DAWN PUTNAM SESSION

Putnam 1965 — Papers A and B

Putnam 120 Archive

Contents

Putnam 1965 — Paper A **5**
1965 A1 5
1965 A2 7
1965 A3 10
1965 A4 13
1965 A5 15
1965 A6 18

Putnam 1965 — Paper B **23**
1965 B1 23
1965 B2 26
1965 B3 28
1965 B4 32
1965 B5 35
1965 B6 37

Putnam 1965 — Paper A

1965 A1

1965 A1. Let ABC be a triangle with angle $A <$ angle $C < 90^\circ <$ angle B. Consider the bisectors of the external angles at A and B, each measured from the vertex to the opposite side (extended). Suppose both of these line-segments are equal to AB. Compute the angle A.

Fields: Euclidean geometry, angle chasing, isosceles triangles **Position:** A1 (opener) **Answer:** Angle $A = 12^\circ$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You draw the triangle, mark the external bisector from A meeting line BC (extended) at X, the external bisector from B meeting line AC (extended) at Y, and you are told $AX = AB = BY$. Two segments have been pinned to a known length, in a figure that is otherwise built entirely from bisectors – the reflex is to make this quantitative. You know a formula (or can derive one via the Sine Rule in triangle ABX) for the length of a bisector from a vertex to the opposite side in terms of the triangle's sides and half-angles, something in the shape $AX = \dfrac{2 \cdot AB \cdot AC \cdot \sin(\text{half-angle})}{|AB - AC|}$ for the internal case, and a cousin formula for the external one. You set $AX = AB$ and $BY = AB$, intending to turn each into a trigonometric equation in A, B, C (with $A + B + C = 180^\circ$ available to eliminate one variable), and then solve the resulting system. This is not a dead end – it is provable this way – but you are now carrying two derived length formulas, each itself proved via the Sine Rule in a sub-triangle, plus the angle-sum constraint, and the algebra to isolate A from two simultaneous trigonometric equations in three unknowns is long, easy to mis-cancel, and gives no insight into why the answer comes out to a clean 12°.

The model chain. Do not compute a length. The condition $AX = AB$ is not data for a formula; it is exactly the hypothesis of the isosceles triangle theorem applied to triangle ABX, and $BY = AB$ is the same theorem applied to triangle ABY. Read off the consequences for free.

First, locate X and Y correctly. Since $\angle B > \angle C$, the external bisector from A meets line BC at a point X with B between X and C (the alternative, C between B and X, would force $\angle B \leq \angle C$, contradicting the hypothesis). Symmetrically, since $\angle A < \angle C$, the external bisector from B meets line AC at a point Y with C between A and Y.

Triangle ABY has $AB = BY$, so it is isosceles with apex B; its base angles are equal: $\angle BAY = \angle BYA$. Since Y lies on ray AC beyond C, $\angle BAY$ is the same angle as $\angle BAC = A$, so $\angle BYA = A$ too. Now look at the straight line through A, B extended to a point Z beyond B: the exterior angle of triangle ABY at B, namely $\angle ZBY$, equals the sum of the two remote interior angles, $\angle BAY + \angle BYA = A + A = 2A$. But BY is the external bisector at B, which by definition bisects the exterior angle of the original triangle at B (the angle between ray BZ and ray BC) into two equal halves $\angle ZBY = \angle YBC$; hence $\angle ZBC = 2\angle ZBY = 4A$.

Triangle ABX has $AX = AB$, so it is isosceles with apex A; its base angles at B and X are equal: $\angle ABX = \angle AXB$. Ray BX is opposite to ray BC (since B lies between X and C), and ray

BZ is opposite to ray BA (since B lies between A and Z); the angle between two rays equals the angle between their two opposite rays, so $\angle ABX = \angle ZBC = 4A$. Combined with the base-angle equality, $\angle AXB = \angle ABX = 4A$.

The apex angle of the same isosceles triangle ABX is $\angle XAB$, and since AX is the external bisector at A, it makes an angle with AB equal to half the exterior angle at A: $\angle XAB = \frac{1}{2}(180^\circ - A) = 90^\circ - \frac{A}{2}$.

Now sum the angles of triangle ABX:

$$\left(90^\circ - \frac{A}{2}\right) + 4A + 4A = 180^\circ.$$

Solving, $90^\circ + \frac{15A}{2} = 180^\circ$, so $A = 12^\circ$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains notice the same two facts – $AX = AB$ makes triangle ABX isosceles, $BY = AB$ makes triangle ABY isosceles – in the same first glance at the figure. They split immediately after, on what those facts are *for*: the natural attempt treats "$AX = AB$" as a length equation to be generated from a bisector-length formula and then solved; the model treats it as the hypothesis of the isosceles triangle theorem, whose conclusion (equal base angles) is already an angle statement requiring no formula and no algebra to extract.

> *$AX = AB$ is not a number I need to compute a formula for and then set equal to something – it is literally the defining condition of an isosceles triangle, and the isosceles triangle theorem hands me two equal angles for free the instant I recognise the shape. Since everything else in this problem – the angles A, B, C, the bisector constructions – is already angular, I should stay in angles and never touch a length formula at all.*

Layer 3 THE WHY

The reason the angle-chase is not a shortcut but the *forced* reading of the hypothesis is that $AX = AB$ is precisely, word for word, the defining condition of the isosceles triangle theorem: two sides of triangle ABX are equal, and the theorem that answers "what follows from two sides being equal" is a statement about angles ($\angle ABX = \angle AXB$), not a statement that produces or requires a length formula. A length formula for a bisector is itself proved by an application of the Sine Rule inside a sub-triangle; reaching for it here means proving a fact you do not need (a closed form for AX) in order to discard it again the moment you set it equal to AB and cancel. The isosceles triangle theorem skips that detour because its conclusion is already phrased in the same currency as the hypothesis: both are angle statements about the same triangle.

The remaining engine of the proof – the exterior angle theorem ($\angle ZBY = \angle BAY + \angle BYA$) and the equality of vertical angles ($\angle ABX = \angle ZBC$) – is available for the same underlying reason the problem is entirely stated in angles: it is a triangle-and-bisector configuration, and every fact about such a configuration that does not depend on the actual side lengths is, by construction, an angle fact. The two collinearity choices (B between X and C; C between A and Y) are themselves forced, not assumed, by the given angle order $A < C < 90^\circ < B$: an external bisector's foot lands on whichever side of the vertex the strict angle inequality forces it to, and getting this wrong flips every subsequent angle relation. The final equation $\left(90^\circ - \frac{A}{2}\right) + 4A + 4A = 180^\circ$ is simply "the three angles of a triangle sum to 180°" applied to triangle ABX – the single most axiomatic fact in the subject – once every angle of that triangle has been expressed in terms of A alone by the chase above.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "these line-segments are equal to AB" attached to bisector constructions. On seeing it, write down immediately: "triangle ABX and triangle ABY are isosceles," before touching any formula.

2. That cue suggests the isosceles triangle (base-angle) theorem together with the exterior angle theorem and vertical angles – a pure angle chase – rather than the Sine Rule or a memorised bisector-length formula.

3. Hesitation occurs right after spotting the isosceles triangles, when the instinct is to reach for a bisector-length formula to make the equality quantitative. Reinforce against it by asking: "is this equal-length condition secretly an equal-angle condition in disguise?" – whenever the two equal sides share a vertex with a known role (apex of an isosceles triangle), the answer is yes.

4. The trap is getting the collinearity order wrong – assuming C is between B and X, or A between C and Y – which silently flips the exterior/interior angle relations and produces a self-consistent but wrong equation. The given strict angle inequalities ($A < C < 90^\circ < B$) must be used explicitly to fix the order before any angle is named.

5. Recognise this faster next time by checking, as soon as a problem states "this segment equals this side" inside a construction built purely from bisectors (no other lengths ever appear), whether the two equal segments share an endpoint – if so, that endpoint is the apex of an isosceles triangle and the base angles are available immediately.

6. Decision rule: when a construction is entirely angular (bisectors, no computed lengths required in the answer) and the hypothesis states a length equality between a constructed segment and an existing side, treat the equality as the hypothesis of the isosceles triangle theorem and resolve the whole problem with angle chasing (base angles, exterior angle theorem, vertical angles, angle sum) – do not import the Sine Rule or a bisector-length formula, which reintroduces the algebra the equal-length condition was designed to make unnecessary.

1965 A2

1965 A2. Show that, for any positive integer n,

$$\sum_{r=0}^{\lfloor n/2 \rfloor} \left[\frac{n-2r}{n} \binom{n}{r} \right]^2 = \frac{1}{n} \binom{2n-2}{n-1},$$

where $\lfloor x \rfloor$ means the greatest integer not exceeding x, and $\binom{n}{r}$ is the binomial coefficient "n choose r," with the convention $\binom{n}{0} = 1$.

Caveat. The book's OCR of this problem's display is badly corrupted (rendered as "[(nf/2] jn- 2r (") t'= ~(2n- 2) 1 / r-o 1 n r ~n n- 1"). The statement above was reconstructed from the surviving fragments (the summation index running to $\lfloor n/2 \rfloor$, the factor $(n-2r)/n$ times $\binom{n}{r}$ squared, the greatest-integer convention note, and the final binomial coefficient $\binom{2n-2}{n-1}$ divided by n), then independently verified two ways: (i) direct algebraic derivation below, matching the OCR'd official solution's stated final line "$= \frac{1}{n}\binom{2n-2}{n-1}$"; (ii) numerical check for $n = 1, \ldots, 7$, which matches the closed form exactly in every case. The reconstruction is treated as settled on this evidence, not merely plausible.

Fields: combinatorics, binomial coefficient identities **Position:** A2 **Answer:** The identity holds for every positive integer n; both sides equal $1, 1, 2, 5, 14, 42, 132, \ldots$ for $n = 1, 2, 3, \ldots$ (Catalan-flavoured), collapsing via the factorial identity $r\binom{n}{r} = n\binom{n-1}{r-1}$ and Vandermonde's convolution to $\frac{1}{n}\binom{2n-2}{n-1}$.

Layer 1 THE WHAT — surface analysis

The natural attempt. You expand the square inside the sum: $\left[\frac{n-2r}{n}\binom{n}{r}\right]^2 = \frac{(n-2r)^2}{n^2}\binom{n}{r}^2$, and note that summing over $r = 0$ to $\lfloor n/2 \rfloor$ is exactly half of summing over $r = 0$ to n (the

summand is symmetric under $r \leftrightarrow n-r$, and any middle term at $r = n/2$ contributes zero anyway since $n - 2r = 0$ there). So the target is $\frac{1}{2}T$, where $T = \sum_{r=0}^{n}(n-2r)^2\binom{n}{r}^2$. You expand $(n-2r)^2 = n^2 - 4nr + 4r^2$ and split T into three sums: $n^2\sum\binom{n}{r}^2 - 4n\sum r\binom{n}{r}^2 + 4\sum r^2\binom{n}{r}^2$. The first sum is the classical $\binom{2n}{n}$. For the second and third you reach for generating functions: $\sum_r \binom{n}{r}^2$ is the coefficient of x^n in $(1+x)^n(1+x)^n$, and to bring down a factor of r or r^2 you differentiate. You form $f(x) = (1+x)^n$, note $xf'(x) = \sum r\binom{n}{r}x^r$, and try to extract $\sum r\binom{n}{r}^2$ as $[x^n]\,\big(xf'(x)\big)f(x)$ – a product of two different power series, one differentiated and one not, whose coefficient extraction now requires you to track which side of the product contributed the r and which the $n-r$, and to differentiate again (product rule, second derivative) to reach the r^2 term. The algebra is doable but every step is a bookkeeping hazard: mis-assign which factor carries the derivative and you silently get $\sum r\binom{n}{r}\binom{n}{n-r}$ instead of $\sum r\binom{n}{r}^2$, which happen to be equal here but are not obviously so without a further symmetry argument you have not yet made explicit.

The model chain. Do not differentiate a generating function; the identity is stated entirely in binomial coefficients, so resolve the r-weighting with the one factorial identity that turns "r copies of $\binom{n}{r}$" back into a plain binomial coefficient:

$$r\binom{n}{r} = n\binom{n-1}{r-1}, \qquad r^2\binom{n}{r} = n\Big[(r-1)\binom{n-1}{r-1} + \binom{n-1}{r-1}\Big] = n(n-1)\binom{n-2}{r-2} + n\binom{n-1}{r-1}.$$

Apply this inside $T = \sum_{r=0}^{n}(n-2r)^2\binom{n}{r}^2$ term by term:

$$\sum_r r\binom{n}{r}^2 = n\sum_r \binom{n-1}{r-1}\binom{n}{r} = n\binom{2n-1}{n-1},$$

$$\sum_r r^2\binom{n}{r}^2 = n^2\sum_r \binom{n-1}{r-1}^2 = n^2\binom{2n-2}{n-1},$$

each step an instance of Vandermonde's convolution $\sum_r \binom{p}{r-1}\binom{q}{r} = \binom{p+q}{q-1}$ (or its symmetric restatement). Substituting,

$$T = n^2\binom{2n}{n} - 4n\cdot n\binom{2n-1}{n-1} + 4\cdot n^2\binom{2n-2}{n-1} = n^2\left[\binom{2n}{n} - 4\binom{2n-1}{n-1} + 4\binom{2n-2}{n-1}\right].$$

Now use two applications of Pascal's rule to collapse the bracket. First, $\binom{2n}{n} = \binom{2n-1}{n-1} + \binom{2n-1}{n} = 2\binom{2n-1}{n-1}$ (the two summands are equal by symmetry, since $(2n-1) - n = n-1$). Substituting,

$$\binom{2n}{n} - 4\binom{2n-1}{n-1} = -2\binom{2n-1}{n-1}.$$

Second, $\binom{2n-1}{n-1} = \binom{2n-2}{n-1} + \binom{2n-2}{n-2}$, so the bracket becomes

$$-2\binom{2n-2}{n-1} - 2\binom{2n-2}{n-2} + 4\binom{2n-2}{n-1} = 2\binom{2n-2}{n-1} - 2\binom{2n-2}{n-2}.$$

Finally, $\binom{2n-2}{n-2} = \frac{n-1}{n}\binom{2n-2}{n-1}$ (a one-line ratio of factorials), so the bracket equals $2\binom{2n-2}{n-1}\left[1 - \frac{n-1}{n}\right] = \frac{2}{n}\binom{2n-2}{n-1}$. Hence

$$T = n^2\cdot\frac{2}{n}\binom{2n-2}{n-1} = 2n\binom{2n-2}{n-1}.$$

Recall T was defined as the sum *before* dividing by n^2 and *before* halving for the range $r = 0$ to $\lfloor n/2\rfloor$; the required sum is $\dfrac{1}{2n^2}T = \dfrac{1}{2n^2}\cdot 2n\binom{2n-2}{n-1} = \dfrac{1}{n}\binom{2n-2}{n-1}$, exactly the right-hand side.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains reach the same decomposition $T = n^2\sum\binom{n}{r}^2 - 4n\sum r\binom{n}{r}^2 + 4\sum r^2\binom{n}{r}^2$ in the same first move – that split is not where they differ. They split on how to evaluate $\sum r\binom{n}{r}^2$ and $\sum r^2\binom{n}{r}^2$: the natural attempt reaches for a generating function and differentiates a product of two power series; the model chain reaches for the single factorial identity $r\binom{n}{r} = n\binom{n-1}{r-1}$ and reduces each weighted sum to a plain Vandermonde convolution.

