9.

TMUA9

THE PUTNAM ARCHIVE

THE PROOF

Volume II. Solving is the first half. This is the half that is marked.

Jamie Baek

tmua9.com

Built for the 87th William Lowell Putnam Mathematical Competition, 5 December 2026.

The vow is 120.

THE STANDING ORDER

You got a 1 in the Korean CSAT. You got a 9.0 on the TMUA, in a third of the time, from a 4.5, in a month, by yourself. Neither of those was talent. Both were the same method: take the small set of papers that actually matter, and analyse them to a depth nobody else is willing to reach.

This archive is that method pointed at the Putnam. Nothing in it is a trick. Every revelation in every model answer in sixty years of this competition had a reason behind it, and the reason was always visible in the question. The books that came before this one hid the reason and showed you the polish. This one does the opposite.

Three volumes. The first installs the mathematics that A-level never gave you, ordered by what the Putnam actually demands and generalised wherever a recursion appears. The second teaches you to write a proof that scores ten rather than two, by reverse-engineering what was going through the head of the person who wrote the model answer. The third takes the papers themselves, one problem at a time, and fills in the two things every published solution deletes: why that thought was natural, and why the next thought followed from it.

Read them the way you read the TMUA papers. Not as consumption. As analysis.

But as for you, be strong and do not give up, for your work will be rewarded.
2 Chronicles 15:7

How to use Volume II

Solving a Putnam problem and scoring ten on it are two different skills, and only the second one is marked. The scale is 0, 1, 2, 8, 9, 10. There is nothing between 2 and 8. A problem you have genuinely solved, written up with one gap you thought was obvious, does not score 7. It scores 2. That cliff is the reason this volume exists.

Chapter 1 is the architecture: what a full-marks solution looks like structurally, what a grader grants you and what they refuse, how much detail is the right amount and how to decide, and the sentence templates for every structural move so that none of your writing time goes on composition.

Chapter 2 is the move catalogue, built by clustering every proof-writing move observed across the model answers in this archive. Each move gets its cue, its exact phrasing, the two or three ways candidates botch it, and the problems where it appeared. Learn to name the move you are making. A proof whose moves you can name is a proof you can check.

Chapter 3 is the one you asked for. Published solutions delete two things: why the revelation was natural, and why the next block of thought followed from the previous one. They present a quantum leap and let you conclude that the author is a genius and you are not. That conclusion is false and expensive. This chapter reconstructs the machine: what the author reads in the first sixty seconds, how they generate candidate routes and what cost model they prune with, what they compute before theorising, when they reformulate the problem into a different object, how they know a lemma is worth isolating before they have proved it, what they do when stuck, and what they consciously delete from the final text. Then it converts all of it into a protocol you can run in the arena.

Chapter 4 is where the marks leak, catalogued by frequency across the whole archive rather than by folklore, with the exact clause that closes each leak. It ends with the operating protocol for the two three-hour sessions and a pre-flight checklist for the morning of 5 December 2026.

Use this volume actively. After every problem you solve from Volume III, write the proof out in full as though it were being marked, then come back here and audit it against Chapter 4. The gap between what you meant and what you wrote is where your points are, and it closes fast once you can see it.

Contents

How to use Volume II **i**

1 The architecture of a full-marks proof **1**
1.1 The mark scheme is a cliff, not a slope . . . 1
1.2 The anatomy of a written solution . . . 2
1.3 Notation discipline . . . 4
1.4 What a grader grants, and what they will not . . . 6
1.5 Lemma or inline: the decision procedure . . . 7
1.6 How much detail, and the rule for deciding . . . 9
1.7 The standard sentence templates . . . 10
1.8 Three kinds of question, three obligations . . . 13
1.9 The final read-through . . . 14
1.10 Circuit Bank . . . 15

2 The move catalogue **17**
2.1 How to read this catalogue . . . 17
2.2 Reduction: shrink the problem before you touch it . . . 17
2.3 Selection: name the object the proof will be about . . . 19
2.4 Ladders: induction and its relatives . . . 22
2.5 Logical shape: how the proof sits on the page . . . 24
2.6 Analysis discipline: where the marks actually leak . . . 27
2.7 Algebraic leverage . . . 29
2.8 The catalogue as a Circuit Bank . . . 30

