposes – τ solves for f pointwise, $x \mapsto 1 - x$ closes the integral.

Caveat. I re-derived $\tau^3 = \mathrm{id}$, the closed form $2f(x) = \arctan x + \arctan \frac{1}{1-x} - \arctan \frac{x-1}{x}$, the six-term arctan sum collapsing to $3\pi/2$, and the final value $3\pi/8$ independently by direct computation, and every step matches Kedlaya and Ng's stated solution exactly, so this chain is fully verified, not merely paraphrased.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is a functional equation of the form $f(x) + f(\tau(x)) = g(x)$ for a specific rational substitution τ; the first thing to write down is $\tau(\tau(x))$ and $\tau(\tau(\tau(x)))$, to check whether τ has finite order.

2. Finding $\tau^3 = \mathrm{id}$ suggests substituting $x, \tau(x), \tau^2(x)$ into the equation to get three linear equations in three unknown values of f, then eliminating by add-subtract-add to isolate $2f(x)$.

3. Hesitation occurs when deciding how to get from a pointwise formula for $f(x)$ to the integral $\int_0^1 f(x)\,dx$; the reinforcement is to ask which substitution actually maps $[0,1]$ to itself – here $x \mapsto 1-x$, not τ – and to use that one for the integral, keeping the two substitutions conceptually separate.

4. The trap is integrating the original functional equation directly and substituting $u = \tau(x)$, since τ maps $(0,1)$ outside itself; this produces an integral of f over an unrelated interval instead of closing the loop.

5. Recognisable faster next time by checking the order of any substitution appearing in a functional equation as the very first step, before attempting to integrate or manipulate the equation in any other way.

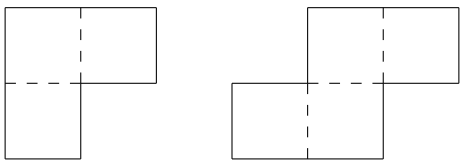
6. Reusable rule: for $f(x) + f(\tau(x)) = g(x)$ with τ of finite order n, solve the resulting $n \times n$ linear system for $f(x)$ pointwise first; only then, if an integral or sum over a specific domain is wanted, look for a *separate* symmetry of that domain (often $x \mapsto a - x$ or $x \mapsto 1/x$) to close the final computation – do not assume τ itself is the right map for that last step.

▷ **Cue:** a functional equation $f(x) + f(\tau x) = g(x)$ for a specific rational substitution τ $\longrightarrow$ **Move:** compute $\tau^2(x)$, $\tau^3(x)$, etc. to find τ's order; if finite, substitute along the orbit to get a closed linear system for $f(x)$

▷ **Cue:** needing $\int_0^1 f(x)\,dx$ after solving for f pointwise via one substitution $\longrightarrow$ **Move:** look for a second, domain-preserving symmetry (e.g. $x \mapsto 1-x$) to pair $f(x)$ with f(its image) and close the integral, rather than reusing the original substitution

2016 A4

2016 A4. Consider a $(2m-1) \times (2n-1)$ rectangular region, where m and n are integers such that $m, n \geq 4$. This region is to be tiled using tiles of the two types shown:



(The dotted lines divide the tiles into 1×1 squares.) The tiles may be rotated and reflected, as long as their sides are parallel to the sides of the rectangular region. They must all fit within the region, and they must cover it completely without overlapping.

What is the minimum number of tiles required to tile the region?

Fields: combinatorics: tilings, colouring invariants, construction matched to a bound **Position:** A4 — the middle of the paper; the bound is a one-line colouring once seen, the construction is fiddly but finite

Answer: mn tiles.

Layer 1 THE WHAT — surface analysis

The natural attempt. The two tiles are an L-tromino (three squares, a 2×2 block with one corner missing) and an S-tetromino (four squares, two dominoes offset by one). The strong student's first instinct is area. The region has area $(2m-1)(2n-1) = 4mn - 2m - 2n + 1$. Every tile has area at most 4, so at least $\lceil (4mn - 2m - 2n + 1)/4 \rceil = mn - \lfloor (2m + 2n - 1)/4 \rfloor$ tiles are needed. That is a lower bound, and it is the wrong one: it is smaller than mn by roughly $(m+n)/2$.

So the student tries to build a tiling using as many S-tetrominoes as possible, hoping to reach the area bound. Draw a 7×7 square ($m = n = 4$). The area bound says $\lceil 49/4 \rceil = 13$ tiles. Try to place S-tetrominoes: they pack into diagonal staircases, and the staircases keep colliding with the edges of the square, where an odd side length leaves ragged single cells that only an L-tromino can pick up. After ten minutes with a pencil you have tilings with 16 tiles and no idea whether 13, 14 or 15 is achievable. That is the wall. Area counts cells; it cannot see that the *shape* of the region — odd by odd — forbids the S-tetromino from being useful near the boundary.

The honest position after the natural attempt: a construction with 16 tiles for 7×7, a lower bound of 13, and a growing suspicion that the answer is mn because $16 = 4 \cdot 4$. What is missing is a lower-bound argument that sees shape rather than area.

The model chain. *Coordinates.* Label the unit squares (i, j) with $1 \le i \le 2m-1$ (column) and $1 \le j \le 2n-1$ (row). Call a square *red* if i and j are both odd. There are m odd values of i and n odd values of j, so exactly mn red squares.

Lower bound. Each tile covers at most one red square.

- The L-tromino lies inside a 2×2 block. A 2×2 block contains two consecutive columns (exactly one odd) and two consecutive rows (exactly one odd), hence exactly one red square. So the tromino covers at most one.
- The S-tetromino, in either orientation, consists of two dominoes lying in two consecutive lines (rows, say), each domino covering two consecutive cells of its line. Only one of the two lines has odd index, and within it the domino covers two consecutive positions, only one of which is odd. So at most one red square.

Every red square must be covered, so at least mn tiles are needed.

Counting what equality forces. If exactly mn tiles are used, with t_3 trominoes and t_4 tetrominoes, then $t_3 + t_4 = mn$ and $3t_3 + 4t_4 = (2m-1)(2n-1)$, giving

$$t_4 = (2m-1)(2n-1) - 3mn = mn - 2m - 2n + 1 = (m-2)(n-2) - 3.$$

For $m = n = 4$ this is $t_4 = 1$, $t_3 = 15$. (It is also why the hypothesis $m, n \ge 4$ is there: for $m = n = 3$ it would be negative.)

Strips. A $2 \times (2k-1)$ strip with $k \ge 2$ is tiled by k tiles. Index its cells by row $\in \{1, 2\}$ and column $1, \ldots, 2k-1$. Take the L-tromino $\{(1,1), (2,1), (1,2)\}$ [row, column]; then for $i = 1, \ldots, k-2$ the S-tetromino $\{(2, 2i), (2, 2i+1), (1, 2i+1), (1, 2i+2)\}$; and finally the L-tromino $\{(2, 2k-2), (2, 2k-1), (1, 2k-1)\}$. These are disjoint and cover all $4k-2$ cells: $3 + 4(k-2) + 3 = 4k-2$.

The 7×7 square with 16 tiles. Split it into a pinwheel of four 3×4 rectangles around the centre cell, cut each into two 3×2 (or 2×3) rectangles, tile seven of those eight with two L-trominoes each, and absorb the centre cell into the eighth. Explicitly, with equal letters marking one tile

(rows top to bottom):

$$\begin{array}{ccccccc} a & a & c & c & e & e & f \\ a & b & D & c & e & f & f \\ b & b & D & D & g & g & h \\ m & m & n & D & g & h & h \\ m & n & n & i & i & k & k \\ o & o & p & i & j & k & l \\ o & p & p & j & j & l & l \end{array}$$

Here D is the single S-tetromino (it contains the centre cell) and the other fifteen letters are L-trominoes. I checked this array cell by cell: 16 tiles, all legal shapes, each covering exactly one red square.

Assembly. For general $m, n \geq 4$, place the 7×7 square in a corner. The remaining $2m - 8$ columns beside it, of height 7, split into $m - 4$ strips 7×2, each a $2 \times (2 \cdot 4 - 1)$ strip needing 4 tiles. The remaining $2n - 8$ rows, of full width $2m - 1$, split into $n - 4$ strips $2 \times (2m - 1)$, each needing m tiles. Total:

$$16 + 4(m - 4) + m(n - 4) = mn.$$

(I also machine-checked the assembled tilings for all $4 \leq m, n \leq 8$.) So the minimum is exactly mn.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths start by asking "what does a tile use up?" At (†) — the choice of *which resource to count* — the natural attempt counts area, a resource every cell consumes equally. The model counts a sparse set of cells that no tile can use twice. Area gives $\approx mn - (m + n)/2$; the sparse set gives mn exactly.

> *Area is losing me $(m + n)/2$ tiles, and that loss scales with the perimeter. So the obstruction lives on the boundary, in the oddness of the sides. I want cells spaced so far apart that no single tile can reach two of them — a lattice with spacing 2 — and $(2m - 1) \times (2n - 1)$ holds exactly $m \times n$ of those.*

> *Every tile lives in two consecutive rows (or two consecutive columns), and in each of those lines it occupies at most two consecutive cells. So if I mark every other cell in every other row, starting from the corner, one tile can reach at most one mark. Now count the marks.*

Layer 3 THE WHY

Why is the odd–odd lattice the natural thing to count, rather than a trick someone happened to know? Drive to what a lower bound for a tiling *is*. A lower bound on the number of tiles is always of the form "there is a quantity Q that the region has, and each tile can absorb at most q of it; so at least Q/q tiles". Area is one choice of Q. The question you must ask is: which choice of Q makes q as small as possible relative to Q? The optimal Q is concentrated where the tiles are weakest.

Now read the problem's phrasing. The sides are $2m - 1$ and $2n - 1$: odd, and specifically *one less than an even number*. The tiles all fit in a 2×3 box, so their "reach" is 2 in one direction. The question "how many cells can I choose, pairwise at distance ≥ 2 in some coordinate, so that no tile of width 2 spans two of them?" has the answer "every other cell, starting from a corner", and $2m - 1$ was chosen precisely so that "every other" gives m and not $m - 1$ or $m - \frac{1}{2}$. The number $2m - 1$ is the problem announcing its own colouring. Once you read $2m - 1$ as "m odd positions", the answer mn is visible in the statement before any tile is placed.

There is also a second "of course". The lower bound is tight only if every tile covers *exactly* one red square, and this forces the tile counts: $t_4 = (m - 2)(n - 2) - 3$. For the 7×7 square that

means 15 trominoes and one tetromino. This is where the construction stops being trial and error. With one tetromino and a centre cell that must be covered, the design question becomes "cover everything except the centre with trominoes, then fix the centre with the tetromino". Trominoes pair into 2×3 rectangles; $48 = 8 \times 6$; and the classical pinwheel of four 3×4 rectangles around a 1×1 hole in a 7×7 square exists because $7 = 3 + 4$. Kedlaya and Ng say "cover all but the center square with eight 2×3 rectangles" without saying why the centre: it is because the equality count tells you there is exactly one odd-man-out tile, and the pinwheel is the standard way to leave exactly the centre uncovered.

And the reduction to 7×7 plus strips is the counting identity read backwards: $mn = 16 + 4(m-4) + m(n-4)$ is only possible because each $2 \times (2k-1)$ strip contains exactly k red squares and is tiled with exactly k tiles. Everything in the construction is the bound, respected locally.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the side lengths $2m-1$, $2n-1$ in a minimum-tiling question. Write down immediately: "odd positions along a side of length $2m-1$: exactly m." Then look at the tiles' bounding boxes (2×2, 2×3).

2. The cue suggests a colouring lower bound: mark a set of cells no tile can cover twice, and count them. Then build a construction in which every tile covers exactly one marked cell.

3. Hesitation occurs after the area bound, when it visibly falls short of mn and you start drawing. What reinforces it away: noticing that the gap between the area bound and the conjectured answer grows like the perimeter, which says the obstruction is a boundary/parity effect and area will never see it. Second hesitation: the 7×7 construction. Reinforce it with the equality count (15 trominoes, 1 tetromino), which turns a search into a design.

4. The trap is area — a correct but weak bound that invites hours of fruitless tiling in pursuit of a number (about 13 for 7×7) that is unachievable. A secondary trap is to overlook that the S-tetromino must be checked in *both* orientations for the "at most one red square" claim; the argument is symmetric under swapping rows and columns, which is why one sentence covers both.

5. Recognise it faster by the phrase "minimum number of tiles" with small tiles and a region whose dimensions are of the form $2m-1$ or $km-1$. The minimum is then usually the size of the sparsest lattice the region contains, spaced by the tiles' reach.

