DAWN PUTNAM SESSION

Putnam 1989 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1989 — Paper A **5**
1989 A1 5
1989 A2 7
1989 A3 9
1989 A4 12
1989 A5 14
1989 A6 17

Putnam 1989 — Paper B **23**
1989 B1 23
1989 B2 25
1989 B3 28
1989 B4 31
1989 B5 33
1989 B6 36

Putnam 1989 — Paper A

1989 A1

1989 A1. How many primes among the positive integers, written as usual in base 10, are alternating 1's and 0's, beginning and ending with 1?

Fields: number theory, algebraic factorisation **Position:** A1 — the opener **Answer:** Exactly one: 101. Every other such number $N_k = 1010\cdots101$ (with k ones) equals $\dfrac{(10^k-1)(10^k+1)}{99}$, and for $k \geq 3$ this splits into two integer factors each greater than 1; $N_1 = 1$ is not prime.

Layer 1 THE WHAT — surface analysis

The natural attempt. You write out the first few candidates: $1, 101, 10101, 1010101, 101010101, \ldots$ You know 101 is prime. You test 10101: digit sum 3, so $3 \mid 10101$ and it is composite. You test 1010101: digit sum 4, not divisible by 3; not by 11 (alternating sum of digits is $1-0+1-0+1-0+1 = 4$); you try 7 and find $1010101 = 7 \cdot 144300.14\ldots$ — no; you try 101 and find $1010101 = 101 \cdot 10001$. Composite. Then 101010101: digit sum 5; you try 101 again and it fails; you try 41, 271 ...

This is where the natural attempt stalls. Every individual number turns out to be composite, but each one fell to a different divisor (3, then 101, then something you have not yet found — in fact $101010101 = 41 \cdot 271 \cdot 9091$), and a pattern of "they keep being composite" is not a proof for all lengths. You have been doing trial division, one case at a time, on an infinite family. The honest student at this point either guesses "one" and writes an incomplete argument, or starts hunting for divisibility rules for each residue of the length — a case analysis with no evident end.

The model chain. Let N_k be the number with k ones separated by single zeros, so N_k has $2k-1$ digits. Its ones sit in the places $10^0, 10^2, 10^4, \ldots, 10^{2k-2}$, so

$$N_k = \sum_{j=0}^{k-1} 100^j = \frac{100^k - 1}{99} = \frac{(10^k-1)(10^k+1)}{99}.$$

$N_1 = 1$ is not prime and $N_2 = 101$ is prime. Now let $k \geq 3$.

Case k odd. Then $10^k - 1$ is divisible by 9 (as $10 \equiv 1 \pmod 9$) and $10^k + 1$ is divisible by 11 (as $10 \equiv -1 \pmod{11}$ and k is odd). Hence

$$N_k = \frac{10^k-1}{9} \cdot \frac{10^k+1}{11},$$

a product of two integers. Since $k \geq 3$, $\frac{10^k-1}{9} \geq 111 > 1$ and $\frac{10^k+1}{11} \geq 91 > 1$. So N_k is composite. (Check: $10101 = 111 \cdot 91$.)

Case k even. Then $10^k - 1 = 100^{k/2} - 1$ is divisible by $100 - 1 = 99$, so

$$N_k = \frac{10^k-1}{99} \cdot (10^k+1),$$

again a product of two integers. Since $k \geq 4$, $\frac{10^k-1}{99} \geq \frac{9999}{99} = 101 > 1$, and $10^k + 1 > 1$. So N_k is composite.

Hence the only prime of the required form is 101, and the answer is **one**.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths begin by listing $1, 101, 10101, \ldots$ They split at the next line. The natural attempt treats each listed number as an integer to be factored; the model treats the list as a single formula, $N_k = \sum_{j<k} 100^j$, and factors the formula. Once $N_k = (100^k - 1)/99$ is on the page, the difference of two squares $100^k - 1 = (10^k - 1)(10^k + 1)$ is forced, and the only remaining work is distributing the 99.

> *I am being asked about infinitely many numbers, so I cannot test them one by one. These are geometric series in disguise — a* 1 *every two places is powers of* 100*. Let me write N_k as a closed form and factor the closed form, not the number.*

> $100^k - 1$ *is a difference of two squares,* $(10^k - 1)(10^k + 1)$*. The* $99 = 9 \cdot 11$ *in the denominator has to be shared between the two factors; how it shares depends on the parity of k, so I split on parity.*

Layer 3 THE WHY

Why is "write it as a geometric series" the natural move rather than a trick? Go back to what a base-10 numeral *is*: the string $d_m \cdots d_1 d_0$ means $\sum d_i 10^i$. The problem tells you the digit string exactly — ones in the even places, zeros in the odd places — so the definition of base-10 notation hands you $\sum_j 10^{2j}$ without any inspiration at all. The phrasing "written as usual in base 10" is the problem-setter pointing at the definition. A question about the *digits* of a family of numbers is a question about a polynomial in 10 evaluated at a fixed pattern, and the only question is what that polynomial is.

Then why does factoring appear? Primality for a whole family cannot be settled by trial division, because there are infinitely many members and the divisors keep changing. It can only be settled by an identity valid for every k at once — an algebraic factorisation $N_k = A(k) \cdot B(k)$ with both factors > 1. The geometric-series closed form $\frac{x^k-1}{x-1}$ is exactly the object that has such factorisations: $x^k - 1$ factors whenever k factors, and $x^{2m} - 1 = (x^m - 1)(x^m + 1)$ always. Here $x = 100$ is itself a square, 10^2, so $100^k - 1$ is a difference of two squares for *every* k, not only for even k. That is the structural reason this family is almost entirely composite: its base is a perfect square. (Compare repunits $\frac{10^k-1}{9}$, base 10, not a square — there the question of which are prime is open, and nobody sets it as an A1.)

The final step — deciding how $99 = 9 \cdot 11$ divides the two factors — is the only place care is needed, and it is forced by the arithmetic of 10 modulo 9 and 11. $10 \equiv 1 \pmod 9$ means $9 \mid 10^k - 1$ always. $10 \equiv -1 \pmod{11}$ means $11 \mid 10^k + 1$ exactly when k is odd, and $11 \mid 10^k - 1$ exactly when k is even. So the parity split is not a choice; it is where 11 lives. The one thing you must not forget is that "product of two integers" proves compositeness only if both factors exceed 1 — and that check is what separates $k = 2$ (where $\frac{10^2-1}{99} = 1$, and indeed 101 is prime) from $k \geq 4$.

Layer $+\alpha$ BACKTESTING

1. The first cue is "alternating 1's and 0's" together with "how many primes": a digit pattern over an unbounded family. Write down $N_k = \sum_{j=0}^{k-1} 100^j = \frac{100^k-1}{99}$ immediately.

2. It suggests algebraic factorisation of the closed form (difference of two squares), rather than testing individual members.

3. Hesitation comes after the factorisation, when you notice the 99 in the denominator and worry that the factors are not integers. Reinforce it away by reducing 10 modulo 9 and modulo 11: $10 \equiv 1$ and $10 \equiv -1$ tell you at once which factor each prime divides.

4. Two traps. First, trial division case by case, which never terminates. Second, writing "$N_k = \frac{(10^k-1)(10^k+1)}{99}$, a product, so composite" without checking that both integer factors exceed 1 — that argument would wrongly declare 101 composite.

5. Whenever a problem describes a number by its digit pattern and asks a question for all lengths, the digit pattern is a geometric (or polynomial) sum in the base; get the closed form before doing anything numerical.

6. Reusable rule: to show an infinite family of integers is composite, find an identity $N = A \cdot B$ valid across the family and verify $A, B > 1$ outside finitely many small cases, which you check by hand. Look first for $x^{2m} - 1 = (x^m - 1)(x^m + 1)$ and, more generally, $x^{ab} - 1$ divisible by $x^a - 1$.

▷ **Cue:** a number is described by a repeating digit pattern (e.g. $1010\cdots01$, $111\cdots1$, 1001001) and the question is about all lengths $\longrightarrow$ **Move:** write it as a geometric series in a power of the base and find the closed form $\frac{x^k-1}{x-1}$

▷ **Cue:** a closed form $\frac{x^k-1}{x-1}$ with x a perfect square, or $x^k - 1$ with k composite $\longrightarrow$ **Move:** factor by difference of squares or by $x^a - 1 \mid x^{ab} - 1$

▷ **Cue:** an algebraic factorisation over a denominator like 99 or $9 \cdot 11$ $\longrightarrow$ **Move:** reduce the base modulo each prime factor of the denominator to decide which factor absorbs it; split on parity of the exponent if $10 \equiv -1$

▷ **Cue:** you conclude "composite" from a product of two integers $\longrightarrow$ **Move:** check both factors exceed 1, and handle the small cases where one factor equals 1 separately

1989 A2

1989 A2. Evaluate $\displaystyle\int_0^a \int_0^b e^{\max\{b^2x^2, a^2y^2\}}\, dy\, dx$ where a and b are positive.

Fields: calculus, multiple integrals **Position:** A2 — routine once the region is cut correctly **Answer:** $\dfrac{e^{a^2b^2} - 1}{ab}$. Cut the rectangle along its diagonal $bx = ay$; on each triangle the integrand depends on one variable only, and the inner integral supplies exactly the factor needed to integrate te^{t^2}.

Layer 1 THE WHAT — surface analysis

The natural attempt. The integral is written with y inside, so you start there: fix $x \in [0, a]$ and integrate $e^{\max\{b^2x^2, a^2y^2\}}$ over $y \in [0, b]$. The maximum switches at $a^2y^2 = b^2x^2$, i.e. at $y = bx/a$, which lies in $[0, b]$. So

$$\int_0^b e^{\max\{\cdot\}}\, dy = \int_0^{bx/a} e^{b^2x^2}\, dy + \int_{bx/a}^b e^{a^2y^2}\, dy = \frac{bx}{a} e^{b^2x^2} + \int_{bx/a}^b e^{a^2y^2}\, dy.$$

The first piece integrates in x without trouble. The second piece is the wall: $\int e^{a^2y^2}\, dy$ has no elementary antiderivative, so you cannot evaluate the inner integral, and you are left with $\int_0^a \int_{bx/a}^b e^{a^2y^2}\, dy\, dx$. Many students stop here, or try substitutions in y that go nowhere. The ones who push on realise they have to swap the order of integration on that piece — which is the model chain, reached late and by the back door.

The model chain. Write $R = [0, a] \times [0, b]$ and $f(x, y) = e^{\max\{b^2x^2, a^2y^2\}}$. Because $a, b, x, y \geq 0$, we have $b^2x^2 \geq a^2y^2$ if and only if $bx \geq ay$. The diagonal $bx = ay$ runs from $(0, 0)$ to (a, b) and cuts R into two triangles:

$$T_1 = \{(x, y) \in R : ay \leq bx\}, \qquad T_2 = \{(x, y) \in R : bx \leq ay\},$$

which overlap only along the diagonal (a set of area zero). On T_1, $f = e^{b^2x^2}$; on T_2, $f = e^{a^2y^2}$.

On T_1 integrate y first: for fixed $x \in [0, a]$, y runs over $[0, bx/a]$ (and $bx/a \le b$). So

$$\iint_{T_1} f = \int_0^a \frac{bx}{a}\, e^{b^2x^2}\, dx = \frac{b}{a}\Big[\frac{e^{b^2x^2}}{2b^2}\Big]_0^a = \frac{e^{a^2b^2}-1}{2ab}.$$

On T_2 integrate x first: for fixed $y \in [0, b]$, x runs over $[0, ay/b]$. So

$$\iint_{T_2} f = \int_0^b \frac{ay}{b}\, e^{a^2y^2}\, dy = \frac{a}{b}\Big[\frac{e^{a^2y^2}}{2a^2}\Big]_0^b = \frac{e^{a^2b^2}-1}{2ab}.$$

The integrand is continuous on the compact rectangle, so Fubini's theorem allows either order on each piece. Adding,

$$\int_0^a \int_0^b e^{\max\{b^2x^2, a^2y^2\}}\, dy\, dx = \frac{e^{a^2b^2}-1}{ab}.$$

(Symmetry check: interchanging $(x, a) \leftrightarrow (y, b)$ maps the problem to itself and $T_1 \leftrightarrow T_2$, which is why the two halves agree.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains find the switching line $y = bx/a$. They split at what they do with it. The natural attempt uses it only as a limit inside the given order of integration, $\int dx \int dy$, and so meets $\int e^{a^2y^2}\, dy$. The model uses it to cut the *region* into two pieces first, and then chooses the order of integration separately on each piece, integrating out the variable that does *not* appear in the integrand.

> *The order $dy\,dx$ was chosen by whoever wrote the problem, not by me. The* max *splits the rectangle into two regions; on each, the integrand depends on only one variable. I will integrate out the other variable first on each piece, so the inner integral is just a length.*

> *Integrating out the missing variable gives a length proportional to the remaining variable — $\frac{bx}{a}$ — and $x\, e^{b^2x^2}$ is exactly the derivative pattern of $e^{b^2x^2}$. That is why the answer is elementary.*

Layer 3 THE WHY

Kedlaya's style of solution for this problem is two lines: "split along the diagonal; each piece is $\frac{e^{a^2b^2}-1}{2ab}$." Here is where the split comes from, and why the order of integration is not a detail.

Start with what a max is. $\max\{u, v\}$ is a function defined piecewise: it equals u on one region and v on the other. Any integral of a piecewise-defined function begins, by definition of how you integrate such a thing, with partitioning the domain into the regions where each formula holds. So cutting R along $\{b^2x^2 = a^2y^2\}$ is not a trick; it is the first line forced by the word "max". The only structural fact needed is that on the positive quadrant $b^2x^2 = a^2y^2$ is the straight line $bx = ay$, and that this line is the rectangle's own diagonal — the constants a, b inside the max were chosen to match the limits of integration precisely so that this happens.

Now the order. A double integral over a region is one number; Fubini's theorem says you may compute it as iterated integrals in either order, as long as the integrand is continuous (here it is). So the $dy\,dx$ in the problem is a notation, not an instruction. On T_1 the integrand is $g(x) = e^{b^2x^2}$, independent of y. If you integrate y first, the inner integral is $\int g(x)\, dy = g(x) \cdot (\text{length of the } y\text{-slice})$ — you have integrated a constant. If you integrated x first, you would be integrating $e^{b^2x^2}$ in x, which has no elementary antiderivative. The general principle: *integrate first with respect to the variable the integrand does not depend on.* That converts the inner integral into a measurement of the region, and all the difficulty of the integrand is postponed to a one-variable integral where the region has contributed a helpful factor.

And the factor is helpful for a reason. The slice of a triangle with a vertex at the origin has length proportional to its distance from the origin: here $\frac{b}{a}x$. So you end up integrating $x\,e^{b^2x^2}$, which is $\frac{1}{2b^2}\frac{d}{dx}e^{b^2x^2}$. This is the same mechanism as polar coordinates in the Gaussian integral, where the Jacobian r turns $\int e^{-r^2}$ into $\int re^{-r^2}\,dr$. Triangles with a vertex at the origin behave like sectors: the width grows linearly, and linear growth is what e^{t^2} needs.