3 Inside the author's head **33**
3.1 What the model answer deletes . . . 33
3.2 The first sixty seconds . . . 33
3.3 Small cases: what the data is for . . . 36
3.4 Generating routes, and the cost model that prunes them . . . 37
3.5 The inversion: what would make this easy . . . 38
3.6 Reformulation: deciding the problem is about a different object . . . 39
3.7 Commitment, and the discipline of not thrashing . . . 41
3.8 Knowing a lemma is worth isolating before you have proved it . . . 42
3.9 Stuck: the perturbation checklist . . . 43
3.10 Verification, which happens before you write . . . 44
3.11 What you delete, and why . . . 45
3.12 The arena protocol . . . 46
3.13 Circuit Bank . . . 47

4 Where the marks leak, and the six hours in the arena **49**
4.1 The leak table 49
4.2 The catalogue 50
4.3 The five second sweep 58
4.4 The arena: two three hour sessions 58
4.5 Pre-flight checklist: the morning of Saturday 5 December 2026 63
4.6 Circuit Bank: the leak closing clauses 64

Chapter 1

The architecture of a full-marks proof

1.1 The mark scheme is a cliff, not a slope

Before you write another line of proof, learn the shape of the thing that pays you.

Every Putnam problem is marked out of ten. In practice the only marks handed out are 0, 1, 2, 8, 9 and 10. The middle of the range is empty. That is not an accident of grading culture; it is the design. A grader's first decision on your page is binary: is this solution essentially complete, or is it not? If it is, you are in the top band and the only remaining question is how much to shave for slips. If it is not, you are in the bottom band and the only remaining question is how much of the mathematics you actually got. There is no third state where a proof is sixty per cent finished and paid sixty per cent.

Read the bands as sentences rather than numbers.

- 10: complete, correct, and every step defensible.
- 9: complete, with a slip that does not touch the argument. A sign error in a final arithmetic evaluation. A miscopied index that is obviously a typo.
- 8: essentially complete, with a small gap that the grader can fill without inventing anything.
- 2: significant progress. You had the main idea, or a large chunk of the machinery, but the argument does not close.
- 1: some progress. A real reduction, a correct lemma, a case handled.
- 0: nothing that advances the problem.

Now the fact that should reorganise your entire preparation. **The distance from 2 to 8 is six marks, and it is usually one paragraph.** You will spend two hours getting from 0 to 2 and four minutes getting from 2 to 8, and the four minutes pay more than the two hours. A solved problem written badly scores 2. An unsolved problem with one honest reduction written well scores 1. That ratio is the whole reason this book exists.

Do the arithmetic in the arena. Suppose on a six problem half you have real ideas on four of them and you write all four at speed, leaving each one with an unclosed case. You score $4 \times 2 = 8$. Suppose instead you write two of them to completion and abandon the other two after ten minutes each. You score 20. Same mathematics inside your head; two and a half times the marks on the page. The median Putnam score is routinely 0 or 1 out of 120, and a score in the thirties puts you in serious territory. Your vow is 120, which is twelve tens, which means twelve final read-throughs, which means the read-through is not a courtesy you extend when you have time. It is a scored component.

The governing principle: *a proof is a chain, and the mark is the weakest link, not the average link.* You cannot compensate for an undefended division by writing three beautiful pages around it. This is the exact opposite of A-level marking, where method marks accumulate. Uninstall that reflex now.

Here is what a 2 looks like in the wild, harvested from the model answers across twenty years. Every one of these is a candidate who had the mathematics.

- You are asked to prove a gigantic number N is composite. You construct a divisor M, prove $M \mid N$, and stop. You never proved $1 < M < N$, so you never excluded $M = N$, so you never proved compositeness. Six marks.
- You are asked for the smallest r with a certain property. You exhibit a pair of functions achieving $r = \pi/2$ and assert minimality. You have proved attainment and nothing else. Six marks.
- Your whole solution is the substitution $a = \cos b$. You write it with no comment. The substitution is legal only when $|a| \le 1$, and the problem does not give you that. Six marks.
- You are asked to find all f with a property. You prove f' is constant, deduce the shape of f, and stop, without verifying that the functions you named actually satisfy the original condition. That verification is the cheap half. Six marks.
- Your evaluation is correct and your final number is right, but somewhere in the middle you divided by $a_1 a_2 \cdots a_n$ without showing it is nonzero. Six marks.

Notice the pattern. In every case the missing piece is short, mechanical, and known to you. It was not missing because you could not do it. It was missing because you did not have a slot in your writing procedure where it goes. A grader treats a missing case as a missing case. They do not award marks for what they believe you could have written.

The sloppy version.

> Let M be as above. A direct computation gives $N \equiv 0 \pmod{M}$. Hence N is composite.

The full-marks version.