6. The reusable decision rule: **for a lower bound on the number of pieces, count a resource that each piece can use at most once, and choose the resource where the pieces are weakest (usually a sparse lattice or the boundary), not the one they all consume equally (area). If the bound is tight, the equality case tells you the composition of any optimal configuration — build the construction from that composition, not by search.**

▷ **Cue:** a tiling or covering problem asks for the minimum number of pieces and the area bound is visibly too weak $\longrightarrow$ **Move:** find a set of cells pairwise too far apart for one piece to cover two; the number of such cells is the lower bound

▷ **Cue:** a side length is written $2m-1$ (or $km-1$) in a grid problem $\longrightarrow$ **Move:** read it as "exactly m positions $\equiv 1 \pmod 2$" and colour the odd–odd lattice; the answer is often $m \cdot n$

▷ **Cue:** a lower bound N is conjectured to be tight and pieces come in two sizes $\longrightarrow$ **Move:** solve the two linear equations (count and area) for the numbers of each piece; the equality case dictates the construction's composition

▷ **Cue:** a 7×7 (or any $(a+b) \times (a+b)$) square must be covered leaving only the centre $\longrightarrow$ **Move:** use the pinwheel of four $a \times b$ rectangles around a 1×1 hole ($7 = 3 + 4$)

▷ **Cue:** a construction must work for all $m, n \geq m_0$ $\longrightarrow$ **Move:** solve the base case $m_0 \times n_0$ by hand, then extend with strips that each respect the bound exactly, so the counts add

2016 A5

2016 A5. Suppose that G is a finite group generated by the two elements g and h, where the order of g is odd. Show that every element of G can be written in the form

$$g^{m_1}h^{n_1}g^{m_2}h^{n_2}\cdots g^{m_r}h^{n_r}$$

with $1 \leq r \leq |G|$ and $m_1, n_1, m_2, n_2, \ldots, m_r, n_r \in \{-1, 1\}$. (Here $|G|$ is the number of elements of G.)

Fields: group theory (finite groups, generated subgroups), combinatorics (pigeonhole, shortest paths in a Cayley digraph) **Position:** A5 — hard position, but the proof is short once the problem is split into "some length" and "length at most $|G|$" **Answer:** A proof: the four blocks $g^{\pm 1}h^{\pm 1}$ generate G (using that g has odd order), and a shortest representation has length at most $|G|$ by pigeonhole on prefixes.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student reads "every element can be written as a word" and starts writing words. Call the four allowed blocks $x_{\pm\pm} = g^{\pm 1}h^{\pm 1}$. Any word in the problem's form is a product $x_1x_2\cdots x_r$ of blocks. The natural goal: produce g and h themselves as such products, since they generate G.

Try g. A block always contains exactly one $g^{\pm 1}$ and one $h^{\pm 1}$, so a product of blocks has alternating letters and starts with g and ends with h. You cannot write g as a word of that shape — the literal word g ends in g. Try g^2: $gh \cdot h^{-1}g$ is g^2, but $h^{-1}g$ is not a block. Try $(gh)(g^{-1}h)^{-1}$: that is $ghh^{-1}g = g^2$, but it uses an *inverse* of a block, which the form does not allow. So in the free world of words the student hits a wall: blocks can only be multiplied, never inverted, and the obvious combinations all need an inverse.

The second part of the natural attempt is to worry about the bound $r \leq |G|$. The student computes small cases — say $G = \mathbb{Z}/3$ — and sees short words suffice, but has no mechanism for a general bound. Even if generation were settled, they would be trying to bound the length of a specific explicit word, which depends on how G is presented. That is a second wall, and it is the same wall in disguise: the student is trying to build words, when the problem is really about the set of elements words can reach.

The model chain. Write a *representation of s of length r* for a choice of $m_1, n_1, \ldots, m_r, n_r \in \{\pm 1\}$ with $s = g^{m_1}h^{n_1}\cdots g^{m_r}h^{n_r}$, $r \geq 1$. Let $S \subseteq G$ be the set of elements having a representation of some length.

Step 1: S is a subgroup. Concatenating representations gives a representation of the product, so S is closed under multiplication, and $S \neq \emptyset$ (it contains gh). For $a \in S$ let $k \geq 1$ be the order of a in the finite group G. Then $a^k = e$, so $e \in S$ (as a^k is a product of k copies of a representation of a); and $a^{-1} = a^{k-1}$ if $k \geq 2$, while $a^{-1} = a = e$ if $k = 1$. Either way $a^{-1} \in S$. So S is a subgroup.

Step 2: $S = G$. $gh \in S$ and $g^{-1}h \in S$, so by Step 1

$$(gh)(g^{-1}h)^{-1} = ghh^{-1}g = g^2 \in S.$$

Let the order of g be $2k+1$. Then $(g^2)^{k+1} = g^{2k+2} = g$, so $g \in S$, hence $g^{-1} \in S$, and $h = g^{-1}(gh) \in S$. Since S is a subgroup containing the generators, $S = G$.

Step 3: the length bound. Fix $s \in G$ and take a representation of s of the smallest possible length r. Suppose $r > |G|$. Define the prefixes

$$s_i = g^{m_1}h^{n_1}\cdots g^{m_i}h^{n_i} \qquad (0 \leq i \leq r-1), \quad s_0 = e.$$

These are $r > |G|$ elements of G, so $s_i = s_j$ for some $0 \leq i < j \leq r-1$. Deleting blocks $i+1, \ldots, j$ gives

$$s = s_i\left(g^{m_{j+1}}h^{n_{j+1}}\cdots g^{m_r}h^{n_r}\right) = s_j\left(g^{m_{j+1}}h^{n_{j+1}}\cdots g^{m_r}h^{n_r}\right),$$

a representation of length $r-(j-i)<r$. It still has length ≥ 1, because $j \leq r-1$ means block r survives. This contradicts minimality, so $r \leq |G|$.

Why the prefixes stop at $r-1$. If one used $s_1, \ldots, s_r$ instead, the coincidence could be $s_i = s_r$, and deleting blocks $i+1, \ldots, r$ could leave length 0 when $i = 0$ — which the problem forbids ($r \geq 1$). Stopping at $r-1$ guarantees the last block is kept. I checked the statement by brute force on random two-generator subgroups of S_3, S_4, S_5 with g of odd order: in every case the shortest representation had length at most $|G|$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that the blocks $g^{\pm 1}h^{\pm 1}$ are the real alphabet. They split at (†): "I need g^2 but that needs the inverse of a block." At (†) the natural attempt keeps looking for a word; the model changes the object from *words* to the *set S of reachable elements*, and observes that in a finite group a set closed under multiplication is automatically closed under inverses.

> *I can't invert a block inside the word, but I don't need to: in a finite group, a^{-1} is a positive power of a. Anything closed under multiplication is a subgroup. So inverses are free, and I can use $(gh)(g^{-1}h)^{-1} = g^2$ without apology.*

> *Odd order is there for exactly one reason: to get g back from g^2. If g has order $2k+1$, then $g = (g^2)^{k+1}$.*

> *The bound $|G|$ is the number of vertices. A shortest path can't visit a vertex twice, so I should look at the sequence of partial products and apply pigeonhole, not try to bound an explicit word.*

Layer 3 THE WHY

Why is passing from words to the set S natural rather than a trick? Go back to what "can be written in the form" means. It is a statement about the image of a map from words to G. Images are sets; the question "is the image all of G?" is a question about a set, and the only tool that shows a subset of a group is everything is "it is a subgroup containing generators". So the moment the problem says "show every element can be written", the first object to define is the set of elements that can, and the first thing to ask is whether it is a subgroup. That is the circle-and-centre "of course" here: *every* "every element can be written as" problem in a group reduces to closure properties of the set of writable elements.

The second "of course" is finiteness. The problem says *finite* group, and the one thing finiteness does for a subset closed under multiplication is to give inverses: the powers $a, a^2, a^3, \ldots$ must repeat, so $a^k = e$ for some $k \geq 1$, and $a^{-1} = a^{k-1}$. In an infinite group this fails — the positive powers of 2 in $(\mathbb{Q}^\times, \cdot)$ are closed under multiplication and contain no inverses. So the word "finite" is in the statement to license exactly the inverse the natural attempt was missing.

Third: why g^2, and why odd order? Look at the four blocks. Each has one $g^{\pm 1}$ and one $h^{\pm 1}$. Comparing two blocks that share the h-part cancels the h: $(gh)(g^{-1}h)^{-1} = g^2$. The exponent is 2 because the two blocks differ by g^1 versus g^{-1}. So blocks give you g^2, never g directly, and g^2 generates $\langle g \rangle$ precisely when the order of g is odd. Kedlaya and Ng's remark makes the obstruction concrete: if G has an index-2 subgroup H containing neither g nor h, then every block $g^{\pm 1}h^{\pm 1}$ lies in H (two elements of the non-trivial coset multiply into H), and all words stay in H. The smallest example: $G = \mathbb{Z}/2$, $g = h = 1$; every block is 0. Odd order of g kills this, because an element of odd order cannot lie outside an index-2 subgroup.

Fourth: the bound $|G|$. Kedlaya and Ng introduce the prefixes $s_0, \ldots, s_{r-1}$ without motivation. Here is where they come from. A word in the blocks is a walk in the Cayley digraph on G whose edges are right-multiplication by the four blocks. A representation of s is a walk from e to

s of length r. A shortest walk between two vertices never repeats a vertex, so it has at most $|G|$ vertices after the start — that is the whole bound. The prefixes are the vertices the walk visits. Stopping at s_{r-1} rather than s_r is the one place where care is needed: the problem demands $r \geq 1$, so for $s = e$ the "walk" must be a genuine cycle, not the empty walk, and keeping the last edge is how that is enforced.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "every element of G can be written in the form". Write down: "Let S be the set of elements that can be so written; show S is a subgroup containing g and h." Then separately: "Bound the length of a *shortest* representation."

2. Method: closure $\Rightarrow$ subgroup (using finiteness for inverses); a commutator-like combination of two blocks to isolate a pure power of one generator; a parity/order argument to recover the generator; pigeonhole on partial products for the length.

3. Hesitation occurs at "I need an inverse and the form forbids it." What reinforces it away is the finiteness reflex: in a finite group, closure under multiplication implies closure under inverses. A second hesitation is at the length bound; reinforce with "shortest path in a graph with $|G|$ vertices has length at most $|G|$".

4. Two traps. (i) Trying to write explicit words for g and h, which cannot work directly because every word starts with $g^{\pm 1}$ and ends with $h^{\pm 1}$. (ii) In the pigeonhole step, using prefixes $s_1, \ldots, s_r$ and deleting down to length 0 when $s = e$; the problem requires $r \geq 1$.

5. Recognise it faster by the combination "finite group" + "every element is a product of the following restricted shapes" + "with at most $|G|$ factors". That trio is always: subgroup generated by the shapes, then shortest path in the Cayley digraph.

6. The reusable decision rule: **to show every element of a structure is reachable, define the set of reachable elements and prove it is closed under the operations that generate the structure; in a finite group, multiplicative closure already gives a subgroup. To bound how many steps are needed, look at a shortest route and apply pigeonhole to the states it passes through — a shortest route never revisits a state.**

▷ **Cue:** "show every element of the finite group G can be written as a product of elements of the following kind" $\longrightarrow$ **Move:** let S be the set of such products; prove S is closed under multiplication, hence (finiteness) a subgroup; show it contains generators

▷ **Cue:** you need an inverse but only products are allowed, in a finite group $\longrightarrow$ **Move:** use $a^{-1} = a^{\mathrm{ord}(a)-1}$

▷ **Cue:** allowed pieces differ only in the sign of one generator's exponent $\longrightarrow$ **Move:** divide one piece by another to isolate the square of that generator; recover the generator from its square using odd order

▷ **Cue:** "with at most $|G|$ factors" (or at most the number of states) $\longrightarrow$ **Move:** take a shortest representation, list its prefixes, pigeonhole a repeat, excise the loop

▷ **Cue:** a representation must have length at least 1 $\longrightarrow$ **Move:** pigeonhole on prefixes $s_0, \ldots, s_{r-1}$, not $s_1, \ldots, s_r$, so the excision never empties the word

2016 A6

2016 A6. Find the smallest constant C such that for every real polynomial $P(x)$ of degree 3 that has a root in the interval $[0, 1]$,

$$\int_0^1 |P(x)|\, dx \leq C \max_{x \in [0,1]} |P(x)|\,.$$

Fields: analysis (polynomial inequalities, numerical quadrature), linear optimisation over a convex set
Position: A6 — the hardest slot on the paper, but the difficulty is entirely in seeing which tool applies; the verification is a page **Answer:** $C = \frac{5}{6}$, attained by $P(x) = 4x^3 - 8x^2 + 5x$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student normalises: scale so that $\max_{[0,1]} |P| = 1$ and try to maximise $\int_0^1 |P|$. The absolute value is unpleasant, so they consider cases by the number of roots of P in $[0,1]$. In the simplest case, one root $a \in [0,1]$ and P of one sign on each side, they write $P(x) = (x-a)(bx^2+cx+d)$ and try to optimise over a, b, c, d subject to $|P| \le 1$ on $[0,1]$.

This goes badly fast. The constraint $|P| \le 1$ is infinitely many inequalities; the student has to guess where the maximum of $|P|$ is attained (an interior critical point? an endpoint? both?), and each guess is a separate calculus problem with a messy system from $P' = 0$. Test cases give numbers: $P(x) = x^3$ gives $\frac{1}{4}$; $P(x) = x$ (not degree 3, but a limit) gives $\frac{1}{2}$; a bump like $P(x) = 4x(1-x)$ adjusted to degree 3 gives about $\frac{2}{3}$. A careful student might try cubics that stay near their maximum over a long stretch and creep towards 0.8. But without a proof of an upper bound, each example is only a lower bound for C, and the optimisation over four parameters with an infinite family of constraints is a wall. Many students guess a Chebyshev-type extremal and get the wrong number.