Layer $+\alpha$ BACKTESTING

1. The first cue is $\max\{\cdot,\cdot\}$ in the integrand. On seeing it, write down the set where the two arguments are equal — here $bx = ay$, the diagonal of the rectangle — and sketch the two triangles.

2. It suggests splitting the region and, on each piece, choosing the order of integration that integrates out the absent variable first.

3. Hesitation occurs when you meet $\int e^{a^2y^2}\,dy$ and think you need a special function. Reinforce it away with the rule: a non-elementary inner integral is a signal to swap the order, not to find a cleverer antiderivative.

4. The trap is obeying the printed order $dy\,dx$ over the whole rectangle. It gets half the answer quickly and then leaves an integral of $e^{a^2y^2}$ that cannot be done in that order.

5. Any double integral whose integrand is e^{t^2}, $\sin(t^2)$, $\frac{\sin t}{t}$ or similar non-elementary function of a single variable is almost certainly designed to be done in the order where that variable is integrated last, with the region supplying a factor of t.

6. Reusable rule: for piecewise integrands (max, min, $|\cdot|$, $\lfloor\cdot\rfloor$), first partition the domain along the switching set; then, independently on each piece, integrate first in whichever variable the local formula ignores.

▷ **Cue:** max, min, or $|\cdot|$ inside a double integral $\longrightarrow$ **Move:** find the switching curve, sketch it on the region, and split the integral along it

▷ **Cue:** the inner integral of an iterated integral has no elementary antiderivative $\longrightarrow$ **Move:** swap the order of integration (Fubini) so that the troublesome variable is integrated last

▷ **Cue:** on a piece of the region the integrand depends on only one variable $\longrightarrow$ **Move:** integrate out the other variable first; the inner integral becomes the length of a slice

▷ **Cue:** integrand e^{ct^2} over a triangle or sector with a vertex at the origin $\longrightarrow$ **Move:** the slice length is proportional to t, producing te^{ct^2}, which integrates exactly

1989 A3

1989 A3. Prove that if

$$11z^{10} + 10iz^9 + 10iz - 11 = 0,$$

then $|z| = 1$. (Here z is a complex number and $i^2 = -1$.)

Fields: complex numbers, polynomials **Position:** A3 — mid-paper; short once the right grouping is seen, a wall of inequalities otherwise **Answer:** Every root has $|z| = 1$. Group the equation as $z^9(11z + 10i) = 11 - 10iz$ and take moduli; a direct computation gives $|11z + 10i|^2 - |11 - 10iz|^2 = 21(|z|^2 - 1)$, so $|z| > 1$ and $|z| < 1$ each make the two sides unequal in modulus.

Layer 1 THE WHAT — surface analysis

The natural attempt. The target is a statement about $|z|$, so you take moduli. The obvious split is "big terms against small terms": $11z^{10} + 10iz^9 = 11 - 10iz$, hence

$$|z|^9\,|11z + 10i| = |11 - 10iz|.$$

Now you try to show $|z| > 1$ is impossible with the triangle inequality. Put $r = |z| > 1$. The left side is at least $r^9(11r - 10)$; the right side is at most $11 + 10r$. You would need $r^9(11r - 10) >$

$11 + 10r$, but at $r = 1$ the left side is 1 and the right side is 21, so the inequality fails for every r close to 1 — exactly the region where the claim has content. Swapping which side you bound above and which below does not help: the lower bound $|11 - 10iz| \geq 10r - 11$ is negative near $r = 1$. The coefficients 11 and 10 are too close for the triangle inequality to separate anything; it throws away the phases, and the phases are the whole problem.

A second natural try is to look for roots: $z = \pm 1$ and $z = \pm i$ do not work ($g(1) = 20i$, $g(i) = -22 - 20 = -42$, and so on), and a degree-10 polynomial will not factor by inspection. At this point many students either give up or reach for Rouché's theorem, which also needs a strict modulus comparison on a circle — the same comparison they could not make.

The model chain. Let $g(z) = 11z^{10} + 10iz^9 + 10iz - 11$. If $g(z) = 0$ then

$$z^9(11z + 10i) = 11 - 10iz. \tag{$*$}$$

Write $z = a + bi$ with a, b real. Then $11z + 10i = 11a + i(11b + 10)$ and $11 - 10iz = (11 + 10b) - 10ai$, so

$$|11z + 10i|^2 = 121(a^2 + b^2) + 220b + 100, \qquad |11 - 10iz|^2 = 100(a^2 + b^2) + 220b + 121,$$

and therefore

$$|11z + 10i|^2 - |11 - 10iz|^2 = 21\big(|z|^2 - 1\big). \tag{$\dagger$}$$

Case $|z| > 1$. By $(\dagger)$, $|11z + 10i| > |11 - 10iz| \geq 0$, so $11z + 10i \neq 0$, and $(*)$ gives $|z|^9 = |11 - 10iz|/|11z + 10i| < 1$, contradicting $|z| > 1$.

Case $|z| < 1$. By $(\dagger)$, $|11 - 10iz| > |11z + 10i|$. But $(*)$ gives $|11 - 10iz| = |z|^9|11z + 10i| \leq |11z + 10i|$, a contradiction.

Hence every root satisfies $|z| = 1$. (Numerically, all ten roots have modulus 1 to machine precision, as they must.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains write $z^9(11z + 10i) = 11 - 10iz$ and take moduli. They split at the next line. The natural attempt *estimates* $|11z + 10i|$ and $|11 - 10iz|$ separately with the triangle inequality. The model *computes* their squared moduli exactly and subtracts, getting $(\dagger)$: the difference is $21(|z|^2 - 1)$, with the troublesome $220b$ cancelling.

> *The two sides of the equation have the same coefficients,* 11 *and* 10*, in swapped roles. So* $|11z + 10i|$ *and* $|11 - 10iz|$ *should be equal exactly when* $|z| = 1$*. I will not estimate them; I will compute both squared moduli and subtract, and the cross terms should cancel.*

> *The coefficients are palindromic up to conjugation and sign:* $11, 10i, \ldots, 10i, -11$*. A polynomial like that has roots symmetric under* $z \mapsto 1/\bar{z}$*, so the unit circle is where the problem lives, and the right tool compares the two "halves" of the polynomial on either side of the circle.*

Layer 3 THE WHY

The official solutions (Gleason–Kedlaya et al.) give three proofs: an intermediate-value count of sign changes of $22\cos 5\theta + 20\cos 4\theta$ on the circle, the modulus computation above, and Rouché's theorem. None says why one should split the polynomial as $z^9(11z + 10i)$ against $11 - 10iz$. Here is where that comes from.

Start with the definition of what is being claimed. $|z| = 1$ means $\bar{z} = 1/z$. So the statement "all roots lie on the unit circle" is really a statement about how g behaves under $z \mapsto 1/\bar{z}$. Compute it: conjugating the coefficients of g and reversing their order gives $-11, -10i, 0, \ldots, 0, -10i, 11$, which is $-g$. Equivalently $z^{10}\,\overline{g(1/\bar{z})} = -g(z)$. So g is *self-inversive*:

its roots come in pairs $z, 1/\bar{z}$, one inside the circle and one outside, unless they sit on the circle. The problem is asking you to rule out the pairs, i.e. to show nothing is inside.

Now the split is forced. A self-inversive polynomial is a "high half" plus the reflected "low half". Here the low half is $h(z) = 10iz - 11$ and the high half is $11z^{10} + 10iz^9 = z^9(11z + 10i)$, and $11z + 10i$ is precisely $-z\,\overline{h(1/\bar{z})}$ — the reflection of h. On the unit circle, reflection preserves modulus, so $|11z + 10i| = |10iz - 11|$ there; that is why the constants must be the same pair $11, 10$ in swapped positions. Off the circle the two moduli separate in a controlled way, and $(\dagger)$ measures the separation exactly: $|11z + 10i|^2 - |11 - 10iz|^2 = (11^2 - 10^2)(|z|^2 - 1)$. The factor $21 = 11^2 - 10^2$ is positive because the constant term of h beats its linear term, i.e. h's only root $11/(10i)$ lies *outside* the unit circle. That single fact — the low half has no root inside — is what the whole proof rests on, and it is also exactly the hypothesis Rouché needs in Solution 3.

This is also why the triangle inequality fails. The identity $(\dagger)$ is an identity of the form $|\alpha z + \beta|^2 - |\bar{\beta}z + \bar{\alpha}|^2 = (|\alpha|^2 - |\beta|^2)(|z|^2 - 1)$; its proof needs the cross terms $2\,\mathrm{Re}(\alpha\bar{\beta}z)$ on both sides to cancel. The triangle inequality replaces each cross term by its worst case independently, so it cannot see the cancellation. Whenever an argument depends on two phase-dependent quantities being equal on a curve, you must compute, not bound.

Finally, the sign-change proof tells the same story from the other side. Substituting $z = iw$ and dividing by $-w^5$ gives $11(w^5 + w^{-5}) + 10(w^4 + w^{-4})$, real on $|w| = 1$ because self-inversive polynomials become real trigonometric polynomials on the circle after this normalisation. Then $22\cos 5\theta$ dominates $20\cos 4\theta$ (again $11 > 10$), forcing ten sign changes, hence ten roots on the circle, hence all of them.

Layer $+\alpha$ BACKTESTING

1. The first cue is the coefficient pattern $11, 10i, \ldots, 10i, -11$: the list reads the same backwards up to conjugation and a sign. On seeing it, write down $z^{10}\overline{g(1/\bar{z})} = -g(z)$ and the grouping $z^9(11z + 10i) = 11 - 10iz$.

2. It suggests comparing $|11z + 10i|$ with $|11 - 10iz|$ as functions of $|z|$, by exact computation of squared moduli, and then arguing by cases $|z| > 1$, $|z| < 1$.

3. Hesitation occurs after taking moduli, when you must decide whether to estimate or compute. Reinforce it away with the rule: if the claimed equality case is a whole curve ($|z| = 1$), inequalities that lose phase information cannot be sharp there, so compute $|\cdot|^2$ exactly.

4. The trap is the triangle inequality. It is the reflex for modulus problems, it is sharp only when phases align, and with coefficients 11 and 10 it fails in precisely the annulus near $|z| = 1$ where the statement is non-trivial.

5. Any polynomial whose coefficient list is its own conjugate reversal up to a unimodular factor is self-inversive; claims about roots on the unit circle for such polynomials reduce to comparing the top half with the reflected bottom half. Recognise it by reading the coefficients from both ends.

6. Reusable rule: to prove a quantity is exactly 1, split the equation into two pieces that are equal in modulus on the unit circle, compute $|A|^2 - |B|^2$ exactly as a multiple of $(|z|^2 - 1)$, and let the sign of that multiple exclude both $|z| < 1$ and $|z| > 1$.

▷ **Cue:** a polynomial whose coefficients read the same backwards up to conjugation and a constant factor of modulus 1 $\longrightarrow$ **Move:** recognise it as self-inversive: roots come in pairs $z, 1/\bar{z}$, so show there are none inside the unit circle

▷ **Cue:** prove every root has $|z| = 1$ $\longrightarrow$ **Move:** group as $z^k A(z) = B(z)$ with $|A| = |B|$ on the circle, take moduli, and argue by the cases $|z| < 1$, $|z| > 1$

▷ **Cue:** need to compare $|\alpha z + \beta|$ with $|\bar{\beta}z + \bar{\alpha}|$ $\longrightarrow$ **Move:** compute the squared moduli exactly: the difference is $(|\alpha|^2 - |\beta|^2)(|z|^2 - 1)$

▷ **Cue:** triangle-inequality estimates fail near the expected equality case $\longrightarrow$ **Move:** switch from bounding to exact computation of $|\cdot|^2$, so the cross terms can cancel

1989 A4

1989 A4. If α is an irrational number, $0 < \alpha < 1$, is there a finite game with an honest coin such that the probability of one player winning the game is α? (An honest coin is one for which the probability of heads and the probability of tails are both $\frac{1}{2}$. A game is finite if with probability 1 it must end in a finite number of moves.)

Fields: probability, number representations **Position:** A4 — mid-paper; short once the binary expansion of α is in view, a dead end if you think in terms of fixed numbers of tosses **Answer:** Yes, for every real $\alpha \in [0,1]$ (irrationality is not needed). Write $\alpha = 0.b_1b_2b_3\ldots$ in binary, toss the coin to generate digits $d_1d_2\ldots$ of a uniform random number X, and stop at the first n with $d_n \neq b_n$; player 1 wins if $d_n = 0$, $b_n = 1$. Player 1 wins with probability $\sum b_n 2^{-n} = \alpha$, the game stops at toss n with probability 2^{-n}, and the expected length is 2.

Layer 1 THE WHAT — surface analysis

The natural attempt. You start with what a coin can do in a fixed number of tosses. After N tosses there are 2^N equally likely outcomes, and any rule that declares a winner after exactly N tosses gives player 1 a probability of the form $k/2^N$. So bounded games produce only dyadic rationals. That is a real theorem, and it is the first wall: an irrational α cannot come from a game of bounded length.

Next you remember the standard trick for $\frac{1}{3}$: toss twice, HH means player 1 wins, HT or TH means player 2 wins, TT means start again. The game ends with probability 1 (it continues forever with probability $\lim(1/4)^k = 0$), and player 1 wins with probability $\frac{1/4}{3/4} = \frac{1}{3}$. The same rejection idea gives every rational p/q: toss N times with $2^N \geq q$, read the result as an integer in $[0, 2^N)$, player 1 wins on $[0, p)$, player 2 on $[p, q)$, restart on $[q, 2^N)$. So now you have every rational.

And here you stall. The rationals are dense, so you try to "approach" α: play the game for p_1/q_1, then…what? A sequence of separate games does not have a limiting winning probability unless the games are welded together, and a mixture $\sum c_k \cdot (\text{game } k)$ needs its own randomisation of the weights c_k. Some students conclude the answer is no, on the grounds that every construction they know outputs a rational. The attempt fails because it treats each game as a finished object with a finite description, when the problem's definition of "finite" explicitly permits unbounded length.

The model chain. Write $\alpha = \sum_{n\geq 1} b_n 2^{-n}$ with each $b_n \in \{0,1\}$. (If α is irrational this expansion is unique; for a dyadic rational pick either one.) Toss the coin repeatedly and let $d_n = 1$ if toss n is heads and $d_n = 0$ if tails.

Rule. Stop at the first n with $d_n \neq b_n$. If $d_n = 0$ (so $b_n = 1$) player 1 wins; if $d_n = 1$ (so $b_n = 0$) player 2 wins.

Finiteness. The game is still running after n tosses exactly when $d_k = b_k$ for $k = 1, \ldots, n$, an event of probability 2^{-n}. These events decrease to "never stops", whose probability is therefore $\lim 2^{-n} = 0$. So the game is finite in the problem's sense.

Winning probability. Player 1 wins at toss n exactly when $d_1 \ldots d_{n-1} = b_1 \ldots b_{n-1}$ and $d_n = 0$ and $b_n = 1$. For fixed n this has probability $2^{-(n-1)} \cdot \frac{1}{2} \cdot b_n = b_n 2^{-n}$. The events for different n are disjoint, so by countable additivity

$$\Pr(\text{player 1 wins}) = \sum_{n\geq 1} b_n 2^{-n} = \alpha.$$

Player 2 wins with probability $\sum(1 - b_n)2^{-n} = 1 - \alpha$, consistent with finiteness.