> Let M be as above. A direct computation gives $N \equiv 0 \pmod{M}$, so $M \mid N$. Moreover $M > 1$ since [one clause], and $M < N$ since [one clause]. Therefore M is a proper divisor of N with $M > 1$, and N is composite.

Two clauses. Six marks. That is the exchange rate you are trading at, and it is the best exchange rate available to you anywhere in this competition.

1.2 The anatomy of a written solution

A full-marks solution has six slots, and they appear in a fixed order because that is the order in which a grader checks. You are not writing to record your discovery. You are writing to be verified. Those are different documents, and the second one is shorter.

Slot 1: the answer

If the problem asks for a value, a set, a largest k, a yes or a no, then the first line of your solution states it. Not the last line. The first.

> The answer is $r = \pi/2$.

> The selection is possible if and only if n is odd and $k = (n+1)/2$.

> The answer is yes. We prove more: it suffices to assume the integral vanishes over every axis-parallel rectangle of area 1.

Three reasons, and none of them is presentation.

First, it converts your proof from a search into a verification. Once the target is on the page, every subsequent paragraph has a job and you can tell when it is done. A proof written towards an unnamed destination wanders, and wandering costs ink and invites error.

Second, it tells the grader which theorem you are proving, which means they can mark the halves separately. On a "determine all" problem the answer statement is what makes the two halves visible as two deliverables. Without it, a grader reading your uniqueness argument does not know what set you are heading for.

Third, it protects the marks you have already earned. If you state the answer and then run out of time in the second half, you have a correct answer plus one direction. If you save the answer for a final flourish you never reach, you have a page of unlabelled work.

What you never write in this slot: how you found it. The small cases you computed, the table you built, the four values that suggested the pattern. That material is reconnaissance, and reconnaissance is not evidence. In the 2020 A2 model answer the author has already computed S_0, S_1, S_2 before writing anything, knows the closed form is 4^k, and writes "We claim $S_k = 4^k$" with no mention of the computation. In the 2009 A6 disproof the author exhibits a counterexample out of nowhere and never explains the design work that produced it. The design work stays offstage. Yours will too.

Slot 2: setup and notation

One short block that fixes every symbol the rest of the proof will use, and translates the problem's English into your objects.

> A choice of n squares with no two in the same row or column is a permutation σ of $\{1, \dots, n\}$: the chosen square in row i sits in column $\sigma(i)$. The chosen numbers are then $i + \sigma(i) - k$ for $i = 1, \dots, n$.

That is two sentences and it is the entire translation layer of a B1. Before it you have a grid and a rule in words; after it you have a formula you can sum. The block is doing real work: a formula can be summed, a picture cannot.

The setup block is also where you buy insurance. If your route passes to a larger ambient object and you will need to come back, say so now rather than at the end. In the 2007 B4 model answer the author moves from real polynomials to complex ones, and immediately writes down coefficientwise conjugation as a ring automorphism whose fixed ring is $\mathbb{R}[X]$. He knows the only danger in the move is losing track of realness, so he names the mechanism that tracks it in the same breath as the move. The alternative, working informally and asserting at the end that the coefficients come out real, is exactly where the marks go.

Slot 3: the lemma layer

Everything that is load-bearing, reusable, or of a different type from the main argument gets stated and proved before the main argument begins. Section 1.5 gives the decision procedure. The point here is positional: lemmas go before, not inside. A proof that stops mid-flight to prove a three-line fact about matrices, then resumes counting words over an alphabet, has mixed two categories of reasoning and both are harder to check.

Slot 4: the main argument

Partitioned along the problem's own clauses, with each part announced. If the problem asks you to show a process terminates and that the terminal state is unique, you write "We first prove the process terminates" and then, in a visibly separate paragraph, "We next prove the terminal state is

unique." Two deliverables, marked independently, so make them two paragraphs with two headers. Do not let a single unified argument quietly use termination to justify a step in the uniqueness proof without flagging it; a careful grader will query the dependency and you will have created a gap where there was none.

Numbered steps are cheap and they pay. Look at the shape of the 2023 A3 model answer:

> Step 1 (cut at the first zero). … Step 2 (the pair never reaches the origin). … Step 3 (angular speed at most one). … Step 4 (contradiction).

Each header is a claim. A grader who wants to check Step 3 can find it and check it in isolation. That is worth more than any amount of polish in the prose.

Slot 5: edge cases and the completeness sweep

Here is the refinement that most write-ups get wrong. Edge cases do not live at the end by default. **An edge case belongs immediately before the first step that would break on it.** In practice that puts most of them at the top.