The model chain. *Step 1: reduce to $P(0) = 0$ and $P \ge 0$.* Let $0 \le r_1 < \cdots < r_k \le 1$ be the roots of P in $[0,1]$ ($k \ge 1$ by hypothesis). They cut $[0,1]$ into intervals $I_0 = [0, r_1]$, $I_i = [r_i, r_{i+1}]$, $I_k = [r_k, 1]$ (some possibly degenerate), each having a root of P at one endpoint at least, and P of constant sign inside. Suppose we know

$$(*) \qquad \int_0^1 Q(t)\,dt \le \tfrac{5}{6} \max_{[0,1]} Q$$

for every polynomial Q of degree at most 3 with $Q(0) = 0$ and $Q \ge 0$ on $[0,1]$. For an interval $I = [u, v]$ of length $\ell = v - u > 0$ with root at u, put $Q(t) = \pm P(u + \ell t)$ with the sign chosen so $Q \ge 0$ (if the root is at v, use $Q(t) = \pm P(v - \ell t)$). Then $\int_I |P| = \ell \int_0^1 Q$ and $\max_I |P| = \max_{[0,1]} Q$, so $(*)$ gives $\int_I |P| \le \frac{5}{6}\,\ell(I) \max_I |P| \le \frac{5}{6}\,\ell(I) \max_{[0,1]} |P|$. Summing over the intervals, whose lengths add to 1:

$$\int_0^1 |P| \le \tfrac{5}{6} \max_{[0,1]} |P|.$$

Step 2: prove $(*)$ by Simpson's rule. For any polynomial Q of degree at most 3,

$$\int_0^1 Q(t)\,dt = \tfrac{1}{6}\big(Q(0) + 4Q(\tfrac{1}{2}) + Q(1)\big).$$

Both sides are linear in Q, so it suffices to check $1, t, t^2, t^3$: the left sides are $1, \frac{1}{2}, \frac{1}{3}, \frac{1}{4}$; the right sides are $\frac{1}{6}(1+4+1) = 1$, $\frac{1}{6}(0+2+1) = \frac{1}{2}$, $\frac{1}{6}(0+1+1) = \frac{1}{3}$, $\frac{1}{6}(0+\frac{1}{2}+1) = \frac{1}{4}$. With $Q(0) = 0$ and $Q(\frac{1}{2}), Q(1) \le M := \max_{[0,1]} Q$,

$$\int_0^1 Q \le \tfrac{1}{6}(0 + 4M + M) = \tfrac{5}{6}M.$$

Step 3: sharpness. Equality needs $Q(0) = 0$, $Q(\frac{1}{2}) = Q(1) = M$, and (since $\frac{1}{2}$ is interior and $Q \le M$) $Q'(\frac{1}{2}) = 0$. With $M = 1$, write $Q(t) = 1 + (t - \frac{1}{2})^2(at + b)$; then $Q(0) = 0$ gives $b = -4$ and $Q(1) = 1$ gives $a = 4$:

$$Q(t) = 1 + 4(t-1)(t-\tfrac{1}{2})^2 = 4t^3 - 8t^2 + 5t.$$

Check: $Q'(t) = 12t^2 - 16t + 5 = (2t-1)(6t-5)$, and $Q(0) = 0$, $Q(\frac{1}{2}) = 1$, $Q(\frac{5}{6}) = \frac{25}{27}$, $Q(1) = 1$. So Q rises from 0 to 1, dips to $\frac{25}{27} > 0$, rises to 1: it is nonnegative with maximum 1 on $[0,1]$, has degree 3 and a root at $0 \in [0,1]$, and $\int_0^1 Q = 1 - \frac{8}{3} + \frac{5}{2} = \frac{5}{6}$. Hence no $C < \frac{5}{6}$ works,

and $C = \frac{5}{6}$ is the answer. (A random search over 2×10^5 cubics with a root in $[0,1]$ gave ratios up to 0.8306, below and close to $\frac{5}{6} \approx 0.8333$, consistent with this.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains normalise $\max|P| = 1$ and split $[0,1]$ at sign changes. They split at (†): "now bound $\int_0^1 Q$ for a nonnegative cubic." At (†) the natural attempt parametrises Q and tries to optimise; the model writes the integral as a fixed linear combination of three point values, and bounds each value by M.

> *The integral of a cubic is a linear function of its four coefficients, so it is a linear combination of any four point values — and Simpson's rule says three suffice: $\frac{1}{6}Q(0) + \frac{4}{6}Q(\frac{1}{2}) + \frac{1}{6}Q(1)$. I don't need to know where the maximum is; every point value is at most the maximum.*

> *The hypothesis "has a root in $[0,1]$" is there to give me a zero I can put at an endpoint. Then the $Q(0)$ term vanishes and the bound is $\frac{4}{6} + \frac{1}{6}$, not 1.*

> *Equality in a bound made of point values forces those point values. So the extremal cubic is pinned down by $Q(0) = 0$, $Q(\frac{1}{2}) = Q(1) = 1$, $Q'(\frac{1}{2}) = 0$ — four conditions, four coefficients.*

Layer 3 THE WHY

Why is quadrature the natural move rather than a remembered trick? Drive to the structure of the two sides. The left side, $\int_0^1 Q$, is a *linear* functional on the four-dimensional space of cubics. The right side, $\max_{[0,1]} Q$, is a bound on *point values*. The only way a linear functional is controlled by point values is by writing it as a combination of point values. Linear algebra guarantees that you can: evaluation at any four distinct points is a basis of the dual space of cubics, so $\int_0^1$ is some combination of four evaluations. So the question is never "is there a quadrature formula?" — there always is — but "which nodes make the weights best?" You want a node at the known root (so its weight is multiplied by 0) and positive weights elsewhere (so that bounding each value by M is legitimate). Simpson's rule, with nodes $0, \frac{1}{2}, 1$ and weights $\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$, is exact for cubics by the symmetry of the nodes about $\frac{1}{2}$ (odd powers of $t - \frac{1}{2}$ integrate to zero), and it has all weights positive and a node at 0. It is the formula the situation asks for.

Kedlaya and Ng's first solution writes $\frac{1}{4}a + \frac{1}{3}b + \frac{1}{2}c = \frac{2}{3}P(\frac{1}{2}) + \frac{1}{6}P(1)$ as if pulled from nowhere. Here is where it comes from: it is Simpson's rule with $P(0) = 0$ substituted, discovered by asking "which combination of $P(\frac{1}{2})$ and $P(1)$ equals the integral of $ax^3 + bx^2 + cx$?" — three unknown coefficients, two point values, one free choice killed by $P(0) = 0$. Their remark phrases the same thing geometrically: the set S of cubics with $P(0) = 0$ and $0 \le P \le 1$ is convex, $\int$ is linear, so the maximum is on a supporting plane, and the two constraints $P(\frac{1}{2}) \le 1$, $P(1) \le 1$ are the faces of S that the plane touches.

The reduction step also deserves its "of course". The integrand is $|P|$, and absolute value is only linear on intervals of constant sign. The roots of P are exactly where the sign may change. So cutting at roots is forced by the absolute value, and every piece so obtained has a root at an endpoint — which is exactly the property Simpson's rule wants. The hypothesis "has a root in $[0,1]$" guarantees that even the end pieces $[0, r_1]$ and $[r_k, 1]$ have a root at one endpoint. Without it the answer is 1: $P(x) = 1 + \varepsilon x^3$ has ratio tending to 1.

Finally, why must the extremal have an interior maximum at $\frac{1}{2}$ touching M? Because equality in Simpson's bound requires $Q(\frac{1}{2}) = M$, and an interior point where a function attains its maximum is a critical point. That is Fermat's theorem, and it is the fourth condition that pins down the cubic.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $\int_0^1 |P|$ against $\max |P|$ with P of degree 3. Write down: "integral of a cubic = weighted sum of point values (Simpson: $\frac{1}{6}, \frac{4}{6}, \frac{1}{6}$ at $0, \frac{1}{2}, 1$)." And from "has a root in $[0,1]$": "move the root to an endpoint."

2. The cue suggests exact quadrature plus a reduction: split at roots, affinely rescale each piece to $[0,1]$ with the root at 0, apply Simpson, bound each point value by the maximum.

3. Hesitation occurs at the absolute value ("how many sign changes? which case?"). What reinforces it away: the splitting argument handles every case at once because each piece has a root at an endpoint and the bound scales with length. A second hesitation is at sharpness; reinforce with "equality forces the point values, plus $Q' = 0$ at the interior node".

4. The trap is optimising over coefficients subject to $|P| \leq 1$ on $[0,1]$ — an infinite family of constraints that invites casework on where the maximum occurs. A second trap is guessing a Chebyshev polynomial as extremal; Chebyshev polynomials extremise the sup norm for a given leading coefficient, which is not this problem.

5. Recognise it faster by any inequality comparing an integral of a low-degree polynomial with its values or maximum. The first thing to try is an exact quadrature rule whose nodes include the points where you know something (roots, endpoints).

6. The reusable decision rule: **when a linear functional (integral, average, derivative at a point) must be bounded by pointwise information on a finite-dimensional space, write the functional as a combination of point evaluations, choosing the nodes where you have information and seeking positive weights. The bound is then the sum of the weights at the free nodes, and equality forces the extremal.**

▷ **Cue:** $\int_0^1 P$ for P of degree at most 3 $\longrightarrow$ **Move:** use Simpson's rule, exact for cubics: $\frac{1}{6}\big(P(0) + 4P(\frac{1}{2}) + P(1)\big)$

▷ **Cue:** an integral of $|P|$ with P a polynomial $\longrightarrow$ **Move:** cut at the roots, so that on each piece $|P| = \pm P$ and each piece has a root at an endpoint

▷ **Cue:** a hypothesis "has a root in the interval" $\longrightarrow$ **Move:** rescale so the root is an endpoint and the corresponding quadrature term vanishes

▷ **Cue:** an upper bound built from point values $Q(x_i) \leq M$ must be shown sharp $\longrightarrow$ **Move:** impose $Q(x_i) = M$ at the nodes and $Q'(x_i) = 0$ at interior nodes; solve the resulting linear system for the extremal

▷ **Cue:** an inequality between a linear functional and a max over a compact convex set of polynomials $\longrightarrow$ **Move:** the extremum lies on a supporting hyperplane; the active constraints are the point evaluations with positive weight

Putnam 2016 — Paper B

2016 B1

2016 B1. Let $x_0, x_1, x_2, \ldots$ be the sequence such that $x_0 = 1$ and for $n \geq 0$, $x_{n+1} = \ln(e^{x_n} - x_n)$ (as usual, the function ln is the natural logarithm). Show that the infinite series $x_0 + x_1 + x_2 + \cdots$ converges and find its sum.

Fields: analysis (sequences, telescoping series) **Position:** B1 – the paper's opener, but the exponentiate-first move is the whole problem **Answer:** The sum converges to $e - 1$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The recursion $x_{n+1} = \ln(e^{x_n} - x_n)$ is defined by composing two named functions, exp then subtract-x_n then ln, and the natural first move is to understand the sequence (x_n) itself: is it positive, is it monotonic, does it converge? Since $e^x - x$ has derivative $e^x - 1 > 0$ for $x > 0$ and equals 1 at $x = 0$, the function $e^x - x$ is strictly increasing on $x > 0$ and exceeds 1 there; by induction $x_n > 0$ for all n (starting from $x_0 = 1 > 0$, if $x_n > 0$ then $e^{x_n} - x_n > 1$ so $x_{n+1} = \ln(e^{x_n} - x_n) > 0$). Also $e^{x_n} - x_n < e^{x_n}$, so $x_{n+1} < x_n$: the sequence is positive and decreasing, hence converges to some limit $L \geq 0$. Passing to the limit in the recursion, $L = \ln(e^L - L)$, i.e. $e^L = e^L - L$, forcing $L = 0$. So $x_n \to 0$.

This establishes that the sequence itself converges to 0 – but the problem asks about the *series* $\sum x_n$, and $x_n \to 0$ is not remotely enough for that: $1/n \to 0$ too, and $\sum 1/n$ diverges. The natural next move is to try to pin down the *rate* at which $x_n \to 0$, using a Taylor expansion. For small x, $e^x - x = 1 + x + \frac{x^2}{2} + O(x^3) - x = 1 + \frac{x^2}{2} + O(x^3)$, so $x_{n+1} = \ln\left(1 + \frac{x_n^2}{2} + O(x_n^3)\right) \approx \frac{x_n^2}{2}$. This says x_n decays doubly-exponentially once it is small (roughly $x_n \sim C \cdot 2^{1-2^n}$ for large n), which is easily enough to make $\sum x_n$ converge by comparison with a geometric-type series. That settles convergence – but it gives no way to find the *exact value* of the sum: the Taylor estimate is only asymptotic, it loses information from every early term, and there is no route from "x_n roughly halves its exponent each step eventually" to a closed-form total. The natural attempt can prove convergence but is structurally incapable of finding the sum.