Bonus. The game stops at toss n with probability 2^{-n}, so the expected number of tosses is $\sum n2^{-n} = 2$, whatever α is.

Interpretation: $X = \sum d_n 2^{-n}$ is uniform on $[0,1]$, and the rule is "player 1 wins iff $X < \alpha$", decided at the first moment the digits make the comparison certain. Kedlaya–Poonen–Vakil phrase it exactly this way.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that N fixed tosses give only multiples of 2^{-N}. They split at the next line. The natural attempt reads this as "the coin produces rationals, so I need a limiting process over games". The model reads it as "N tosses give the first N binary digits of a uniform number, so unboundedly many tosses give *all* of them", and compares that number with α digit by digit.

> *A fixed number of tosses gives me $k/2^N$. That is a truncated binary expansion. The problem explicitly allows games of unbounded length, so I should use all the digits: the coin is a uniform random real in $[0,1]$, written out one binary digit at a time, and player 1 wins if it lands below α.*

> *I never need to know the whole random number. I only need to know which side of α it is on, and that is settled at the first digit where it disagrees with α.*

Layer 3 THE WHY

Why is the binary expansion the natural object and not a trick? Go back to what the problem gives you. A fair coin produces one bit with probability $\frac{1}{2}$ each way. A sequence of independent fair bits *is*, by definition of binary notation, a real number $X \in [0,1]$, and the probability that its first m digits equal a given string is 2^{-m}, which is the length of the dyadic interval those digits describe. So $\Pr(X \in I) = \text{length}(I)$ for every dyadic interval, hence X is uniform. The coin is not a device for making rationals; it is a device for sampling a uniform real, slowly. Once you see that, "win with probability α" translates to "win iff $X < \alpha$", because for a uniform X that event has probability α by the definition of uniform.

Why the stopping rule is forced: comparison of two numbers in the same base is lexicographic — you scan until the first differing digit, and that digit decides. That is the same algorithm you use to decide that $0.3141\ldots < 0.32$. Applied to X and α, the first differing digit n decides the winner and no later toss can change it. With probability 1 such a digit exists, since the only way it fails to exist is $X = \alpha$ digit for digit, and a single point has probability 0.

Why the problem defines "finite". The parenthetical is the hint the whole problem hangs on. If "finite" meant "bounded number of tosses", the answer would be no, by the $k/2^N$ argument above: an irrational is never a dyadic rational. The setters therefore wrote the definition as "ends with probability 1", which is precisely the licence to use infinitely many digits. Reading the definition as a clue, not as legal boilerplate, is the move. The word "irrational" is a decoy in the same spirit: it steers you towards thinking rationals are easy and irrationals special, when the construction treats every real identically.

Why the natural attempt stalls. The rejection-sampling game for p/q is a correct and useful circuit, but it depends on a finite list of outcomes and a restart. An irrational has no finite description to put into that list. The fix is not to stack infinitely many rational games; it is to notice that the rejection game itself is secretly comparing X with a number whose binary expansion is eventually periodic. Remove the periodicity and the same comparison works for any expansion.

Kedlaya–Poonen–Vakil's solution begins by defining $X = \sum d_n/2^n$ and asserting it is uniform. Here is where that comes from: it is the only object that uses every toss and turns "probability" into "length", which is what α is.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the definition in the parenthesis: a game is finite if it ends with probability 1, not in a bounded number of moves. On seeing it, write down: "bounded

games give only $k/2^N$; so the game must be allowed to run arbitrarily long" and then $\alpha = 0.b_1b_2b_3\ldots_2$.

2. Binary expansion plus a uniform random variable: the coin tosses are the binary digits of $X \sim U[0,1]$, and the event $X < \alpha$ has probability α.

3. Hesitation occurs when you realise you cannot finish generating X. Reinforce it away with: to decide an inequality between two reals you only need the first digit where they differ, and that digit arrives at a finite time with probability 1.

4. Two traps. First, the word "irrational", which suggests irrationals need special treatment; they do not. Second, a strict reading of "finite" as "bounded", which gives the wrong answer "no" through a correct argument about $k/2^N$.

5. Whenever a problem asks for an arbitrary real probability from a source of discrete fair randomness, think of the source as producing the digits of a uniform real, and of the target as a threshold on it.

6. Reusable rule: to realise a real number α as an outcome of an unbounded process, expand α in the process's natural base and decide at the first digit where the process diverges from the expansion. The same rule settles lexicographic comparisons, dyadic interval arguments, and stopping-time constructions.

▷ **Cue:** a fair coin and a target probability that is not of the form $k/2^N$ ⟶ **Move:** treat the tosses as binary digits of a uniform $X \in [0,1]$ and let player 1 win iff $X < \alpha$

▷ **Cue:** a problem defines "finite" as "ends with probability 1" ⟶ **Move:** the intended construction has unbounded length; bounded versions are the trap

▷ **Cue:** need to decide $X < \alpha$ while X is revealed one digit at a time ⟶ **Move:** stop at the first digit where X and α differ; that digit decides

▷ **Cue:** a rational probability p/q from a fair coin ⟶ **Move:** rejection sampling: toss N times with $2^N \geq q$, restart on outcomes in $[q, 2^N)$

▷ **Cue:** the problem restricts to irrational values ⟶ **Move:** check whether the construction works for all reals; the restriction may be a decoy

1989 A5

1989 A5. Let m be a positive integer and let $\mathcal{G}$ be a regular $(2m+1)$-gon inscribed in the unit circle. Show that there is a positive constant A, independent of m, with the following property. For any points p inside $\mathcal{G}$ there are two distinct vertices v_1 and v_2 of $\mathcal{G}$ such that

$$\big|\,|p - v_1| - |p - v_2|\,\big| < \frac{1}{m} - \frac{A}{m^3}.$$

Here $|s - t|$ denotes the distance between the points s and t.

Fields: geometry, inequalities, pigeonhole principle **Position:** A5 — late in the paper, but short once the target is read correctly; the difficulty is entirely in seeing what the A/m^3 is asking for **Answer:** Any A with $0 < A \leq 1/18$ works (the official solution notes any $A < \pi^2/32$ works for large m). The $2m+1$ distances from p to the vertices lie in an interval of length at most $w = 2\cos\frac{\pi}{4m+2}$, the longest diagonal; pigeonhole on the $2m$ gaps gives a gap at most $w/(2m) = \frac{1}{m}\cos\frac{\pi}{4m+2} < \frac{1}{m} - \frac{1}{18m^3}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You see "two distinct vertices" and "difference of distances less than something", and pigeonhole comes to mind: there are $2m+1$ distances, and you want two of them close together. You bound the range. Every vertex is on the unit circle and p is inside it, so each distance $|p - v|$ lies in $[0,2]$. Sort the distances $d_1 \leq d_2 \leq \cdots \leq d_{2m+1}$. The $2m$ consecutive gaps add up to $d_{2m+1} - d_1 \leq 2$, so some gap is at most $2/(2m) = 1/m$.

And there you stop, one hair short. You have $\le 1/m$; the problem wants $< \frac{1}{m} - \frac{A}{m^3}$. The natural response is to think the pigeonhole was too crude and to start computing. Put $p = (x, y)$ and the vertices at angles $2\pi k/(2m+1)$; then $|p - v_k|^2 = 1 + x^2 + y^2 - 2x\cos\frac{2\pi k}{2m+1} - 2y\sin\frac{2\pi k}{2m+1}$, and you try to locate two indices k whose distances are provably close. That leads into a case analysis on where p sits relative to the vertices, with square roots of trigonometric expressions to compare — a genuine Solution 2 exists along these lines and proves a far stronger bound, but it is a page of careful estimates and it is not what the problem is asking for. Most students who go down this road run out of time in the case analysis.

The model chain. Let $n = 2m + 1$ and let the vertices be $v_k = e^{2\pi ik/n}$, $k = 0, \dots, n-1$.

Step 1: the longest diagonal. Two vertices v_j, v_k subtend a central angle $2\pi\ell/n$ with $\ell \in \{1, \dots, m\}$ (take the shorter arc), so $|v_j - v_k| = 2\sin\frac{\pi\ell}{n}$. This is increasing in ℓ on $1 \le \ell \le m$ because $\pi\ell/n < \pi/2$ there. The maximum is at $\ell = m$:

$$w = 2\sin\frac{\pi m}{2m+1} = 2\sin\left(\frac{\pi}{2} - \frac{\pi}{2(2m+1)}\right) = 2\cos\frac{\pi}{4m+2}.$$

Step 2: the range of the distances. Let $d_1 \le \cdots \le d_n$ be the distances from p to the vertices, with $d_1 = |p - v_a|$, $d_n = |p - v_b|$, $a \ne b$. By the triangle inequality,

$$d_n - d_1 = |p - v_b| - |p - v_a| \le |v_b - v_a| \le w.$$

Step 3: pigeonhole. $\sum_{i=1}^{2m}(d_{i+1} - d_i) = d_n - d_1 \le w$, and the $2m$ summands are non-negative, so some $d_{i+1} - d_i \le w/(2m)$. The two vertices realising d_i and d_{i+1} are distinct, and

$$\big||p - v_1| - |p - v_2|\big| \le \frac{w}{2m} = \frac{1}{m}\cos\frac{\pi}{4m+2}.$$

Step 4: the constant. Let $x = \pi/(4m+2)$, so $0 < x \le \pi/6$. Then $1 - \cos x = 2\sin^2(x/2)$, and on $0 < y < \pi/2$ we have $\sin y > 2y/\pi$ (the sine is concave there, so its graph lies above the chord from $(0,0)$ to $(\pi/2, 1)$). With $y = x/2$ this gives $\sin(x/2) > \frac{1}{4m+2}$, so

$$1 - \cos x > \frac{2}{(4m+2)^2} = \frac{1}{2(2m+1)^2} \ge \frac{1}{2(3m)^2} = \frac{1}{18m^2},$$

using $2m + 1 \le 3m$ for $m \ge 1$. Hence

$$\frac{w}{2m} = \frac{1}{m} - \frac{1 - \cos x}{m} < \frac{1}{m} - \frac{1}{18m^3},$$

and $A = 1/18$ works for every $m \ge 1$. (A numerical check over $1 \le m \le 39$ with random p confirms the chain; the true smallest gap is far below this bound, of order $\frac{2\pi^2}{3n^2}$.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains sort the distances and pigeonhole the gaps. They split at the bound on the range. The natural attempt writes "each distance lies in $[0, 2]$", so the range is at most 2. The model writes $d_n - d_1 \le |v_b - v_a| \le w$: the range of the distances is bounded by the longest *diagonal of the polygon*, not by the diameter of the circle. Everything after that line is bookkeeping.

> *Pigeonhole with range* 2 *gives* $1/m$ *exactly. The target is* $1/m$ *minus a tiny correction of order* $1/m^3$*, which is* $\frac{1}{m}$ *times something of order* $\frac{1}{m^2}$*. So I don't need a new method; I need the range to be* 2 *minus something of order* $1/m^2$*. What is a little less than* 2 *here?*

> *The polygon has an odd number of vertices, so no two vertices are diametrically opposite. The longest diagonal falls short of 2 by about $\frac{\pi^2}{4n^2}$ — that is exactly an order-$1/m^2$ shortfall. And the triangle inequality says two distances from p can never differ by more than the distance between the two vertices.*

Layer 3 THE WHY

Kedlaya, Poonen and Vakil's Solution 1 opens with "the greatest distance between two vertices of $\mathcal{G}$ is $w = 2\cos\frac{\pi}{4m+2}$" and never says why one would look at it. Here is where that comes from.

Start from the shape of the target. The right-hand side $\frac{1}{m} - \frac{A}{m^3}$ is not an arbitrary expression: it is a leading term plus a correction two orders smaller. When a problem hands you "main term minus a lower-order term", it is telling you the main term comes from a crude argument and the correction from a second-order refinement of that same argument. So the first job is to find the crude argument that produces exactly $\frac{1}{m}$. Pigeonhole produces it: $2m+1$ numbers, $2m$ gaps, range 2, so $\frac{2}{2m} = \frac{1}{m}$. The $2m$ in the denominator is why the problem is phrased with $2m+1$ vertices and with m rather than n — the statement is written in the natural variables of the pigeonhole count.

Now ask where an $\frac{A}{m^3} = \frac{1}{m}\cdot\frac{A}{m^2}$ correction can come from. Pigeonhole has only two inputs: the number of gaps and the length of the range. The number of gaps is fixed at $2m$. So the range must be $2 - O(1/m^2)$ rather than 2. The signature of an $O(1/m^2)$ shortfall in geometry is $1-\cos\theta$ with θ of order $1/m$: a chord that subtends almost, but not quite, a straight angle. That points straight at the vertex configuration, and at the defining property of an *odd* regular polygon: since $v \mapsto -v$ is not a symmetry of the vertex set when n is odd, no vertex is antipodal to another, and the nearest thing to an antipode is off by half a step, an angle of π/n. That is the $\frac{\pi}{4m+2} = \frac{\pi}{2n}$ in w.

Why is the longest diagonal, rather than anything about p, the right bound for the range? Because a distance from p is a 1-Lipschitz function of the vertex: $\big||p-u| - |p-v|\big| \le |u-v|$ is the triangle inequality in the triangle puv. So the spread of the n distances is controlled entirely by the spread of the vertices, whatever p is. This has a consequence worth noticing: the hypothesis "p inside $\mathcal{G}$" is never used. The model chain works for every point in the plane. The natural attempt used "p inside the circle" to get $d \le 2$, and that use is exactly what cost it the correction term — it bounded each distance separately, when the problem is about differences of distances.

Finally, why does the constant need care at small m? The expansion $\cos x = 1 - \frac{x^2}{2} + O(x^4)$ gives $\frac{w}{2m} = \frac{1}{m} - \frac{\pi^2}{32m^3} + o(m^{-3})$, which proves the claim for large m only. The official solution patches the finitely many small cases by shrinking A, which is legitimate because $\frac{w}{2m} < \frac{1}{m}$ for every m. The model chain above avoids the patch by using the concavity bound $\sin y > 2y/\pi$, which holds on the whole range, so a single explicit $A = 1/18$ covers all m at once. The worst case is $m = 1$, where $m^2(1-\cos\frac{\pi}{6}) = 1 - \frac{\sqrt{3}}{2} \approx 0.134$; the ratio rises towards $\pi^2/32 \approx 0.308$ as $m \to \infty$.

> **Caveat.** The official Solution 2 and the "stronger result" in Kedlaya–Poonen–Vakil (a gap below π^2/n^2, improved to the sharp $\frac{2\pi^2}{3n^2}$) are quoted, not re-derived, here. A random-sampling check for $1 \le m \le 39$ gives worst observed gaps of about $0.66\,\pi^2/n^2$, consistent with the sharp constant, but that is evidence, not proof. The model chain above does not depend on them.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the form of the bound, $\frac{1}{m} - \frac{A}{m^3}$: a leading term with a correction two orders down. On seeing it, write down the crude argument that gives the leading term exactly: sort the $2m+1$ distances, $2m$ gaps, range ≤ 2, some gap $\le \frac{1}{m}$.