> *The factor of r multiplying $\binom{n}{r}$ is not a differentiation cue – it's a combinatorial identity. $r\binom{n}{r} = n\binom{n-1}{r-1}$ turns "r copies of* choose *r* from *n" into "n copies of* choose $r-1$ from $n-1$," *which is again a plain binomial coefficient I can convolve with $\binom{n}{r}$ using Vandermonde – no derivative, no product rule, no risk of mis-assigning which power series contributed which factor.*

Layer 3 THE WHY

The identity $r\binom{n}{r} = n\binom{n-1}{r-1}$ follows directly from the factorial definition $\binom{n}{r} = \frac{n!}{r!(n-r)!}$: multiply both sides by r and cancel one factor of r from $r!$, giving $\frac{n!}{(r-1)!(n-r)!} = n \cdot \frac{(n-1)!}{(r-1)!(n-r)!} = n\binom{n-1}{r-1}$ – nothing beyond the definition of the binomial coefficient is used. It also has a one-line combinatorial reading: "r times the number of r-subsets of an n-set" counts pairs (subset, marked element of the subset); by choosing the marked element first (n ways) and then the remaining $r-1$ elements from what is left ($\binom{n-1}{r-1}$ ways), the same count is $n\binom{n-1}{r-1}$. Both readings stay entirely inside the combinatorial world the problem is posed in.

This is exactly why it is the forced move and not a lucky substitution: the identity to be proved is stated purely in the language of binomial coefficients, with no exponential, trigonometric, or explicit polynomial variable anywhere in sight. A generating function is a tool for extracting coefficients *from* a polynomial or power series; reaching for one here means manufacturing a variable x that the problem never asked for, differentiating with respect to it, and then discarding it again once the coefficient is extracted – a round trip through continuous machinery to answer a question that is already, natively, discrete. The factorial identity, by contrast, never leaves the discrete world: it converts an r-weighted binomial coefficient into an unweighted one of adjacent parameters, which is exactly the input Vandermonde's convolution (itself proved by a counting argument, choosing r objects from a union of two disjoint sets in two different ways) is built to consume. Once every weighted term has been converted, the rest of the proof is two applications of Pascal's rule – $\binom{n}{r} = \binom{n-1}{r} + \binom{n-1}{r-1}$, the recursive definition of the binomial coefficient itself – which is as close to an axiom as this subject has. The final ratio $\binom{2n-2}{n-2} = \frac{n-1}{n}\binom{2n-2}{n-1}$ is likewise nothing but the factorial definition again, confirming that at no point in the model chain does the argument leave the combinatorial register the problem was stated in.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the appearance of $r\binom{n}{r}$ (or $r^2\binom{n}{r}$) inside a sum of squared binomial coefficients. On seeing it, write down immediately: $r\binom{n}{r} = n\binom{n-1}{r-1}$.
2. That identity suggests converting every r-weighted binomial coefficient into a plain binomial coefficient with shifted parameters, then applying Vandermonde's convolution $\sum_r \binom{p}{r}\binom{q}{s-r} = \binom{p+q}{s}$ to each resulting sum.
3. Hesitation occurs exactly when the eye lands on the factor r attached to $\binom{n}{r}$ and the reflex fires "differentiate the generating function." Reinforce against it by asking: is this r counting something (an element of a chosen subset), or is it acting as a continuous parameter? In a purely combinatorial identity it is always the former.
4. The trap is the generating-function route's bookkeeping: differentiating $(1+x)^n$ and forming a product of two different power series (one differentiated, one not) to extract $\sum r\binom{n}{r}^2$

requires tracking which factor of the product contributed the r and which contributed the plain $\binom{n}{n-r}$ – a mis-assignment silently produces the wrong sum, one that happens to equal the right answer here but is not obviously equal without a further symmetry argument.

5. Recognise this faster next time by checking, as soon as an identity to prove involves only binomial coefficients (no exponential/trigonometric function, no explicit polynomial variable x), that a factorial-identity-and-Vandermonde route exists by default, and reach for it before reaching for calculus.

6. Decision rule: when an identity is stated purely in the language of binomial coefficients, solve it purely in that language – convert every weighted coefficient ($r\binom{n}{r}$, $r^2\binom{n}{r}$, etc.) into a sum of plain binomial coefficients via $r\binom{n}{r} = n\binom{n-1}{r-1}$, then finish with Vandermonde's convolution and Pascal's recursion; reserve generating-function calculus for identities already stated in terms of a polynomial variable.

1965 A3

1965 A3. Show that, for any sequence $a_1, a_2, \cdots$ of real numbers, the two conditions

$$\lim_{n\to\infty} \frac{e^{ia_1} + e^{ia_2} + \cdots + e^{ia_n}}{n} \text{ exists} \tag{A}$$

and

$$\lim_{n\to\infty} \frac{e^{ia_1} + e^{ia_2} + \cdots + e^{ia_{n^2}}}{n^2} \text{ exists} \tag{B}$$

are equivalent.

Caveat. The book's OCR of this statement is severely corrupted (rendered as fragments like "e<i4•1 + e<iao) + ... + eCI4.)" and "e<i4,) + e<ia<> + ... + e<ia.")"), with the two limit expressions barely legible and no visible indication of what index the second limit runs to. The reconstruction above – condition (A) the Cesàro-type mean of e^{ia_k} over the first n terms, condition (B) the same mean but restricted to the first n^2 terms – is fixed by the surviving solution text, which explicitly substitutes $m = n^2 + k$ with $0 \le k \le 2n$ and uses the bound $|S(t+k) - S(t)| \le k$ for $S(t) = e^{ia_1} + \cdots + e^{ia_t}$; this scale (gap of size $O(n)$ between consecutive perfect squares) only makes sense if (B)'s index runs over n^2. The reconstructed statement was independently re-derived and verified in full below (both directions), and its proof structure and final bound $2k/m \le 4/n$ match the OCR'd solution's surviving arithmetic exactly, so the reconstruction is treated as settled.

Fields: analysis, sequences of complex numbers, Cesàro-type limits **Position:** A3 **Answer:** Conditions (A) and (B) are equivalent: (A)⇒(B) is immediate because a subsequence of a convergent sequence converges to the same limit, and (B)⇒(A) follows by squeezing the general term $S(m)/m$ against its nearest perfect-square anchor $S(n^2)/n^2$ using the bound $|S(t+k) - S(t)| \le k$.

Layer 1 THE WHAT — surface analysis

The natural attempt. Write $c_k = e^{ia_k}$ and $S(t) = c_1 + c_2 + \cdots + c_t$, so (A) says $S(n)/n$ converges and (B) says $S(n^2)/n^2$ converges. The direction (A)⇒(B) is free: $S(n^2)/n^2$ is literally the subsequence of $S(n)/n$ obtained by restricting to perfect-square indices, and a subsequence of a convergent sequence converges to the same limit. The real content is (B)⇒(A), and here the natural instinct, since $S(n)/n$ looks like a Cesàro average, is to reach for a Cesàro- or Stolz-type convergence theorem: something of the flavour "if the averages along a well-behaved subsequence converge, and the terms are controlled, the full average converges." You try to set

this up directly – perhaps attempting to bound $S(n) - S(n-1) = c_n$ and control how $S(n)/n$ drifts from one n to the next, hoping to invoke a named theorem about Cesàro means or to run an epsilon-argument comparing $S(n)/n$ and $S(n+1)/(n+1)$ term by term. But this stalls: the hypothesis (B) only ever tells you anything at the sparse indices $1, 4, 9, 16, \ldots$; it says nothing at all about $S(n) - S(n-1)$ for a general n, and no amount of massaging a Cesàro-mean theorem produces control at every index from information given only at an index set of density $\to 0$. You are trying to force a theorem built for dense, consecutive-index control onto a hypothesis that only supplies sparse anchors, and the fit is not there.

The model chain. Do not look for a named theorem; build the one inequality the sparse hypothesis actually supports. First record two facts that come for free from $|c_k| = |e^{ia_k}| = 1$: (i) $|S(t)| \le t$ for every t (triangle inequality on t unit vectors), and (ii) $|S(t+k) - S(t)| = |c_{t+1} + \cdots + c_{t+k}| \le k$ for any $t \ge 0$, $k \ge 0$ (triangle inequality on k unit vectors).

Assume (B): $S(n^2)/n^2 \to L$ as $n \to \infty$, for some complex number L. Given a general positive integer m, let $n = \lfloor \sqrt{m} \rfloor$, so $n^2 \le m < (n+1)^2$, and write $m = n^2 + k$ with $0 \le k \le 2n$ (since $(n+1)^2 - n^2 = 2n + 1$). Compare $S(m)/m$ to $S(n^2)/n^2$:

$$\left| \frac{S(m)}{m} - \frac{S(n^2)}{n^2} \right| = \left| \frac{S(m) - S(n^2)}{m} + S(n^2) \left(\frac{1}{m} - \frac{1}{n^2} \right) \right| \le \frac{|S(m) - S(n^2)|}{m} + |S(n^2)| \left(\frac{1}{n^2} - \frac{1}{m} \right),$$

using $m \ge n^2$ so that $1/n^2 - 1/m \ge 0$. By fact (ii), $|S(m) - S(n^2)| = |S(n^2 + k) - S(n^2)| \le k$. By fact (i), $|S(n^2)| \le n^2$. So the right side is at most

$$\frac{k}{m} + n^2 \left(\frac{1}{n^2} - \frac{1}{m} \right) = \frac{k}{m} + \frac{m - n^2}{m} = \frac{k}{m} + \frac{k}{m} = \frac{2k}{m}.$$

Since $k \le 2n$ and $m \ge n^2$, this is at most $\dfrac{4n}{n^2} = \dfrac{4}{n}$. As $m \to \infty$, necessarily $n = \lfloor \sqrt{m} \rfloor \to \infty$, so $4/n \to 0$, forcing

$$\left| \frac{S(m)}{m} - \frac{S(n^2)}{n^2} \right| \to 0.$$

Since $S(n^2)/n^2 \to L$ by hypothesis (B), and the difference above vanishes, $S(m)/m \to L$ as well. This is exactly condition (A), with the same limit L.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree instantly that (A)⇒(B) is free – a convergent sequence's subsequence converges to the same limit, no work needed. They split entirely on the harder direction, (B)⇒(A): the natural attempt looks for a pre-packaged theorem (Cesàro-mean or Stolz-type) that promises to upgrade convergence on a subsequence to convergence of the whole sequence; the model chain builds the bridging inequality by hand, using only the one property every c_k has by definition ($|c_k| = 1$) to bound how far S can move between the sparse anchor points.

> *I don't have control over S at every index n – only at the perfect squares, where (B) hands me a limit. So I cannot invoke a theorem that implicitly needs consecutive differences $S(n) - S(n-1)$; the only fact I actually get for free from $|e^{ia_k}| = 1$ is that S cannot move by more than 1 per step, and that is precisely the tool that lets me anchor any index m to the nearest perfect square below it and bound the drift.*

Layer 3 THE WHY

The reason the anchoring argument is forced, not merely convenient, is that the hypothesis (B) is information supplied on an extremely sparse index set – the perfect squares, which have density $\to 0$ among the positive integers – while the conclusion (A) is wanted at every index.

Any bridge from a sparse set to the full sequence has to be built entirely out of facts that hold at every index, and the single universal fact available here comes directly from the definition of the objects in the problem: $e^{i\theta}$ lies on the unit circle for every real θ, by the definition of the complex exponential (Euler's formula), so $|c_k| = 1$ for every k, with no exceptions and no need for any further hypothesis on the a_k. The triangle inequality applied to k such unit vectors in a row is what turns that pointwise fact into the step-bound $|S(t+k) - S(t)| \leq k$ – again nothing beyond the definition of the modulus and the triangle inequality, no theorem imported from elsewhere.

The scale of the problem is what makes the bound close the argument rather than merely bound it loosely: the gap between consecutive perfect squares, $(n+1)^2 - n^2 = 2n+1$, is $O(n)$, while the anchor itself is n^2 – so the ratio of gap to anchor, $O(n)/n^2 = O(1/n)$, vanishes as $n \to \infty$. This is not a coincidence to be checked after the fact; it is the reason the problem chose n^2 (rather than, say, 2^n, where the analogous gap 2^n is the same order as the anchor 2^n and no such squeeze would close). Once this scale relationship is seen, the rest is arithmetic: the error introduced by moving from $S(n^2)/n^2$ to $S(m)/m$ is bounded by (step bound)/(anchor size) plus (a rearranged difference-of-reciprocals term that telescopes to the same order), both of which are forced to zero by the same $O(1/n)$ scale. A Cesàro- or Stolz-type theorem, by contrast, is designed to convert control of *consecutive* differences into control of the mean; it is the wrong tool exactly because the hypothesis here never supplies consecutive differences at all, only values at a sparse anchor set, and no theorem can manufacture information the hypothesis does not contain.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that a limit hypothesis is stated on an explicit, sparse subsequence (here: at every perfect square index n^2) while the conclusion is wanted at every index. On seeing it, write down immediately the gap between consecutive terms of that sparse index set: $(n+1)^2 - n^2 = 2n+1$.

2. That gap suggests bounding how much the quantity in question ($S(t)$, the partial sum) can change over exactly that gap, using whatever uniform per-step bound the problem's own definitions supply – here, $|c_k| = 1$ giving $|S(t+k) - S(t)| \leq k$ – and then squeezing a general index between its two nearest sparse anchors.

3. Hesitation occurs right when the sequence $S(n)/n$ is recognised as a Cesàro-type average and the instinct is to reach for a named convergence theorem (Stolz–Cesàro, Cauchy criterion) that implicitly assumes control at every consecutive index. Reinforce against it by asking: does my hypothesis actually give me consecutive differences $S(n) - S(n-1)$, or only differences at a sparse set of points? Here it is only the latter, and no theorem can supply what the hypothesis withholds.

4. The trap is assuming that convergence (or boundedness) of S at perfect squares says anything at all about the individual increments $c_n = S(n) - S(n-1)$ – it does not, since the a_n between consecutive squares are entirely unconstrained by (B). Trying to control those increments directly, rather than the aggregate drift over the whole gap, is a dead end.

5. Recognise this faster next time by checking, whenever a limit hypothesis is stated along a subsequence indexed by n^p for some fixed $p \geq 2$ rather than by n itself, the size of the gap $(n+1)^p - n^p = O(n^{p-1})$ against the anchor size n^p: if the ratio $O(n^{p-1})/n^p = O(1/n) \to 0$, a squeeze argument through the nearest anchor is very likely to close the gap between the subsequence and the full sequence.

6. Decision rule: to upgrade a limit hypothesis from a sparse subsequence (n_k) to the full sequence, find a uniform per-step bound on how much the quantity of interest can change (typically a triangle-inequality consequence of a modulus-one, boundedness, or Lipschitz-type fact built into the problem's own definitions), and squeeze the general term between its two nearest subsequence anchors using that bound – reach for this construction, not a named Cesàro/Stolz theorem, whenever the hypothesis is silent about consecutive differences and

only supplies information at a sparse index set whose gaps are asymptotically smaller than the indices themselves.

1965 A4

1965 A4. At a party, assume that no boy dances with every girl but each girl dances with at least one boy. Prove that there are two couples gb and $g'b'$ which dance whereas b does not dance with g' nor does g dance with b'.

Caveat. Identity check: the MAA volume's solution under 1965 A-4 takes a boy b with a maximal number of partners, a girl g' he does not dance with, a partner b' of g', and finds a partner g of b not dancing with b'; that proves exactly this statement. The OCR text of the statement is clean. The party is tacitly finite (otherwise "a boy with a maximal number of partners" need not exist, and the statement can fail: boys $b_1, b_2, \ldots$ with b_k dancing with girls $g_1, \ldots, g_k$ satisfy both hypotheses and have nested partner sets, so no two couples of the required kind exist). The analysis below assumes finitely many boys and girls.