The model chain. Exponentiate the defining equation. Since $x_{n+1} = \ln(e^{x_n} - x_n)$, raising e to both sides gives

$$e^{x_{n+1}} = e^{x_n} - x_n,$$

which rearranges to

$$x_n = e^{x_n} - e^{x_{n+1}}.$$

This is an exact identity, true for every $n \geq 0$, expressing the very term being summed as a difference of consecutive values of the auxiliary sequence $y_n := e^{x_n}$. Summing it from $n = 0$ to N telescopes completely:

$$x_0 + x_1 + \cdots + x_N = \sum_{n=0}^{N} (e^{x_n} - e^{x_{n+1}}) = e^{x_0} - e^{x_{N+1}} = e - e^{x_{N+1}}.$$

It remains only to know $x_{N+1} \to 0$ as $N \to \infty$, which is exactly the monotone-convergence argument above ($x_n > 0$, decreasing, limit L satisfies $e^L = e^L - L \Rightarrow L = 0$) – now used only to control the tail term $e^{x_{N+1}} \to e^0 = 1$, not to estimate a decay rate. Taking $N \to \infty$,

$$\sum_{n=0}^{\infty} x_n = e - 1.$$

Both convergence and the exact value drop out of the same single line.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains begin identically: establish $x_n > 0$, decreasing, hence convergent to a limit L, and solve $L = \ln(e^L - L)$ to get $L = 0$. They split immediately after that, at the question of what to do *next* with a sequence known only to tend to 0. The natural attempt asks "how fast does $x_n \to 0$?" and reaches for a Taylor expansion of the recursion. The model instead goes back to the defining equation itself, before taking limits, and asks what algebraic identity it hides – specifically, it undoes the ln by exponentiating, which turns the additive quantity being summed, x_n, into a difference of two terms of a new sequence e^{x_n}.

> *I am being asked to sum a sequence, and the recursion for that sequence is wrapped in a* ln *– before I do anything else with it, I should undo that* ln *by exponentiating, because a* ln *sitting on the outside of a recursion is very often there to be cancelled, and cancelling it is free information I have not yet used.*

> *Once I have $e^{x_{n+1}} = e^{x_n} - x_n$, I notice that x_n – exactly the term I need to sum – is isolated on one side, equal to a difference $e^{x_n} - e^{x_{n+1}}$ of consecutive terms of a new sequence. A sum of consecutive differences telescopes; I should sum this identity over n before doing anything else, rather than trying to estimate the size of x_n term by term.*

Layer 3 THE WHY

The instinct to estimate the decay rate of x_n is not wrong in general – it is exactly the right move if the question were only "does $\sum x_n$ converge?". It is the wrong move here because the question also demands the exact value, and an asymptotic estimate is lossy by construction: it discards the precise relationship between x_n and x_{n+1} in exchange for an inequality that only holds "eventually" and only up to a constant. The recursion, by contrast, is not an inequality; it is an equation, true exactly and for every single n starting at $n = 0$, and an equation which is exact should be manipulated exactly before it is ever approximated. Exponentiating loses nothing because ln and exp are exact inverses on the relevant domain; a Taylor expansion loses information at every step because it is inherently an approximation. This is the general reason to prefer algebraic manipulation of an exact recursion over asymptotic analysis whenever the target is a specific closed-form value rather than mere existence of a limit.

The deeper structural fact is this: a series $\sum a_n$ has closed form exactly when a_n can be exhibited as $g(n) - g(n+1)$ for some sequence g, because then partial sums collapse by the definition of a telescoping sum – there is no other mechanism by which an infinite sum of positive terms yields a clean closed-form answer like $e - 1$. So once the problem asks for an *exact value*, not just "does it converge", the right question to ask of the recursion is not "how big is x_n" but "can x_n be written as a difference of consecutive terms of some auxiliary sequence built from the recursion itself" – and the recursion, read literally, is already offering exactly that shape (x_{n+1} defined in terms of e^{x_n}) the moment the ln is stripped off. Kedlaya and Ng's solution begins by exponentiating for precisely this reason: it is not an isolated trick but the direct consequence of asking the right question of an exact recursion when an exact answer is demanded.

Caveat. I re-derived the telescoping identity $x_n = e^{x_n} - e^{x_{n+1}}$ and the limit computation $L = 0$ independently, and both match Kedlaya and Ng's solution exactly (sum $= e - 1$); the Taylor-expansion decay-rate analysis in the natural attempt is my own addition to illustrate the trap and is a standard local expansion, not a claim requiring external verification.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the ln sitting on the outside of the recursion, $x_{n+1} = \ln(\cdots)$; on seeing a defining recursion wrapped in ln or another invertible function, write down what the recursion looks like after inverting it, before doing any analysis on x_n itself.

2. That exponentiated form, $e^{x_{n+1}} = e^{x_n} - x_n$, suggests isolating the summand: rearrange so the term you need to sum stands alone on one side, equal to a difference of two terms of an auxiliary sequence.

3. Hesitation occurs right after proving $x_n \to 0$ monotonically, because that result feels like it should be the main content of the problem; the reinforcement is to notice the question asked for an exact sum, not just convergence, and a monotone-limit argument alone never produces an exact value.

4. The trap is treating the monotone-convergence argument as the finish line and reaching for asymptotic/Taylor estimates to "finish" – this proves convergence but structurally cannot produce the closed form, and time spent computing the decay rate is time spent on a road that does not reach the answer.

5. Recognisable faster next time by treating "recursion defined via an invertible function (ln, exp, a bijective substitution) wrapped around an expression in x_n" as an instant cue to invert first and look for the summand isolated as a difference.

6. Reusable rule: when asked for the exact sum of a sequence given by a recursion, look for a way to write the general term as $g(x_n) - g(x_{n+1})$ using the recursion's own defining equation (often by undoing whatever function wraps it) before estimating growth or decay rates – exact questions want exact (telescoping) manipulations, not asymptotic ones.

▷ **Cue:** a recursion defined as $x_{n+1} = f(\text{invertible function of } x_n)$, e.g. $\ln(\cdots)$ or $\arctan(\cdots)$, and an exact sum is requested $\longrightarrow$ **Move:** invert the outer function first (exponentiate, take tangent, etc.) to expose an exact algebraic identity before estimating growth

▷ **Cue:** the exact value (not just convergence) of $\sum x_n$ is requested for a recursively-defined sequence $\longrightarrow$ **Move:** search for $x_n = g(x_n) - g(x_{n+1})$ so the sum telescopes, rather than bounding decay rate

2016 B2

2016 B2. Define a positive integer n to be *squarish* if either n is itself a perfect square or the distance from n to the nearest perfect square is a perfect square. For example, 2016 is squarish, because the nearest perfect square to 2016 is $45^2 = 2025$ and $2025 - 2016 = 9$ is a perfect square. (Of the positive integers between 1 and 10, only 6 and 7 are not squarish.) For a positive integer N, let $S(N)$ be the number of squarish integers between 1 and N, inclusive. Find positive constants α and β such that $\lim_{N\to\infty} \frac{S(N)}{N^\alpha} = \beta$, or show that no such constants exist.

Fields: number theory / analysis (lattice-point counting, Riemann sums) **Position:** B2 – second problem on the paper, but the density argument requires real care with the boundary between consecutive squares

Answer: $\alpha = \frac{3}{4}, \beta = \frac{4}{3}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The definition is checkable one integer at a time, so the natural first move is to compute: list $n = 1, \ldots, 20$, find the nearest square to each, test whether the gap

is itself a square, and count. This reproduces the given fact (only 6, 7 fail below 10) and can be pushed further by hand or by a short mental table up to $N = 50$ or so. From such data a student might try to guess α by comparing $S(N)$ at two values of N, e.g. estimating $\log(S(N_2)/S(N_1))/\log(N_2/N_1)$ – essentially curve-fitting a power law to a handful of noisy data points. This can produce a plausible-looking α close to $3/4$, but it is not a proof and it is not even reliable numerically: $S(N)$ is not smooth, it jumps unevenly as N crosses regions near large squares versus regions far from any square, and no finite amount of tabulation pins down α exactly, let alone β. The deeper problem is structural: squarishness is defined relative to *whichever* square is nearest, and the natural attempt never organises the integers by which square is nearest to them – it treats each n as an isolated case, so there is no way to turn the count into a sum that can be estimated by an integral.

The model chain. Organise the integers not by their own size but by which perfect square is nearest to them. Fix $n \geq 1$ and ask: which integers m have n^2 as their nearest perfect square? The next square up is $(n+1)^2 = n^2 + 2n + 1$; the midpoint between n^2 and $(n+1)^2$ is $n^2 + n + \frac{1}{2}$, so every integer $m \leq n^2 + n$ is at least as close to n^2 as to $(n+1)^2$. Symmetrically, every integer $m \geq n^2 - n + 1$ is at least as close to n^2 as to $(n-1)^2$. So the block of integers whose nearest square is n^2 is

$$\{\, n^2 - n + 1,\ n^2 - n + 2,\ \ldots,\ n^2 + n \,\},$$

a block of exactly $2n$ consecutive integers, and these blocks partition the positive integers as n ranges over $1, 2, 3, \ldots$ (up to the harmless ambiguity exactly at a midpoint, which affects at most one integer per boundary and is negligible for an asymptotic count).

Within the block for n, write $m = n^2 + d$ with d ranging over the $2n$ values $-(n-1), \ldots, n$. The integer m is squarish exactly when $d = 0$ or $|d|$ is a perfect square. On the positive side, $d = k^2$ is available for $k = 1, \ldots, \lfloor \sqrt{n} \rfloor$ (need $k^2 \leq n$). On the negative side, $d = -k^2$ is available for $k = 1, \ldots, \lfloor \sqrt{n-1} \rfloor$ (need $k^2 \leq n - 1$, since $d \geq -(n-1)$). Including $d = 0$, the number of squarish integers in the block for n is

$$1 + \lfloor \sqrt{n} \rfloor + \lfloor \sqrt{n-1} \rfloor \approx 2\sqrt{n}.$$

Now sum over blocks. Since $S(N)$ is (up to boundary effects of size $O(\sqrt{N})$, negligible against the main term) the sum of this block-count for $n = 1, \ldots, \lfloor \sqrt{N} \rfloor$,

$$S(N) \approx \sum_{n=1}^{\sqrt{N}} 2\sqrt{n} \approx \int_0^{\sqrt{N}} 2\sqrt{x}\, dx = \frac{4}{3}\left(\sqrt{N}\right)^{3/2} = \frac{4}{3}N^{3/4},$$

using the standard fact that a sum $\sum f(n)$ approximates $\int f(x)\, dx$ with error controlled by the variation of f, which is routine to make rigorous with upper/lower Riemann-sum bounds here since $\sqrt{x}$ is monotonic. Hence

$$\lim_{N \to \infty} \frac{S(N)}{N^{3/4}} = \frac{4}{3},$$

so $\alpha = \frac{3}{4}$, $\beta = \frac{4}{3}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both approaches start from the same definition and can verify it on small cases. They split at the very first structural choice: how to organise the count. The natural attempt indexes by n itself (the number being tested). The model indexes by *which square is nearest to* n – i.e. it inverts the roles, treating n^2 as the primary variable and asking how many integers belong to its neighbourhood, before ever asking whether a specific n is squarish.

> *Testing squarishness integer by integer will never tell me the exponent α, because α is a statement about how a* count *grows, and a count only becomes tractable once I know how many integers I'm counting per unit of the variable that actually controls the problem – here, that variable is which square is nearest, not the integer itself.*

> *The block of integers nearest to n^2 has size exactly $2n$ (it runs from just past the midpoint below n^2 to the midpoint above), and within that block of size $2n$, roughly $2\sqrt{n}$ of the offsets d are perfect squares (in size at most n, so up to $\sqrt{n}$ in absolute value, times two for sign) – this is a self-similar sub-count of exactly the shape "how many perfect squares are there up to X", which I already know how to sum via an integral.*

Layer 3 THE WHY

The reorganisation is forced by the structure of the definition itself: "nearest perfect square to n" is a relation between n and the set of squares, and any relation of this kind partitions the domain into blocks attached to the object being compared against (here, the squares n^2), not into blocks attached to the object being tested. Trying to count squarish integers by testing n one at a time throws away the fact that the squares $1, 4, 9, 16, \ldots$ get sparser as they grow, so the neighbourhoods around them get *wider* as they grow – and it is precisely this widening (block size $2n$, growing without bound) combined with a shrinking *relative* density of squares within each block (about $2\sqrt{n}$ out of $2n$, a fraction $\sim 1/\sqrt{n} \to 0$) that produces a nontrivial power law at all. Neither fact is visible from a per-n table; both are immediate once the integers are grouped by nearest square.