2. It suggests keeping the pigeonhole and sharpening its one soft input, the length of the range, by an amount of order $1/m^2$. The oddness of $2m+1$ then says where the shortfall lives: the longest diagonal $2\cos\frac{\pi}{4m+2} < 2$.

3. Hesitation comes right after reaching "$\leq \frac{1}{m}$": it feels like the method failed and a new one is needed. Reinforce it away by dividing the target by the bound you got: $\left(\frac{1}{m} - \frac{A}{m^3}\right)\Big/\frac{1}{m} = 1 - \frac{A}{m^2}$. That tells you precisely what the refinement must deliver, and that it is small.

4. Two traps. The first is bounding each distance by 2 separately, which cannot see that the vertices themselves are never 2 apart. The second is abandoning pigeonhole for coordinates and trigonometric case analysis — zombie maths that proves more than is asked at many times the cost. A minor third: proving the constant only asymptotically and forgetting small m.

5. Whenever a target is "crude bound minus lower-order term", the problem is almost always the crude argument with one input tightened. Recognise it by the exponents: a correction of relative size $1/m^2$ smells of $1-\cos(c/m)$, i.e. a near-straight angle or a near-diameter chord.

6. Reusable rule: to force two of N quantities close together, pigeonhole the sorted gaps against the range, and bound the range by the Lipschitz constant times the spread of the underlying objects — differences of distances are bounded by distances between the points, not by bounds on the individual distances.

▷ **Cue:** a bound of the shape $\frac{1}{m} - \frac{A}{m^3}$, main term minus a correction two orders smaller $\longrightarrow$ **Move:** find the crude argument giving exactly $\frac{1}{m}$, then sharpen its softest input by a relative $O(1/m^2)$

▷ **Cue:** N real quantities, show two of them differ by less than $L/(N-1)$ $\longrightarrow$ **Move:** sort them and pigeonhole the $N-1$ consecutive gaps against the range L

▷ **Cue:** difference of distances $|p-u| - |p-v|$ from a common point $\longrightarrow$ **Move:** bound it by $|u-v|$ (triangle inequality): the spread of distances is controlled by the spread of the targets, for any p

▷ **Cue:** a regular polygon with an odd number of vertices $\longrightarrow$ **Move:** no two vertices are antipodal; the longest diagonal is $2\cos\frac{\pi}{2n}$, short of the diameter by about $\frac{\pi^2}{4n^2}$

▷ **Cue:** a Taylor expansion proves an inequality only for large m $\longrightarrow$ **Move:** replace it by a global inequality such as $\sin y > \frac{2y}{\pi}$ on $(0, \frac{\pi}{2})$ or $\sin y < y$, and get one explicit constant for every m

1989 A6

1989 A6. Let $\alpha = 1 + a_1 x + a_2 x^2 + \cdots$ be a formal power series with coefficients in the field of two elements. Let

$$a_n = \begin{cases} 1 & \text{if every block of zeros in the binary expansion of } n \text{ has an even number of zeros in the block} \\ 0 & \text{otherwise.} \end{cases}$$

(For example, $a_{36} = 1$ because $36 = 100100_2$ and $a_{20} = 0$ because $20 = 10100_2$.) Prove that $\alpha^3 + x\alpha + 1 = 0$.

Fields: algebra (formal power series over $\mathbb{F}_2$), combinatorics on binary digits, functional equations
Position: A6 — the last problem of the morning paper, and it looks it: an unfamiliar ring, a digit condition, a cubic. It collapses to half a page once one structural fact about characteristic 2 is used
Answer: A proof, not a number. The digit rules $a_{2m+1} = a_m$, $a_{4m} = a_m$, $a_{4m+2} = 0$ say exactly that

$\alpha = x\alpha(x^2) + \alpha(x^4)$; in characteristic 2 this is $\alpha = x\alpha^2 + \alpha^4$, i.e. $\alpha(\alpha^3 + x\alpha + 1) = 0$, and α is invertible because its constant term is 1.

Layer 1 THE WHAT — surface analysis

The natural attempt. You read "prove $\alpha^3 + x\alpha + 1 = 0$" and you do what you do with any identity between power series: compare coefficients. Over $\mathbb{F}_2$ the coefficient of x^n in α^3 is

$$[x^n]\,\alpha^3 = \sum_{i+j+k=n} a_i a_j a_k \pmod 2,$$

so you must show that this equals $[x^n](x\alpha + 1) = a_{n-1} + [n = 0]$ for every n. You check small cases. The sequence starts

$$a_0, a_1, \ldots, a_{15} = 1, 1, 0, 1, 1, 0, 0, 1, 0, 1, 0, 0, 1, 0, 0, 1,$$

and for $n = 0, 1, 2, 3$ the identity holds (for instance $[x^2]\alpha^3 = 3a_0^2 a_2 + 3a_0 a_1^2 = 0 + 1 = 1 = a_1$). Encouraged, you look for a general argument, and you hit the wall. The triple convolution has no reason to respect binary digits: when you add $i + j + k$, carries scramble the blocks of zeros of the summands, and the condition "every block of zeros has even length" does not survive addition in any way you can track. You try to split $\alpha^3 = \alpha^2 \cdot \alpha$ and count pairs, you try induction on n with the cubic as a recursion for a_n (the cubic does determine α uniquely among series with constant term 1, but the recursion it gives expresses a_n through a sum of products of earlier terms, not through digits). Twenty minutes later you have a table of verified coefficients and no proof. This is the honest default, and it is a dead end: not because it is too long, but because addition of indices is the wrong operation to pair with a condition on digits.

The model chain. Write $\mathbb{F}_2[[x]]$ for the ring of formal power series with coefficients in $\mathbb{F}_2 = \{0, 1\}$. By the statement, $a_0 = 1$.

Step 1: squaring is substitution. For any $\beta = \sum b_n x^n \in \mathbb{F}_2[[x]]$,

$$\beta^2 = \sum_n b_n^2 x^{2n} + 2\sum_{i<j} b_i b_j x^{i+j} = \sum_n b_n x^{2n} = \beta(x^2),$$

because $2 = 0$ in $\mathbb{F}_2$ and $b^2 = b$ for $b \in \{0, 1\}$. Applying this twice, $\beta^4 = \beta(x^4)$. (For power series the cross-term computation is legitimate coefficient by coefficient: each coefficient of β^2 is a finite sum.) Hence

$$\alpha^4 = \sum_m a_m x^{4m}, \qquad x\alpha^2 = \sum_m a_m x^{2m+1}.$$

Step 2: three digit rules. Write the binary expansion of 0 as the empty string, consistent with $a_0 = 1$. Every $n \geq 0$ is of exactly one of the forms $2m + 1$, $4m$, $4m + 2$ with $m \geq 0$.

- $n = 2m + 1$: the expansion of n is that of m followed by the digit 1. Appending a 1 creates no zeros and does not merge or lengthen any block, so the blocks of zeros of n are those of m: $a_{2m+1} = a_m$. (For $m = 0$: $n = 1 = 1_2$, no blocks, $a_1 = 1 = a_0$.)

- $n = 4m$, $m \geq 1$: the expansion is that of m followed by 00. If m ends in 1 this adds a new final block of length 2; if m ends in a block of k zeros, that block becomes length $k + 2$, same parity. All other blocks are unchanged. So $a_{4m} = a_m$; and for $m = 0$ the equation reads $a_0 = a_0$.

- $n = 4m + 2$: the expansion ends $\ldots 10$, so its final block of zeros has length exactly 1, which is odd: $a_{4m+2} = 0$.

Step 3: compare coefficients of $\alpha^4 + x\alpha^2 + \alpha$. Let b_n be its coefficient of x^n. By Step 1, α^4 contributes only at $n \equiv 0 \pmod 4$ and $x\alpha^2$ only at odd n. So:

$$b_{2m+1} = a_m + a_{2m+1} = 0, \qquad b_{4m} = a_m + a_{4m} = 0, \qquad b_{4m+2} = a_{4m+2} = 0,$$

by Step 2. Therefore

$$\alpha^4 + x\alpha^2 + \alpha = \alpha\,(\alpha^3 + x\alpha + 1) = 0.$$

Step 4: cancel α. Since α has constant term 1, write $\alpha = 1 + \gamma$ with $\gamma \in x\,\mathbb{F}_2[[x]]$; then $\alpha^{-1} = 1 + \gamma + \gamma^2 + \cdots$ is a well-defined power series (the coefficient of x^n involves only $\gamma^0, \ldots, \gamma^n$), and $\alpha\alpha^{-1} = 1$. Multiplying by α^{-1} gives $\alpha^3 + x\alpha + 1 = 0$. ■

(Verified by machine: the identity holds for the first 2000 coefficients, and the three digit rules hold on the same range.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start at the same line: "compare coefficients". They split on *which* identity to compare. The natural attempt compares coefficients of α^3 against $x\alpha + 1$ and is forced into a triple convolution. The model first multiplies the target by α, to get $\alpha^4 + x\alpha^2 + \alpha = 0$ (†), and compares coefficients there. At (†) every power of α is a power of 2, every such power is a substitution $x \mapsto x^{2^k}$, and no convolution ever appears.

> *I'm in characteristic* 2. *Squaring is free here:* $(a+b)^2 = a^2 + b^2$, *so* α^2 *is* α *with* x *replaced by* x^2. *Cubes are not free. So I want an equation in which every power of* α *is* 1, 2 *or* 4 — *multiply the cubic by* α *and see what I get.*

> *The condition on* n *is about binary digits, so I should read* n *from its last digits: odd, a multiple of* 4, *or* 2 *mod* 4. *Replacing* x *by* x^2 *appends a* 0 *to every exponent; multiplying by* x *after that appends a* 1. *The series is going to satisfy an equation that says "append a digit".*

> α *starts with* 1, *so it is invertible. Throwing away a factor of* α *costs nothing.*

Layer 3 THE WHY

Kedlaya, Poonen and Vakil's solution opens: "It suffices to prove that $\alpha^4 + x\alpha^2 + \alpha = 0$." That sentence is the entire problem, and it is stated without a word of motivation. Here is where it comes from.

Start with the defining fact of the coefficient field. In $\mathbb{F}_2$, $1 + 1 = 0$. The binomial expansion of $(u+v)^2$ is $u^2 + 2uv + v^2$, and the middle term is killed. So squaring is additive, $(u+v)^2 = u^2 + v^2$, and by induction it passes through any finite sum, hence through every coefficient of a power series. Combined with $b^2 = b$ for $b \in \{0, 1\}$, this says $\alpha^2 = \alpha(x^2)$. This is not a trick; it is the single most important structural fact about characteristic p (the Frobenius map $u \mapsto u^p$ is a ring homomorphism), and any problem that says "the field of two elements" is inviting you to use it. The question to ask on reading "coefficients in $\mathbb{F}_2$" is: *what becomes linear that is usually not?* The answer is squaring.

Now look at the target, $\alpha^3 + x\alpha + 1$. The exponent 3 is not a power of 2, so Frobenius cannot touch α^3 directly — that is why coefficient comparison produced a convolution. The cheapest fix is to change the exponents: multiply by α and they become $4, 2, 1$, all powers of 2. A polynomial of the form $c_2Y^4 + c_1Y^2 + c_0Y$ is called *additive* (or linearised) in characteristic 2, because it respects addition. Multiplying by α is legitimate because α is a unit, and it is visible that α is a unit from the statement's own first words, "$\alpha = 1 + \cdots$": the problem put the constant term 1 there on purpose.

Why should an additive equation be the right thing to test against a digit condition? Because the three operations it produces are exactly digit operations on exponents. $x^n \mapsto x^{2n}$ appends a 0 to n in binary; $x^n \mapsto x^{2n+1}$ appends a 1; $x^n \mapsto x^{4n}$ appends 00. So $\alpha(x^2)$, $x\alpha(x^2)$, $\alpha(x^4)$ are the generating functions of "a read after appending 0, 1, 00". A condition on blocks of zeros is a condition checkable by reading digits one at a time with a finite memory ("am I inside a block of zeros, and is its length so far odd or even?"). Reading from the right, appending 1 resets nothing that matters, appending 00 preserves every parity, and appending a single 0 after a 1 creates an odd block that can only be repaired by more zeros. That is Step 2, and the equation $\alpha = x\alpha(x^2) + \alpha(x^4)$ is Step 2 written in generating-function language. Frobenius then turns this functional equation into the polynomial equation $\alpha = x\alpha^2 + \alpha^4$. The problem was never "compute α^3"; it was "notice that the digit rule and the equation are the same statement".

One subtlety the official solution passes over: a_0. Written as the single digit 0, the number 0 has a block of one zero and would get $a_0 = 0$. The statement overrides this by writing $\alpha = 1 + a_1 x + \cdots$. The consistent convention is that 0 has the empty expansion, and with it the rule $a_{4m} = a_m$ is valid at $m = 0$ and $a_{2m+1} = a_m$ is valid at $m = 0$. A proof that ignores $n = 0$ and $n = 1$ has a hole exactly there.

Caveat. The background fact that a power series over a finite field of characteristic p is algebraic over $\mathbb{F}_p(x)$ if and only if its coefficient sequence is p-automatic (Christol's theorem, 1979; Christol–Kamae–Mendès France–Rauzy 1980) is quoted as context for the reusable rule, not proved here. The proof of the problem does not use it.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "coefficients in the field of two elements". On seeing it, write down $\beta^2 = \beta(x^2)$ for every power series β over $\mathbb{F}_2$, and underneath it the three exponent maps $n \mapsto 2n, 2n+1, 4n$ as "append 0, append 1, append 00".

2. It suggests converting the required identity into one where every power of α is a power of 2 (here, by multiplying by the unit α), so that each term is a substitution, and then reading off coefficient equality as a statement about the last binary digits of n.

3. Hesitation comes at the cubic: α^3 resists Frobenius and the student reaches for the convolution. Reinforce it away with a single question: "can I change this equation into an equivalent one whose exponents are all $1, 2, 4, \ldots$?" Multiplying by α is the first thing to try, and the constant term 1 in the statement guarantees you can undo it.

4. The trap is coefficient comparison on the cubic itself — zombie maths that is correct at every step and cannot finish, because index addition destroys the digit structure. The minor trap is the edge case $n = 0$: forgetting that the statement defines $a_0 = 1$ and that 0 must be read as the empty expansion.

5. Whenever a sequence is defined by a property of its base-p digits and appears as the coefficients of a series over $\mathbb{F}_p$, expect a functional equation in $\beta(x^p)$, $\beta(x^{p^2}), \ldots$, and expect Frobenius to turn it into a polynomial equation. The cue is the pairing "digits of n" with "characteristic p".

6. Reusable rule: to prove an algebraic identity for a digit-defined series over $\mathbb{F}_p$, first derive the digit recursion (a_{pm+r} in terms of a_m, or a_{p^2m+r} in terms of a_m), write it as a functional equation, apply $\beta(x^p) = \beta^p$, and compare with the target after multiplying the target by whatever unit makes its exponents powers of p. Example: Thue–Morse $t_{2m} = t_m$, $t_{2m+1} = 1 + t_m$ gives $T = (1+x)T^2 + \frac{x}{1+x^2}$, i.e. $(1+x)^3T^2 + (1+x)^2T + x = 0$ over $\mathbb{F}_2$ (machine-checked to 400 terms).