Fields: combinatorics (bipartite graphs, extremal principle) **Position:** A4 — mid-paper; short once the conclusion is read correctly **Answer:** True. Let $N(b)$ be the set of girls boy b dances with. The conclusion says precisely that some two boys have incomparable sets $N(b) \not\subseteq N(b')$ and $N(b') \not\subseteq N(b)$. If all the sets were nested, the largest would be the union, which is every girl — contradicting the hypothesis that no boy dances with every girl.

Layer 1 THE WHAT — surface analysis

The natural attempt. Draw it. Boys on one side, girls on the other, an edge for each pair that dances. You want two edges gb and $g'b'$ with the "cross" pairs bg' and $b'g$ both missing: a picture of two disjoint dances with no crossover.

The forward-construction instinct is to build the configuration one piece at a time. Take any boy b. He does not dance with every girl, so there is a girl g' he misses. She dances with somebody, say b'; and $b' \neq b$ because b does not dance with g'. Now you want a partner g of b whom b' does not dance with. Is there one? Not necessarily. Perhaps b' dances with every partner of b and with g' as well.

So you patch. If that happens, replace b by b' and start again: b' also fails to dance with somebody, that girl has a partner b'', and so on. Each time the construction fails, the new boy's partner set strictly contains the old one's (it contains all of the old set plus g'). A strictly increasing chain of subsets of a finite set of girls must stop, so eventually the construction succeeds.

That is a correct proof, and it is worth noticing that it is correct. But it costs a loop, a termination argument, and a reader who has to trust that "and so on" has been handled. Under exam pressure most people who start this way either stop at the first failure ("perhaps b' dances with all of b's partners — then I'm stuck") or write the loop without stating why it ends.

The model chain. Assume the party is finite and write $N(b)$ for the set of girls boy b dances with. The hypotheses are:

(H1) $N(b) \neq G$ for every boy b, where G is the set of all girls;

(H2) every girl lies in some $N(b)$, that is, $\bigcup_b N(b) = G$.

There is at least one girl (otherwise (H1) fails: a boy vacuously dances with every girl of an empty set), so by (H2) there is at least one boy.

Choose a boy b for whom $|N(b)|$ is as large as possible; there are finitely many boys, so such a b exists. By (H1) there is a girl $g' \notin N(b)$. By (H2) there is a boy b' with $g' \in N(b')$; since $g' \notin N(b)$, $b' \neq b$.

Suppose every girl in $N(b)$ also danced with b'. Then $N(b) \cup \{g'\} \subseteq N(b')$, and $g' \notin N(b)$, so $|N(b')| \geq |N(b)| + 1$, contradicting the maximality of $|N(b)|$. Hence there is a girl $g \in N(b)$ with $g \notin N(b')$.

Now gb dance ($g \in N(b)$), $g'b'$ dance ($g' \in N(b')$), b does not dance with g' ($g' \notin N(b)$), and g does not dance with b' ($g \notin N(b')$). Also $g \neq g'$ and $b \neq b'$, so these are two genuinely different couples. ■

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the first three moves: a boy b, a girl g' he misses, her partner b'. They split at the choice of b. The natural attempt takes *any* boy and has to repair the construction when b' turns out to dominate b. The model takes b to be the boy with the most partners, so that "b' dominates b" is impossible from the outset and the repair loop never runs.

> *The only way my construction can fail is if b' dances with everyone b dances with, and then some. So I'll pick b so that nobody can dance with strictly more girls than he does. Then the failure case is dead before it starts.*

> *Four conditions, two "do" and two "does not". Written in terms of partner sets, that is $g \in N(b) \setminus N(b')$ and $g' \in N(b') \setminus N(b)$: the two sets are incomparable. So I'm being asked to show the partner sets are not all nested.*

Layer 3 THE WHY

Why is the maximal boy the natural choice rather than a trick? Look at what the loop in the natural attempt actually did. Every time it failed, it moved to a boy with a strictly bigger partner set. A process that only ever moves upward in a finite ordering, and succeeds unless it can move upward, is a process that succeeds at the top. So start at the top. The extremal choice is not a separate idea from the loop; it is the loop with the iterations deleted. That is the general content of the extremal principle: whenever a failed attempt hands you a strictly "bigger" object of the same kind, pick the biggest object at the start.

The second critical thought goes deeper, down to what the conclusion *means*. The problem states four conditions on four people, which is designed to look like a search. Translate each condition into set language and the search collapses into one sentence: $N(b)$ and $N(b')$ are incomparable under inclusion. The negation is then equally short: for every two boys, one partner set contains the other. A family of finite sets in which every two are comparable is a chain, and a finite chain has a largest member. The largest member contains every other member, so it equals their union. By (H2) the union is all of G. So the boy owning the largest set dances with every girl, and (H1) forbids that.

Seen this way, the two hypotheses are the two halves of one contradiction. (H2) says the union of the partner sets is everything; (H1) says no single partner set is everything. For a chain these cannot both hold, because in a chain the union *is* a member. So the problem is the statement "a finite family of sets whose union is not a member of the family cannot be a chain". The MAA solution's "boy with a maximal number of girls" is the top of the would-be chain, and the step "otherwise b' would have more partners than b" is the comparability failing exactly where it must.

This also explains why the maximum should be taken by *size* rather than by inclusion. By inclusion, a boy b with $N(b)$ maximal (not contained in any other) also works: if every partner of b danced with b', then $N(b) \subsetneq N(b')$. Either ordering does the job. Size is chosen because it is a number, and choosing the boy with the largest number is a move nobody has to justify.

Finally, the finiteness caveat is not pedantry; it is where the proof lives. With infinitely many boys dancing with $\{g_1\}, \{g_1, g_2\}, \{g_1, g_2, g_3\}, \ldots$ the partner sets form a chain whose union is every girl, yet no boy dances with every girl. The chain has no top, and the argument has nothing to stand on. If you ever write an extremal argument, the first line must say why the extreme exists.

Layer $+\alpha$ BACKTESTING

1. The cue is a conclusion with a pattern of "does" and "does not" across two pairs. Write down, at once, $N(b)$ for each boy and translate: $g \in N(b) \setminus N(b')$, $g' \in N(b') \setminus N(b)$.

2. Either the extremal principle (take the boy with most partners) or contradiction via chains (if no two partner sets are incomparable they are nested, so the largest is the union). Both are three lines.

3. Hesitation comes after choosing b, g', b' when you realise b' might dance with all of b's partners. The reinforcement is the question "what would that make b'?" — the answer, "a boy with strictly more partners", tells you which b to have chosen.

4. The trap is starting with an arbitrary boy and either getting stuck at the domination case or iterating without proving termination. A secondary trap is forgetting to say the party is finite, which is exactly what makes the maximum exist.

5. When a combinatorial conclusion lists several "is related" and "is not related" conditions, rewrite it as a statement about neighbourhood sets. Two neighbourhoods, one element in each difference, is incomparability.

6. If your construction fails only when some object is strictly bigger (or better) than the one you chose, choose the biggest one at the start and state why it exists. If you need to show a family of sets is not a chain, compare its union with its members.

▷ **Cue:** a conclusion of the form "$x \sim y$, $x' \sim y'$, but $x \not\sim y'$ and $x' \not\sim y$" $\longrightarrow$ **Move:** write neighbourhood sets $N(x)$, $N(x')$ and read the conclusion as "$N(x)$ and $N(x')$ are incomparable"

▷ **Cue:** a construction that fails only when the next object strictly dominates the current one $\longrightarrow$ **Move:** start from an extremal object (maximum size, maximal under inclusion) so the failure case is impossible

▷ **Cue:** hypotheses "no single set is everything" and "the union of the sets is everything" $\longrightarrow$ **Move:** a finite chain contains its union, so the sets cannot be nested

▷ **Cue:** an extremal argument over a possibly infinite family $\longrightarrow$ **Move:** state finiteness (or compactness) first; without it the extreme may not exist and the claim may be false

1965 A5

1965 A5. In how many ways can the integers from 1 to n be ordered subject to the condition that, except for the first integer on the left, every integer differs by 1 from some integer to the left of it?

Caveat. Identity check: the MAA volume's solution under 1965 A-5 conjectures 2^{n-1}, shows every admissible arrangement ends in 1 or n, and counts by induction; that is a solution of this statement. The only OCR damage is a stray bullet after "condition that", read here as a comma. The answer 2^{n-1} was checked by exhaustive enumeration for $n \leq 8$.

Fields: combinatorics (enumeration, bijections) **Position:** A5 — late in the paper, but one structural observation reduces it to a binary choice **Answer:** 2^{n-1}. At every stage the integers already written form a block of consecutive integers, so each new integer extends the block by one at the bottom or at the top. Read from the right, the arrangement is determined by $n-1$ choices "current minimum or current maximum".

Layer 1 THE WHAT — surface analysis

The natural attempt. List small cases. $n = 1$: (1). $n = 2$: $(1,2), (2,1)$. $n = 3$: the six orders of $1,2,3$; $(1,3,2)$ fails because 3 differs by 1 from nothing to its left, and $(3,1,2)$ fails the same way at 1. The four survivors are

$$(1,2,3),\ (2,1,3),\ (2,3,1),\ (3,2,1).$$

$n = 4$ gives 8 after a patient check of 24 orders. The pattern $1,2,4,8$ suggests 2^{n-1}.

Now prove it. The instinct is to count by the first entry, since the first entry is the only unconstrained one. Start with k. The next entry must be $k-1$ or $k+1$. After that, the third entry must differ by 1 from k or from the second entry, and the options depend on which of $k \pm 1$ you took. You write a tree for $k = 3$, $n = 5$ and it branches in a way that is hard to summarise, because at each stage you are asking "which integers differ by 1 from *something* already written, and are not yet written?", and that set looks as though it could be large.

Push a little further and the tree tells you its secret: the answer to that question is never more than two integers. But unless you stop and prove that, the tree remains a tree, and the count remains a conjecture.

The model chain. Call an ordering $(a_1, \ldots, a_n)$ of $\{1, \ldots, n\}$ *admissible* if each a_i with $i \geq 2$ differs by 1 from some a_j with $j < i$.

Claim. In an admissible ordering, for every i the set $S_i = \{a_1, \ldots, a_i\}$ is a block of consecutive integers $\{p, p+1, \ldots, q\}$.

Proof. Induction on i. $S_1 = \{a_1\}$ is a block. Suppose $S_{i-1} = \{p, \ldots, q\}$. The new entry a_i is not in S_{i-1} (the entries are distinct), and it differs by 1 from some element of $\{p, \ldots, q\}$, so $a_i \in \{p-1, \ldots, q+1\}$. Not being in $\{p, \ldots, q\}$, it is $p-1$ or $q+1$, and S_i is again a block. □

Conversely, if every S_i is a block then the ordering is admissible: a_i is $p-1$ or $q+1$ for the block $S_{i-1} = \{p, \ldots, q\}$, and so differs by 1 from p or from q. So admissible orderings are exactly those in which every initial segment is a block.

Counting. Read the ordering from the right. Since $S_n = \{1, \ldots, n\}$ and S_{n-1} is a block of $n-1$ integers inside it, a_n is 1 or n: the last entry is the minimum or the maximum of the whole set. In general a_i is the minimum or the maximum of the block S_i, because $S_{i-1} = S_i \setminus \{a_i\}$ must still be a block.

Now build an admissible ordering backwards. Choose $a_n \in \{\min S_n, \max S_n\}$: two options. That fixes S_{n-1}, a block of $n-1$ integers; choose a_{n-1} as its minimum or its maximum: two options. Continue down to S_2, a block of two integers, where a_2 is one of them: two options. Finally a_1 is the one remaining integer. Every sequence of $n-1$ choices produces an ordering whose initial segments are all blocks (each S_i is a block by construction), and distinct choice sequences give distinct orderings (the first position where the choices differ receives different integers, since a block of at least two integers has distinct minimum and maximum). So the number of admissible orderings is

$$2^{n-1}. \qquad \blacksquare$$

Forward count, as a check. Counting from the left instead, the first entry k is free, and each later entry extends the block downward or upward; exactly $k-1$ steps go down and $n-k$ go up, in any order among the $n-1$ steps. That gives $\binom{n-1}{k-1}$ orderings starting with k, and $\sum_{k=1}^{n} \binom{n-1}{k-1} = 2^{n-1}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains list the small cases and conjecture 2^{n-1}. They split at the question "what can the next entry be?". The natural attempt answers it entry by entry, inside a growing tree. The model answers it once, for all stages, by describing the *set* of entries written so far — it is always an interval — and the tree collapses to two branches per step.

> *I keep asking what the next number can be. Let me instead ask what the set of numbers written so far looks like. In every example it has no gaps. If that's always true, the next number can only be one below the bottom or one above the top.*

> *The answer is a power of* 2*, so somewhere there is a binary choice made* $n-1$ *times. The interval grows by one each step, down or up: that's the choice.*

Layer 3 THE WHY

Why look at the set of written entries rather than the entries themselves? Because of the form of the condition. "Differs by 1 from *some* integer to the left" is a condition on the new entry relative to a set, not relative to any particular earlier entry. The earlier entries are relevant only through the set they form; their order is irrelevant to what may come next. When a rule refers to "some earlier element", the state of the process is the set of earlier elements, and the first job is to describe what that set can look like. Here the description is immediate: a set of integers grown from one point by adding neighbours can never skip a number, because a number is only ever added next to one already present. That is the same reason a connected region grown by adding adjacent squares stays connected; on a line, "connected" means "interval".

Once the state is an interval $\{p, \ldots, q\}$, the rule's options are forced: the integers within 1 of the interval and outside it are $p-1$ and $q+1$, at most two of them. The problem has turned into a walk: two endpoints, each step moves one of them outward. The count 2^{n-1} is then not a coincidence to be verified but a reading of the structure, provided you count in the direction where both options always exist. Going forward, the options can be blocked: if $p = 1$ you cannot go down. That is why the forward count needs the binomial sum. Going backward, the last entry is the min or the max of the current interval, and an interval with at least two elements always has two distinct extremes. No blocking, no case analysis, a clean 2 at every step.

The MAA solution reaches the same fact by a different road: it shows by induction that an admissible n-arrangement ends in 1 or n, then uses the reflection $a \mapsto n+1-a$ to show half end in n, and identifies those with $(n-1)$-arrangements. Each step is correct, but "ends in 1 or n" is stated and proved without the reason it is true. The reason is the interval claim: the first $n-1$ entries form an interval of length $n-1$ inside $\{1, \ldots, n\}$, so the missing entry is an end. Once you see intervals, the MAA's induction, reflection and deletion are all one sentence.

Note the direction of travel in the model: from "the answer is 2^{n-1}" to "find the binary choice". A power of 2 in a small-case table is itself a cue. It says the objects are in bijection with binary strings of length $n-1$, and the proof's job is to find the bit. Asking "what is being chosen $n-1$ times between two options?" is a far better search than "how do I set up an induction?".

Layer $+\alpha$ BACKTESTING

1. The cue is "from some integer to the left of it". Write down the set S_i of entries so far and draw it on a number line for the $n = 4$ examples: it is always an interval.

2. State space description plus bijection: prove each S_i is an interval, then encode each admissible ordering by the sequence of "extend down / extend up" moves (or, backwards, "remove min / remove max").

3. Hesitation comes when counting forward and the number of choices seems to depend on where you are (at an end of $\{1, \ldots, n\}$ one option disappears). Reinforce by counting from the other end, where both options always exist, or by accepting the binomial sum.