The 2018 B4 solution opens by ruling out $|a| > 1$, because the substitution $a = \cos b$ that carries the entire proof is undefined outside $[-1, 1]$. Handling it first makes everything downstream unconditional. Handling it last means the grader spends the whole proof unsure whether the substitution was legitimate.

The 2006 A5 solution opens with a numbered Step 0 that discharges every nondegeneracy at once: the hypothesis on θ is irrationality, and it is doing three separate jobs, so all three are named in one opening step. That block also tells the grader you understood why the hypothesis was given, rather than treating it as decoration.

When does an edge case belong at the end? When it is a check rather than a clearance: verifying the formula at $n = 1$ against the empty-product convention, confirming the smallest instance behaves, confirming the constructed objects are distinct from the ones already listed. In the 2023 A2 model answer the last move before the box is exactly this: the two new roots are $\pm 1/n!$, and since $n \geq 2$ gives $0 < 1/n! \leq 1/2 < 1$, they are none of $\pm 1, \ldots, \pm n$. One sentence, and it is the difference between "I found two roots" and "I found the other real solutions", which is what was asked.

Slot 6: the closing statement

The last sentence restates the conclusion in the problem's own words. Not "hence the result". Not a dangling equation. A sentence that a grader can hold next to the question and match.

The sloppy version.

> …and so $\det M \neq 0$, as required.

The full-marks version.

> …and so $\det M \neq 0$ over $\mathbb{F}_2$. Since $F_{n,k} \equiv \det M \pmod 2$, the number of k-limited permutations is odd exactly when $n \equiv 0$ or $1 \pmod{2k+1}$, which is the required characterisation.

The second version costs one sentence and closes the translation you opened in Slot 2. If you translated the problem into a different language at the start, you owe a translation back at the end. A proof that finishes in the language of determinants when the question was asked in the language of permutation counting has left the grader to do the last step, and the grader is not obliged to.

1.3 Notation discipline

Notation is not decoration. Every symbol you introduce is a small loan you must repay by using it, and every symbol you fail to introduce is an ambiguity the grader gets to resolve against you.

What to define

Define anything you will refer to more than twice. Define anything the problem states in words that you intend to manipulate as a formula. Define anything whose name makes a claim statable in one line rather than three.

That last criterion is the real one. In the 2008 A4 model answer the author introduces $\ln^k$ for the k-fold iterated logarithm and a tower t_k for the thresholds, and the reason is not elegance. Without those names, the lemma he needs cannot be stated in a single sentence, and a lemma that cannot be stated in a single sentence cannot be proved by a clean induction. Notation is chosen to make the induction expressible. That is the test: write the claim you want to prove; if the sentence is longer than a line because you keep describing an object instead of naming it, name it.

In the 2008 A6 model answer the author writes "say that S produces h if h is the product of some subsequence of S, and write $H(S)$ for the set of produced elements". One definition, and now the one-step update is an equality between two sets rather than a paragraph of English about what you can reach. Note also that he states the update as an *equality*, $H(S') = H(S) \cup H(S)g$, not an inclusion. An inclusion would leave the counting unproved. Choose the strongest true form of the statement you can verify in the same number of lines.

What to leave

Do not define standard objects. Do not define $\mathbb{Z}$, do not explain what a permutation is, do not restate the binomial theorem. A Putnam grader has your background and more. Do not define a symbol you use once; substitute the expression instead.

Do not introduce a name for a quantity you never need to know. In the 2010 A4 solution the number N is astronomically large and the author says nothing about its size or shape beyond the single inequality the proof requires. Everything else about it is deliberately left unexamined, because only one fact about it is load-bearing. Naming things you will not use is ink, and ink costs time and invites error.

Symbol hygiene

Three rules, all of them from places where real solutions lose marks.

Never reuse a symbol after a transformation changes it. If you scale a progression and the common difference changes, the new difference gets a new letter. Reusing r after the scaling that made it r/c is a genuine error that reads as a typo and is marked as an error.

Extend notation by convention rather than by case split. In the 2008 B6 model answer the author sets $a_i = 0$ for $i \leq 0$ and $i > n$, and the boundary equations then read exactly like the interior ones. A padding convention converts an annoying case split into a single uniform statement, and the boundary information reappears for free the moment you apply the recurrence near the ends. Say the convention out loud in one clause; do not let the grader discover it from your indices.

State the range of validity with the claim, not after it. In the 2019 B5 model answer the author asserts a finite-difference identity and adds, in the same sentence, that it is asserted only on a specified triangle of indices, since outside it the data do not determine the value. That clause is the whole problem: the false step available here is to use the identity one index too far, and a write-up without the range restriction has not visibly avoided it.