Once grouped this way, the count "$1 + \lfloor\sqrt{n}\rfloor + \lfloor\sqrt{n-1}\rfloor$ squarish integers per block" is itself just "how many perfect squares are $\leq n$", counted on each side of 0 – the same primitive fact used to convert $\sum_n(\text{something} \sim \sqrt{n})$ into an integral $\int \sqrt{x}\,dx$. So the whole problem reduces twice to the same elementary estimate (Riemann-sum for $\sqrt{x}$), once to size the blocks' squarish content and once to sum that content over blocks – and the exponent $3/4$ is exactly what falls out of integrating $\sqrt{x}$ over a range of width $\sqrt{N}$: $\int_0^{\sqrt{N}} \sqrt{x}\,dx$ has degree $\frac{3}{2}$ in $\sqrt{N}$, i.e. degree $\frac{3}{4}$ in N. The $3/4$ is not a number to be pattern-matched from data; it is the sum of one $1/2$ (density of squares, $\sqrt{\cdot}$) and one $1/2$ (width of range, $\sqrt{N}$), halved by the integration, and it appears the moment the two nested square-root estimates are composed correctly.

> **Caveat.** I re-derived the block structure (size $2n$, boundary at $n^2 \pm n$), the per-block squarish count $1 + \lfloor\sqrt{n}\rfloor + \lfloor\sqrt{n-1}\rfloor$, and the resulting exponent $\alpha = 3/4$, $\beta = 4/3$ independently, and this matches Kedlaya and Ng's solution. I have not independently re-derived their Riemann-sum error bounds line by line (the passage from $S(N) \approx \int 2\sqrt{x}\,dx$ to a fully rigorous two-sided bound with explicit error terms); the mechanism – sandwiching $S(N)$ between upper and lower Riemann sums for the monotone function $\sqrt{x}$ – is standard and I am confident it goes through, but I flag the epsilon-management as the one part taken on trust from the official solution rather than checked digit-by-digit myself.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "nearest perfect square", which names a relation to a specific other set (the squares) rather than a property of n alone; on seeing a definition phrased as "distance to the nearest member of set T", write down the plan to partition the domain by which element of T is nearest, before testing any individual case.

2. That partition suggests a two-level count: first, how many integers fall in the neighbourhood of a given $t \in T$ (block size), and second, within that neighbourhood, how many satisfy the

extra condition (here, offset is a perfect square) – multiply/sum these rather than testing membership directly.

3. Hesitation occurs at deciding the exact cutoff between consecutive blocks (does the midpoint round up or down, and does it matter); the reinforcement is to note that a single boundary integer per block is $O(\sqrt{N})$ total against a main term of order $N^{3/4}$, hence provably negligible, so the boundary convention can be fixed by convenience and need not be agonised over.

4. The trap is trying to guess α numerically from small values of $S(N)$; the counting function is not smooth enough at small N for curve-fitting to be reliable, and no amount of tabulation substitutes for identifying the two nested square-root estimates that actually produce the exponent.

5. Recognisable faster next time by treating any "distance to nearest special number" counting problem as an instant cue for the two-step block decomposition, and by recalling that composing two independent $\sqrt{\cdot}$-type densities (density of squares among integers; width of a neighbourhood) through an integral typically yields an exponent that is a simple fraction built from $1/2$'s.

6. Reusable rule: to count integers up to N satisfying a condition phrased relative to a sparser reference set (squares, primes, etc.), partition $[1, N]$ into blocks attached to each reference element, count the condition's frequency within a typical block as a function of the block's index, then sum that block-count over blocks via a Riemann-sum/integral estimate to extract the growth exponent – do not attempt to estimate the exponent from raw tabulated data.

▷ **Cue:** a definition phrased as "distance from n to the nearest element of a sparser set T (squares, primes, ...)" $\longrightarrow$ **Move:** partition the domain into blocks by which element of T is nearest, before testing individual n

▷ **Cue:** counting a two-parameter condition (block index and offset within block) where both parameters independently contribute a square-root-type density $\longrightarrow$ **Move:** sum the per-block count via a Riemann sum/integral to read off the growth exponent directly

2016 B3

2016 B3. Suppose that S is a finite set of points in the plane such that the area of triangle $\triangle ABC$ is at most 1 whenever A, B, and C are in S. Show that there exists a triangle of area 4 that (together with its interior) covers the set S.

Fields: combinatorial geometry (extremal triangle, homothety) **Position:** B3 – third problem, the paper's easiest true "construction" problem but the containment argument is the whole content **Answer:** No numerical answer; the covering triangle is the anticomplementary (medial-image) triangle of a maximum-area triple, scaled to area 4.

Layer 1 THE WHAT — surface analysis

The natural attempt. The bound "every triangle formed from S has area at most 1" is a global constraint on all $\binom{|S|}{3}$ triples, and the natural first move is to try to exploit it directly: find the triangle $ABC \subseteq S$ of maximum area (it exists since S is finite), and try to cover S by scaling ABC up. A natural guess is to scale ABC by a factor of 3 from its centroid, since a factor of 3 turns area 1 into area 9 – too big – or, aiming directly for area 4, to scale by a factor of 2 from the centroid. The scaling-by-2-from-centroid triangle does have area $4 \cdot (\text{area } ABC) \leq 4$, which matches the target, so this looks promising. But checking containment is where the attempt stalls: to show every point $P \in S$ lies inside this centroid-scaled triangle, one wants to bound the position of P using the area constraint $[PBC] \leq [ABC]$ (and the two analogous inequalities), and the centroid-scaled triangle's sides are not, in general, positioned so that this area inequality

translates directly into "P is on the correct side of that specific side". The scaling factor 2 is right, but scaling from the *centroid* does not put the triangle's sides where the area argument needs them to be, and the student is left trying to verify containment by coordinates or ad hoc case-work that does not close cleanly.

The model chain. Let $A, B, C \in S$ be a triple maximising the area of $\triangle ABC$ (exists since S is finite); write $[XYZ]$ for area, so $[ABC] \leq 1$ by hypothesis. Define A', B', C' to be the points for which A, B, C are respectively the midpoints of $B'C'$, $C'A'$, $A'B'$ – that is, $A'B'C'$ is the *anticomplementary triangle* of ABC (equivalently, the image of ABC under the homothety centred at the centroid with ratio -2). This triangle has sides twice as long as ABC's and is similar to it, so

$$[A'B'C'] = 4[ABC] \leq 4.$$

(If the inequality is strict, scale $A'B'C'$ up slightly about its own centroid until its area is exactly 4; this can only enlarge the region and so does not affect containment.)

The key property of this specific construction, not shared by an arbitrary scaling, is that each side of $A'B'C'$ passes through a vertex of ABC and is parallel to the opposite side of ABC: in particular, side $B'C'$ passes through A and is parallel to BC. Take any $P \in S$. Since ABC has maximum area among all triples from S, $[PBC] \leq [ABC]$. Both triangles PBC and ABC share the base BC, so area is proportional to the distance from the apex to line BC; hence the distance from P to line BC is at most the distance from A to line BC. Equivalently, P lies in the half-plane bounded by the line through A parallel to BC – exactly the line containing $B'C'$ – that contains B and C. Applying the same argument to $[PCA] \leq [ABC]$ and $[PAB] \leq [ABC]$ places P in the analogous half-planes bounded by lines $C'A'$ and $A'B'$. The intersection of these three half-planes is precisely the closed triangular region bounded by $A'B'C'$ (each half-plane is bounded by one side of $A'B'C'$ and contains the opposite vertex of ABC, which is exactly the defining property of a triangle as an intersection of three half-planes). Hence $P \in \triangle A'B'C'$. Since $P \in S$ was arbitrary, $S \subseteq \triangle A'B'C'$, a triangle of area 4 (after the harmless rescaling above) covering S.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the first move: take a maximum-area triple A, B, C from S, since $[ABC] \leq 1$ and maximality is the only way to extract a single triangle that dominates every other triple's area. They split at the choice of *how* to blow that triangle up to area 4. The natural attempt scales by 2 from the centroid – correct on area alone – and then tries to retrofit a containment proof. The model instead asks what specific enlarged triangle makes the area-maximality inequality $[PBC] \leq [ABC]$ translate directly into a half-plane condition, and constructs the anticomplementary triangle precisely because its sides run through the vertices of ABC parallel to the opposite sides.

> *I don't want to pick a scaling factor first and then hunt for a containment proof that fits it – I want to work backwards from the inequality I actually have, $[PBC] \leq [ABC]$ for every $P \in S$, and ask what geometric statement that inequality is really making, then choose the enlarged triangle to be exactly the region that statement describes.*

> *Two triangles sharing a base have areas proportional to their apex's distance from the line containing that base – so "$[PBC] \leq [ABC]$" is not really a statement about area at all, it's a disguised statement about distance from the line BC; the moment I see two triangles compared that share an edge, I should convert the area inequality into a distance-from-a-line inequality, because distance inequalities are half-planes, and half-planes are exactly what I need to build a containing triangle.*

Layer 3 THE WHY

The reason scaling from the centroid fails to yield an easy containment proof, while the

anticomplementary construction succeeds, is that area-maximality is fundamentally a *base-sharing* comparison: the hypothesis $[PBC] \leq [ABC]$ only becomes a clean geometric constraint on P once P and A are compared as apexes of triangles sharing the exact same base BC, since area of a triangle with fixed base is (by the definition of area as $\frac{1}{2} \cdot \text{base} \times \text{height}$) linear in the height alone. Any enlarged triangle whose sides are not aligned with BC, CA, AB forces you to re-derive that relationship in new coordinates for no reason; a triangle whose sides *are* parallel to BC, CA, AB and pass through the opposite vertices converts the area hypothesis into a half-plane condition for free, because "distance from P to line BC is at most distance from A to line BC" is now visibly the statement "P is on the near side of the line through A parallel to BC" – and that line is, by construction, a side of the covering triangle. The anticomplementary triangle is not a lucky guess; it is the unique (up to the harmless final rescaling) triangle whose three sides are simultaneously parallel to BC, CA, AB and pass through A, B, C respectively, which is exactly the structural requirement the area-maximality hypothesis needs in order to become three half-plane statements rather than one messy area computation. This is also why the factor is forced to be 2 (not chosen to hit "area 4" as a target): the homothety of ratio -2 about the centroid is the *unique* homothety sending A, B, C to points of which A, B, C are midpoints, so "area exactly 4" is a consequence of the construction, not an input to it.

Caveat. I re-derived the half-plane argument coordinate-free (placing line BC as $y = 0$ and $A = (a, h)$, so $B'C'$ is $y = h$, and $[PBC] \leq [ABC]$ becomes $|y_P| \leq h$) and confirmed the three half-planes intersect exactly in $\triangle A'B'C'$; this matches Kedlaya and Ng's argument. I take the elementary fact that an anticomplementary triangle's sides are parallel to the medial triangle's opposite sides as a standard, checkable synthetic-geometry fact (verified by the coordinate computation above) rather than re-deriving it from first principles inside this analysis.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "area of triangle ABC is at most 1 whenever $A, B, C \in S$" – a bound on every triple simultaneously; on seeing a uniform bound over all triples from a finite set, the first thing to write down is "take the extremal (here, maximum-area) instance", since finiteness guarantees it exists and it alone can dominate all the others.

2. That extremal triple suggests the method of comparing every other point $P \in S$ against the fixed extremal triangle via the same measurement (area), which in turn suggests converting each of the three resulting area inequalities into a geometric constraint on P's position.

3. Hesitation occurs at choosing how to enlarge ABC to area 4: many scalings achieve the right area, and it is tempting to fix the scaling first; the reinforcement is to work backwards from the inequality $[PBC] \leq [ABC]$ and ask what region of the plane that inequality actually describes, then pick the enlargement to be exactly that region.

4. The trap is treating "quadruple the area" as purely a numerical target (any scale factor 2 homothety achieves it) and only afterward attempting to prove containment; the correct construction is forced by the containment argument, and the area value 4 falls out as a consequence rather than being the design goal.

5. Recognisable faster next time by treating "two triangles share a base/edge and their areas are compared" as an instant cue to rewrite the comparison as a distance-from-a-line (half-plane) statement, since area with a fixed base is linear in height.

6. Reusable rule: to build a region containing a finite set S from a bound on every k-tuple drawn from S, take an extremal k-tuple, then convert each remaining comparison inequality into a half-plane (or similar elementary region) constraint by comparing against a shared reference (a shared base, a shared vertex, etc.), and let the containing region be the intersection of those elementary constraints, rather than choosing an enlargement by its size alone.

▷ **Cue:** a uniform bound holds for every k-tuple drawn from a finite set S (e.g. every triangle

from S has area ≤ 1) $\longrightarrow$ **Move:** take an extremal (maximum, here) k-tuple; it exists by finiteness and dominates every comparison

▷ **Cue:** two triangles compared by area share a common base or vertex $\longrightarrow$ **Move:** rewrite the area inequality as a distance-from-a-line inequality, i.e. a half-plane condition, since area with fixed base is linear in height

2016 B4

2016 B4. Let A be a $2n \times 2n$ matrix, with entries chosen independently at random. Every entry is chosen to be 0 or 1, each with probability $1/2$. Find the expected value of $\det(A - A^t)$ (as a function of n), where A^t is the transpose of A.