4. The trap is building the branching tree of "which integers are within 1 of something written" without noticing that the set written is an interval, and hence that the tree has at most two branches per node. A second trap is conjecturing 2^{n-1} from small cases and then proving it by an induction whose key lemma (the arrangement ends in 1 or n) is asserted rather than proved.

5. When a rule constrains the next element relative to "some earlier element", stop tracking the sequence and track the set. When a small-case table shows powers of 2, look for an independent binary choice made at each step.

6. Describe the reachable states before counting paths. If the states are intervals (or connected sets) grown one step at a time, the process is a walk on its endpoints, and the count is a product of per-step option counts. Count in the direction where the number of options is constant.

▷ **Cue:** "differs by 1 from some integer already placed" ⟶ **Move:** the placed set is always an interval; the next entry is one below its minimum or one above its maximum

▷ **Cue:** a condition referring to "some earlier element" rather than a specific one ⟶ **Move:** the state of the process is the set of earlier elements; describe what that set can be

▷ **Cue:** a small-case table reading $1, 2, 4, 8, \ldots$ ⟶ **Move:** look for a binary choice made at each of $n-1$ steps and build the bijection with binary strings

▷ **Cue:** a forward count whose number of options varies with position ⟶ **Move:** count backwards from the end, where the options may be uniform

1965 A6

1965 A6. In the plane with orthogonal Cartesian coordinates x and y, prove that the line whose equation is $ux + vy = 1$ will be tangent to the curve $x^m + y^m = 1$ (where $m > 1$) if and only if $u^n + v^n = 1$ and $m^{-1} + n^{-1} = 1$.

Caveat. Identity check: the MAA volume's solution under 1965 A-6 computes the tangent $x_0^{m-1}x + y_0^{m-1}y = 1$ at a point of $x^m + y^m = 1$ and verifies $u^n + v^n = 1$ with $n = m/(m-1)$; it solves this statement. The OCR renders the exponents as "$X''' + Y'''$" and "$u^\bullet + v^\bullet$"; the reading $x^m + y^m = 1$, $u^n + v^n = 1$ is confirmed by that solution. **The problem as set is not true without restriction**, and the MAA solution says so: "The problem is not well set being true only under rather heavy restrictions on the x, y, u and v. For example, all is in order if they are nonnegative." This analysis proves the statement for the first-quadrant arc $x, y \geq 0$ and lines with $u, v \geq 0$, and gives explicit counterexamples to the unrestricted reading, verified numerically: for $m = 3$ the tangent at $(-1, 2^{1/3})$ has $u^{3/2} + v^{3/2} = 3$; for $m = 3/2$ the line $-x + 2^{1/3}y = 1$ satisfies $u^3 + v^3 = 1$ but touches no point of the curve.

Fields: analysis / geometry (tangent lines, implicit curves, conjugate exponents) **Position:** A6 — last on the paper, but the difficulty is in reading the statement, not in the calculus **Answer:** With $n = m/(m-1)$ (the unique n with $\frac{1}{m} + \frac{1}{n} = 1$), and $x, y, u, v \geq 0$: the line $ux + vy = 1$ is tangent to the arc $x^m + y^m = 1$, $x, y \geq 0$, if and only if $u^n + v^n = 1$; the point of tangency is $(u^{n/m}, v^{n/m})$. Without sign restrictions the statement fails in both directions.

Layer 1 THE WHAT — surface analysis

The natural attempt. Tangency means: the line meets the curve at a point (x_0, y_0), and has the curve's slope there. Differentiate $x^m + y^m = 1$ implicitly: $mx^{m-1} + my^{m-1}y' = 0$, so $y' = -x_0^{m-1}/y_0^{m-1}$. The line $ux + vy = 1$ has slope $-u/v$. So you need

$$ux_0 + vy_0 = 1, \qquad x_0^m + y_0^m = 1, \qquad \frac{u}{v} = \frac{x_0^{m-1}}{y_0^{m-1}}.$$

Three equations, and you want to eliminate x_0, y_0. The third says (x_0, y_0) is proportional to $(u^{1/(m-1)}, v^{1/(m-1)})$: put $x_0 = \lambda u^{1/(m-1)}$, $y_0 = \lambda v^{1/(m-1)}$ with $\lambda > 0$. Write $n = m/(m-1)$, so that $u \cdot u^{1/(m-1)} = u^n$ and $\left(u^{1/(m-1)}\right)^m = u^n$. Then

$$\lambda\,(u^n + v^n) = 1, \qquad \lambda^m(u^n + v^n) = 1.$$

Dividing, $\lambda^{m-1} = 1$, so $\lambda = 1$ and $u^n + v^n = 1$. And $n = m/(m-1)$ is the same as $\frac{1}{m} + \frac{1}{n} = 1$.

This works, and it is the honest default. Its costs: it divided by v and by y_0 (so the tangents at $(1,0)$ and $(0,1)$ need separate treatment), it took real $(m-1)$th roots of u and v without asking whether they are positive, and it proved only one direction — the converse, that every (u,v) with $u^n + v^n = 1$ gives a tangent, was never addressed. The sign question is not a technicality: it is where the problem, as printed, is false.

The model chain. **Setting.** Let $m > 1$ and $n = m/(m-1)$, so $n > 1$, $\frac{1}{m} + \frac{1}{n} = 1$, and $(m-1)(n-1) = 1$. Let

$$C = \{(x,y) : x \geq 0,\ y \geq 0,\ x^m + y^m = 1\},$$

and consider lines $ux + vy = 1$ with $u, v \geq 0$. (A line not through the origin has exactly one equation of the form $ux + vy = 1$, so (u,v) is determined by the line.)

Tangent at a point of C. Let $F(x,y) = x^m + y^m$ on $x, y \geq 0$. Because $m > 1$, F has continuous partial derivatives $F_x = mx^{m-1}$, $F_y = my^{m-1}$ there (including on the axes), and at $(x_0, y_0) \in C$ they are not both zero, since $x_0^m + y_0^m = 1$. So the tangent line to C at (x_0, y_0) is

$$x_0^{m-1}(x - x_0) + y_0^{m-1}(y - y_0) = 0, \quad \text{i.e.} \quad x_0^{m-1}x + y_0^{m-1}y = x_0^m + y_0^m = 1.$$

(At $(1,0)$ this is the vertical line $x = 1$, correctly.)

Only if. If $ux + vy = 1$ is tangent to C at (x_0, y_0), then by uniqueness of the coefficients $u = x_0^{m-1}$ and $v = y_0^{m-1}$, both ≥ 0. Since $(m-1)n = m$,

$$u^n + v^n = x_0^{(m-1)n} + y_0^{(m-1)n} = x_0^m + y_0^m = 1.$$

If. Let $u, v \geq 0$ with $u^n + v^n = 1$. Put $x_0 = u^{n-1}$, $y_0 = v^{n-1}$ (note $n - 1 = 1/(m-1)$). Then $x_0, y_0 \geq 0$, $x_0^m = u^{(n-1)m} = u^n$ because $(n-1)m = n$, so $x_0^m + y_0^m = u^n + v^n = 1$ and $(x_0, y_0) \in C$. Also $x_0^{m-1} = u^{(n-1)(m-1)} = u$ and likewise $y_0^{m-1} = v$. So the tangent to C at (x_0, y_0) is $ux + vy = 1$. ■

Why the restriction is needed. Take $m = 3$, so $n = \frac{3}{2}$, and let the curve be the whole real curve $x^3 + y^3 = 1$. The point $(-1, 2^{1/3})$ lies on it, and the tangent there is $x + 2^{2/3}y = 1$. Then $u^{3/2} + v^{3/2} = 1 + 2 = 3 \neq 1$: a tangent line violating the condition. In the other direction take $m = \frac{3}{2}$, so $n = 3$; here $x^{3/2}$ is real only for $x \geq 0$, so the curve is C. The pair $(u,v) = (-1, 2^{1/3})$ satisfies $u^3 + v^3 = 1$, but every tangent to C has $u = x_0^{1/2} \geq 0$, so $-x + 2^{1/3}y = 1$ is not tangent. Both examples appear, in general form, in the MAA solution.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains compute the tangent from the gradient (mx^{m-1}, my^{m-1}). They split at the next line. The natural attempt matches *slopes*, $u/v = x_0^{m-1}/y_0^{m-1}$, and then has to solve for an unknown scale λ. The model writes the tangent line in the normalised form $x_0^{m-1}x + y_0^{m-1}y = 1$ — the same right-hand side as $ux + vy = 1$ — and reads off u and v with no scale to solve for, which also makes the converse a one-line substitution.

> *The given line is normalised with right-hand side* 1. *If I write the tangent with right-hand side* 1 *as well, I can compare coefficients directly instead of comparing slopes and then chasing a constant.*

> *Before proving an "if and only if" about u^n with n possibly fractional, I have to know what signs u and v can have. Let me test the statement on a point with a negative coordinate.*

Layer 3 THE WHY

Why is n, defined by $\frac{1}{m} + \frac{1}{n} = 1$, the exponent that appears? Start from the case you know. For $m = 2$ the curve is the unit circle and the line $ux + vy = 1$ is at distance $1/\sqrt{u^2 + v^2}$ from

the origin; it is tangent exactly when that distance is 1, i.e. $u^2 + v^2 = 1$. So for $m = 2$ the set of tangent lines, recorded as points (u, v), is again the unit circle, and $n = 2$. The problem claims the same for every m, except that the curve of (u, v)'s is the n-curve rather than the m-curve. That is a statement about *duality*: record each tangent line $ux + vy = 1$ by its coefficient point (u, v); the tangent lines of one curve trace out another curve, and here the curve $x^m + y^m = 1$ is sent to $u^n + v^n = 1$.

Where does the conjugate exponent come from? From the coefficients of the tangent: $u = x_0^{m-1}$. To turn a relation among the x_0^m into a relation among the u's you must raise u to the power that turns $m - 1$ into m, namely $\frac{m}{m-1}$. That is all n is. Every property used in the proof — $(m-1)n = m$, $(n-1)m = n$, $(m-1)(n-1) = 1$ — is one rearrangement of $\frac{1}{m} + \frac{1}{n} = 1$, and the proof is nothing but those identities applied to exponents.

There is a second, inequality-based reason, and it explains the geometry. Young's inequality (weighted AM–GM, from concavity of log) says that for $a, b \geq 0$,

$$ab \leq \frac{a^m}{m} + \frac{b^n}{n}, \qquad \text{with equality iff } a^m = b^n.$$

Apply it to ux and vy with $x, y, u, v \geq 0$, $x^m + y^m = 1$, $u^n + v^n = 1$:

$$ux + vy \leq \frac{x^m + y^m}{m} + \frac{u^n + v^n}{n} = \frac{1}{m} + \frac{1}{n} = 1.$$

So the whole arc C lies on the origin side of the line $ux + vy = 1$, and touches it exactly where $x^m = u^n$ and $y^m = v^n$, i.e. at $(u^{n/m}, v^{n/m})$. That is tangency in its original Greek sense — a line that meets the curve and does not cross it — and it is proved with no derivatives at all. The condition $\frac{1}{m} + \frac{1}{n} = 1$ is exactly what makes the two right-hand sides add up to 1. This is Hölder's inequality in its smallest case, and it is the reason the unit balls of the ℓ^m and ℓ^n norms are dual to each other.

Why does the statement break for negative values? Both arguments above use a^m and b^n as *nonnegative* quantities that can be raised to further powers freely: $(x_0^{m-1})^n = x_0^m$ is the rule $(a^r)^s = a^{rs}$, which the MAA solution notes holds for $a \geq 0$. For negative a and fractional exponents that rule fails or the powers are undefined, and the counterexamples in the model chain are exactly where it fails. The setter meant the first-quadrant picture; a solver should say so explicitly and not prove a false statement.

Layer $+\alpha$ BACKTESTING

1. The cue is a line in the normalised form $ux + vy = 1$ and a curve $F(x, y) = 1$. Write down the tangent at (x_0, y_0) in the same normalised form, $F_x(x_0, y_0)\, x + F_y(x_0, y_0)\, y = \text{const}$ divided through by the constant; here $x_0^{m-1}x + y_0^{m-1}y = 1$.

2. Compare coefficients, then use exponent identities derived from $\frac{1}{m} + \frac{1}{n} = 1$. For the converse, invert: $x_0 = u^{n-1}$, $y_0 = v^{n-1}$. Alternatively, Young's inequality proves both directions at once.

3. Hesitation arises over fractional powers of possibly negative numbers. Resolve it by testing a point with a negative coordinate before writing the proof, and then stating the sign restriction.

4. The trap is the statement itself: as printed it is false unless $x, y, u, v \geq 0$. A second trap is proving only "tangent $\Rightarrow u^n + v^n = 1$" and never constructing the point of tangency for the converse. A third is dividing by v or y_0 and losing the tangents at $(1, 0)$ and $(0, 1)$.

5. When a problem pairs exponents m, n with $\frac{1}{m} + \frac{1}{n} = 1$, expect Young's or Hölder's inequality, and expect a duality between an object built from m and one built from n. When a statement involves real powers with non-integer exponents, check signs first.

6. To find the tangent lines of $F = 1$ in the form $ux + vy = 1$, write the tangent in that same form and read off (u, v) as a function of the point; the set of tangent lines is the image of the curve under that map. Before proving any "if and only if", test it on the edge of its domain.

- ▷ **Cue:** a line given as $ux + vy = 1$ and a question about tangency $\longrightarrow$ **Move:** write the tangent at (x_0, y_0) with right-hand side 1 and compare coefficients
- ▷ **Cue:** exponents m, n with $\frac{1}{m} + \frac{1}{n} = 1$ $\longrightarrow$ **Move:** use $(m-1)n = m$ and $(m-1)(n-1) = 1$; think Young/Hölder: $ab \le a^m/m + b^n/n$
- ▷ **Cue:** a curve lies on one side of a line and meets it at one point $\longrightarrow$ **Move:** prove tangency by an inequality with its equality case rather than by derivatives
- ▷ **Cue:** non-integer powers of variables whose signs are not stated $\longrightarrow$ **Move:** test a negative value before proving anything; restrict to the nonnegative case explicitly if needed

Putnam 1965 — Paper B

1965 B1

1965 B1. Evaluate

$$\lim_{n\to\infty} \underbrace{\int_0^1 \int_0^1 \cdots \int_0^1}_{n} \cos^2\Big(\frac{\pi}{2n}(x_1 + x_2 + \cdots + x_n)\Big)\, dx_1\, dx_2 \cdots dx_n.$$

Fields: Real analysis, multiple integrals, symmetry arguments **Position:** B1 (opener) **Answer:** The limit is $\frac{1}{2}$ – and in fact the n-fold integral equals $\frac{1}{2}$ exactly, for every n, not merely in the limit.

Caveat. The MAA OCR drops the displayed formula for this problem entirely – the scanned page reads only "B-1. Evaluate" with a blank where the integral should be. The statement above is reconstructed by working backward from the book's own solution to "B-1," which begins "The change of variables $x \mapsto 1 - x$ yields $\int \cdots \cos^2\{\frac{\pi}{2n}(x_1 + \cdots + x_n)\}\, dx_1 \cdots dx_n = \int \cdots \sin^2\{\cdots\} \cdots$" and concludes "Each of these expressions ... must equal $\frac{1}{2}$. The limit is also $\frac{1}{2}$." This pins the integrand, the domain, and the scaling factor $\pi/2n$ uniquely; I solved the reconstructed problem independently (below) and it reproduces the book's solution line for line, which is the strongest available check that the reconstruction is correct.