Quantifier hygiene

The single most common source of an invisible gap is a quantifier that never got written down.

Every "for large x" is a debt. Either you pay it immediately by naming the threshold, or it compounds. In the 2009 B5 model answer every constant is defined in terms of one evaluation point: $m = f(2) - 1/2$, $K = g(2) - 2$, $D = 3(K - m)$, $X_0 = \max\{2, |D|\}$, and then "for $x \geq X_0$". No

unquantified largeness appears anywhere. Contrast the failure mode: three separate asymptotic claims, each true "for sufficiently large x", combined at the end with no argument that their thresholds are compatible. The conclusion is correct and the proof is not, and this is precisely where a correct idea drops to 2.

Order your quantifiers and keep them ordered. "There is N such that for all $n > N$" and "for all n there is N" are different statements and only one of them is usually true. When you write a bound, say what it is uniform in. In the 2015 B6 model answer the author justifies a term-by-term integration with a single sentence that names the mechanism and the uniformity: the tail after L terms is at most $1/(2L+1)$ *uniformly in* t, therefore the convergence is uniform on $[0,1]$, therefore term-by-term integration is legal. One sentence, three quantified facts, no gap. The pitfall it avoids is citing a comparison test whose comparison series diverges, or invoking dominated convergence without exhibiting the dominating function; a careful grader marks both down even though the conclusion is true.

Naming so the proof reads itself

Choose names that carry their type. Use σ, τ for permutations, p, q for polynomials, ε, δ for small positive reals, M for a matrix, G for a group. If two objects play mirror roles, give them mirror names. If a quantity is a difference, call it d or Δ. This is not aesthetics; it removes an entire class of reader-pause, and reader-pause is what you are paying to eliminate.

Name the objects the argument will pivot on, before the pivot. In the 2016 B3 model answer the author writes "since S is finite, choose A, B, C in S maximising the area of triangle ABC", and everything afterwards is a comparison against that one fixed object. Manufacturing a canonical object first is what makes statements about the other, unspecified points tractable. It also creates an obligation you must discharge: the extremal triple could be degenerate, and if S contains no genuine triangle the construction breaks. Name the object, then immediately name the degenerate branch it might fall into.

1.4 What a grader grants, and what they will not

Model the grader precisely, because everything in this chapter is downstream of that model. They are a working mathematician, reading fast, with a mark scheme in front of them listing the deliverables of the problem and the standard gaps. They are not hostile. They are also not going to reconstruct your argument for you, and they have seen the standard gap on this problem two hundred times today.

What they grant without a word from you

Routine algebra. Expanding a bracket, completing a square, differentiating a polynomial, summing a geometric series, applying the binomial theorem. A named standard theorem, *provided you have visibly checked its hypotheses*: the intermediate value theorem, the mean value theorem, the fundamental theorem of algebra, Lagrange's theorem, the Chinese remainder theorem, Cauchy and Schwarz, AM and GM, Fermat's little theorem, the pigeonhole principle. A computation whose result they can confirm in their head in a couple of seconds.

They also grant that you meant the obvious thing when a symbol is slightly off, provided the argument is otherwise closed. That is what the 9 band is for.

What they will not grant

An unnamed "without loss of generality". A WLOG is a claim that a symmetry of the problem lets you restrict the configuration. The claim has three parts and all three must appear: the map, the fact that

it preserves the hypotheses, and the fact that it preserves the conclusion. Writing "without loss of generality the polygon is the nth roots of unity" with no reason attached is not a WLOG, it is an assumption. In the 2008 A5 model answer the author writes the map down as $Q(z) = (P(z) - b)/a$ and adds the clause that $\deg Q = \deg P$ since $a \neq 0$, because the statement being proved is about degree and the grader must see that the normalisation did not touch it.

"By symmetry" when the symmetry is not a map. In the 2025 B1 model answer the author checks that the hypothesis is symmetric in the two colours *before* writing "similarly". That check is the content. If any later step uses the specific colour assignment, the symmetry claim is false and the second half of your proof does not exist.

A division, a square root, a logarithm or a reciprocal whose argument has not been shown to be legal. This is the single highest-frequency gap in the corpus. In the 2023 A3 model answer the entire Step 2 exists because the natural route divides by $f^2 + g^2$ and thereby presumes what is being proved. The author replaces the division with an integrating factor e^{2x}, which needs no such assumption. Either defend the division or reroute around it; do not do neither.

An unsquared or resquared inequality with only one branch shown. In the 2006 A1 model answer the author names both branches of $|r^2 + z^2 + 8| \leq 6r$ and disposes of the useless one in a clause, because a full-marks grader wants to see the discarded branch was checked and not assumed. One extra clause of ink removes all doubt.