Fields: probability (linearity of expectation, independence); linear algebra (Leibniz expansion of the determinant, skew-symmetric matrices); enumerative combinatorics (perfect matchings) **Position:** B4 — mid-paper; the answer is short once you refuse to compute any determinant **Answer:** $\mathbb{E}\det(A - A^t) = \dfrac{(2n)!}{4^n\, n!} = \dfrac{(2n-1)!!}{2^n}$, i.e. (number of perfect matchings of $2n$ points) $\times\, 2^{-n}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The strong student starts with small cases, which is right. For $n = 1$,

$$A - A^t = \begin{pmatrix} 0 & x \\ -x & 0 \end{pmatrix}, \qquad x = a_{12} - a_{21},$$

so $\det = x^2$. Here x is $-1, 0, 1$ with probabilities $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$, hence $\mathbb{E}x^2 = \frac{1}{2}$. Good.

For $n = 2$ the matrix is 4×4 skew-symmetric with six independent above-diagonal entries $x_{12}, x_{13}, x_{14}, x_{23}, x_{24}, x_{34}$. The student now expands the determinant — by cofactors along the first row, since that is what one does with a 4×4 — and after a page arrives (correctly) at

$$\det = (x_{12}x_{34} - x_{13}x_{24} + x_{14}x_{23})^2.$$

Expanding the square and taking expectations: every cross term such as $x_{12}x_{34}x_{13}x_{24}$ contains some x_{ij} to the first power and averages to 0; the three squares each give $\frac{1}{2} \cdot \frac{1}{2} = \frac{1}{4}$. So $\mathbb{E}\det = \frac{3}{4}$.

Now $n = 3$: a 6×6 cofactor expansion is out of the question by hand. The student has $\frac{1}{2}, \frac{3}{4}$ and tries to guess. $\frac{1}{2}, \frac{3}{4}$ could be $1 - 2^{-n}$, could be $\frac{n+1}{4}$, could be $\frac{(2n-1)!!}{2^n}$. Two data points do not distinguish them, and the third data point costs a 6×6 determinant. This is the wall: the method generates data at a cost growing like the determinant itself, and the answer cannot be pinned without either a third value or a structural argument. The $n = 2$ computation did contain the structural argument — "cross terms vanish, squares give $\frac{1}{4}$ each, there are three squares" — but it was buried under the cofactor expansion rather than lifted out of it.

The model chain. Write $M = A - A^t$, so $M_{ij} = a_{ij} - a_{ji}$. Then $M_{ii} = 0$, $M_{ji} = -M_{ij}$, and the $\binom{2n}{2}$ variables $\{M_{ij} : i < j\}$ are *independent* (they are built from disjoint pairs of entries of A), each taking the values $-1, 0, 1$ with probabilities $\frac{1}{4}, \frac{1}{2}, \frac{1}{4}$. In particular $\mathbb{E}M_{ij} = 0$ and $\mathbb{E}M_{ij}^2 = \frac{1}{2}$.

Step 1: expand, and take expectation term by term. By the Leibniz formula,

$$\det M = \sum_{\sigma \in S_{2n}} \operatorname{sgn}(\sigma) \prod_{i=1}^{2n} M_{i\,\sigma(i)},$$

a finite sum, so by linearity $\mathbb{E}\det M = \sum_\sigma \operatorname{sgn}(\sigma)\, E_\sigma$ where $E_\sigma = \mathbb{E}\prod_i M_{i\sigma(i)}$.

Step 2: most E_σ vanish. Decompose σ into disjoint cycles.

- If σ has a fixed point i, the product contains $M_{ii} = 0$, so $E_\sigma = 0$.
- If σ has a cycle $(i_1\, i_2\, \ldots\, i_m)$ with $m \geq 3$, then the unordered pair $\{i_1, i_2\}$ occurs in the product exactly once: within the cycle, the pairs $\{i_k, i_{k+1}\}$ are distinct because $m \geq 3$ means $\sigma(\sigma(i_1)) \neq i_1$; other cycles use disjoint indices. So the product is $\pm M_{i_1 i_2} \cdot Q$ with Q a product of variables independent of $M_{i_1 i_2}$ (write $M_{i_1 i_2}$ as $\pm$ the variable with the smaller index first). Hence $E_\sigma = \pm\, \mathbb{E}M_{i_1 i_2} \cdot \mathbb{E}Q = 0$.

So only σ whose cycles are *all transpositions* survive: fixed-point-free involutions, i.e. perfect matchings of $\{1, \ldots, 2n\}$.

Step 3: the survivors. If $\sigma = (i_1\, j_1)(i_2\, j_2)\cdots(i_n\, j_n)$, then $\operatorname{sgn}\sigma = (-1)^n$ and

$$\prod_i M_{i\sigma(i)} = \prod_{k=1}^{n} M_{i_k j_k} M_{j_k i_k} = \prod_{k=1}^{n} (-M_{i_k j_k}^2) = (-1)^n \prod_{k=1}^{n} M_{i_k j_k}^2.$$

The n pairs are distinct, so by independence $E_\sigma = (-1)^n \left(\frac{1}{2}\right)^n$, and $\operatorname{sgn}(\sigma) E_\sigma = 2^{-n}$: every survivor contributes the same positive amount.

Step 4: count. The number of perfect matchings of $2n$ points is $(2n-1)(2n-3)\cdots 3 \cdot 1 = (2n-1)!!$ (the partner of the smallest unmatched point can be chosen in $2n-1$, then $2n-3$, ...ways), which equals $\frac{(2n)!}{2^n n!}$. Therefore

$$\mathbb{E}\det(A - A^t) = \frac{(2n)!}{2^n n!} \cdot \frac{1}{2^n} = \frac{(2n)!}{4^n\, n!}.$$

Check: $n = 1$ gives $\frac{1}{2}$, $n = 2$ gives $\frac{3}{4}$, $n = 3$ gives $\frac{15}{8}$. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the first line: "det of a random matrix — I need $\mathbb{E}$ of a polynomial in independent variables." They split on the next: the natural attempt *evaluates the polynomial first and averages second* (compute det for a generic 4×4, then take $\mathbb{E}$); the model *averages first, term by term, and evaluates nothing*. At the Leibniz expansion, linearity of expectation turns one impossible computation into $(2n)!$ trivial ones, and independence plus $\mathbb{E}M_{ij} = 0$ kills almost all of them before any arithmetic happens.

> *I am asked for an expected value, not a determinant. Expectation is linear and the determinant is a sum of products. So I never need to know what* det *M is for any particular M; I need to know, for each monomial in the Leibniz sum, whether its average is zero.*

> *Each variable M_{ij} has mean zero and they are independent. So a monomial survives averaging only if every variable in it appears to an even power. In a product $\prod_i M_{i\sigma(i)}$, the pair $\{i, j\}$ can appear at most twice — once as M_{ij}, once as M_{ji} — and it appears twice exactly when σ swaps i and j. So only permutations made entirely of swaps survive.*

> *My $n = 2$ answer was a sum of three equal squares and 3 is the number of ways to pair up four points. That is not a coincidence to be guessed from; it is the count of surviving permutations, and I should prove that directly for general n.*

▷ **Cue:** "find the expected value of" a polynomial expression (determinant, trace of a power, permanent) in independent random entries ⟶ **Move:** expand into monomials and apply linearity of expectation to each monomial; do not compute the expression itself

- ▷ **Cue:** independent random variables each with mean zero $\longrightarrow$ **Move:** a monomial has nonzero expectation only if every variable appears to power at least 2; list the index structures for which that happens
- ▷ **Cue:** a matrix built as $A - A^t$, $A + A^t$, or with a symmetry forced on it $\longrightarrow$ **Move:** identify the genuinely independent variables (one per unordered pair) before anything else; the diagonal is often forced to zero
- ▷ **Cue:** a small-case answer that is a sum of k identical terms $\longrightarrow$ **Move:** interpret k combinatorially — it is almost always the number of surviving index configurations — and prove the count in general

Layer 3 THE WHY

Kedlaya and Ng's first solution begins by writing $\det(A - A^t)$ as the Leibniz sum and taking expectation term by term. Here is where that comes from, and why it is not a choice but the only road.

Start with the definition. The determinant is *defined* (or, from A-level onwards, characterised) as the alternating multilinear form, and the Leibniz formula $\sum_\sigma \operatorname{sgn}\sigma \prod_i M_{i\sigma(i)}$ is that definition written out. Cofactor expansion is a recursive way of organising the same sum for hand computation of a *specific* matrix. When the matrix is not specific — when its entries are random variables — the recursive organisation buys nothing, and the flat sum is what the definition gives you. So "Leibniz" is not a trick; it is the determinant with its packaging removed.

Now expectation. The one property of $\mathbb{E}$ that holds with no hypotheses at all is linearity: $\mathbb{E}\sum = \sum\mathbb{E}$ for a finite sum. The one property that needs a hypothesis is multiplicativity: $\mathbb{E}(XY) = \mathbb{E}X \cdot \mathbb{E}Y$ *when* X, Y *are independent*. The problem hands you independence ("chosen independently") in its first sentence. So the tools on the table are exactly: split a sum, and split a product of independent pieces. A determinant is a sum of products. The problem is engineered so that both tools apply, and nothing else is needed.

Then why do the terms die? Because of $\mathbb{E}M_{ij} = 0$, and that is where the phrase "$A - A^t$" earns its place in the question. If you asked for $\mathbb{E}\det A$, each entry has mean $\frac{1}{2}$, nothing dies, and you get $\sum_\sigma \operatorname{sgn}\sigma \cdot 2^{-2n} = 0$ for a different reason. Subtracting the transpose makes each variable symmetric about 0 — so $\mathbb{E}M_{ij} = 0$ — and couples M_{ij} to M_{ji}, so the only way a variable can appear twice in a Leibniz monomial is through a transposition. Those two structural facts together select the perfect matchings. The combinatorics of the answer is dictated by the construction $A - A^t$; that is the axiom-level reason the double factorial appears.

It is worth seeing why the sign works out uniformly, since a mixture of signs would have made the answer a cancellation problem. A transposition contributes $M_{ij}M_{ji} = -M_{ij}^2$, a factor of -1, and also contributes -1 to $\operatorname{sgn}\sigma$. The two signs cancel pairwise. So every surviving term is *positive*, and the expected determinant is a pure count. (This is the shadow of a deeper fact: for skew-symmetric M of even size, $\det M = \operatorname{Pf}(M)^2$, a perfect square — which is the $n = 2$ formula the natural attempt found. Bhargava's second solution in the official text starts from the Pfaffian and gets the same count from $\mathbb{E}\operatorname{Pf}^2$. You do not need the Pfaffian, but it explains why the natural attempt saw a square.)

Finally, the rule that should have fired on reading the problem. *Evaluate-then-average* is zombie maths when the question is about an average: it computes an object far richer than what is asked, and throws nearly all of it away. *Average-then-evaluate* computes only what survives. The $n = 2$ cofactor expansion is a full page; the Leibniz argument for all n is five lines.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the pair "expected value" + "entries chosen independently". On seeing it, write $\det M = \sum_\sigma \operatorname{sgn}\sigma \prod_i M_{i\sigma(i)}$ and, beside it, the two facts $\mathbb{E}M_{ij} = 0$ and $\mathbb{E}M_{ij}^2 = \frac{1}{2}$ for the independent variables M_{ij}, $i < j$.

2. Linearity of expectation over the Leibniz monomials, followed by a survival analysis: which index patterns σ give a monomial in which every variable appears to an even power.

3. Hesitation occurs at "I can't expand a $2n \times 2n$ determinant". It is reinforced away by the reminder that you are not going to expand it; you are going to average each of its $(2n)!$ terms, and nearly all of them average to zero for a reason you can state in one line. A second hesitation is at the cycle-of-length-≥ 3 case: make sure the pair $\{i_1, i_2\}$ really appears only once, which uses $m \geq 3$.

4. The trap is the small-case-and-guess route: $\frac{1}{2}, \frac{3}{4}$ fit several formulas, and the third data point costs a 6×6 determinant. A subtler trap is forgetting that M_{ij} and M_{ji} are the *same* random variable up to sign, so treating all $4n^2$ entries of M as independent (which would kill even the transpositions and give the wrong answer 0).

5. Any time a question asks for $\mathbb{E}$ of an algebraic object built from independent mean-zero pieces, the answer is a count of "pairings" — index configurations in which every variable is used exactly twice. Recognise it by asking, before computing anything: *what is the mean of a single entry?* If it is zero, you are counting pairings.

6. Reusable decision rule: *when asked for an average of a complicated function, average its simplest additive pieces instead; use linearity first, independence second, and evaluate only what survives.* This covers expected determinants and traces of random matrices, expected numbers of fixed points or cycles, moment calculations, and any "expected value of a sum over configurations" problem.

2016 B5

2016 B5. Find all functions f from the interval $(1, \infty)$ to $(1, \infty)$ with the following property: if $x, y \in (1, \infty)$ and $x^2 \leq y \leq x^3$, then $(f(x))^2 \leq f(y) \leq (f(x))^3$.

Fields: functional equations and inequalities; real analysis (logarithmic change of variables); number theory (irrationality of $\log 3 / \log 2$, density of irrational rotations) **Position:** B5 — hard; the answer is guessable in a minute, the uniqueness proof is the whole problem **Answer:** Exactly the power functions $f(x) = x^c$ with $c > 0$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The guess is immediate and the student makes it: $f(x) = x^c$ works for every $c > 0$, because $x^2 \leq y \leq x^3$ gives $x^{2c} \leq y^c \leq x^{3c}$ (raising to a positive power preserves order), and $x^c > 1$ for $x > 1$. The student also notices $c \leq 0$ fails the codomain.