Layer 1 THE WHAT — surface analysis

The natural attempt. The object in front of you is an n-fold integral of a function that depends on the x_i only through their sum $s = x_1 + \cdots + x_n$, followed by a limit as $n \to \infty$. The most direct way in is to actually compute the integral for fixed n and then take the limit, and the standard tool for integrating a cosine of a linear combination of independent variables is to pass to a complex exponential, because $\cos\theta = \mathrm{Re}(e^{i\theta})$ turns a sum in the exponent into a *product* over the variables:

$$\int_0^1 \cdots \int_0^1 e^{i\frac{\pi}{n}s}\, dx_1 \cdots dx_n = \left(\int_0^1 e^{i\pi x/n}\, dx\right)^n,$$

since $e^{i\pi s/n} = \prod_j e^{i\pi x_j/n}$ and the integral factors variable by variable. Writing $\cos^2\theta = \frac{1+\cos 2\theta}{2}$ first, the problem reduces to evaluating $\left(\int_0^1 e^{i\pi x/n} dx\right)^n$ and taking its real part.

The single integral is elementary: $\int_0^1 e^{i\pi x/n} dx = \frac{n}{i\pi}\left(e^{i\pi/n} - 1\right)$. Writing $e^{i\pi/n} - 1 = 2i\sin(\pi/2n)\, e^{i\pi/2n}$ (a standard half-angle manipulation), this becomes $\frac{2n}{\pi}\sin(\pi/2n)\, e^{i\pi/2n}$. Raising to the nth power,

$$\left(\int_0^1 e^{i\pi x/n} dx\right)^n = \left(\frac{2n}{\pi}\sin\frac{\pi}{2n}\right)^n e^{i\pi/2}.$$

Now you need $\lim_{n\to\infty}\left(\frac{2n}{\pi}\sin\frac{\pi}{2n}\right)^n$. Since $\sin u = u - \frac{u^3}{6} + O(u^5)$ with $u = \pi/2n$, the base is $1 - \frac{\pi^2}{24n^2} + O(n^{-4})$, so the whole bracket is $\left(1 - \frac{c}{n^2} + \cdots\right)^n$ for a constant $c = \pi^2/24$. This

is where the natural attempt becomes fragile: it is tempting to treat this like the classical limit $(1 + a/n)^n \to e^a$, but the correction here is $O(1/n^2)$, not $O(1/n)$ – the exponent n is being applied to a base that is 1 plus something of order $1/n^2$, and $n \cdot (1/n^2) \to 0$, so the bracket tends to $e^0 = 1$, not to some nontrivial exponential. Carrying this out correctly gives $\left(\int_0^1 e^{i\pi x/n} dx\right)^n \to 1 \cdot e^{i\pi/2} = i$, whose real part is 0. So the cosine term vanishes in the limit and $I_n = \frac{1}{2} + \frac{1}{2}\,\mathrm{Re}(\cdots) \to \frac{1}{2} + 0 = \frac{1}{2}$. This chain is completely legitimate and does reach the right answer – but it computes the exact value of I_n for every n via a genuinely delicate asymptotic expansion, when the limit in the question turns out not to need $n \to \infty$ at all.

The model chain. Write I_n for the integral. Apply the substitution $y_i = 1 - x_i$ to every variable simultaneously; this is measure-preserving on $[0,1]^n$ (it is just a relabelling of each coordinate axis), so it leaves the value of the integral unchanged while transforming the integrand. Since $\sum y_i = n - \sum x_i$,

$$\frac{\pi}{2n}\sum y_i = \frac{\pi}{2} - \frac{\pi}{2n}\sum x_i,$$

and $\cos^2\left(\frac{\pi}{2} - \theta\right) = \sin^2\theta$. Hence the substitution turns the integrand $\cos^2(\cdots)$ into $\sin^2(\cdots)$, while the value of the integral (just a relabelled dummy variable) cannot change:

$$I_n = \int\cdots\int \cos^2\left(\frac{\pi}{2n}\sum x_i\right) dx = \int\cdots\int \sin^2\left(\frac{\pi}{2n}\sum x_i\right) dx =: J_n.$$

Now add the two integrals, using $\cos^2\theta + \sin^2\theta = 1$ pointwise on the entire domain:

$$I_n + J_n = \int_0^1\cdots\int_0^1 1\, dx_1\cdots dx_n = 1.$$

Since $I_n = J_n$, this forces $2I_n = 1$, i.e. $I_n = \frac{1}{2}$ – exactly, for every positive integer n, with no asymptotics required at all. In particular the limit as $n \to \infty$ is $\frac{1}{2}$.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains correctly reduce the problem via $\cos^2\theta = \frac{1+\cos 2\theta}{2}$ (or, equivalently, both recognise that the integrand depends only on $\sum x_i$). They split at what to do with that reduced integral. The natural attempt treats "evaluate then take a limit" as two separate jobs and reaches for the general-purpose machine – complex exponentials – to get an exact closed form for I_n, planning to send $n \to \infty$ afterwards. The model chain asks a different question first: is there a symmetry of the box $[0,1]^n$ that acts on this particular integrand in a useful way – and finds that $x_i \mapsto 1 - x_i$ swaps $\cos^2 \leftrightarrow \sin^2$ exactly, collapsing "integrate, then take a limit" into "add two copies of the same unknown integral and solve a one-line equation for it," with n never appearing in the answer at all.

> *I don't actually need a formula for I_n as a function of n – I need its limit. If I can find a symmetry that pins down I_n to a single value independent of n in one step, the limit is that value trivially, and I never have to touch an asymptotic expansion at all.*

> *The integrand is a function of $\sum x_i$ alone, and the domain $[0,1]^n$ has an obvious symmetry that acts on that sum in a predictable way: reflecting every coordinate $x_i \mapsto 1 - x_i$ sends $\sum x_i \mapsto n - \sum x_i$, which is exactly the reflection $\theta \mapsto \frac{\pi}{2} - \theta$ that swaps* $\cos^2$ *and* $\sin^2$*. A substitution that turns my integrand into a different, complementary integrand over the* same *domain is a signal to add the two and use* $\cos^2 + \sin^2 = 1$*, not to compute either one directly.*

▷ **Cue:** an integral of a symmetric function of several variables over a symmetric domain (here, an integrand depending only on $x_1 + \cdots + x_n$ over the cube $[0,1]^n$), especially when the

answer is asked for only in a limit $\longrightarrow$ **Move:** look for a domain symmetry (reflection, permutation, complementation) that maps the integrand to a recognisably related one – here $x_i \mapsto 1 - x_i$ turning $\cos^2 \to \sin^2$ – add the two resulting equal integrals and use a pointwise identity ($\cos^2 + \sin^2 = 1$) to solve for the value in one line, before reaching for series expansions or asymptotics

▷ **Cue:** a limit of a quantity that a symmetry argument shows is exactly constant in the parameter being sent to infinity $\longrightarrow$ **Move:** prove the exact, n-independent value first; the limit is then immediate and no asymptotic estimate (Taylor expansion, L'Hopital, dominant-term analysis) is needed at all

Layer 3 THE WHY

The natural attempt fails to be the efficient path – though it is not wrong – because it commits to *computing* before asking whether the quantity needs to be computed at all. Complex exponentials are the right tool when you actually need the value of I_n as a function of n; here the problem only asks for a limit, which is a much weaker demand, and a weaker demand should be met with a correspondingly lighter argument if one is available. The half-angle identity $2i \sin(\pi/2n) e^{i\pi/2n} = e^{i\pi/n} - 1$ used in the natural attempt is itself already circling the actual content of the symmetry – the factor $e^{i\pi/2n}$ is precisely the "reflection about the midpoint $\frac{1}{2}$" that the model chain exploits directly, but the natural attempt buries it inside a magnitude-and-phase computation instead of using it as a substitution.

Why is $x_i \mapsto 1 - x_i$ the natural move once you see it? Because the domain $[0, 1]$ has exactly one nontrivial measure-preserving involution – reflection about its midpoint $\frac{1}{2}$ – and the problem's constant $\pi/2n$ is precisely calibrated so that this reflection acts on the phase $\frac{\pi}{2n} \sum x_i$ as a shift by $\frac{\pi}{2}$, the exact amount needed to swap cosine and sine. This is not a coincidence to be marvelled at; it is what "the coefficient in front of $\sum x_i$ is $\pi/2n$, and there are n variables each ranging over an interval of length 1" is *for*: $\frac{\pi}{2n} \cdot n \cdot 1 = \frac{\pi}{2}$ is exactly the phase swept out by $\sum x_i$ ranging from its minimum (0) to its maximum (n), and a linear phase that sweeps out exactly $\pi/2$ over the full range of the domain is the signature of a problem built for the reflection trick. Recognising "the total range of the phase is $\pi/2$, a quarter period of $\cos^2$" as the cue – rather than treating $\pi/2n$ as an arbitrary constant to be pushed through a limit – is what the book's one-line solution is doing, and it is worth naming explicitly rather than admiring silently.

The identity $\cos^2\theta + \sin^2\theta = 1$ is doing all the remaining work, and it is worth being honest about why it applies pointwise inside a multiple integral: it holds for *every* point $(x_1, \ldots, x_n)$ in the cube, not merely on average, so integrating it over the cube gives exactly 1 (the volume of the cube) with no approximation anywhere. The model chain therefore has no error term to control, in contrast to the natural attempt's $O(1/n^2)$ correction that has to be shown to vanish under the nth power – a step that is a well-known trap (see Layer $+\alpha$).

Layer $+\alpha$ BACKTESTING

1. The first useful cue is that the integrand depends on the n variables only through their sum $\sum x_i$, over the symmetric cube $[0, 1]^n$. On seeing it, write down: "does a coordinate symmetry of the cube act simply on this sum?" before reaching for any integration technique.

2. That cue suggests checking the reflection $x_i \mapsto 1 - x_i$ (the only nontrivial measure-preserving involution of $[0, 1]$ applied coordinatewise), and checking what it does to the phase $\frac{\pi}{2n} \sum x_i$ – here, a shift by exactly $\pi/2$, converting $\cos^2$ to $\sin^2$.

3. The hesitation point is right after reducing via $\cos^2\theta = \frac{1+\cos 2\theta}{2}$: it feels natural to keep going and actually integrate the cosine term via complex exponentials, because that is a reliable, general method you already trust. It is reinforced away by asking explicitly, before computing anything: "do I actually need I_n for every n, or only its limit?" – and noticing that a symmetry argument, if one exists, answers the weaker question in one line.

4. The trap is the $n \to \infty$ in the problem statement: it invites you to think asymptotically from the first line, when in fact the exact value is n-independent and the limit is a red herring

designed to make the symmetry solution look more surprising than it is. A second, more technical trap sits inside the natural attempt itself: mistaking the $O(1/n^2)$ correction to $\left(\frac{2n}{\pi}\sin\frac{\pi}{2n}\right)^n$ for an $O(1/n)$ correction would send the limit to $e^{-\pi^2/24}$ instead of 1 – an easy sign of not tracking the order of the Taylor remainder carefully enough.

5. To recognise this faster: whenever a multi-dimensional integral's integrand is a function of a single linear statistic (a sum, or more generally $a \cdot x$) over a box, and the box has an obvious coordinatewise symmetry, check what that symmetry does to the statistic before setting up any computation – it is often designed to relate the integral to itself or to a complementary integral.

6. Decision rule: when asked for a limit (or any weaker fact) about a quantity defined by an integral or sum with a manifest symmetry, look for an identity that pins the quantity down exactly and independently of the parameter first (via a symmetry-induced pairing plus a pointwise algebraic identity), and only fall back on direct computation and asymptotics if no such identity is found – solving the exact problem is frequently easier than estimating the limit of an approximate one.

1965 B2

1965 B2. In a round-robin tournament with n players $P_1, P_2, \ldots, P_n$ (where $n > 1$), each player plays one game with each of the other players and the rules are such that no ties can occur. Let w_r and l_r be the number of games won and lost, respectively, by P_r. Show that

$$\sum_{r=1}^{n} w_r^2 = \sum_{r=1}^{n} l_r^2.$$

Fields: Combinatorics, algebraic identities **Position:** B2 (mid-paper) **Answer:** The identity holds for every $n > 1$ and every possible tournament outcome; $\sum w_r^2 = \sum l_r^2$ always.

Layer 1 THE WHAT — surface analysis

The natural attempt. You want to compare $\sum w_r^2$ and $\sum l_r^2$, and the most direct route into "w_r^2" is to write w_r itself as a sum of indicators over the games P_r plays: $w_r = \sum_{s \neq r} [P_r \text{ beats } P_s]$, where $[\cdot]$ is 1 if true and 0 otherwise. Squaring,

$$w_r^2 = \sum_{s \neq r} \sum_{t \neq r} [P_r \text{ beats } P_s] \cdot [P_r \text{ beats } P_t],$$

and summing over r gives a triple sum over ordered triples (r, s, t) with $s, t \neq r$. You would then want to do the identical expansion for l_r^2 and show the two triple sums agree. This is possible in principle – split each triple sum into the "diagonal" terms $s = t$ (which just recover $\sum_r w_r = \sum_r l_r$, itself needing a short argument) and the "off-diagonal" terms $s \neq t$ (which require a case analysis on how P_s and P_t each fare against P_r, and whether s, t beat each other, none of which is actually relevant to the claim). This is a completely legitimate way to grind out the identity, but it is drowning a two-line algebraic fact in bookkeeping over ordered triples of players – the indicator expansion is exactly the right tool for extracting w_r from raw game outcomes, and exactly the wrong tool for comparing w_r^2 against l_r^2, because once you square, you have thrown away the one clean relationship between w_r and l_r that the problem is built around, before you ever use it.

The model chain. Every player P_r plays exactly $n - 1$ games (one against each other player), and each game is won or lost by P_r with no ties, so

$$w_r + l_r = n - 1 \qquad \text{for every } r.$$

Also, every game played has exactly one winner and one loser, so summing over all games, the total number of wins across all players equals the total number of losses across all players (both count the same set of games, once each):

$$\sum_{r=1}^{n} w_r = \sum_{r=1}^{n} l_r.$$

Now use the difference-of-squares factorisation on each term before summing:

$$\sum_{r=1}^{n} w_r^2 - \sum_{r=1}^{n} l_r^2 = \sum_{r=1}^{n} (w_r^2 - l_r^2) = \sum_{r=1}^{n} (w_r - l_r)(w_r + l_r).$$

Since $w_r + l_r = n - 1$ is the *same constant for every* r, it factors out of the sum:

$$\sum_{r=1}^{n} (w_r - l_r)(w_r + l_r) = (n-1) \sum_{r=1}^{n} (w_r - l_r) = (n-1) \left(\sum_r w_r - \sum_r l_r \right) = (n-1) \cdot 0 = 0.$$

Hence $\sum w_r^2 = \sum l_r^2$. ■

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same raw material – w_r and l_r as counts of games won and lost by each player – and both, in some form, need the two basic facts $w_r + l_r = n - 1$ and $\sum w_r = \sum l_r$. They split at how w_r^2 is handled. The natural attempt expands w_r^2 down to the level of individual games (an indicator sum over pairs of opponents), squaring *before* relating w_r to l_r. The model chain instead applies $a^2 - b^2 = (a-b)(a+b)$ to $w_r^2 - l_r^2$ *as a pair*, algebraically, before ever touching what w_r and l_r mean in terms of individual games – so it only ever needs the two summary facts about them, never the game-by-game structure.