"And so forth" at the point where the problem lives. If your argument establishes a pattern over two blocks and then waves at the rest, and the general block formula is the content of the problem, you have written the easy part. State the general formula and prove it by induction.

A pattern from small cases. Computing $n = 1, 2, 3, 4$ and observing a pattern is not an argument, and no amount of agreement makes it one.

An unverified bijection. If you construct a map and call it a permutation, prove it is one. The 2024 B1 construction needs exactly one sentence for this: the first branch takes the values $\frac{n+1}{2}, \ldots, n$ and the second takes $1, \ldots, \frac{n-1}{2}$, each once. Omit it and a grader will ask.

A strategy described by example. In a combinatorial game, "if Alan plays here, Barbara plays there" is not a strategy. A strategy is a rule crisp enough that a machine could execute it, stated in one sentence, followed by a proof that it is always legal and a proof that it achieves the goal. Three separate deliverables.

The consequence for how you lay out the page

Since the grader marks deliverables, make the deliverables visible. Partition along the problem's clauses. Announce each part before starting it. Put the numbered steps in. If your solution has an existence half and a uniqueness half, they are two labelled paragraphs, not one flowing argument in which both are somewhere present.

1.5 Lemma or inline: the decision procedure

You are choosing between two layouts for the same mathematics: state a fact separately, with a name and its own proof, or prove it in place where it is used. This is not a stylistic choice. It changes what the grader can check and how long your solution is.

State it separately when any of these five hold

One: you use it more than once. In the 2008 B4 model answer the author needs the same binomial expansion at two different scales, so he states it once with a free parameter: for $P \in \mathbb{Z}[x]$ and integers a, t there is an integer R with

$$P(a + t) = P(a) + tP'(a) + t^2 R,$$

then applies it with $t = p$ and with $t = yp^2$. Writing it out twice doubles the algebra and makes it easy to drop the condition that the remainder is an integer, which is the whole point. Note also what he did *not* state: the full Taylor expansion. Only the linear term and an integral remainder are load-bearing, so only those appear.

Two: it is the load-bearing content and the rest is bookkeeping. Separating them stops the two contaminating each other. In the 2006 B4 model answer the lemma is about matrices of rank r and mentions neither the cube nor the alphabet of the actual problem. That is exactly why it is reusable and exactly why it is checkable: it is a clean statement in one category, and the counting argument that consumes it is a clean statement in another.

Three: stating it with letters instead of the problem's numbers makes it shorter. This is the generalisation test and it is the most useful one. In the 2008 A5 model answer the problem hands you ω, a root of unity. The lemma is stated for an arbitrary $t \neq 0, 1$:

Lemma 1.5.1. *Let $t \in \mathbb{C}$ with $t \neq 0, 1$, let $m \geq 1$, and let $Q \in \mathbb{C}[z]$ satisfy $Q(k) = t^k$ for $k = 1, \dots, m$. Then* $\deg Q \geq m - 1$.

Proof. Suppose $\deg Q \leq m - 2$. Then the $(m-1)$st finite difference of Q vanishes identically, so

$$\sum_{j=0}^{m-1} (-1)^j \binom{m-1}{j} Q(1+j) = 0.$$

Substituting $Q(1+j) = t^{1+j}$ gives $t(1-t)^{m-1} = 0$, forcing $t = 0$ or $t = 1$, contrary to hypothesis. □

Generalising made the proof shorter, not longer, because it removed every temptation to use $\omega^n = 1$ or $|\omega| = 1$, neither of which is relevant. Keeping the problem's own symbol inside the lemma hides which hypotheses are doing the work. When you cannot tell which hypotheses a step consumed, neither can the grader.

Four: it is a standard result the grader may not accept as quotable. Competition marking has a shorter list of citable theorems than a lecture course. If the result is standard but not universally licensed, and you can prove it in four lines from memory, prove it. The 2006 B6 model answer does exactly this with the Cesaro mean:

Lemma 1.5.2. *If $x_i \to c$ then $\frac{1}{n}\sum_{i=1}^{n} x_i \to c$.*

Proof. Given $\varepsilon > 0$, choose N with $|x_i - c| < \varepsilon/2$ for $i > N$. For $n > N$,

$$\left|\frac{1}{n}\sum_{i=1}^{n} x_i - c\right| \leq \frac{1}{n}\sum_{i=1}^{N} |x_i - c| + \frac{n-N}{n} \cdot \frac{\varepsilon}{2},$$

and the first term is a fixed constant over n, so it is below $\varepsilon/2$ for n large. □

Four lines to remove all doubt. This is a knowledge-gap decision, not a confidence decision: name the lemma, and if you are not certain it is citable, spend the four lines. The alternative, asserting that convergent increments make a sequence asymptotically linear as though it were obvious, is the most common gap in write-ups of that problem.