▷ **Cue:** coefficients in $\mathbb{F}_2$ (or $\mathbb{F}_p$), a power of a series in the statement $\longrightarrow$ **Move:** use Frobenius: $\beta^2 = \beta(x^2)$, $\beta^p = \beta(x^p)$; squaring is a substitution, not a convolution

▷ **Cue:** an identity with a power of α that is not a power of p, in characteristic p $\longrightarrow$ **Move:** multiply by a unit (often α itself) so every exponent becomes a power of p, then compare

coefficients

- ▷ **Cue:** a sequence defined by a condition on the binary digits of n $\longrightarrow$ **Move:** split n by its last digits ($2m+1$, $4m$, $4m+2$, …) and find a_n in terms of a_m
- ▷ **Cue:** $x^n \mapsto x^{2n}$, x^{2n+1}, x^{4n} in a generating function $\longrightarrow$ **Move:** read them as "append 0", "append 1", "append 00" to the binary expansion of the exponent
- ▷ **Cue:** a power series with constant term 1, and an equation you want to divide by it $\longrightarrow$ **Move:** it is a unit: $(1+\gamma)^{-1} = \sum \gamma^k$ with $\gamma \in x\,R[[x]]$; cancel freely

Putnam 1989 — Paper B

1989 B1

1989 B1. A dart, thrown at random, hits a square target. Assuming that any two parts of the target of equal area are equally likely to be hit, find the probability that the point hit is nearer to the center than to any edge. Express your answer in the form $\dfrac{a\sqrt{b}+c}{d}$, where a, b, c, d are integers.

Fields: geometric probability, coordinate geometry (parabolas), integration **Position:** B1 – opening problem of the second session **Answer:** $\dfrac{4\sqrt{2}-5}{3} \approx 0.2190$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You put the square at $[-1,1]^2$, centre at the origin, total area 4. "Nearer to the centre than to any edge" means nearer to the centre than to every edge, i.e. nearer than to the closest edge. The distance from (x,y) to the edges $x=\pm1$, $y=\pm1$ is $1-|x|$, $1-|y|$, so the region is

$$R=\Big\{(x,y):\ \sqrt{x^2+y^2}<1-|x| \text{ and } \sqrt{x^2+y^2}<1-|y|\Big\}.$$

Squaring, $\sqrt{x^2+y^2}<1-x$ becomes $y^2<1-2x$: a parabola opening left, with focus at the centre and directrix the edge $x=1$. Four such parabolas, one per edge; R is the curvy "cushion" inside all four. Now you try to integrate over the whole cushion at once, or over a quarter of it (the part with $x,y\ge0$). In the quarter, the boundary is made of two parabolic arcs, $y^2=1-2x$ and $x^2=1-2y$, meeting at a corner you have to locate. Solve simultaneously: subtracting gives $y^2-x^2=2(y-x)$, so $x=y$ or $x+y=2$; the relevant solution is $x=y$ with $x^2+2x-1=0$, $x=\sqrt{2}-1$. Then you set up

$$\int_0^{\sqrt{2}-1}(\text{top}-\text{bottom})\,dx+\int_{\sqrt{2}-1}^{1/2}(\dots)\,dx,$$

with top boundary $y=(1-x^2)/2$ on the first piece and $y=\sqrt{1-2x}$ on the second. It can be done – but the second piece needs $\int\sqrt{1-2x}\,dx$ between surd limits, the first a polynomial, and the two must be split at $\sqrt{2}-1$ without a sign slip. It is a quarter-page of surd arithmetic with three places to go wrong.

The model chain. Use the full symmetry of the square: its eight symmetries (four rotations, four reflections) preserve the centre and permute the edges, so they map R to itself. The lines $y=0$, $x=0$, $y=\pm x$ cut the square into eight congruent triangles, each mapped onto every other by some symmetry, so $P(R)=\text{area}(R\cap T)/\text{area}(T)$ for any one of them. Take

$$T=\{(x,y):0\le y\le x\le1\},\qquad \text{area}(T)=\tfrac{1}{2}.$$

In T the nearest edge is $x = 1$ (since $1 - x \le 1 - y$ and $1 - x \le 1 + y$, $1 + x$), so on T the condition is the single inequality $x^2 + y^2 < (1 - x)^2$, i.e.

$$x < \frac{1 - y^2}{2}.$$

Integrate in horizontal strips, i.e. with respect to x first: for fixed y, the strip of T is $y \le x \le 1$, and R keeps $y \le x < (1 - y^2)/2$. This is non-empty iff $y < (1 - y^2)/2$, i.e. $y^2 + 2y - 1 < 0$, i.e. $0 \le y < \sqrt{2} - 1 =: t$. Hence

$$\text{area}(R \cap T) = \int_0^t \Big(\frac{1 - y^2}{2} - y \Big) dy = \frac{t}{2} - \frac{t^3}{6} - \frac{t^2}{2}.$$

With $t^2 = 3 - 2\sqrt{2}$ and $t^3 = t \cdot t^2 = (\sqrt{2} - 1)(3 - 2\sqrt{2}) = 5\sqrt{2} - 7$,

$$\text{area}(R \cap T) = \frac{3(\sqrt{2} - 1) - (5\sqrt{2} - 7) - 3(3 - 2\sqrt{2})}{6} = \frac{4\sqrt{2} - 5}{6}.$$

Dividing by $\frac{1}{2}$, the probability is $\dfrac{4\sqrt{2} - 5}{3}$. (A Monte Carlo run of 10^6 darts gives 0.2189 against the exact 0.21895.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree on the region: inside all four parabolas $y^2 < 1 - 2x$ and its rotations. They split at the choice of the piece to integrate over. The natural attempt takes a quarter (cut by the axes), where two different parabolas form the boundary and the integral must be split at their intersection. The model takes an eighth (cut also by the diagonals), where only one edge is ever nearest, so only one parabola appears; and it integrates in x first, so the integrand is a polynomial in y with no square root.

> *The square has eight symmetries, not four. If I also cut along the diagonals, then in each triangle I know which edge is nearest, and "nearer than any edge" collapses to "nearer than that one edge" – one parabola, no intersection to split at.*

> *The parabola is $x = (1 - y^2)/2$ when solved for x and $y = \sqrt{1 - 2x}$ when solved for y. I will integrate in whichever variable makes it a polynomial, so the strips run horizontally.*

Layer 3 THE WHY

Why is the diagonal cut natural and not a trick? Unpack the phrase "nearer to the centre than to *any* edge". It is a statement about a minimum: $|p| < \min(d_1, d_2, d_3, d_4)$ where d_i are the distances to the four edges. A minimum of several functions is piecewise, and the pieces are the regions where each function wins. For the four edges of a square, the region where edge $x = 1$ is nearest is the triangle bounded by the two diagonals and that edge. So the diagonals are not decoration; they are the boundaries of the "which edge wins" regions – the lines along which the definition of the minimum changes formula. Any time a condition contains "min" or "max" or "any" over several objects, the first thing to draw is where each object wins. The axes, by contrast, are irrelevant to the minimum: cutting along them leaves each quarter with two competing edges, which is exactly why the quarter-integral needs a split.

This is the same structural fact as the Blueprint's circle: the geometry of the definition names the objects of interest. Here the definition is "distance to the nearest edge", and the nearest-edge regions are cut by the diagonals, therefore you cut by the diagonals.

The second decision – integrate in x first – comes from the other definition in play. The locus of points equidistant from a point (the centre) and a line (the edge) is a parabola with that focus

and directrix. Its equation is quadratic in the coordinate *parallel* to the directrix (y) and linear in the coordinate *perpendicular* to it (x), because the directrix distance is $1 - x$ and only the focus distance brings in y^2. Solving for the linear variable is free; solving for the quadratic one costs a square root. So the parabola is a polynomial graph $x = g(y)$, and horizontal strips are the cheap direction. This is the cheapness principle in its purest form: before writing an integral, ask which variable the boundary is a function *of*, polynomially.

Finally, the answer's shape is a confirmation cue. The problem tells you the answer is $(a\sqrt{b}+c)/d$. A single surd means a single irrational intersection point, and $t = \sqrt{2}-1$ – where the parabola meets the diagonal $y = x$ – is it. If your computation were producing $\sqrt{3}$ or nested roots, you would know you had set up the region wrongly.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "nearer to the centre than to any edge" – a comparison with the nearest of several objects. On seeing it, write down: $|p| < \min_i d_i(p)$, and draw the regions where each d_i is the minimum.

2. It suggests symmetry reduction to one fundamental region – here one of eight triangles – in which the minimum is a single known function, followed by one integral over that region.

3. Hesitation occurs after drawing the four parabolas, when the cushion looks like it needs a piecewise description. What reinforces it away is asking which symmetries the problem has; a square has eight, and using all of them removes the piecewise boundary.

4. The trap is the quarter-square decomposition, which looks natural because axes are what you always draw, and which forces a split integral with a square root in one piece. A secondary trap is misreading "any edge" as "some edge" (nearer than to at least one edge), which gives a much larger region.

5. Whenever a condition involves the nearest (or farthest) of several symmetric objects, the fundamental domain of the symmetry group is exactly the region where one object is nearest. Cut by the lines that decide "who wins", not by the coordinate axes.

6. The reusable rule: reduce by the full symmetry group before integrating; then, for each boundary curve, integrate first in the variable in which the curve is polynomial.

▷ **Cue:** a condition compares distance to a point with distance to the nearest of several lines or objects $\longrightarrow$ **Move:** draw the regions where each object is nearest; restrict to one of them, where the condition is a single inequality

▷ **Cue:** a region or probability on a square, regular polygon, or cube $\longrightarrow$ **Move:** use the full symmetry group (eight for a square), not only reflections in the axes, and compute on one fundamental piece

▷ **Cue:** a locus is equidistant from a point and a line $\longrightarrow$ **Move:** it is a parabola with that focus and directrix; it is polynomial in the coordinate parallel to the directrix, so integrate that way

▷ **Cue:** the answer is promised in the form $(a\sqrt{b}+c)/d$ $\longrightarrow$ **Move:** expect exactly one irrational intersection point; locate it first and use it as a check

1989 B2

1989 B2. Let S be a non-empty set with an associative operation that is left and right cancellative ($xy = xz$ implies $y = z$, and $yx = zx$ implies $y = z$). Assume that for every a in S the set $\{a^n : n = 1, 2, 3, \ldots\}$ is finite. Must S be a group?

Fields: abstract algebra (semigroups, groups), pigeonhole principle **Position:** B2 – early second-session problem **Answer:** Yes: S is necessarily a group.

Layer 1 THE WHAT — surface analysis

The natural attempt. A group needs three things beyond associativity: an identity, inverses, and closure (closure is given). You know that in a finite group every element has finite order,

$a^k = e$. So you try to find e first: "the identity should be some power of a". You write a^n for a moment, then hit the problem that you do not have an e to write $a^k = e$ against. You try defining e as "the element with $ea = a$" and then need to show it exists, is the same for every a, and works on the other side. Often the attempt drifts into looking for a counterexample instead – the natural numbers under addition are cancellative and not a group, so "perhaps the answer is no" – until you notice that $\{n \cdot 1\}$ is infinite there, so the finiteness hypothesis kills it. Now you have neither a proof nor a counterexample, and half the time is gone.

The model chain. The answer is yes.

Step 1: a repeat. Fix $a \in S$ (possible since $S \neq \emptyset$). The map $n \mapsto a^n$ sends the infinite set $\{1, 2, 3, \dots\}$ into a finite set, so by pigeonhole there are integers $1 \leq m < n$ with $a^m = a^n$. Put $k = n - m \geq 1$, so $a^m = a^{m+k} = a^m a^k = a^k a^m$ (associativity lets us regroup powers freely).

Step 2: $e = a^k$ is a two-sided identity for all of S. Let $x \in S$ be arbitrary. Then

$$a^m x = a^{m+k} x = a^m \, (a^k x),$$

and left cancellation by the element a^m gives $x = a^k x$. Likewise

$$x a^m = x a^{m+k} = (x a^k) \, a^m,$$

and right cancellation by a^m gives $x = x a^k$. So $e := a^k$ satisfies $ex = xe = x$ for every $x \in S$.

Step 3: e does not depend on a. Any two two-sided identities e, e' satisfy $e = ee' = e'$. So every element's construction produces the same e.

Step 4: inverses. Let $b \in S$. Repeating Steps 1–2 for b gives $j \geq 1$ with b^j an identity, hence $b^j = e$ by Step 3. The element $c = b^{2j-1}$ exists because $2j - 1 \geq 1$, and

$$bc = cb = b^{2j} = b^j b^j = ee = e.$$

(Using b^{2j-1} rather than b^{j-1} avoids the case $j = 1$, where b^0 is not an element we have.)

Associativity is given, closure is given, and we have an identity and inverses: S is a group.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths agree that the identity should be a power of a and that finiteness of $\{a^n\}$ is the only lever. They split at the line $a^m = a^{m+k}$. The natural attempt wants to write $a^k = e$ and cannot, because it has no e yet. The model reads $a^m = a^m \cdot a^k$ as "multiplying by a^k did nothing to a^m", multiplies both sides by an arbitrary x, and cancels a^m to get "multiplying by a^k does nothing to x".

> *I cannot write $a^k = e$ because there is no e. But $a^m = a^m a^k$ already says a^k acts like an identity on a^m. If I put an arbitrary x next to a^m, cancellation should let me strip a^m off and keep the "acts like an identity" part on x.*

> *Pigeonhole gives me a repeat $a^m = a^n$, not $a^k = e$. I should use the repeat as it is, and let cancellation turn it into the identity.*

Layer 3 THE WHY

Why is this the natural reading of $a^m = a^{m+k}$? Start from what an identity *is*: an element e with $ex = x = xe$ for all x. That is an equation of the shape "something times x equals x". The only equation the hypotheses hand you is $a^m = a^m a^k$, which already has that shape with $x = a^m$. The defect is that it holds for one particular x. So the question is precisely: how do I transport "a^k fixes a^m" to "a^k fixes everything"? Cancellation is the only hypothesis left unused, and cancellation is literally a tool for removing a common factor. Put x on the far side, so that

a^m becomes a common left factor of $a^m x$ and $a^m(a^k x)$, and cancel it. The step is not inspired; it is the one move available that uses the remaining hypothesis on the remaining gap.

It helps to see why the finite case is familiar. The textbook proof that a finite cancellative semigroup is a group says "left multiplication by a is injective, hence a bijection of the finite set S". That argument needs S finite, and here S may be infinite – only each cyclic sub-semigroup $\{a^n\}$ is finite. The model chain respects exactly that: all pigeonholing happens inside $\{a^n\}$, and cancellation carries the conclusion out to arbitrary x. When a hypothesis is local (each element separately) and the conclusion is global (a single identity for all), look for the step that turns local into global; here it is the cancellation of a^m against arbitrary x, followed by uniqueness of identities.

Uniqueness deserves a word, because it is where a careless proof leaks. Step 2 constructs an identity from one a. Nothing yet says b's construction gives the same element; if it did not, "$b^j = e$" would be meaningless. The two-line argument $e = ee' = e'$ closes this, and it is the reason Step 3 is written out.