> *I am being asked to compare two sums of squares, $\sum w_r^2$ and $\sum l_r^2$ – but I already know a simple linear relationship between w_r and l_r individually ($w_r + l_r = n - 1$) and between their totals ($\sum w_r = \sum l_r$). Rather than expanding each square back down to the games that produced it, I should factor $w_r^2 - l_r^2$ as $(w_r - l_r)(w_r + l_r)$ first, so that the linear facts I already have can be plugged in directly.*

> *The quantity $w_r + l_r$ is not just some expression – it is the constant $n - 1$, the same for every single player, because everyone plays everyone else exactly once. A term that would otherwise vary with r inside a sum, but is secretly constant, should be pulled out in front of the summation sign immediately; that is what turns the whole right-hand side into $(n-1)$ times a sum I already know is zero.*

- ▷ **Cue:** a sum of squared quantities $\sum a_r^2$ is to be compared with another $\sum b_r^2$, and a_r, b_r satisfy a simple linear relation (individually, and/or in total) $\longrightarrow$ **Move:** apply $a_r^2 - b_r^2 = (a_r - b_r)(a_r + b_r)$ term by term before expanding either a_r or b_r into their finer-grained definitions; if $a_r + b_r$ (or $a_r - b_r$) turns out to be constant across the index, pull it out of the sum and reduce the problem to the simpler linear fact
- ▷ **Cue:** each of n objects is compared against every other exactly once (a round-robin structure), producing two counts per object (e.g. wins/losses, in-degree/out-degree) that sum to the fixed constant $n - 1$ $\longrightarrow$ **Move:** record "$\text{count}_1 + \text{count}_2 = n - 1$, the same for every object" as an immediately available algebraic fact, and look for places (difference of squares, generating sums) where a quantity constant across the index can be factored out of a summation

Layer 3 THE WHY

The natural attempt fails to be efficient because it treats "w_r" as fundamentally a description of a set of games, and reaches for that description the moment a nonlinear operation (squaring)

is applied – as if nonlinearity always demands going back to first principles. But the problem supplies two facts about w_r and l_r that are already at exactly the right level of abstraction to answer a question about $\sum w_r^2$ versus $\sum l_r^2$: a per-player linear relation ($w_r + l_r = n - 1$) and a global linear relation ($\sum w_r = \sum l_r$). The identity $a^2 - b^2 = (a - b)(a + b)$ exists precisely to convert a quadratic difference into a product of two linear expressions – so it is the tool that meets linear information exactly where it lives, without demanding any information (like which specific players beat which) that the problem does not need and was never asked to use.

Why is $w_r + l_r = n - 1$ the fact that matters, rather than, say, w_r or l_r individually? Because a round-robin tournament's defining structural feature is that *every* player plays *every other* player exactly once – this is what makes $w_r + l_r$ constant across r in the first place, rather than an arbitrary number that happens to vary from player to player. Once you see that $(w_r - l_r)(w_r + l_r)$ has a factor that does not depend on r at all, summing over r becomes trivial: constants pull outside sums, and what remains, $\sum(w_r - l_r)$, is answered by the second fact, that total wins equal total losses – which is itself just "every game has one winner and one loser," restated as a sum. Both ingredients are one-line consequences of the words "round-robin, no ties," read for their algebraic content rather than their combinatorial content; the model solution's economy comes entirely from noticing that this problem does not actually require knowing *who* beat *whom*, only these two aggregate constraints.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the phrase "each player plays one game with each of the other players," which pins down $w_r + l_r = n - 1$ as a constant independent of r. On seeing it, write that identity down immediately, before looking at what is being asked about w_r, l_r.

2. That cue, combined with the target expression $\sum w_r^2 = \sum l_r^2$ (a comparison of two sums of squares), suggests factoring the difference $w_r^2 - l_r^2 = (w_r - l_r)(w_r + l_r)$ rather than expanding either square from its game-by-game definition.

3. The hesitation point is right after writing $w_r = \sum_{s \neq r}[\text{beats}]$: expanding an unfamiliar quantity into indicator sums over games feels like solid, unimpeachable progress, because it is always available and always correct. It is reinforced away by asking, before expanding: "what algebraic identity exists specifically for comparing two sums of squares?" – the answer, difference of squares, is a standing tool that should be tried before any case-by-case expansion.

4. The trap is that squaring an indicator sum *feels* like it demands returning to the indicators to make sense of the square (cross terms, diagonal terms, and so on) – but the problem only ever needed $w_r + l_r$ and $\sum(w_r - l_r)$, both linear, and factoring first avoids ever generating a cross term.

5. To recognise this faster: whenever a problem compares $\sum a_r^k$ against $\sum b_r^k$ for some power $k > 1$, and the problem's hypotheses give you a linear relationship between a_r and b_r (individually or in aggregate), try factoring $a_r^k - b_r^k$ into linear factors first (difference of squares for $k = 2$, difference of cubes for $k = 3$, etc.) before expanding a_r or b_r into anything finer-grained.

6. Decision rule: when a problem gives you a fact at a certain level of abstraction (here: a linear relation between two per-object quantities, plus a linear relation between their totals) and asks a question one algebraic degree higher (here: quadratic), reach first for the factorisation that converts the higher-degree question into the lower-degree one you already have an answer to, rather than dropping down to a more granular description of the objects involved.

1965 B3

1965 B3. Prove that there are exactly three right-angled triangles whose sides are integers while the area is numerically equal to twice the perimeter.

Fields: Number theory, Pythagorean triples, Diophantine equations **Position:** B3 (mid-paper) **Answer:** Exactly three such triangles exist: $(9, 40, 41)$, $(10, 24, 26)$, and $(12, 16, 20)$.

Layer 1 THE WHAT — surface analysis

The natural attempt. The condition "area equals twice the perimeter" is easy to write down for a right triangle with legs a, b and hypotenuse c: $\frac{1}{2}ab = 2(a + b + c)$. The most immediately available strategy, given that you already know a good supply of small Pythagorean triples, is to test them one by one: $(3, 4, 5)$ gives area 6, perimeter 12, $2\times$ perimeter $= 24 \neq 6$; $(6, 8, 10)$ gives area 24, perimeter 24, $2\times$ perimeter $= 48 \neq 24$; $(5, 12, 13)$ gives area 30, perimeter 30, needs 60; $(9, 12, 15)$ gives area 54, perimeter 36, needs 72; $(8, 15, 17)$ gives area 60, perimeter 40, needs 80; $(12, 16, 20)$ gives area 96, perimeter 48, needs 96 – a hit. Continuing through $(7, 24, 25)$, $(20, 21, 29)$, $(10, 24, 26)$ (another hit, area 120, perimeter 60, needs 120), $(9, 40, 41)$ (area 180, perimeter 90, needs 180, a third hit), and onward through triples with ever-larger hypotenuses, you start to suspect the three you have found are all of them.

This search is completely valid as a way to *find* candidates, and it does locate all three answers. But it cannot, by itself, *prove* there are exactly three, because the list of Pythagorean triples is infinite, and nothing about testing triples up to some hypotenuse bound tells you that no larger triple satisfies the condition – you would need either an a priori bound on the hypotenuse, or an argument that the condition becomes impossible to satisfy once the legs are large enough. Checking triples by hand does not supply that bound; it just keeps not finding a fourth hit, which is evidence, not proof. "Prove that there are exactly three" is a claim about completeness over an infinite search space, and finite case-checking, however far it is carried, does not by itself close that space off.

The model chain. Every primitive Pythagorean triple (one with $\gcd(a, b, c) = 1$) has the form

$$a = m^2 - n^2, \quad b = 2mn, \quad c = m^2 + n^2,$$

for integers $m > n \geq 1$ with $\gcd(m, n) = 1$ and m, n of opposite parity; every Pythagorean triple is k times a primitive one, for some positive integer k. So write the general triple as

$$a = k(m^2 - n^2), \qquad b = 2kmn, \qquad c = k(m^2 + n^2).$$

Compute both sides of the condition $ab = 4(a + b + c)$ (equivalent to $\frac{1}{2}ab = 2(a + b + c)$) in terms of k, m, n:

$$ab = k(m^2 - n^2) \cdot 2kmn = 2k^2mn(m - n)(m + n),$$

$$a + b + c = k(m^2 - n^2) + 2kmn + k(m^2 + n^2) = k\left(2m^2 + 2mn\right) = 2km(m + n).$$

The condition $ab = 4(a + b + c)$ becomes

$$2k^2mn(m - n)(m + n) = 4 \cdot 2km(m + n) = 8km(m + n).$$

Both sides are positive, so $m, n, m + n$ are all nonzero and can be cancelled: divide by $2km(m + n)$ to get

$$k\,n\,(m - n) = 4.$$

This is now a finite problem: $k, n, m - n$ are positive integers whose product is 4. But the parametrization requires m, n to have opposite parity, which forces $m - n$ to be *odd*. The only odd divisor of 4 is 1, so $m - n = 1$ is forced, and then $kn = 4$, giving exactly three factorizations: $(k, n) \in \{(4, 1), (2, 2), (1, 4)\}$, with $m = n + 1$ in each case:

- $n = 1$, $m = 2$, $k = 4$: $\gcd(2,1) = 1$, opposite parity. $a = 4(4-1) = 12$, $b = 4 \cdot 2 \cdot 2 \cdot 1 = 16$, $c = 4(4+1) = 20$.
- $n = 2$, $m = 3$, $k = 2$: $\gcd(3,2) = 1$, opposite parity. $a = 2(9-4) = 10$, $b = 2 \cdot 2 \cdot 3 \cdot 2 = 24$, $c = 2(9+4) = 26$.
- $n = 4$, $m = 5$, $k = 1$: $\gcd(5,4) = 1$, opposite parity. $a = 25 - 16 = 9$, $b = 2 \cdot 5 \cdot 4 = 40$, $c = 25 + 16 = 41$.

Each of $\gcd(m,n) = 1$ and opposite parity is verified directly in every case, so all three parameter triples do correspond to genuine Pythagorean triples, and every solution of $kn(m-n) = 4$ with the required parity has been enumerated exhaustively (three positive-integer factorizations of 4, one of which, $m - n = 2$, is excluded by parity, leaving exactly these three). Hence there are exactly three right triangles with integer sides and area equal to twice the perimeter: $(12,16,20)$, $(10,24,26)$, $(9,40,41)$. ■

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from the same equation, $\frac{1}{2}ab = 2(a+b+c)$, and both correctly identify small Pythagorean triples as the objects to test. They split at how the infinitude of Pythagorean triples is handled. The natural attempt tests specific triples one at a time, forever leaving open the possibility that some untested, larger triple also works. The model chain instead substitutes the *general parametric formula* for all Pythagorean triples into the condition *before* testing anything, turning "search an infinite set of triples for a property" into "solve a single equation, $kn(m-n) = 4$, in the finitely many parameters that generate that infinite set" – at which point the infinitude of triples has been absorbed into finitely many divisor choices, and completeness is automatic.

> *Checking triples one at a time can only ever produce a list of hits, never a certificate that the list is exhausted, because the space of Pythagorean triples is infinite. If I want to* prove *exactly three exist, I need to work with the object that generates* all *Pythagorean triples at once – the (m,n,k) parametrization – and turn the area/perimeter condition into an equation in m,n,k themselves, so that "how many triples satisfy this" becomes "how many integer solutions does this equation have," which I can actually bound.*

> *After substituting the parametrization, the condition collapsed to $kn(m-n) = 4$ – a product of three positive integers equal to 4, which has only finitely many factorizations no matter how large m,n,k might have looked before I imposed the constraint. This is the moment the infinite search space becomes finite: I should actively watch for a substitution that turns "for which of infinitely many objects does property X hold" into "how many ways can a fixed small integer be factored," because the latter is always checkable by hand.*

▷ **Cue:** a problem asks to find or count all integer-sided right triangles (Pythagorean triples) satisfying some additional numerical condition, especially phrased "prove there are exactly k..." $\longrightarrow$ **Move:** substitute the general parametrization $a = k(m^2 - n^2)$, $b = 2kmn$, $c = k(m^2 + n^2)$ ($\gcd(m,n) = 1$, $m > n$, opposite parity, $k \geq 1$) into the condition before testing any specific triples; this converts an open-ended search over an infinite family into a Diophantine equation in the finitely-appearing parameters k, m, n, which can then be solved by factoring a fixed integer

▷ **Cue:** a Diophantine equation reduces, after algebraic simplification, to a product of positive-integer unknowns equalling a fixed small constant (e.g. $k\,n\,(m-n) = 4$) $\longrightarrow$ **Move:** enumerate the finitely many factorizations of that constant directly, and check any side

conditions (parity, coprimality, ordering) on each factorization individually, rather than searching the original unknowns

Layer 3 THE WHY

The natural attempt fails to prove completeness because "check triples in increasing order until you stop finding hits" has no stopping criterion built into it – it is testing a property against an infinite set, one element at a time, and no finite amount of that testing logically forecloses the rest of the set. The only way to assert "exactly three" with a right to the word *exactly* is to replace "search over triples" with "solve an equation with a provably finite solution set," and the tool for doing that here is already sitting inside the definition of a Pythagorean triple: every one of them, without exception, is generated by exactly one triple of parameters (k, m, n) satisfying $\gcd(m, n) = 1$ and opposite parity. This generating formula is not a shortcut invented for this problem; it *is* what a Pythagorean triple is, structurally, and any condition imposed on (a, b, c) is therefore automatically a condition on (k, m, n) – the only question is whether that condition, once rewritten, is easy to solve.

Why does $ab = 4(a + b + c)$ collapse so cleanly? Because both ab and $a + b + c$, when expressed in m, n, k, are built from the same two factors $m^2 - n^2 = (m - n)(m + n)$ and $2mn$, so their ratio is forced to simplify: $a + b + c = k \cdot 2m(m + n)$ contains the factor $(m + n)$ that also appears inside $ab = 2k^2mn(m - n)(m + n)$, and once that common factor of $m(m + n)$ cancels from both sides, what remains, $kn(m - n) = 4$, is linear in each of k, n, and $(m - n)$ separately even though the original condition was quadratic in a, b. This is the general phenomenon behind why "area equals a linear multiple of the perimeter" problems for Pythagorean triples always reduce to small, finite Diophantine conditions: area is degree 2 in the legs, perimeter is degree 1, but in the (m, n, k) parametrization the legs themselves are degree 2 in (m, n) and degree 1 in k, so "area / perimeter is bounded" becomes "a specific product of parameters is bounded" – structurally guaranteed to be a finite search once you have switched to the parameters that generate the triples rather than the triples themselves. The parity condition (that m, n have opposite parity, hence $m - n$ is odd) is not a technicality to be checked afterwards – it is what cuts the three abstract factorizations of 4 (as 1×4, 2×2, 4×1 along the $n(m - n)$ split, further times k) down correctly, ruling out $m - n = 2$ and $m - n = 4$ before they can produce spurious "solutions" that are not actually valid triple-generators.

Layer $+\alpha$ BACKTESTING

1. The first useful cue is the word "exactly" in "prove that there are exactly three" combined with the fact that the underlying objects (Pythagorean triples) form an infinite family. On seeing it, write down: a finite case search by itself cannot prove completeness here; I need a parametrization of the whole infinite family and an equation in its parameters.

2. That cue suggests substituting the standard (m, n, k) generating formula for Pythagorean triples into the given condition ($ab = 4(a + b + c)$) before testing any numbers, so the condition becomes an equation in m, n, k rather than a property to be checked triple by triple.

3. The hesitation point is right after finding the first hit, $(12, 16, 20)$, by trial: it feels like real progress and invites continuing the same search a bit further to "make sure." It is reinforced away by asking explicitly: "what would make me certain there is no fourth triple, no matter how large?" – trial search cannot answer that question in principle, so it is a signal to switch to the parametrized equation.