Then the converse. The natural move is to exploit special values: put $y = x^2$ to get $f(x)^2 \leq f(x^2) \leq f(x)^3$; put $y = x^3$ to get $f(x)^2 \leq f(x^3) \leq f(x)^3$. Iterating, $f(x^{2^k}) \geq f(x)^{2^k}$ and $f(x^{3^k}) \leq f(x)^{3^k}$. The student hopes to squeeze $f(x^t)$ between $f(x)^t$ from both sides for all t, and then argue $\log f(x) / \log x$ is constant.

The wall arrives quickly. Every inequality obtained relates f at one point to f at a point whose *exponent* is between 2 and 3 times larger (or smaller) — a factor-of-$3/2$ slack at every step. Compose k steps and the slack compounds to $(3/2)^k$. The student has lots of inequalities and none of them tight: they bound $\log f(x^t) / \log f(x)$ between $2^a 3^{-b}$ and $3^a 2^{-b}$, and those two numbers get *further* apart as a, b grow. It looks as though composition can only lose information, and the student concludes (wrongly) that there may be non-power solutions, or tries to prove continuity or monotonicity first and gets nowhere — neither is given, and neither is needed.

The model chain. **Step 0: the power functions work.** Done above; $c > 0$ is forced by $f > 1$.

Step 1: take logarithms twice. For $s > 0$ put $g(s) = \log f(e^s)$. Since $f > 1$, $g : (0, \infty) \to (0, \infty)$. With $x = e^s$, $y = e^t$ ($s, t > 0$), the hypothesis reads: if $2s \leq t \leq 3s$ then $2g(s) \leq g(t) \leq 3g(s)$. Now for $u \in \mathbb{R}$ put $h(u) = \log g(e^u)$, so $h : \mathbb{R} \to \mathbb{R}$. With $s = e^u$, $t = e^v$ and $a = \log 2$, $b = \log 3$:

$$v - u \in [a, b] \implies h(v) - h(u) \in [a, b]. \tag{H}$$

The goal becomes: $h(v) - h(u) = v - u$ for all u, v. Then $h(u) = u + C$, $g(s) = e^C s$, and $f(y) = \exp(g(\log y)) = y^{e^C}$, a power function with exponent $e^C > 0$.

Step 2: compose (H). Reading (H) with u, v swapped: if $v - u \in [-b, -a]$ then $h(v) - h(u) \in [-b, -a]$. *Claim:* for integers $p, q \geq 0$, not both 0,

$$v - u \in I_{p,q} := [\, pa - qb,\ pb - qa \,] \implies h(v) - h(u) \in I_{p,q}. \tag{$\mathrm{H}_{p,q}$}$$

Proof. Let $d = v - u \in I_{p,q}$ and write $d = pa - qb + \lambda(p+q)(b-a)$ with $\lambda \in [0,1]$. Put $\delta = a + \lambda(b-a) \in [a,b]$ and $\eta = b - \lambda(b-a) \in [a,b]$; then $p\delta - q\eta = d$. Walk from u to v by p steps of $+\delta$ followed by q steps of $-\eta$. By (H) and its swapped form, h changes by an amount in $[a,b]$ on each forward step and in $[-b,-a]$ on each backward step; adding, $h(v) - h(u) \in [pa - qb,\ pb - qa]$. ■

Step 3: only one endpoint of $I_{p,q}$ matters. Fix u, v and $d = v - u$. The interval $I_{p,q}$ is wide, width $(p+q)(b-a)$, but if $d \in I_{p,q}$ then ($\mathrm{H}_{p,q}$) gives in particular

$$h(v) - h(u) \leq pb - qa \qquad \text{and} \qquad h(v) - h(u) \geq pa - qb.$$

So I want an upper endpoint $pb - qa$ just above d with the lower endpoint below d, and separately a lower endpoint just below d with the upper endpoint above d.

Step 4: density. $\theta = b/a = \log 3/\log 2$ is irrational: $\theta = m/n$ with positive integers would give $3^n = 2^m$, impossible by parity. Hence the fractional parts $\{p\theta\}$, $p = 1, 2, 3, \ldots$, are dense in $[0,1)$ and every open subinterval contains infinitely many of them (Kronecker; Dirichlet's pigeonhole proof suffices). Fix $\varepsilon > 0$.

Upper bound. $pb - qa = a(p\theta - q)$. Choose p with $\{p\theta\}$ in the open interval $\big(\{d/a\}, \{d/a\} + \varepsilon/a\big)$ (shrink ε so this lies inside $[0,1)$), and take $q = \lfloor p\theta \rfloor - \lfloor d/a \rfloor$. Then $d < pb - qa < d + \varepsilon$. Infinitely many p qualify; as $p \to \infty$, $q \sim p\theta \to \infty$, so $q \geq 0$ and

$$pa - qb = a(p - q\theta) = a\big(p - p\theta^2\big) + O(1) \to -\infty,$$

because $q = p\theta + O(1)$ and $\theta > 1$. So for large such p, $d \in I_{p,q}$ and $h(v) - h(u) < d + \varepsilon$.

Lower bound. Symmetrically, $pa - qb = -a(q\theta - p)$; choose q large with $\{q\theta\}$ in the open window $\big(\{-d/a\}, \{-d/a\} + \varepsilon/a\big)$ (again shrinking ε) and put $p = \lfloor q\theta \rfloor - \lfloor -d/a \rfloor$, so that $q\theta - p \in (-d/a,\ -d/a + \varepsilon/a)$, i.e. $d - \varepsilon < pa - qb < d$. Now $p = q\theta + O(1) \to \infty$ and $pb - qa = a(p\theta - q) = a(q\theta^2 - q) + O(1) \to +\infty$, so $d \in I_{p,q}$ for large such q. Then $h(v) - h(u) > d - \varepsilon$.

As ε was arbitrary, $h(v) - h(u) = v - u$ for all u, v, and by Step 1 $f(x) = x^c$ with $c = e^{h(0)} > 0$. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that composing the hypothesis is the only source of information. They split at the moment the composed inequality appears as

$$2^a 3^{-b} \leq \frac{\log f(x^t)}{\log f(x)} \leq 3^a 2^{-b}.$$

The natural attempt reads this as a *pair* of bounds whose ratio $(3/2)^{a+b}$ explodes, and concludes composition is lossy. The model reads it as *two separate one-sided bounds*, each attached to a different composite: a composite whose *upper* endpoint happens to land near the target, and another whose *lower* endpoint does. The width of the window is irrelevant; only its edges are used, and irrationality lets the edges land anywhere.

> *The condition relates exponents multiplicatively — y is between x^2 and x^3 — and the conclusion is about exponents. Multiplicative in the exponent means additive after two logarithms. Let me take* log *of* log *and see the hypothesis as "if $v-u$ is in $[\log 2, \log 3]$ then so is $h(v)-h(u)$". That is a statement about differences, and differences add.*

> *Composing p forward steps and q backward steps gives an interval of possible displacements that gets wider, not narrower. But I do not need the whole interval. If the upper end of the interval lies just above $v-u$ and $v-u$ is still inside, then $h(v)-h(u)$ is at most that upper end — which is almost $v-u$.*

> *The upper ends are the numbers $p\log 3 - q\log 2$. If $\log 3/\log 2$ were rational these would live on a lattice and I would be stuck with a gap; it is irrational, so they are dense, and I can put an upper end anywhere I like.*

- ▷ **Cue:** a functional inequality whose hypothesis is "y between x^α and x^β" (multiplicative in exponents) ⟶ **Move:** substitute twice, $h(u) = \log\log f(e^{e^u})$, to turn it into "if $v-u \in [\log\alpha, \log\beta]$ then $h(v)-h(u) \in [\log\alpha, \log\beta]$"
- ▷ **Cue:** a condition of the form "displacement in an interval ⇒ increment in the same interval" ⟶ **Move:** compose it: p forward and q backward steps give the Minkowski-sum interval $[pa-qb,\ pb-qa]$ for both
- ▷ **Cue:** bounds that get wider under iteration ⟶ **Move:** stop using both ends; use one endpoint of each iterate and choose the iterate so that endpoint lands near the target
- ▷ **Cue:** integer combinations $p\beta - q\alpha$ with β/α irrational ⟶ **Move:** they are dense in $\mathbb{R}$ (Kronecker), so a one-sided endpoint can be placed within any ε of any target

Layer 3 THE WHY

Kedlaya and Ng's solution begins by defining $g(x) = \log f(e^x)$ and then $h(x) = \log g(e^x)$. Here is where the double logarithm comes from.

Read the hypothesis at the level of definitions. "$x^2 \le y \le x^3$" says: the *exponent* of y (relative to any fixed base) is between 2 and 3 times the exponent of x. So the natural variable is $s = \log x$, and in it the hypothesis says "$t/s \in [2,3]$". The conclusion is similarly "$g(t)/g(s) \in [2,3]$". Ratios in $[2,3]$ are awkward to compose — they multiply — and the one operation that turns multiplication into addition is the logarithm. Taking it gives "$v-u \in [\log 2, \log 3] \Rightarrow h(v)-h(u) \in [\log 2, \log 3]$". The two logarithms are not a trick; each one is forced by the observation that the thing being compared is a ratio. After both, the target $f(x) = x^c$ has become $h(u) = u + C$ — "h has increments equal to the increments of its argument" — and the hypothesis is visibly a weak form of that. The problem has become: does "increments in $[a,b]$ whenever the step is in $[a,b]$" force "increment equals step"?

Why should it? Test the analogous statement with $a = b$: "$v-u = a \Rightarrow h(v)-h(u) = a$" says only that $h(u)-u$ is a-periodic, which does *not* force linearity. Now test a window with rational endpoints, $[a,b] = [1,2]$. The step function $h(u) = \lfloor u \rfloor$ satisfies it: if $v-u \in [1,2]$ then $\lfloor v \rfloor - \lfloor u \rfloor$ is an integer strictly between $v-u-1 \ge 0$ and $v-u+1 \le 3$, so it is 1 or 2. So the condition does *not* force linearity in general, and whatever proof works must use something specific to $\log 2$ and $\log 3$. It must also use the ends of the window: the floor function fails to be linear only by jumping, and a jump is detected only by a displacement landing exactly where it matters. That is why the proof uses endpoints. The composite statements ($H_{p,q}$) are *tight at their endpoints*: when $v-u = pb-qa$ exactly, the upper bound says $h(v)-h(u) \le v-u$ with no slack. So composites are not lossy at all at the edges; they are lossy only in the middle. The natural attempt looked at the middle.

And then irrationality. The endpoints available are $\{pb-qa\}$. If b/a were rational, say both a, b integer multiples of some δ, all endpoints would lie on $\delta\mathbb{Z}$, a displacement strictly between

lattice points could never be pinned to within less than the gap, and indeed $h(u) = \delta\lfloor u/\delta \rfloor$ would be a genuine non-linear solution (the same floor-difference argument as for $[1, 2]$ shows it). That $\log 3/\log 2$ is irrational is the precise reason the answer is only the power functions, and the proof of that irrationality is the parity of 3^n versus 2^m. So the chain of "of course" reaches the axioms: two logs because the hypothesis is about ratios of exponents; composition because that is the only thing a local condition can do; endpoints because composites are tight exactly there; density because $2^m \neq 3^n$.

A note on what the proof does not need. No continuity, monotonicity or measurability of f is assumed or proved; the density argument works pointwise for each fixed pair u, v. A student who spends time trying to prove f is increasing first is solving a problem that is not in the way.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "$x^2 \leq y \leq x^3$" paired with "$f(x)^2 \leq f(y) \leq f(x)^3$" — the same shape on both sides, multiplicative in the exponent. On seeing it, write $g(s) = \log f(e^s)$, then $h(u) = \log g(e^u)$, and the translated hypothesis (H).

2. Composition of an interval-to-interval condition along a path, then a density argument (irrational rotation) to choose composites whose one endpoint approaches the target.

3. Hesitation occurs when the composed bounds are seen to widen. It is reinforced away by testing one composite at its endpoint: if $v - u = pb - qa$ exactly, ($\mathrm{H}_{p,q}$) says $h(v) - h(u) \leq v - u$ — a sharp inequality. Seeing one sharp instance tells you composition is not lossy where it matters.

4. Two traps. First, trying to establish regularity (continuity, monotonicity) that is never used. Second, using both endpoints of the same composite, which produces the ever-widening bracket and the false impression that non-power solutions might exist. A smaller trap: forgetting that $f > 1$ is what makes $g > 0$ and hence $\log g$ defined.

5. Faster recognition: whenever a condition is "if the input changes by an amount in a window, the output changes by an amount in the same window", expect the answer to be "output change = input change" as soon as the window's endpoints have an irrational ratio (or here, the log-endpoints). The irrationality is the tell; look for where the problem has hidden it.