Finally, the counterexample instinct, which is worth running to confirm that no hypothesis is decorative. Drop finiteness: $(\mathbb{Z}_{>0}, +)$ is associative and cancellative with no identity. Drop left cancellation: on $\{u, v\}$ define $xy = x$. This is associative, since $(xy)z = x = x(yz)$; every set of powers is $\{a\}$, finite; right cancellation holds, since $yx = zx$ reads $y = z$; but $uv = u = uu$ with $v \neq u$, so left cancellation fails, and the structure is not a group (there is no e with $eu = u$ and $ev = v$ simultaneously, as $ex = e$). The mirror rule $xy = y$ shows right cancellation is needed too. Each hypothesis kills a counterexample, so each must appear in the proof; in the model chain each appears exactly once – finiteness in Step 1, left cancellation for $ex = x$, right cancellation for $xe = x$.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is "the set $\{a^n\}$ is finite". On seeing it, write down: by pigeonhole, $a^m = a^{m+k}$ for some $m \geq 1, k \geq 1$.

2. It suggests: treat a^k as a candidate identity, and use cancellation against an arbitrary x to promote "a^k fixes a^m" to "a^k fixes everything".

3. Hesitation occurs at the moment you want to write $a^k = e$ and there is no e. It is reinforced away by rereading the definition of identity as the equation $ex = x$, and noticing that $a^m a^k = a^m$ is already that equation for one x.

4. Two traps. First, drifting toward "no" because $(\mathbb{N}, +)$ is a cancellative non-group – which forgets that its powers are infinite. Second, importing the finite-group argument ("injective on a finite set is bijective"), which needs S finite; S here may be infinite.

5. Whenever a structure axiom (identity, inverse, fixed point) must be produced from nothing, look for an equation the hypotheses give that already has the axiom's shape for one special element, then find the hypothesis that transports it to all elements.

6. The reusable rule: from a repeat $a^m = a^{m+k}$ in a cancellative setting, a^k is an identity; then inverses are high powers such as a^{2k-1}. More generally, cancellation turns "holds for one element" into "holds for all" by putting an arbitrary element next to the special one and cancelling.

▷ **Cue:** a set of powers $\{a^n\}$, or any orbit under iteration, is finite $\longrightarrow$ **Move:** pigeonhole a repeat $a^m = a^{m+k}$ and work with that equation directly

▷ **Cue:** an identity must be proved to exist and there is no candidate $\longrightarrow$ **Move:** find an equation of the form $u \cdot w = u$ for some u; then w is the candidate, and cancellation extends it

▷ **Cue:** cancellation is a hypothesis and an equation holds for one special element $\longrightarrow$ **Move:** multiply by an arbitrary x so that the special element is a common factor, then cancel it

▷ **Cue:** an inverse is needed as a power but the natural exponent could be 0 $\longrightarrow$ **Move:** shift by a full period: use a^{2k-1} instead of a^{k-1}

▷ **Cue:** a construction depends on a chosen element but should be canonical ⟶ **Move:** prove uniqueness of the constructed object (e.g. $e = ee' = e'$) before using it for other elements

1989 B3

1989 B3. Let f be a function on $[0, \infty)$, differentiable and satisfying

$$f'(x) = -3f(x) + 6f(2x)$$

for $x > 0$. Assume that $|f(x)| \leq e^{-\sqrt{x}}$ for $x \geq 0$ (so that $f(x)$ tends rapidly to 0 as x increases). For n a non-negative integer, define

$$\mu_n = \int_0^\infty x^n f(x)\, dx$$

(sometimes called the nth moment of f).

a) Express μ_n in terms of μ_0.

b) Prove that the sequence $\{\mu_n \frac{3^n}{n!}\}$ always converges, and that the limit is 0 only if $\mu_0 = 0$.

Fields: real analysis (improper integrals, integration by parts), functional-differential equations, infinite products **Position:** B3 – middle of the second session **Answer:** (a) $\mu_n = \dfrac{n!\,\mu_0}{3^n \prod_{k=1}^{n}(1-2^{-k})}$; (b) $\mu_n 3^n/n! \to \mu_0/\prod_{k=1}^{\infty}(1-2^{-k})$, and the product is positive (about 0.2888), so the limit is 0 only if $\mu_0 = 0$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You see a differential equation and try to solve it. $f' = -3f$ alone gives e^{-3x}, so you try $f = \sum_k a_k e^{-c_k x}$. The term $6f(2x)$ turns e^{-cx} into e^{-2cx}, so the exponents must be closed under doubling: $c_k = 3 \cdot 2^k$. Matching coefficients of $e^{-3\cdot 2^k x}$ gives $-3 \cdot 2^k a_k = -3a_k + 6a_{k-1}$, so $a_k = -2a_{k-1}/(2^k - 1)$. That is a genuine family of solutions, and you can even compute its moments. But the problem did not say f is of this form. It says "let f be a function … satisfying"; you would have to prove every such f is in your family – a uniqueness theorem for a delay-type equation with an advanced argument $2x$, which is not a Putnam B3 task. The wall is that solving the equation answers a question nobody asked. You wanted μ_n, not f.

The model chain. Throughout, $n \geq 0$ is an integer.

Everything converges. Since $|f(x)| \leq e^{-\sqrt{x}}$ and, substituting $x = u^2$, $\int_0^\infty x^n e^{-\sqrt{x}} dx = 2\int_0^\infty u^{2n+1} e^{-u} du = 2(2n+1)! < \infty$, each μ_n converges absolutely. By the equation, $|f'(x)| \leq 3e^{-\sqrt{x}} + 6e^{-\sqrt{2x}} \leq 9e^{-\sqrt{x}}$ for $x > 0$, so $\int_0^\infty x^n f'(x)\, dx$ converges absolutely too. Also f is continuous on $[0, \infty)$ (it is differentiable), so f' is continuous on $(0, \infty)$ by the equation, and integration by parts is legitimate on every $[\varepsilon, R] \subset (0, \infty)$.

The moment identity. Multiply the equation by x^n and integrate over $(0, \infty)$. On the right,

$$\int_0^\infty x^n f(2x)\, dx = \frac{1}{2^{n+1}} \int_0^\infty u^n f(u)\, du = \frac{\mu_n}{2^{n+1}} \qquad (u = 2x),$$

so the right side is $-3\mu_n + 6\mu_n/2^{n+1} = -3(1 - 2^{-n})\mu_n$. On the left, for $n \geq 1$,

$$\int_\varepsilon^R x^n f'(x)\, dx = \left[x^n f(x)\right]_\varepsilon^R - n\int_\varepsilon^R x^{n-1} f(x)\, dx.$$

As $R \to \infty$, $R^n f(R) \to 0$ because $|R^n f(R)| \le R^n e^{-\sqrt{R}} \to 0$; as $\varepsilon \to 0^+$, $\varepsilon^n f(\varepsilon) \to 0 \cdot f(0) = 0$ because $n \ge 1$ and f is continuous at 0. So the left side is $-n\mu_{n-1}$, and

$$n\,\mu_{n-1} = 3\big(1 - 2^{-n}\big)\mu_n \qquad (n \ge 1). \tag{$*$}$$

(For $n = 0$ the same computation reads $-f(0) = 0$: every admissible f has $f(0) = 0$. This is not needed below.)

Part (a). For $n \ge 1$, $1 - 2^{-n} \ne 0$, so $(*)$ gives $\mu_n = \dfrac{n}{3(1-2^{-n})}\mu_{n-1}$, and telescoping,

$$\mu_n = \frac{n!}{3^n \prod_{k=1}^{n}(1-2^{-k})}\,\mu_0.$$

Part (b). Hence $\mu_n \dfrac{3^n}{n!} = \dfrac{\mu_0}{P_n}$ with $P_n = \prod_{k=1}^{n}(1-2^{-k})$. The P_n are positive and decreasing (each factor lies in $[\frac{1}{2}, 1)$). For a lower bound, $\ln(1-t) \ge -2t$ on $[0, \frac{1}{2}]$ (the left side is concave, the right side linear, and they agree at $t = 0$ and satisfy $\ln\frac{1}{2} \approx -0.693 \ge -1$ at $t = \frac{1}{2}$), so

$$\ln P_n = \sum_{k=1}^{n} \ln(1-2^{-k}) \ge -2\sum_{k=1}^{n} 2^{-k} > -2.$$

A decreasing sequence bounded below converges, so $P_n \to P \ge e^{-2} > 0$. Therefore $\mu_n 3^n/n! \to \mu_0/P$, which is 0 if and only if $\mu_0 = 0$.

Checks. The hypotheses are satisfiable by a non-zero f: the series family from the natural attempt, $f = a_0 \sum_{k\ge 0} \frac{(-2)^k}{\prod_{j=1}^{k}(2^j-1)} e^{-3\cdot 2^k x}$ with a_0 small, is bounded by a constant times e^{-3x} and satisfies the equation; its value at 0 is 0 by Euler's identity $\sum_k (-1)^k q^{k(k-1)/2} z^k/(q;q)_k = (z;q)_\infty$ at $q = \frac{1}{2}$, $z = 1$, consistent with the $n = 0$ remark. Numerically, for this f the ratio $\mu_n 3^n/n!$ agrees with μ_0/P_n to 15 digits for $n \le 11$, and $P \approx 0.288788$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both paths start from the equation. They split at the first line. The natural attempt asks "what is f?" and starts building solutions. The model asks "what is μ_n?" and applies the operation that produces μ_n – multiply by x^n and integrate – to both sides of the equation, at once, without knowing f.

> *They want the moments, not f. So I apply "multiply by x^n and integrate" to the whole equation. The f' becomes μ_{n-1} by parts, the $f(2x)$ becomes $\mu_n/2^{n+1}$ by substitution, and I get a recursion in the moments.*

> *The recursion has n on one side and 3 on the other, so μ_n grows like $n!/3^n$ times a correction. Dividing by $n!/3^n$ leaves a product $\prod(1-2^{-k})$, and a product whose factors are 1 minus a summable sequence converges to something non-zero.*

Layer 3 THE WHY

Why is "integrate against x^n" natural rather than a trick? Look at the three operations in the equation – differentiation, scaling the argument by 2, and multiplication by constants – and ask what each does to a moment. Scaling: $\int x^n f(2x)dx = 2^{-(n+1)}\mu_n$, a pure multiple. Differentiation: by parts, $\int x^n f' = -n\mu_{n-1}$, a shift in the index. Constants: nothing. So moments turn every operation in this equation into arithmetic on the sequence (μ_n). That is the defining property of a good transform: it diagonalises the operators you are faced with. The monomial x^n is the natural test function precisely because it is homogeneous – $(2x)^n = 2^n x^n$ – and the

argument-doubling in $f(2x)$ is a homogeneity statement. When the equation's operators and the requested quantity are both "compatible with x^n", integrating against x^n is not an idea; it is reading the problem.

Standard write-ups of moment problems begin "multiply by x^n and integrate by parts". Here is where that comes from: the question defines μ_n explicitly and asks for a relation among the μ_n. Any relation among integrals of f must come from the only fact you know about f – the equation – integrated against something. The problem even tells you what to integrate against.

The rigour sits in the boundary terms, and they are where the hypothesis $|f| \le e^{-\sqrt{x}}$ earns its place. At infinity, $x^n f(x) \to 0$ for every n because $e^{-\sqrt{x}}$ beats every polynomial; this is also why all μ_n exist. The bound is much weaker than exponential decay, but strong enough for every moment, so the solver should notice that each use of the bound is a "polynomial times $e^{-\sqrt{x}}$" estimate and nothing more. At 0, the boundary term vanishes because of the factor x^n with $n \ge 1$, not because of anything about $f(0)$. For $n = 0$ there is no such factor, and the identity instead forces $f(0) = 0$ – a genuine consequence worth noticing, not an inconsistency.

Part (b) is a statement about the product $\prod(1 - 2^{-k})$. Why must it be non-zero? Take logarithms: an infinite product of numbers $1 - t_k$ with $t_k \ge 0$ is a sum $\sum \ln(1 - t_k)$, and $\ln(1 - t) \approx -t$ for small t. So the product converges to a positive number exactly when $\sum t_k < \infty$. Here $t_k = 2^{-k}$, geometric, summable. The inequality $\ln(1 - t) \ge -2t$ on $[0, \frac{1}{2}]$ is the rigorous form of "$\approx -t$". The "only if $\mu_0 = 0$" clause is then the statement that the limit is μ_0 divided by a non-zero constant.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the question asks about $\mu_n = \int x^n f$, not about f. On seeing it, write down: "multiply the equation by x^n and integrate from 0 to ∞".

2. It suggests the moment method: integration by parts for f', the substitution $u = 2x$ for $f(2x)$, giving a first-order recursion $n\mu_{n-1} = 3(1 - 2^{-n})\mu_n$.

3. Hesitation occurs at the boundary terms $[x^n f]_0^\infty$ – are they zero? It is reinforced away by checking each end separately: at ∞ use the decay bound against a polynomial, at 0 use the factor x^n with $n \ge 1$.

4. The trap is solving the differential equation. It produces a correct family of solutions but no proof that every f is in it, and it spends time the moment method does not need. A secondary trap is asserting the infinite product is positive without proof; a product of factors less than 1 can tend to 0.

5. Whenever the requested quantity is a transform of the unknown function (moments, Laplace transform, generating function), apply that transform to the whole equation before trying to solve it. Check that the transform turns each operator in the equation into something algebraic.

6. The reusable rule: an integral transform whose kernel is homogeneous (x^n, e^{-sx}) converts differentiation into index or parameter shifts and argument scaling into multiplication; use it to get a recursion, telescope, and control the resulting infinite product by summing logarithms.

▷ **Cue:** the question asks for moments $\int x^n f$ of a function defined by an equation $\longrightarrow$ **Move:** multiply the equation by x^n and integrate over the whole domain

▷ **Cue:** a term $f(cx)$ appears inside an integral $\int x^n f(cx)\,dx$ $\longrightarrow$ **Move:** substitute $u = cx$; it becomes $c^{-(n+1)}\mu_n$

▷ **Cue:** a term f' appears inside $\int x^n f'\,dx$ $\longrightarrow$ **Move:** integrate by parts to $-n\mu_{n-1}$ and check both boundary terms separately

▷ **Cue:** a sequence is normalised by $n!/c^n$ and the limit is asked $\longrightarrow$ **Move:** expect a telescoped product; the normalisation is chosen to cancel the leading growth

▷ **Cue:** an infinite product $\prod(1 - t_k)$ with $t_k \ge 0$ must be shown non-zero $\longrightarrow$ **Move:** take logs and use $\ln(1 - t) \ge -2t$ for $t \le \frac{1}{2}$; positivity follows from $\sum t_k < \infty$

1989 B4

1989 B4. Can a countably infinite set have an uncountable collection of non-empty subsets such that the intersection of any two of them is finite?

Fields: set theory; combinatorics; cardinality **Position:** B4 — middle of the paper; one idea, and the idea is a change of what the countable set *is* **Answer:** Yes. Take the countable set to be the set of finite 0–1 strings; to each infinite 0–1 sequence attach the set of its finite initial segments. There are uncountably many such sets, and two distinct ones share only the prefixes before the first place the sequences differ, a finite set.