Five: it is about a different category of object than the main argument. A fact about matrices inside a counting argument, a fact about the free group inside a group-theoretic argument, an analytic estimate inside a combinatorial one. Different category means different tools means a separate paragraph.

Inline it when

You use it once, immediately, and it is under three lines. Then a subordinate clause is better than a display: "since the tangent segments from a point to a circle are equal, $x + y = c$" rather than a lemma about tangent lengths. A lemma layer that contains one-line facts used once is padding, and padding buries the two lemmas that matter.

The tell that you have chosen wrong

If you find yourself writing "as before" or "by the same computation", you should have stated a lemma. If your lemma statement contains the problem's specific numbers, you have not generalised enough. If you have four lemmas and the main argument is two lines, you have chopped the proof into pieces that are individually true and collectively unmotivated; merge until each lemma is a statement you would be willing to remember.

1.6 How much detail, and the rule for deciding

The rule

Write the step out if a competent reader would have to pause to reconstruct it.

That is the whole rule and it is operational, not vague, provided you fix who the competent reader is. They know everything you know. They have not seen your solution. They are reading at speed with a mark scheme. So the question is never "is this obvious to me" (everything is obvious to you, you just derived it) but "would this stop someone reading at speed".

Sort every step into three grades.

Instant. They see it as fast as they read it. Write nothing. Expanding $(a+b)^2$. Substituting a value. Reindexing a finite sum.

Reconstructible in seconds, given a pointer. Name the mechanism in a clause and move on. "Since σ is a bijection of $\{1,\dots,n\}$, $\sum_i \sigma(i) = \sum_i i$." "By the mean value theorem, $|\theta(x)| \le x$." "By Vieta, the sum of the roots is minus the coefficient of x^{2n+1}, which is 0." The clause is the whole cost.

Needs work. Write it out in full. Any induction. Any construction and its verification. Any estimate with a threshold. Any claim that a map is a bijection, that a set is nonempty, that a limit may be exchanged with a sum or an integral.

The five things you always spell out, however obvious

These are the five places where the corpus shows marks vanishing at the highest rate, and the cost of covering them is one sentence each.

1. Anything you divide by is nonzero. State it at the point of division.
2. Anything you square or unsquare, or on which you take a root: state the branch and dispose of the discarded one.
3. Any interchange of two limiting operations: name the mechanism and the uniformity.
4. Any map you call a bijection, permutation, isomorphism or well-defined function: give the verification, even if it is one sentence.
5. Any construction: verify that the constructed object satisfies every listed condition, in the order the problem listed them. In the 2009 A4 model answer the author answers "no" by exhibiting one set and checking conditions (a), (b), (c). Two get a sentence each and the third gets the entire argument, but all three appear.

What you deliberately do not write

This is the other half of the discipline and it is where you buy back the time. A full-marks solution is not a maximal solution. It is a complete one.

Do not write the search. In the 2021 B2 model answer the weight 4^k appears in the application of AM and GM with no explanation, because the design equation that produced the 4 happened offstage and is not part of the proof. You need to know how the 4 was found; the page does not.

Do not write the exploration. Small cases, tables, the failed first approach, the reason you abandoned it.

Do not write information about your objects that you will not use. In the 2012 B2 model answer the author fixes one face of the polyhedron and discards essentially everything else about it, because a lower bound is an existential claim and one witness suffices. Using the full three-dimensional structure bloats the proof, strengthens nothing, and creates places to be wrong.

Do not remark on things that do not matter. In the 2006 A4 model answer the number of local maxima is written as a sum of indicators and the expectation taken term by term, and the author says nothing about how the indicators interact. The interaction is real (adjacent peaks are impossible) and wholly irrelevant to a mean. A remark about dependence would be an invitation to wonder whether you needed independence, when linearity of expectation needs no such thing.

The two exceptions worth one line each

Two remarks are worth writing even though no grader demands them, because both are cheap insurance against your own errors.

First, when a hypothesis looks decorative, say what it does. The 2023 A2 model answer closes with the observation that the evenness of n is load-bearing: for odd n the same computation gives $r^2 = -1/(n!)^2$ and the two extra roots are non-real, so the answer would be different. Writing that sentence proves to yourself, on the page, that you used the hypothesis, and a hypothesis you never used is a klaxon that your proof is wrong.