6. Reusable decision rule: *when iterated bounds widen, use each iterate for one side only and select iterates so that side lands where you need it; density of an irrational rotation is what lets you select.* And before that: *when a hypothesis compares ratios, take logarithms until it compares differences.*

2016 B6

2016 B6. Evaluate

$$\sum_{k=1}^{\infty} \frac{(-1)^{k-1}}{k} \sum_{n=0}^{\infty} \frac{1}{k2^n + 1}.$$

Fields: series (absolute convergence, rearrangement); self-similarity under $k \mapsto 2k$; integral representations; binary expansion **Position:** B6 — the last problem of the paper; hard to start, short to finish once the structure is seen **Answer:** The sum equals 1.

Layer 1 THE WHAT — surface analysis

The natural attempt. A strong student looks at the inner sum $T(k) = \sum_{n\geq 0} \frac{1}{k2^n+1}$ and tries to evaluate it in closed form. It has no elementary closed form (it is a Lambert-type series in k), and after ten minutes the student accepts that and switches to the standard analytic move: expand $\frac{1}{k2^n+1}$ as a geometric series. For $n = 0$ and $k = 1$ this is $\frac{1}{2}$ and the expansion $\sum_m (-1)^m k^{-m-1} 2^{-n(m+1)}$ diverges, so the $n = 0$ term has to be split off. The student writes

$S = S_1 + S_2$ with

$$S_1 = \sum_k \frac{(-1)^{k-1}}{k(k+1)} = 2\ln 2 - 1, \qquad S_2 = \sum_{k\geq 1}\sum_{n\geq 1}\sum_{m\geq 0} \frac{(-1)^{k-1+m}}{k^{m+2}\,2^{n(m+1)}},$$

justifies absolute convergence, sums over n to get $\frac{1}{2^{m+1}-1}$, sums over k to get an alternating zeta value $(1-2^{-m-1})\zeta(m+2)$, notices the factor $2^{m+1}-1$ cancels, resums over m as a geometric series in $-\frac{1}{2k}$, and obtains $S_2 = \sum_k \frac{1}{k(2k+1)} = 2 - 2\ln 2$. So $S = 1$.

This works — it is Kedlaya and Ng's first solution — but it is three nested sums, two absolute-convergence checks, a zeta-function identity and a lucky cancellation. In the arena it is forty minutes, and at every stage the student does not know whether the next sum will close. The cancellation of $2^{m+1}-1$ looks like good fortune. It is not; it is the doubling structure of the problem leaking through, and the natural attempt never names it.

The model chain. Let $T(k) = \displaystyle\sum_{n=0}^{\infty} \frac{1}{k2^n+1}$ for $k \geq 1$, so that $S = \displaystyle\sum_{k\geq 1} \frac{(-1)^{k-1}}{k} T(k)$.

Step 1: convergence. $0 < T(k) \leq \sum_n \frac{1}{k2^n} = \frac{2}{k}$, so the k-th term of S is at most $\frac{2}{k^2}$ in absolute value. Hence S converges absolutely, and so do the two sums

$$O = \sum_{k \text{ odd}} \frac{T(k)}{k}, \qquad E = \sum_{k \text{ even}} \frac{T(k)}{k},$$

with $S = O - E$ (splitting an absolutely convergent series into two subseries is legitimate).

Step 2: the doubling identity. Separating the $n = 0$ term and re-indexing $n \mapsto n+1$,

$$T(k) = \frac{1}{k+1} + \sum_{n=1}^{\infty} \frac{1}{k2^n+1} = \frac{1}{k+1} + \sum_{n=0}^{\infty} \frac{1}{(2k)2^n+1} = \frac{1}{k+1} + T(2k).$$

Step 3: express E through $O+E$. Every even k is 2ℓ for a unique $\ell \geq 1$, so

$$E = \sum_{\ell\geq 1} \frac{T(2\ell)}{2\ell} = \frac{1}{2}\sum_{\ell\geq 1} \frac{T(\ell) - \frac{1}{\ell+1}}{\ell} = \frac{1}{2}\sum_{\ell\geq 1} \frac{T(\ell)}{\ell} - \frac{1}{2}\sum_{\ell\geq 1} \frac{1}{\ell(\ell+1)} = \frac{1}{2}(O+E) - \frac{1}{2},$$

using $\sum_{\ell\geq 1} \frac{1}{\ell(\ell+1)} = \sum_\ell \left(\frac{1}{\ell} - \frac{1}{\ell+1}\right) = 1$ (telescoping), and splitting the difference of two convergent series termwise.

Step 4: solve. $E = \frac{1}{2}O + \frac{1}{2}E - \frac{1}{2}$ gives $\frac{1}{2}E = \frac{1}{2}O - \frac{1}{2}$, so

$$S = O - E = 1. \qquad \square$$

A second view, same structure. Since $\frac{1}{m} = \int_0^1 x^{m-1}\,dx$, write $\frac{1}{k2^n+1} = \int_0^1 x^{k2^n}dx$. The triple object $\sum_k \sum_n \frac{1}{k} \int_0^1 x^{k2^n}dx$ has nonnegative-majorant sum $\sum_k \sum_n \frac{1}{k(k2^n+1)} \leq 2\sum_k k^{-2} < \infty$, so sums and integral may be interchanged (Tonelli on the absolute values, then Fubini). Summing over k first gives $\ln(1+x^{2^n})$, and

$$\sum_{n\geq 0} \ln\left(1+x^{2^n}\right) = \ln \prod_{n\geq 0}\left(1+x^{2^n}\right) = \ln \frac{1}{1-x} \qquad (0 \leq x < 1),$$

because each nonnegative integer has exactly one binary expansion. Hence, substituting $u = 1-x$,

$$S = \int_0^1 -\ln(1-x)\,dx = \int_0^1 -\ln u\,du = \left[\,u - u\ln u\,\right]_{0^+}^1 = 1,$$

using $u \ln u \to 0$ as $u \to 0^+$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write down $T(k)$ and both notice they cannot evaluate it. They split at the next line. The natural attempt says "then expand it" and descends into power series. The model says "then find what T does under a natural operation on k", and the operation the summand advertises is $k \mapsto 2k$: shifting n by one is the same as doubling k. That gives $T(k) = \frac{1}{k+1} + T(2k)$, and the outer sign $(-1)^{k-1}$ — which only sees the parity of k — is exactly what distinguishes k from $2k$.

> *I cannot compute $T(k)$. So I should stop trying and ask what relation T satisfies. In $k2^n$, raising n by one is the same as doubling k. So $T(k)$ and $T(2k)$ differ by exactly one term: $T(k) = \frac{1}{k+1} + T(2k)$.*

> *The outer sum's sign $(-1)^{k-1}$ depends only on whether k is odd or even, and the even k are exactly the doubles. So I should split the sum into odd and even k, rewrite the even part with $k = 2\ell$, and use the doubling identity to turn it back into a sum over all ℓ. Then I have one linear equation in two unknowns, and the combination I want might be exactly the one it determines.*

> *Before I rearrange anything, check absolute convergence: $T(k) \le 2/k$, so the terms are $O(1/k^2)$. Good — I may split into odd and even freely.*

- ▷ **Cue:** an inner sum over n involving $k \cdot 2^n$ (or $k \cdot r^n$) $\longrightarrow$ **Move:** shifting n is the same as scaling k; write the functional equation $T(k) = (\text{first term}) + T(2k)$
- ▷ **Cue:** an outer sign $(-1)^k$ together with a doubling structure in k $\longrightarrow$ **Move:** split into odd and even k; substitute $k = 2\ell$ in the even part and apply the functional equation to get a linear relation
- ▷ **Cue:** a sum of reciprocals $\frac{1}{\text{linear in } k}$ that will not close $\longrightarrow$ **Move:** replace $\frac{1}{m}$ by $\int_0^1 x^{m-1}\,dx$, sum the series inside the integral, and look for a product or logarithm identity
- ▷ **Cue:** a product $\prod_n (1 + x^{2^n})$ or a sum over 2^n $\longrightarrow$ **Move:** binary expansion: the product is $\sum_{m \ge 0} x^m = \frac{1}{1-x}$

Layer 3 THE WHY

Liang Xiao's solution in the official text (the fourth of four) is the one given above, and Kedlaya and Ng present it without comment on where the split into odd and even k comes from. Here is where it comes from.

Start with the definition of the object. The inner sum ranges over the numbers $k, 2k, 4k, 8k, \ldots$ — the orbit of k under doubling. That is what $k2^n$ *means*: it is not an arbitrary two-parameter expression, it is "k, then double, then double". So the function T is a sum along doubling orbits, and any function defined as a sum along an orbit satisfies a one-step recurrence along that orbit: remove the first element and the rest is the sum from the next element. $T(k) = \frac{1}{k+1} + T(2k)$ is not a trick; it is the definition of T read one step at a time. The same thing is true of $\sum_{n \ge 0} a_n$ versus $\sum_{n \ge 1} a_n$ — the tail of a series is a series of the same kind.

Why is that recurrence useful here, when it would not be for a sum like $\sum_k T(k)/k^2$? Because the *outer* sum has the sign $(-1)^{k-1}$, and that sign is the indicator of parity. The positive integers split uniquely into odd numbers and doubles, and the doubles are the image of the doubling map. So the outer sum naturally splits into "odd k" and "$2\times$(all k)", and the recurrence converts T on the second piece back to T on all k. The alternating sign is the problem telling you which decomposition to use: it is $+$ on the part the doubling map misses and $-$ on its image. That is exactly the form in which one step of a self-similar decomposition closes up into a linear equation.

It is worth seeing why the equation determines exactly $O - E$ and not O or E separately. Write the relation as $E - \frac{1}{2}(O + E) = -\frac{1}{2}$, i.e. $\frac{1}{2}(E - O) = -\frac{1}{2}$. The doubling map sends a sum over all k to a sum over even k with a factor $\frac{1}{2}$ from the $\frac{1}{k}$ weight. The weight $\frac{1}{k}$ is homogeneous of degree -1, and that factor $\frac{1}{2}$ is what makes the coefficients of O and E come out equal and opposite. With a weight $\frac{1}{k^2}$ the relation would be $E = \frac{1}{4}(O + E) - \ldots$, which does not isolate $O - E$. So the problem's $\frac{1}{k}$ and $(-1)^{k-1}$ are matched to each other: that pairing is the design, and the answer 1 is the telescoping sum $\sum \frac{1}{\ell(\ell+1)}$ that the first terms of the orbits leave behind.

The integral view shows the same structure from the other side, and it explains the "lucky" cancellation in the natural attempt. After the substitution $\frac{1}{m} = \int_0^1 x^{m-1} dx$, the alternating $\frac{1}{k}$-weighted sum becomes $\ln(1 + x^{2^n})$, and the sum over the doubling orbit becomes $\ln \prod_n (1 + x^{2^n})$. That product is $\frac{1}{1-x}$ precisely because every integer has one binary expansion — the doubling orbit again, now as the statement "$2^0, 2^1, 2^2, \ldots$ are a basis for the integers in base two". The factor $2^{m+1} - 1$ that cancelled in the power-series route is the same binary structure appearing as $\sum_{n \geq 1} 2^{-n(m+1)}$ meeting $(1 - 2^{-(m+1)})\zeta(m + 2)$; both factors come from separating out the even k. The natural attempt did the odd/even split implicitly, inside a zeta identity, without knowing it. The model does it on purpose on line one.

So: when an inner sum runs over an orbit, find the orbit's recurrence; when an outer weight distinguishes the orbit's image from its complement, split along it. That is the whole problem.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is $k2^n$ inside the inner sum. On seeing it, write $T(k)$ for the inner sum and, underneath, $T(k) = \frac{1}{k+1} + T(2k)$.

2. A self-similarity (functional-equation) method: split the outer sum by the parity of k, substitute $k = 2\ell$ in the even part, apply the recurrence, and solve the resulting linear equation. Alternative: an integral representation turning the double sum into $\int_0^1 -\ln(1 - x)\, dx$ via binary expansion.

3. Hesitation occurs immediately after writing $T(k)$, because it has no closed form and every instinct says a series cannot be summed without one. It is reinforced away by the question: *what does T do under the operation built into its definition?* A second hesitation is at rearranging a conditionally-looking alternating series; it is removed by the one-line bound $T(k) \leq 2/k$, which makes everything absolutely convergent.

4. The trap is the power-series expansion of $\frac{1}{k2^n+1}$: it is always available, it always looks like progress, and it produces a triple sum whose closing depends on an identity for alternating zeta values. It also fails at $n = 0$, $k = 1$, forcing an early split whose meaning is not explained. It works, at four times the cost. A secondary trap is trying to evaluate $T(k)$ itself.

5. Faster recognition: whenever a summand involves $k \cdot r^n$, k^{2^n}, x^{2^n} or any expression where changing n is the same as transforming k, look for a one-step recurrence along the orbit first. Whenever an outer weight or sign depends on a residue of k that the transformation respects (parity for doubling), expect the split by that residue to close up.

6. Reusable decision rule: *if you cannot evaluate an inner object, find the equation it satisfies under the symmetry built into its definition, and let the outer structure solve that equation for you.* Self-similar sums, functional equations for generating functions, and the classical evaluation $\sum(-1)^{k-1}/k = \ln 2$ via $\sum_{\text{all}} - 2\sum_{\text{even}}$ are all instances.