Layer 1 THE WHAT — surface analysis

The natural attempt. The countable set feels like $\mathbb{N}$, so the strong student works in $\mathbb{N}$. First reaction: can you even get infinitely many pairwise almost-disjoint sets? Disjoint ones, yes — the sets $\{2^k m : m \text{ odd}\}$ for $k = 0, 1, 2, \dots$ — but a family of pairwise *disjoint* non-empty subsets of $\mathbb{N}$ is at most countable, since each set contains a different least element. So the finite overlaps must be used.

Next try: arithmetic progressions $\{a + nd\}$. Two progressions with coprime differences meet in infinitely many points (Chinese remainder theorem), so these fail at once. Powers: $P_k = \{k^j : j \geq 1\}$ for k not a perfect power; these are pairwise disjoint beyond 1, but they are indexed by $k \in \mathbb{N}$, so there are only countably many. Every family built from a formula with an integer parameter is countable, because the parameter is.

Here the student usually decides the answer is "no" and looks for a diagonal argument: list the sets $S_1, S_2, \dots$ and build a set almost disjoint from all of them? That proves only that a *countable* almost-disjoint family can be enlarged, which is the opposite of an upper bound. The wall: every attempt to index sets by integers produces countably many, and every attempt to prove "no" produces nothing, because the answer is yes.

The model chain. Since the property "the intersection of any two is finite" is preserved by any bijection, we may choose the countably infinite set to be any countably infinite set we like. Take

$$T = \{\text{finite strings of 0s and 1s}\} \quad (\text{including the empty string}).$$

T is countably infinite: there are 2^ℓ strings of length ℓ, and listing length 0, then length 1, then length 2, … enumerates T.

For each infinite sequence $\omega = (\omega_1, \omega_2, \dots) \in \{0,1\}^{\mathbb{N}}$ let

$$S_\omega = \{\text{empty string},\ \omega_1,\ \omega_1\omega_2,\ \omega_1\omega_2\omega_3,\ \dots\} \subseteq T,$$

the set of finite initial segments of ω. Each S_ω is non-empty (indeed infinite).

Finite intersections. Let $\omega \neq \eta$ and let r be the first index with $\omega_r \neq \eta_r$. A string of length $\ell \geq r$ in S_ω has rth letter ω_r, and one in S_η has rth letter η_r, so no string of length $\geq r$ lies in both. Hence $S_\omega \cap S_\eta$ consists of the r strings of lengths $0, 1, \dots, r-1$ — a finite set.

Uncountably many sets. The map $\omega \mapsto S_\omega$ is injective: ω is recovered from S_ω because the unique string of length ℓ in S_ω is $\omega_1 \cdots \omega_\ell$. And $\{0,1\}^{\mathbb{N}}$ is uncountable by Cantor's diagonal argument: given any list $\omega^{(1)}, \omega^{(2)}, \dots$, the sequence whose kth term is $1 - \omega_k^{(k)}$ differs from the kth listed sequence in place k, so is not on the list. So $\{S_\omega\}$ is an uncountable collection of non-empty subsets of T, any two meeting in a finite set.

Transport to any countably infinite set X through a bijection $T \to X$. The answer is **yes**. □

Same proof, second dress. Take $X = \mathbb{Q}$. For each irrational r, let S_r be the set of decimal truncations $\lfloor 10^k r \rfloor / 10^k$, $k \geq 0$; it is infinite (its elements converge to r without reaching it). For $r \neq s$, pick disjoint open intervals $U \ni r$, $V \ni s$. All but finitely many elements of S_r lie in U and

all but finitely many of S_s lie in V; a common element lies outside U or outside V, so there are finitely many. Distinct r give distinct sets since r is the limit. The irrationals are uncountable.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree that disjoint families are countable and that the finite overlaps must be exploited. They split at the choice of ground set. (N) keeps "countably infinite set" $= \mathbb{N}$ and looks for sets defined by arithmetic formulas with a parameter. (M) says: I am free to relabel, so I choose a countable set whose elements are *finite approximations of uncountably many objects*, and index my sets by those objects.

> *Anything I index by an integer is countable, so the index of my sets has to be something uncountable — a real number, or an infinite* 0–1 *sequence. What countable data does a real number carry? Its finite approximations. So let the countable set be the approximations, and let each real own the set of its own approximations.*

> *Two different infinite sequences agree up to some point and then disagree forever after. So their sets of prefixes share an initial chunk and nothing more. "Finite intersection" is exactly "eventually different".*

- ▷ **Cue:** a question about a countable set where only the cardinality matters $\longrightarrow$ **Move:** choose the most convenient countable set — finite strings, $\mathbb{Q}$, lattice points — and transport by a bijection at the end
- ▷ **Cue:** you need uncountably many objects built inside a countable set $\longrightarrow$ **Move:** index them by reals or infinite 0–1 sequences, and represent each by the countable set of its finite approximations
- ▷ **Cue:** two sets must overlap only finitely $\longrightarrow$ **Move:** make them branches of a tree, or sequences converging to different limits: they share a start and then separate

Layer 3 THE WHY

Why is the relabelling the natural move and not a trick? Because the problem's phrasing tells you only cardinality matters: "a countably infinite set". The only structure a countably infinite set has is its size; any arithmetic structure you bring by calling it $\mathbb{N}$ is baggage you chose. The natural attempt suffers from that baggage: in $\mathbb{N}$, sets you can write down come from formulas, formulas have integer parameters, and integer parameters produce countable families. The diagnosis "all my families are countable because my index set is countable" is the thought that forces a new index set.

Why finite approximations? Go to the definition of uncountable. The standard uncountable sets — $\mathbb{R}$, $\{0,1\}^{\mathbb{N}}$ — are exactly the sets of *infinite* objects each determined by *countably many finite pieces of data*. That is why Cantor's diagonal argument works: an infinite sequence is a limit of finite strings, and there are only countably many strings but uncountably many limits. So the uncountable family you need is sitting inside the proof that uncountable sets exist. The countable set is the set of finite pieces; each infinite object claims the pieces that approximate it.

Why do the intersections come out finite? Unpack "the intersection of any two is finite" for two infinite sets: they must eventually go their separate ways. For approximations to a limit, "eventually separate" is what having different limits means. For prefixes of sequences it is the first disagreement — after the first place two sequences differ, no later prefix of one can be a prefix of the other. So the finiteness condition is the definition of "distinct limit", read in the language of sets. The circle example of the Blueprint has the same shape: the centre is the object of interest because the definition of a circle mentions it; here the limit is the object of interest because the definition of an infinite sequence (a coherent chain of finite prefixes) mentions it.

This is the standard construction of a large *almost disjoint family*, and the tree picture is worth keeping: a countable binary tree has uncountably many branches, and two branches share only the nodes above their split point.

Layer $+\alpha$ BACKTESTING

1. The first cue is "countably infinite set" with no further structure, set against "uncountable collection". On seeing it write: *I may choose the countable set; my index must be uncountable.*
2. Build sets from finite approximations of an uncountable family of limits: prefixes of infinite binary sequences, or rationals converging to irrationals.
3. Hesitation comes after several countable families have failed, when the student starts to believe the answer is "no". Reinforce it away by asking what the index set of the attempted family was: if it was $\mathbb{N}$, the failure says nothing about the problem.
4. The trap is identifying the countable set with $\mathbb{N}$ and searching among arithmetic families (progressions, powers), all of which are countable or overlap infinitely. A secondary trap: a diagonal argument that enlarges countable families feels like progress towards "no" but proves nothing either way.
5. When a problem needs uncountably many objects inside a countable world, look for a tree: countably many nodes, uncountably many branches. Branches that share only finitely many nodes are the default model of "almost disjoint".
6. Reusable decision rule: *if only the size of a set matters, relabel it as whatever set makes the construction visible; and to manufacture uncountably many things from countable material, let each thing be a limit and let its countable material be its approximations.*

1989 B5

1989 B5. Label the vertices of a trapezoid T (quadrilateral with two parallel sides) inscribed in the unit circle as A, B, C, D so that AB is parallel to CD and A, B, C, D are in counterclockwise order. Let s_1, s_2, and d denote the lengths of the line segments AB, CD, and OE, where E is the point of intersection of the diagonals of T, and O is the center of the circle. Determine the least upper bound of $\frac{s_1-s_2}{d}$ over all such T for which $d \neq 0$, and describe all cases, if any, in which it is attained.

Fields: Euclidean geometry; trigonometry **Position:** B5 — late in the paper, but a coordinate computation once the right coordinate is chosen **Answer:** The least upper bound is 2. In fact $\frac{s_1-s_2}{d} = \pm 2\sin\theta$, where θ is the angle between the diagonals; the value 2 is attained exactly when the diagonals are perpendicular, T is not a square, and AB is the longer parallel side.

Layer 1 THE WHAT — surface analysis

The natural attempt. An inscribed trapezoid is isosceles, so put the axis of symmetry on the y-axis. The parallel sides are horizontal chords at heights h and k, with half-lengths $\sqrt{1-h^2}$ and $\sqrt{1-k^2}$. Take $A = (-\sqrt{1-h^2}, h)$ and $C = (\sqrt{1-k^2}, k)$; the diagonal AC crosses the axis at

$$y_E = \frac{h\sqrt{1-k^2} + k\sqrt{1-h^2}}{\sqrt{1-h^2} + \sqrt{1-k^2}}.$$

Then

$$\frac{s_1 - s_2}{d} = \frac{2\left(\sqrt{1-h^2} - \sqrt{1-k^2}\right)\left(\sqrt{1-h^2} + \sqrt{1-k^2}\right)}{\left|h\sqrt{1-k^2} + k\sqrt{1-h^2}\right|} = \frac{2(k^2-h^2)}{\left|h\sqrt{1-k^2} + k\sqrt{1-h^2}\right|}.$$

This is correct, and the difference-of-squares cancellation is a good sign. Now the student maximises a function of two variables on the square $(-1,1)^2$: partial derivatives in h and k,

each with square roots in the denominator, a boundary where the denominator vanishes (where $d = 0$), and a sign to track. The critical-point equations are solvable but ugly, and at the end the student still has to explain *which trapezoids* attain the bound. Many stop at a numerical guess near 2.

The model chain. **Symmetry.** The perpendicular bisector of a chord passes through O. AB and CD are parallel, so they have the same perpendicular bisector ℓ through O, and reflection in ℓ maps each chord to itself, swapping its endpoints: $A \leftrightarrow B$, $C \leftrightarrow D$. It therefore swaps the diagonals AC and BD, so it fixes their intersection E: $E \in \ell$.

Angle coordinates. Take ℓ as the y-axis. Every point of the unit circle is $(\cos, \sin)$ of an angle; measuring from the top, write the endpoints of AB as $(\pm\sin\alpha, \cos\alpha)$ and those of CD as $(\pm\sin\beta, \cos\beta)$, with $\alpha, \beta \in (0, \pi)$ and $\alpha \neq \beta$. So $s_1 = 2\sin\alpha$, $s_2 = 2\sin\beta$. The diagonal AC crosses the axis (if A and C were on the same side of ℓ, then AC and its mirror image BD would lie in opposite open half-planes and could not meet), so, after reflecting the whole figure in ℓ if necessary (which changes none of s_1, s_2, d), $A = (-\sin\alpha, \cos\alpha)$ and $C = (\sin\beta, \cos\beta)$.

Locate E. The point $A + t(C - A)$ has x-coordinate 0 at $t = \frac{\sin\alpha}{\sin\alpha + \sin\beta}$, where

$$y_E = \cos\alpha + \frac{\sin\alpha(\cos\beta - \cos\alpha)}{\sin\alpha + \sin\beta} = \frac{\cos\alpha\sin\beta + \sin\alpha\cos\beta}{\sin\alpha + \sin\beta} = \frac{\sin(\alpha+\beta)}{\sin\alpha + \sin\beta}.$$

So $d = \dfrac{|\sin(\alpha+\beta)|}{\sin\alpha + \sin\beta}$, and $d \neq 0$ means $\alpha + \beta \neq \pi$.

The ratio. Using $\sin^2\alpha - \sin^2\beta = \sin(\alpha+\beta)\sin(\alpha-\beta)$ (expand both sine products; the cross terms cancel and $\sin^2\alpha\cos^2\beta - \cos^2\alpha\sin^2\beta = \sin^2\alpha - \sin^2\beta$),

$$\frac{s_1 - s_2}{d} = \frac{2(\sin\alpha - \sin\beta)(\sin\alpha + \sin\beta)}{|\sin(\alpha+\beta)|} = \frac{2\sin(\alpha+\beta)\sin(\alpha-\beta)}{|\sin(\alpha+\beta)|} = \pm 2\sin(\alpha-\beta),$$

with the sign of $\sin(\alpha+\beta)$. Hence $\frac{s_1 - s_2}{d} \leq 2$ for every admissible T.

Equality. We need $|\sin(\alpha-\beta)| = 1$, i.e. $|\alpha - \beta| = \frac{\pi}{2}$, together with a positive ratio, i.e. $s_1 > s_2$. Now compute the angle between the diagonals:

$$\overrightarrow{AC} \cdot \overrightarrow{BD} = -(\sin\alpha + \sin\beta)^2 + (\cos\beta - \cos\alpha)^2 = 2\cos(\alpha-\beta)\big(\cos(\alpha+\beta) - 1\big),$$

and $\cos(\alpha+\beta) \neq 1$ since $0 < \alpha + \beta < 2\pi$. So the diagonals are perpendicular iff $|\alpha - \beta| = \frac{\pi}{2}$. (More generally $|\sin(\alpha-\beta)| = \sin\theta$ for the angle θ between the diagonals, by the inscribed-angle theorem: that angle is half the sum of the two equal arcs of size $|\alpha - \beta|$ cut off between the diagonals.)

Conversely, every trapezoid with perpendicular diagonals and $d \neq 0$ attains 2 once labelled with AB longer. Such a trapezoid has $s_1 \neq s_2$: $\sin\alpha = \sin\beta$ with $\alpha \neq \beta$ forces $\alpha + \beta = \pi$, i.e. $d = 0$. Example: $\alpha = \frac{2\pi}{3}$, $\beta = \frac{\pi}{6}$ gives $s_1 = \sqrt{3}$, $s_2 = 1$, $d = \frac{\sqrt{3}-1}{2}$, ratio 2.

Answer. The least upper bound is 2. It is attained precisely by the trapezoids whose diagonals are perpendicular, excluding the square (the one case with $d = 0$), labelled so that AB is the longer of the two parallel sides. □

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains use the symmetry axis and put E on it. They split at the choice of coordinate for a point on the circle. (N) uses the height h and pays with $\sqrt{1-h^2}$ everywhere. (M) uses the angle, $(\sin\alpha, \cos\alpha)$, and the same formula for y_E becomes $\sin(\alpha+\beta)/(\sin\alpha + \sin\beta)$ — an addition formula in plain sight.