4. The trap is mistaking a long, unsuccessful further search (checking triples with larger and larger hypotenuses and not finding a fourth hit) for a proof of completeness; absence of a counterexample after finite effort is evidence, not a proof, when the search space is infinite. A second, smaller trap is forgetting the opposite-parity requirement on m, n when enumerating factorizations of $4 = k \cdot n \cdot (m - n)$, which would let $m - n = 2$ or $m - n = 4$ slip in as spurious extra "triples."

5. To recognise this faster: whenever a problem says "prove there are exactly k" integer solutions to a condition involving Pythagorean triples (or any other family with a known complete parametrization – e.g. Gaussian integers, Pell equation solutions), reach immediately for that parametrization and substitute, rather than beginning with trial and error on small cases.

6. Decision rule: to prove finiteness or an exact count of integer objects satisfying a condition, when the objects belong to a family with a known complete generating parametrization, substitute that parametrization into the condition first; this converts an unbounded search into a Diophantine equation whose solution set can be bounded and enumerated explicitly (often by factoring a fixed integer), which is the only way to earn the word "exactly."

1965 B4

1965 B4. Consider the function

$$f(x,n) = \frac{\binom{n}{0} + \binom{n}{2}x + \binom{n}{4}x^2 + \cdots}{\binom{n}{1} + \binom{n}{3}x + \binom{n}{5}x^2 + \cdots}$$

where n is a positive integer. Express $f(x,n+1)$ rationally in terms of $f(x,n)$ and x. Hence, or otherwise, evaluate $\lim_{n\to\infty} f(x,n)$ for suitable fixed values of x. (The symbols $\binom{n}{r}$ represent the binomial coefficients.)

Caveat. The MAA OCR drops the binomial symbols in the displayed fraction. The statement above is reconstructed from the book's own solution, which uses $f(x,n+1) = \dfrac{f(x,n)+x}{f(x,n)+1}$ and $f(x,n) = \sqrt{x}\,\dfrac{(1+\sqrt{x})^n + (1-\sqrt{x})^n}{(1+\sqrt{x})^n - (1-\sqrt{x})^n}$; both hold exactly for the even-index numerator and odd-index denominator shown, and were checked numerically for several x and n.

Fields: Algebra (binomial sums, recurrences), analysis (limits of iterated Möbius maps) **Position:** B4 (upper middle) **Answer:** $f(x,n+1) = \dfrac{f(x,n)+x}{f(x,n)+1}$; for real $x \geq 0$ the limit is $\sqrt{x}$, for real $x < 0$ there is no limit (over $\mathbb{C}$: the limit is the square root of x with positive real part whenever $x \notin (-\infty, 0)$).

Layer 1 THE WHAT — surface analysis

The natural attempt. Write $E_n = \sum_k \binom{n}{2k}x^k$ and $O_n = \sum_k \binom{n}{2k+1}x^k$, so $f(x,n) = E_n/O_n$. Pascal's rule $\binom{n+1}{r} = \binom{n}{r} + \binom{n}{r-1}$ splits each coefficient of E_{n+1} and O_{n+1}:

$$E_{n+1} = \sum_k \binom{n}{2k}x^k + \sum_{k\geq 1} \binom{n}{2k-1}x^k = E_n + xO_n, \qquad O_{n+1} = O_n + E_n.$$

Dividing, $f(x,n+1) = \dfrac{E_n + xO_n}{E_n + O_n} = \dfrac{f(x,n)+x}{f(x,n)+1}$. That part goes through for almost everyone.

Then the standard reflex: "if the limit L exists, pass to the limit in the recursion", giving $L = (L+x)/(L+1)$, so $L^2 = x$, so $L = \sqrt{x}$. Many scripts stop here and write "the limit is $\sqrt{x}$". That is not a proof. It shows only that *if* a limit exists and is not -1, it is one of $\pm\sqrt{x}$. It says nothing about existence, nothing about which sign, and nothing about what "suitable" x means. The student who notices this then tries to prove convergence directly: is the map $g(F) = (F+x)/(F+1)$ a contraction? Its derivative is $(1-x)/(F+1)^2$, which for large x and small F is bigger than 1 in size (at $x = 50$ the iterates go 1, 25.5, 2.85, 13.73, 4.33, . . . , bouncing

on both sides of $\sqrt{50}$). A monotone-convergence argument fails because the sequence oscillates. The attempt stalls with a correct recursion, a correct guess, and no proof.

The model chain. *Recursion.* As above, $E_{n+1} = E_n + xO_n$ and $O_{n+1} = O_n + E_n$, so $f(x, n+1) = \dfrac{f(x,n)+x}{f(x,n)+1}$, with $f(x,1) = \binom{1}{0}/\binom{1}{1} = 1$.

Case $x = 0$. $f(0,n) = \binom{n}{0}/\binom{n}{1} = 1/n \to 0 = \sqrt{0}$.

Case $x > 0$. Put $s = \sqrt{x} > 0$. The fixed points of $g(F) = (F+x)/(F+1)$ are $\pm s$, and

$$g(F) - s = \frac{F + s^2 - sF - s}{F+1} = \frac{(1-s)(F-s)}{F+1}, \qquad g(F) + s = \frac{(1+s)(F+s)}{F+1}.$$

Every $f(x,n)$ is positive (it is a ratio of positive sums), so $f(x,n) + s \neq 0$ and we may divide:

$$r_n := \frac{f(x,n) - s}{f(x,n) + s} \quad \text{satisfies} \quad r_{n+1} = a\, r_n, \qquad a := \frac{1-s}{1+s}.$$

Since $r_1 = (1-s)/(1+s) = a$, we get $r_n = a^n$. Solving $r_n = (f-s)/(f+s)$ for f,

$$f(x,n) = s\,\frac{1+a^n}{1-a^n} = \sqrt{x}\,\frac{(1+\sqrt{x})^n + (1-\sqrt{x})^n}{(1+\sqrt{x})^n - (1-\sqrt{x})^n}.$$

For $s > 0$, $|1-s| < 1+s$, so $|a| < 1$, $a^n \to 0$, and $f(x,n) \to s = \sqrt{x}$. (At $x = 1$, $a = 0$ and $f \equiv 1$.)

Case $x < 0$. Suppose the real sequence $f(x,n)$ is defined for all large n and tends to a real L. If $L \neq -1$, continuity of g at L gives $L = (L+x)/(L+1)$, i.e. $L^2 = x < 0$, impossible. If $L = -1$, the identity $f(x,n+1)\big(f(x,n)+1\big) = f(x,n) + x$ tends to $0 = -1 + x$, so $x = 1$, again impossible. Hence no limit exists. (Explicitly, writing $x = -t^2$, $t > 0$, and $\theta = \arctan t \in (0, \pi/2)$, the closed form becomes $f(x,n) = t\cot(n\theta)$, which oscillates and is undefined whenever $n\theta \in \pi\mathbb{Z}$; for $x = -1$ it runs $1, 0, -1,$ undefined, $1, \ldots$.)

Complex remark. The same computation works for complex $x \notin (-\infty, 0]$ with s the square root having $\operatorname{Re} s > 0$: then $|1-s| < |1+s|$, so $|a| < 1$ and $f(x,n) \to s$. This is the book's closing comment.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains share the recursion and the fixed-point equation $L^2 = x$. They split at the line "$L = (L+x)/(L+1)$, so $L = \sqrt{x}$". The natural attempt treats the fixed points as the answer and then fights the map head-on for convergence. The model uses the two fixed points as *coordinates*: it measures f against both of them at once, through the ratio $(f-s)/(f+s)$, and that ratio turns the nonlinear recursion into multiplication by a constant.

> *Passing to the limit only tells me what the limit would have to be, not that there is one — so I still owe a convergence proof, and I shouldn't try to get it from the derivative of a map that visibly overshoots.*

> *This recursion is a fraction-linear map with two fixed points $\pm\sqrt{x}$; if I look at the distance to one fixed point divided by the distance to the other, the map should turn into multiplication by a fixed number, and then everything is a geometric sequence.*

Layer 3 THE WHY

Why the cross-ratio coordinate, and why is it not a trick? Start with what $g(F) = (F + x)/(F+1)$ is. It is a Möbius (fraction-linear) map: numerator and denominator are both linear in F. The defining property of such maps is that they are determined by what they do to three

points, and a map fixing two points s and $-s$ has one degree of freedom left. The natural way to see that single degree of freedom is to put the two fixed points at 0 and ∞. The coordinate that sends $s \mapsto 0$ and $-s \mapsto \infty$ is $r = (F - s)/(F + s)$; in that coordinate a Möbius map fixing 0 and ∞ must be $r \mapsto ar$, because the only fraction-linear maps fixing 0 and ∞ are scalings. So linearisation is not an inspired guess; it is forced once you ask "what does this map do, stripped of the positions of its fixed points?". The two-line computation of $g(F) \mp s$ in the model chain is that fact verified by hand, with no theory quoted.

There is a second, independent route to the same place, and it is the one the book's closed form comes from. E_n and O_n are the even and odd parts of the binomial expansion of $(1 + y)^n$ with $y^2 = x$:

$$(1 + \sqrt{x})^n = E_n + \sqrt{x}\, O_n, \qquad (1 - \sqrt{x})^n = E_n - \sqrt{x}\, O_n.$$

Adding and subtracting gives E_n and O_n, and the ratio is the closed form. The question's phrasing is the cue here: "$\binom{n}{0} + \binom{n}{2}x + \binom{n}{4}x^2 + \cdots$" is the even-index half of a binomial expansion in which x plays the role of a square. Whenever a sum takes every second binomial coefficient, split $(1 + y)^n$ and $(1 - y)^n$ into even and odd parts. The linear recursion $E_{n+1} = E_n + xO_n$, $O_{n+1} = E_n + O_n$ is the matrix $\begin{pmatrix} 1 & x \\ 1 & 1 \end{pmatrix}$, whose eigenvalues are $1 \pm \sqrt{x}$; the fixed points $\pm\sqrt{x}$ of g and the eigenvalues $1 \pm \sqrt{x}$ of the matrix are the same data, and $a = (1 - \sqrt{x})/(1 + \sqrt{x})$ is the ratio of the eigenvalues. Convergence happens precisely when one eigenvalue strictly dominates in size, which is $|1 - s| < |1 + s|$, i.e. $\operatorname{Re} s > 0$. For negative x the two eigenvalues $1 \pm it$ have equal modulus, neither dominates, and the ratio rotates forever. That is the structural reason for "suitable values of x".

Finally, the sign. The equation $L^2 = x$ has two roots, and a script that writes $\sqrt{x}$ without comment has skipped a step. Positivity of every $f(x, n)$ for $x \geq 0$ settles it, and so does the dominance argument: the limit is the fixed point associated with the larger eigenvalue.

Layer $+\alpha$ BACKTESTING

1. The first cue is the alternating-index binomial sums $\binom{n}{0} + \binom{n}{2}x + \cdots$ over $\binom{n}{1} + \binom{n}{3}x + \cdots$. On seeing it, write $(1 \pm \sqrt{x})^n = E_n \pm \sqrt{x}\, O_n$.

2. Two methods, both short: Pascal's rule for the recursion, and the even/odd split of $(1 \pm \sqrt{x})^n$ (or the fixed-point ratio $(f - s)/(f + s)$) for the closed form and the limit.

3. Hesitation occurs right after $L^2 = x$: the student senses a gap but does not know how to fill it. Reinforce with the rule that a fixed-point equation identifies candidates only, and that for a fraction-linear recursion the gap is closed by the ratio to the two fixed points.

4. The trap is writing $\lim f = \sqrt{x}$ from $L = (L + x)/(L + 1)$ alone, with no existence proof, no choice of sign, and no discussion of negative x. "For suitable fixed values of x" is the examiner telling you that some x fail.

5. Recognise faster: any recursion $u_{n+1} = (\alpha u_n + \beta)/(\gamma u_n + \delta)$ is a ratio of two sequences obeying a 2×2 linear recursion, and its long-run behaviour is decided by the eigenvalues of that matrix.

6. Decision rule: for an iterated fraction-linear map with distinct fixed points p, q, set $r_n = (u_n - p)/(u_n - q)$; then $r_{n+1} = \lambda r_n$ for a constant λ, and the sequence converges to p exactly when $|\lambda| < 1$.

▷ **Cue:** A sum takes every second binomial coefficient, $\binom{n}{0} + \binom{n}{2}x + \binom{n}{4}x^2 + \cdots$ ⟶ **Move:** Write $x = y^2$ and split $(1 + y)^n$ and $(1 - y)^n$ into even and odd parts.

▷ **Cue:** A limit is wanted for $u_{n+1} = (\alpha u_n + \beta)/(\gamma u_n + \delta)$ ⟶ **Move:** Find the two fixed points p, q and track $(u_n - p)/(u_n - q)$, which is geometric.

▷ **Cue:** Passing to the limit in a recursion gives an equation for L ⟶ **Move:** Treat its roots as candidates only; still prove existence and choose the root by sign or dominance.

1965 B5

1965 B5. Consider collections of unordered pairs of V different objects $a, b, c, \cdots, k$. Three pairs such as bc, ca, ab are said to form a triangle. Prove that, if $4E \leq V^2$, it is possible to choose E pairs so that no triangle is formed.

Caveat. The OCR reads "be, ca, ab" and "$4E \sim V2$". These are restored as bc, ca, ab and $4E \leq V^2$; the book's solution ends with "$mn \geq V^2/4 \geq E$", which confirms the inequality's direction. That final line of the book's solution is itself inaccurate for odd V (then $mn = (V^2-1)/4 < V^2/4$); the gap is closed by integrality of E, as in the model chain below.

Fields: Combinatorics (extremal graph theory) **Position:** B5 (upper) **Answer:** Split the objects into two groups of sizes $\lceil V/2 \rceil$ and $\lfloor V/2 \rfloor$ and use only pairs with one object from each group: there are $\lfloor V^2/4 \rfloor \geq E$ such pairs and no three of them form a triangle, so any E of them will do.

Layer 1 THE WHAT — surface analysis

The natural attempt. In graph language: V vertices, choose E edges, no triangle. The first instinct is greedy: keep adding pairs, each time choosing one that does not complete a triangle with two pairs already chosen. With $V = 4$ and $E = 4$ you try ab, bc, cd and then da — fine, a 4-cycle. With $V = 5$ and $E = 6$ you add ab, bc, cd, de, ea (a pentagon, triangle-free) and then look for a sixth pair: every remaining pair is a diagonal such as ac, and ac together with ab, bc is a triangle. Every diagonal of the pentagon closes a triangle with two of its sides, so the greedy process is stuck at 5 pairs although $6 \leq 25/4$. You restart with a different order, succeed by luck, and realise that greedy success depends on early choices you have no rule for.

The second instinct is induction on V: from a good configuration on V objects, add one object and join it to as many old objects as possible without a triangle. The new object may be joined to a set of old objects only if no two of them are already paired. To keep the count growing like $V^2/4$ you need a large set of mutually unpaired old objects — about half of them — at every stage. At that point you are already describing the answer, but the induction has hidden it: you are carrying a structure without having named it, and the bookkeeping of "the largest unpaired set in the current configuration" becomes the whole problem.

The model chain. Call a set of pairs *triangle-free* if no three of its pairs form a triangle. Any subset of a triangle-free set is triangle-free, because a triangle in the subset would be a triangle in the larger set. So it suffices to exhibit one triangle-free set with at least E pairs and then keep any E of them.

Split the V objects into two groups A and B with $|A| = m$, $|B| = n$, $m + n = V$. Let P be the set of all mn pairs with one object in A and one in B.