> *Every point here lives on the unit circle. Its honest coordinate is an angle, not a height: if I use the height I will be carrying $\sqrt{1-h^2}$ around for the whole problem. Let me write the vertices as $(\pm\sin\alpha, \cos\alpha)$.*

> *The numerator is $\sin^2\alpha - \sin^2\beta$ and the denominator is $\sin(\alpha+\beta)$. A difference of squares of sines factors as $\sin(\alpha+\beta)\sin(\alpha-\beta)$ — the denominator is sitting inside the numerator. The ratio is $2\sin(\alpha-\beta)$, and a sine is at most 1.*

> *The problem asks me to describe when the bound is attained, so $\alpha-\beta = \pm\frac{\pi}{2}$ must mean something geometric. What in this figure is an angle equal to $\alpha-\beta$? The angle between the diagonals.*

▷ **Cue:** points on a unit circle, a quantity to optimise ⟶ **Move:** parametrise by angles, $(\cos t, \sin t)$, never by one Cartesian coordinate with a square root

▷ **Cue:** $\sin^2\alpha - \sin^2\beta$ (or $\cos^2\beta - \cos^2\alpha$) next to $\sin(\alpha\pm\beta)$ ⟶ **Move:** factor as $\sin(\alpha+\beta)\sin(\alpha-\beta)$ and cancel

▷ **Cue:** a figure inscribed in a circle with a pair of parallel chords ⟶ **Move:** use the common perpendicular bisector through the centre as an axis of symmetry; it is an isosceles trapezoid

▷ **Cue:** a bound of the form $c\,|\sin(\text{something})|$ and a request to describe equality ⟶ **Move:** translate "something $= \frac{\pi}{2}$" into a perpendicularity in the figure

Layer 3 THE WHY

Why is the angle the right coordinate? Go back to the definition of the circle: the points at distance 1 from O. The one free parameter of such a point is its direction from O, which is an angle; x and y are two coordinates bound by one constraint, and eliminating one of them is what produces the square root. So $\sqrt{1-h^2}$ is not part of the problem — it is a tax paid for choosing a parameter that does not match the object. Once the angle is used, the constraint $x^2+y^2=1$ is absorbed into $\sin^2+\cos^2=1$, and the algebra of the problem becomes the algebra of addition formulas.

Why does the addition formula appear at E? The diagonal AC joins angle α to angle $-\beta$ (measured from the top). The line through two points of a circle carries information about the *sum* and *difference* of their angles — its direction depends on the half-sum, its distance from O on the half-difference. That is why $\sin(\alpha+\beta)$ and $\sin(\alpha-\beta)$ are the natural outputs, and why the problem was set with a difference $s_1 - s_2$ in the numerator: $s_1^2 - s_2^2$ factors into exactly these two.

The natural attempt reached the same cancellation, $2(k^2-h^2)/|\ldots|$, and stalled because in h, k the denominator $h\sqrt{1-k^2} + k\sqrt{1-h^2}$ is not visibly a factor of k^2-h^2. It is — it is $\sin(\alpha+\beta)$ in disguise — but you only see that when the coordinates are angles. That is the whole cost of the wrong parameter: the structure is present and invisible.

Why is perpendicularity the equality case? An angle between two chords of a circle is half the sum of the arcs they cut off; here the symmetric picture makes the two arcs equal, each $|\alpha-\beta|$, so the angle between the diagonals is $|\alpha-\beta|$ (or its supplement). The answer $\pm 2\sin\theta$ is therefore a statement about the figure, not the coordinates: in an inscribed trapezoid, $|s_1 - s_2| = 2d\sin\theta$. There is no official solution for this year in the archive's sources, so this one is constructed here; the point is that the form $2\sin(\alpha-\beta)$ is forced by using the circle's own coordinate, and that its equality case is read off from the inscribed-angle theorem rather than solved for.

Layer $+\alpha$ BACKTESTING

1. The first cue is "inscribed in the unit circle". On seeing it write the vertices as $(\pm\sin\alpha, \cos\alpha)$, $(\pm\sin\beta, \cos\beta)$, using the symmetry of an inscribed trapezoid.

2. Angle parametrisation followed by addition formulas: the intersection of the diagonals is at height $\sin(\alpha+\beta)/(\sin\alpha+\sin\beta)$, and the ratio collapses to $\pm 2\sin(\alpha-\beta)$.

3. Hesitation occurs at the choice of coordinates, and again after the ratio is found, when the student must describe the equality case. Reinforce the first with "what is the natural parameter of a point on a circle?" and the second with "which angle in the figure equals $\alpha-\beta$?"

4. The trap is Cartesian heights with square roots, leading to two-variable calculus. A secondary trap is the sign: the problem asks for $s_1 - s_2$, not its absolute value, so the labelling matters and AB must be the longer side in the equality case; and $d \neq 0$ excludes the square, the one perpendicular-diagonal trapezoid with $E = O$.

5. Whenever the objects are constrained to a circle, sphere or other curve with a standard parametrisation, use that parametrisation first. Square roots of $1-x^2$ in a computation are the signal that you parametrised the wrong way.

6. Reusable decision rule: *parametrise by the object's own degree of freedom; then look for a difference of squares, because a numerator $A^2 - B^2$ over a denominator involving $A+B$-type terms usually cancels to something bounded.*

1989 B6

1989 B6. Let $(x_1, x_2, \ldots x_n)$ be a point chosen at random from the n-dimensional region defined by $0 < x_1 < x_2 < \cdots < x_n < 1$. Let f be a continuous function on $[0,1]$ with $f(1) = 0$. Set $x_0 = 0$ and $x_{n+1} = 1$. Show that the expected value of the Riemann sum

$$\sum_{i=0}^{n}(x_{i+1} - x_i)f(x_{i+1})$$

is $\int_0^1 f(t)P(t)\,dt$, where P is a polynomial of degree n, independent of f, with $0 \leq P(t) \leq 1$ for $0 \leq t \leq 1$.

Fields: probability; integration over a simplex; order statistics **Position:** B6 — the last problem, but not the hardest on this paper: a direct integration works, and one probabilistic reinterpretation makes it a three-line argument **Answer:** $P(t) = 1-(1-t)^n$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The student integrates. The region $\Delta = \{0 < x_1 < \cdots < x_n < 1\}$ has volume $1/n!$, so the density is $n!$. The last term is $(1-x_n)f(1) = 0$, so the sum is $\sum_{k=1}^{n}(x_k - x_{k-1})f(x_k)$, and by linearity it suffices to compute, for each k,

$$I_k = n!\int_\Delta (x_k - x_{k-1})\,f(x_k)\,dx.$$

Fix $x_k = t$ and $x_{k-1} = s < t$ (for $k \geq 2$). The variables $x_1 < \cdots < x_{k-2}$ range over a simplex in $(0,s)$ of volume $s^{k-2}/(k-2)!$, and $x_{k+1} < \cdots < x_n$ over a simplex in $(t,1)$ of volume $(1-t)^{n-k}/(n-k)!$. Hence

$$I_k = n!\int_0^1 f(t)\frac{(1-t)^{n-k}}{(n-k)!}\int_0^t (t-s)\frac{s^{k-2}}{(k-2)!}\,ds\,dt.$$

The inner integral is $\frac{t^k}{(k-1)!} - \frac{t^k}{k(k-2)!} = \frac{t^k}{k!}$. For $k=1$ the factor is $x_1 - 0 = t$, and $t = t^1/1!$ agrees. So

$$I_k = \int_0^1 f(t)\binom{n}{k}t^k(1-t)^{n-k}\,dt, \qquad \sum_{k=1}^{n} I_k = \int_0^1 f(t)\sum_{k=1}^{n}\binom{n}{k}t^k(1-t)^{n-k}\,dt = \int_0^1 f(t)\big(1-(1-t)^n\big)dt$$

by the binomial theorem (the full sum from $k = 0$ is $(t + 1 - t)^n = 1$). This works. It costs a page, the simplex-volume bookkeeping is error-prone (the student who forgets the $k = 1$ case, or writes $(k-1)!$ for $(k-2)!$, gets a wrong polynomial), and nothing in it explains *why* P comes out between 0 and 1.

The model chain. **Reformulate.** Choosing a point uniformly from Δ is the same as choosing n independent uniform points in $(0, 1)$ and sorting them: sorting is an $n!$-to-one, volume-preserving map from the cube (minus a null set) onto Δ. So let $x_1 < \cdots < x_n$ be the sorted values of n independent uniform points, $x_0 = 0$, $x_{n+1} = 1$.

A gap length is a probability. Let U be one more uniform point, independent of the others. Given the x_i, the probability that U lands in $(x_i, x_{i+1}]$ is $x_{i+1} - x_i$. Let S be the right-hand end of the gap containing U, i.e. the least of $x_1, \ldots, x_{n+1}$ that is $\geq U$. Then

$$\sum_{i=0}^{n} (x_{i+1} - x_i)\, f(x_{i+1}) = \mathbb{E}\big[f(S) \mid x_1, \ldots, x_n\big],$$

so the required expected value is $\mathbb{E}[f(S)]$.

Distribution of S. $S = 1$ exactly when no x_i exceeds U, and then $f(S) = f(1) = 0$. For $0 < t < 1$, $S \leq t$ iff $U \leq t$ and at least one x_i lies in $[U, t]$. Conditioning on $U = u \leq t$, the chance that none of the n points lies in $[u, t]$ is $(1 - (t - u))^n$, so

$$\Pr(S \leq t) = \int_0^t \big(1 - (1 - t + u)^n\big) du = t - \frac{1 - (1 - t)^{n+1}}{n + 1}.$$

Differentiating, S has density $1 - (1 - t)^n$ on $(0, 1)$ (plus an atom at 1 of mass $\frac{1}{n+1}$, which $f(1) = 0$ kills). Therefore

$$\mathbb{E}\Big[\sum_{i=0}^{n} (x_{i+1} - x_i) f(x_{i+1})\Big] = \int_0^1 f(t)\, P(t)\, dt, \qquad P(t) = 1 - (1 - t)^n.$$

P is a polynomial of degree exactly n (leading coefficient $(-1)^{n+1}$), independent of f, and for $0 \leq t \leq 1$ we have $0 \leq (1 - t)^n \leq 1$, so $0 \leq P(t) \leq 1$. □

Checks. $n = 1$: $P(t) = t$, and indeed $\mathbb{E}[x_1 f(x_1) + (1 - x_1) f(1)] = \int_0^1 t f(t)\, dt$. $n = 2$: direct integration gives $2t(1 - t) + t^2 = 2t - t^2 = 1 - (1 - t)^2$. As $n \to \infty$, $P \to 1$ on $(0, 1]$ and the expected Riemann sum tends to $\int_0^1 f$, as a Riemann sum should. I also checked the formula by Monte Carlo simulation for $n = 1, 2, 4, 7$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains drop the last term using $f(1) = 0$ and reduce to a sum over gaps. They split at the factor $(x_{i+1} - x_i)$. (N) treats it as an algebraic factor inside an n-fold integral. (M) reads it as a probability: the length of a subinterval of $[0, 1]$ is the chance a uniform point falls in it. The whole Riemann sum becomes one expectation, $\mathbb{E} f(S)$, and the problem becomes "find the density of one random variable".

> *Each term is (length of a gap) times (f at the gap's right end). In $[0, 1]$ a length is a probability. If I throw one more uniform point U, the sum is the average of f at the right end of whichever gap U lands in.*

> *So the answer is $\int f \cdot$ (density of that right end), and P is a density — no wonder it is between 0 and 1. I only need the distribution of one random variable.*

> *Why is $f(1) = 0$ in the hypothesis? Because the last gap's right end is 1, which is an atom, not a density. The condition is there to delete it.*

- ▷ **Cue:** a weighted sum $\sum$(length of subinterval of $[0,1]$) $\cdot\, g(\dots)$ $\longrightarrow$ **Move:** read each length as the probability that an extra uniform point lands there; the sum becomes one expectation
- ▷ **Cue:** a uniform point in $\{0 < x_1 < \cdots < x_n < 1\}$ $\longrightarrow$ **Move:** replace it by n independent uniforms, sorted
- ▷ **Cue:** an expectation to be written as $\int_0^1 f(t)P(t)\,dt$ for all f $\longrightarrow$ **Move:** find a random variable S with $\mathbb{E} = \mathbb{E}f(S)$; then P is its density
- ▷ **Cue:** a hypothesis like $f(1) = 0$ that seems arbitrary $\longrightarrow$ **Move:** ask what term or atom it deletes; that tells you where the boundary lies in the computation

Layer 3 THE WHY

Why is the probabilistic reading natural rather than inspired? Look at the form the answer is asked for: $\int_0^1 f(t)P(t)\,dt$ *for every continuous* f, with P independent of f. A functional of f that is linear, and given by integrating f against a fixed weight, is the definition of "f averaged against a distribution". The problem is announcing that the expected Riemann sum is $\mathbb{E}f(S)$ for some random variable S with density P — and the bound $0 \le P \le 1$ is the statement that the density of S never exceeds that of a uniform variable. Once you read the target that way, your job is to find S, and the Riemann sum tells you what it is: f is evaluated at right endpoints of gaps, weighted by gap length. The only thing that turns a length into a probability is a uniform point. So you throw one.

Why does $1-(1-t)^n$ come out? S lands near t when one of the n sample points sits at t and U falls somewhere in the empty stretch immediately to its left; the computation above turns that into $\frac{d}{dt}\Pr(S \le t) = 1-(1-t)^n$. Read the answer back as a probability: it is the chance that at least one of the n points lies in $[0,t]$. The natural attempt produced the same function in the form $\sum_{k\ge1}\binom{n}{k}t^k(1-t)^{n-k}$, which is that same probability split by the number $k \ge 1$ of points below t. The binomial sum counts by cases; the closed form counts the complement. Both chains are computing one probability.

Why is the direct integration still worth knowing? Because it is the honest fallback, and it is the natural attempt that *succeeds*; the lesson is not that it fails but that it spends a page on bookkeeping which the reinterpretation spends zero lines on, and that it leaves $0 \le P \le 1$ as an unexplained algebraic fact. The problem punishes zombie integration only by time; it rewards the reading by explaining the answer.

There is no official solution for this year in the archive's sources; both chains here were constructed and checked independently (the binomial computation, the distributional computation, small cases by hand and a numerical simulation).

Layer $+\alpha$ BACKTESTING

1. The first cue is the requested form $\int_0^1 f(t)P(t)\,dt$ with P independent of f and $0 \le P \le 1$. On seeing it write: *P is a density; find the random variable.*

2. Probabilistic reinterpretation: an extra uniform point U turns gap lengths into probabilities, and the Riemann sum into $\mathbb{E}f(S)$ with S the right end of U's gap.

3. Hesitation occurs at the start of the n-fold integral, when the student is tempted to integrate term by term without asking what the sum means. Reinforce it away with the question: "what does a length in $[0,1]$ measure?"

4. Two traps. The bookkeeping trap: simplex volumes $s^{k-2}/(k-2)!$ and the special case $k=1$ are where factors get lost. The hypothesis trap: ignoring $f(1) = 0$, which is exactly what removes the atom of S at 1; without it the expectation is $\int fP + \frac{f(1)}{n+1}$, not of the required form.

5. Whenever an expected value is linear in a function f and the answer is to be $\int fP$, look for the random variable whose density is P. Whenever a weight is the length of an interval inside $[0, 1]$, look for the uniform point that makes it a probability.

6. Reusable decision rule: *before integrating a sum of lengths times values over a simplex, reinterpret lengths as probabilities of an auxiliary uniform point; the sum collapses to a single expectation, and the unknown weight is a density you can compute from one event.*