P is triangle-free. A triangle bc, ca, ab involves three distinct objects a, b, c. Each lies in A or B, and there are three objects but only two groups, so two of them lie in the same group. The pair formed by those two is not in P. So no triangle has all three of its pairs in P.

P is large enough. Take $m = \lceil V/2 \rceil$, $n = \lfloor V/2 \rfloor$.

- If V is even, $mn = V^2/4 \geq E$.
- If V is odd, $mn = \dfrac{V+1}{2} \cdot \dfrac{V-1}{2} = \dfrac{V^2-1}{4}$. Now V^2 is odd, and $4E$ is a multiple of 4 with $4E \leq V^2$; as $4E \neq V^2$, we have $4E \leq V^2 - 1$, i.e. $E \leq (V^2-1)/4 = mn$.

In both cases $|P| = mn \geq E$. (If $E \leq 0$ there is nothing to prove.) Choose any E pairs from P; they form no triangle.

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains agree on the translation into "pairs without a triangle". They split at the first move: the natural attempt *builds* a set of pairs one at a time and checks each addition; the model *describes* a whole family of pairs by a rule — "one from each side" — whose triangle-freeness is proved once, for all of them at the same time.

> *Checking triangles pair by pair will never scale; I want a rule for which pairs to take that makes a triangle impossible for a structural reason, and a triangle needs three objects — so if I sort the objects into two bins and only take pairs that cross between bins, some two of any three objects share a bin and their pair is missing.*

> *The bound is $V^2/4 = (V/2)^2$, which is what I get from two halves of size $V/2$ each; the target number is telling me the shape of the construction.*

Layer 3 THE WHY

Why two groups, and why is this not pulled from a hat? Start from what a triangle is: three objects, all three pairs present. To forbid it you must guarantee that among any three objects *some* pair is missing. "Among any three, two are alike" is the pigeonhole principle with three pigeons and two holes, so the moment you want a guaranteed missing pair among every triple, you are led to colour the objects with two colours and leave out the same-colour pairs. That is the axiom-level reason: triangle-free is a statement about triples, and two colours is the largest number of colours for which every triple must repeat a colour.

The second piece of evidence is the number in the hypothesis. $V^2/4$ is $(V/2)\cdot(V/2)$, a product of two numbers summing to V and as equal as possible. A bound of the form "product of two parts" is the fingerprint of a bipartite count: every object in one part paired with every object in the other. By AM–GM, $mn \le \left(\frac{m+n}{2}\right)^2 = V^2/4$, with the maximum at the balanced split, so the hypothesis $4E \le V^2$ is exactly "E is at most the best such product". Reading the hypothesis as a construction hint is the whole solution; a problem that supplies a clean bound like $V^2/4$ in an existence statement is almost always telling you the extremal object.

The problem also punishes a small lapse the book's own solution makes. For odd V, the balanced split gives $(V^2-1)/4$, which is less than $V^2/4$. The claim $mn \ge V^2/4$ is then false, and it is integrality that rescues the argument: E is a whole number, $V^2 \equiv 1 \pmod 4$, so $4E \le V^2$ forces $4E \le V^2 - 1$. Write that line; it is where a careful grader looks.

For context, this is the easy half of Mantel's theorem (1907): a triangle-free graph on V vertices has at most $\lfloor V^2/4 \rfloor$ edges, with equality only for the balanced complete bipartite graph. The problem asks only for the construction, not for the upper bound, and the construction is the half you can derive from the definition of a triangle alone. Turán's theorem extends it: to avoid $r+1$ mutually paired objects, split into r groups and take all cross pairs — the same pigeonhole with r holes.

Layer $+\alpha$ BACKTESTING

1. The first cue is "no triangle" paired with the bound $V^2/4$. On seeing it, write: "two groups of size about $V/2$, all pairs across" and check that three objects must put two in one group.

2. The method is construction by a global rule (a two-colouring), verified by pigeonhole on three objects and counted by $\lceil V/2 \rceil \lfloor V/2 \rfloor$.

3. Hesitation occurs at odd V, where the product falls short of $V^2/4$. Reinforce with integrality: E is an integer and $V^2 \equiv 1 \pmod 4$, so $4E \le V^2$ means $4E \le V^2 - 1$.

4. The trap is greedy addition of pairs, which can stall early (the pentagon on 5 objects blocks every further pair at $5 < 6$), and a sloppy odd case that asserts $mn \ge V^2/4$.

5. Recognise faster: when a problem forbids a small complete configuration on $r+1$ objects, colour the objects with r colours and keep only relations between different colours; a bound that looks like $(V/r)^2\binom{r}{2}$ confirms it.
6. Decision rule: to show a large family avoiding a forbidden local pattern exists, find a global labelling under which the pattern cannot occur by pigeonhole, then count the family and optimise the label sizes; subsets of a pattern-free family stay pattern-free, so "at least E" gives "exactly E".

▷ **Cue:** Forbidden triangle (three objects, all three pairs present) $\longrightarrow$ **Move:** Two-colour the objects and keep only pairs with different colours; among three objects two share a colour.
▷ **Cue:** An existence claim with bound $V^2/4$ $\longrightarrow$ **Move:** Read $V^2/4$ as $(V/2)(V/2)$: split into two balanced halves and take all cross pairs.
▷ **Cue:** A construction count falls short of a real bound only by a fraction $\longrightarrow$ **Move:** Use integrality of the quantity being bounded, e.g. $V^2 \equiv 1 \pmod 4$ turns $4E \le V^2$ into $4E \le V^2 - 1$.

1965 B6

1965 B6. If A, B, C, D are four distinct points such that every circle through A and B intersects (or coincides with) every circle through C and D, prove that the four points are either collinear (all on one line) or concyclic (all on one circle).

Caveat. The OCR reads "(all of one line)"; restored as "(all on one line)", matching the parallel "(all on one circle)". The book's solution under B-6 (perpendicular bisectors of AB and CD, concentric circles, parallel case) solves exactly this statement. Two steps of that solution are asserted without proof — that the bisectors cannot coincide, and that the two circles in the parallel case are disjoint ("clearly") — and both are proved in full below.

Fields: Geometry (circles, loci), proof by contradiction with case split **Position:** B6 (last) **Answer:** Suppose neither collinear nor concyclic; the centres of circles through A, B fill the perpendicular bisector p of AB, those through C, D fill the bisector q of CD; $p \neq q$, and whether p, q meet or are parallel you can exhibit two disjoint circles (concentric, or tangent to a common line from opposite sides), a contradiction.

Layer 1 THE WHAT — surface analysis

The natural attempt. The hypothesis is about *every* circle, so you start testing particular ones. The first circle to try is the circle through A, B and C. It certainly meets every circle through C and D — at C. Useless. You try the circles with diameters AB and CD: the hypothesis says they meet, which gives an inequality $MN \le \frac{1}{2}(AB + CD)$ between the midpoints M, N and the lengths. That is a condition, not a contradiction, and more such circles give more inequalities with no end in sight.

Then you try to argue directly: assume the points are not concyclic, so D is off the circle $\Gamma = ABC$; find a circle through C and D that misses some circle through A and B. The instinct is to make one circle small. But a circle through two fixed points C, D has diameter at least CD — you cannot shrink it away. You try to make one circle huge and "far away", and discover that a huge circle through A, B is nearly the line AB, which still cuts through the region where the circles through C, D live. You are pushing circles around by hand with no description of the family you are choosing from, and each configuration of the four points seems to need its own picture. The wall is the absence of a parameter: "every circle through A and B" has not been turned into something you can quantify over.

The model chain. *Setup.* A circle passes through A and B exactly when its centre O satisfies $OA = OB$, i.e. when O lies on the perpendicular bisector p of AB; its radius is then OA.

Every point of p is the centre of one such circle. Likewise the circles through C and D are the circles centred at points of the perpendicular bisector q of CD. Suppose, for contradiction, that A, B, C, D are neither collinear nor concyclic. Two lines either coincide, or are distinct and parallel, or meet in exactly one point; we rule out each.

Case 1: $p = q$. Lines AB and CD are both perpendicular to p. If they are the same line the four points are collinear, excluded. Otherwise take coordinates with p the y-axis; then $A, B = (\pm u, h_1)$ and $C, D = (\pm v, h_2)$ with $u, v > 0$ and $h_1 \neq h_2$. For $O = (0, k)$,

$$OA^2 = OC^2 \iff u^2 + (h_1 - k)^2 = v^2 + (h_2 - k)^2 \iff 2(h_2 - h_1)k = v^2 + h_2^2 - u^2 - h_1^2,$$

which has a solution k because $h_1 \neq h_2$. That O is equidistant from all four points, so they are concyclic, excluded. Hence $p \neq q$.

Case 2: p and q meet at one point O. Let Γ_1 be the circle centred at O through A (so also through B) and Γ_2 the circle centred at O through C (so also through D). If $OA = OC$ then $\Gamma_1 = \Gamma_2$ contains all four points, excluded. So $OA \neq OC$, and two concentric circles of different radii have no common point, since every point of Γ_1 is at distance OA from O and every point of Γ_2 at distance OC. Thus Γ_1 neither meets nor coincides with Γ_2: contradiction.

Case 3: $p \parallel q$, $p \neq q$. Lines AB and CD are perpendicular to the common direction of p and q, so they are parallel. If they coincide, the points are collinear, excluded. So they are distinct parallel lines. Choose coordinates with AB on $y = h$ and CD on $y = -h$ ($h > 0$), p the line $x = b$, q the line $x = d$, $b \neq d$; so $A, B = (b \pm u, h)$ and $C, D = (d \pm v, -h)$ with $u, v > 0$. Put $k = (u^2 + h^2)/(2h) > 0$ and let Γ_1 be the circle with centre (b, k) and radius k. Then

$$\operatorname{dist}\big((b,k), A\big)^2 = u^2 + (h-k)^2 = u^2 + h^2 - 2hk + k^2 = k^2,$$

so Γ_1 passes through A and (by symmetry in $x = b$) through B. Every point (x, y) of Γ_1 has $y \geq k - k = 0$, with $y = 0$ only at $(b, 0)$: Γ_1 is tangent to the x-axis at $(b, 0)$ from above. In the same way, with $k' = (v^2 + h^2)/(2h)$, the circle Γ_2 with centre $(d, -k')$ and radius k' passes through C, D, lies in $y \leq 0$, and meets the x-axis only at $(d, 0)$. A common point of Γ_1 and Γ_2 would have $y = 0$ and so equal both $(b, 0)$ and $(d, 0)$, forcing $b = d$. So they are disjoint (and distinct): contradiction.

All three cases are impossible, so the four points are collinear or concyclic. (These Γ_1, Γ_2 are the book's circles ABP and CDQ, with $P = (b, 0)$ and $Q = (d, 0)$ on the line midway between AB and CD.)

Layer 2 THE HOW — the bifurcation

Bifurcating point. Both chains start from "assume not collinear, not concyclic, find two circles that miss". They split at the line "find a circle through C and D that misses a circle through A and B". The natural attempt searches for the circles by picture. The model first *parametrises* each family by its centre, so that "every circle through A and B" becomes "every point of the line p", and the problem becomes a question about two lines.

> *A circle is a centre and a radius, and the circles through two fixed points are pinned down by where the centre is — the centre must be equidistant from the two points, so it runs along the perpendicular bisector; I should stop drawing circles and start choosing points on two lines.*

> *Two circles can be kept apart whatever their sizes in only two ways: give them the same centre and different radii, or put them on opposite sides of a line; the first needs p and q to meet, the second is what's left when they are parallel.*

Layer 3 THE WHY

Why the perpendicular bisector? Go to the definition. A circle is the set of points at a fixed distance from a centre. "Passes through A and B" therefore says $OA = r = OB$, and the set of

O with $OA = OB$ is the perpendicular bisector of AB. That is the Blueprint's circle principle applied to a family instead of a single circle: the centre is the object of interest, so the family of circles is the family of centres, which is a line. Once each family is a line, the hypothesis "every member of one family meets every member of the other" is a statement about two lines p and q, and two lines admit exactly three configurations: equal, parallel, crossing. The case split is not a choice; it is the classification of pairs of lines.

Why concentric circles in Case 2? The question in each case is: when do two circles fail to meet? In terms of centres O_1, O_2 and radii r_1, r_2, two circles are disjoint exactly when $O_1O_2 > r_1 + r_2$ (apart) or $O_1O_2 < |r_1 - r_2|$ (one inside the other). The radii are not free — they are forced to be O_1A and O_2C — and the "apart" option is hard to guarantee because circles through fixed points cannot be made small. The "inside" option is easiest when $O_1O_2 = 0$: then disjointness needs only $r_1 \neq r_2$, and $r_1 = r_2$ would mean concyclic. So if p and q share a point, choose it as both centres.

Why tangency in Case 3? With $p \parallel q$ the centres can never coincide, so the "inside" route is gone and you need "apart". Parallel bisectors mean $AB \parallel CD$, so there is a strip between the two chords and a midline m inside it. The cleanest certificate that two sets are apart is a line separating them. A circle through A, B that touches m lies entirely on the AB side of m; the tangent point is where its centre's line p crosses m. Do the same on the other side. Tangent points $(b, 0)$ and $(d, 0)$ differ because $p \neq q$, so the circles cannot share even the one point of m they touch. The book writes "clearly" at this step; the tangency computation $k = (u^2 + h^2)/(2h)$ is what that word was standing in for.

Why must $p \neq q$ be proved? If $p = q$ the configuration is an isosceles trapezoid (or a degenerate one on a line), and such a trapezoid is cyclic — the linear equation for k always has a solution. So $p = q$ is not a separate contradiction; it is itself the concyclic conclusion, and a proof that skips it has skipped a genuine case.

One warning about the converse, which the problem does not ask. It is false: for collinear points A, B, C, D at $0, 1, 2, 3$ on the x-axis, the circles with diameters AB and CD are disjoint. So the hypothesis is strictly stronger than "collinear or concyclic"; only one direction is being proved, and a solution should not claim an equivalence.

Layer $+\alpha$ BACKTESTING

1. The first cue is "every circle through A and B". On seeing it, write: "centres of these circles = perpendicular bisector p of AB".

2. The method is contradiction plus a case split on the relative position of the two bisectors, with an explicit pair of disjoint circles in each case.

3. Hesitation occurs in the parallel case, where concentric circles are unavailable. Reinforce with the separating-line idea: circles tangent to the midline from opposite sides at different points cannot meet.

4. Two traps: testing only particular circles (such as circle ABC, which meets every circle through C at C), and forgetting the case $p = q$, where the points are in fact concyclic and "cannot coincide" needs proof.

5. Recognise faster: any hypothesis quantified over "all circles through two points" is a hypothesis over points of a line; any conclusion "concyclic" is the existence of one point equidistant from all four.

6. Decision rule: when a family of geometric objects is defined by passing through fixed points, parametrise it by the locus of its centre (or other defining datum) before choosing members; then classify the relative position of the loci and handle each configuration with the extreme member.

▷ **Cue:** "Every circle through A and B" $\longrightarrow$ **Move:** Replace by "every centre on the perpendicular bisector of AB", radius forced.

▷ **Cue:** Need two circles that certainly do not meet, and their centres may be chosen at a common point $\longrightarrow$ **Move:** Make them concentric; different radii guarantee disjointness.

▷ **Cue:** Need two circles to be apart, with two parallel chords given $\longrightarrow$ **Move:** Take the midline of the strip and the circles tangent to it from opposite sides.

▷ **Cue:** Conclusion "the four points are concyclic" $\longrightarrow$ **Move:** Find one point equidistant from all four; with symmetry this is a single linear equation.