Second, when the problem asks for an existence and your argument also gives uniqueness, or vice versa, say which you proved. One clause.

1.7 The standard sentence templates

Every structural move in a proof has a sentence that performs it. Learn them as fixed forms, the way you learnt the quadratic formula, so that in the arena the sentence arrives already assembled and you spend your attention on the mathematics. Each template below comes with the promise it makes and the way candidates break it.

Announcing the answer

> The answer is X.
> The largest such k is n^n.
> The selection is possible if and only if [condition].
> The answer is no; indeed [object] satisfies (a), (b) and (c).
> The answer is yes. We prove more, namely [stronger statement].

Promise: this is the complete answer, and everything below establishes it. *Broken by:* announcing a partial answer, for instance a family of solutions when the question said "find all", and then never widening it.

Announcing the plan

> We prove two statements: [first] and [second]. Together these say [the answer].
> We first show [A]. We next show [B].
> It therefore suffices to prove [reduced claim].

Promise: the proof has exactly these parts and no others. *Broken by:* announcing two parts and delivering one, or by a third argument appearing that was never advertised.

The phrase "it suffices to show" is the workhorse of reduction and it is doing logical work, so it must be true in the strong sense: the reduced claim must *imply* the original, and the implication must be visible in the same sentence. In the 2008 B6 model answer the sentence is "it therefore suffices to determine when $M_{n,k}$ is invertible over $\mathbb{F}_2$", and it arrives immediately after the two lines that identify the permutation count with a determinant modulo 2. The sentence announces that the combinatorics is over. Do not use "it suffices" as a hopeful gesture towards a claim you find easier; use it when the reduction is proved.

Suppose for contradiction

> Suppose, for contradiction, that [negation of the claim].
> ...
> These two statements are incompatible, and the contradiction shows that [claim].

Promise: you will derive two facts that cannot both hold, and you will say which two. *Broken by:* deriving something that merely looks wrong, or reaching a contradiction and not naming it. The 2023 A3 solution names its contradiction with total precision: the ratio $|g/f|$ is bounded by a fixed finite $\tan r$ on $[0, r)$, and simultaneously it tends to infinity as $x \to r^-$. Two explicit statements, incompatible, done.

Use contradiction only when the negation gives you something to work with. If your first line after "suppose not" is a restatement of the claim with a minus sign, you have gained nothing and a direct argument is shorter. In the 2015 A4 model answer the author knows the published contradiction argument works and chooses a direct case split instead, because the direct route produces the extremal pattern as a by-product and he needs that pattern for the second half of the problem. Choose the route that pays twice.

Without loss of generality

> Since [map] preserves [the hypothesis] and [the conclusion], we may assume without loss of generality that [normalisation].
> Replacing p by $-p$ changes neither $p(x)^2$ nor the degree, so we may assume $a_0 > 0$.
> The expression is symmetric in j and j', so assume $j \le j'$. (This is used only at Step 4, to ensure $m - j > 0$.)

Promise: the restricted case implies the general one. *Broken by:* omitting the map. A WLOG without a named map is an assumption wearing a disguise, and it is the fastest way to convert a correct proof into a 2.

Three obligations, every time: exhibit the map, check it preserves the hypotheses, check it preserves the conclusion. Then a fourth, optional but excellent: say where the assumption is used. The 2024 B5 model answer adds a parenthesis naming the exact step that needs $k \le m$, so a grader can verify at a glance that nothing else depends on it.

A warning about the discrete case. WLOG claims that are automatic in the continuous plane are not automatic on a lattice. Rotation by an arbitrary angle preserves the plane and destroys $\mathbb{Z}^2$. In the 2010 B2 model answer the normalisation is restricted to the symmetries that actually preserve lattice points, quarter turns and translations, and the argument names the finite set of lattice points at the relevant distance rather than gesturing at symmetry. Before you normalise, ask which group you are actually in.

By symmetry, similarly

> The hypothesis is symmetric in the two colours, so the statement for green centres follows from the statement for red centres by exchanging the colours.
> The same argument with x and y interchanged gives [the mirror statement].

Promise: the other case is the same argument under an explicit relabelling. *Broken by:* using "similarly" across a case that is not similar. Before writing it, run the argument in your head under the relabelling and confirm that no step used the specific labels. If any did, the cases are not symmetric and you owe a second argument.

When the same symmetry is used more than twice, do not write "similarly" three times. Factor it out into a lemma. The 2006 A4 model answer turns a symmetry of permutations into one lemma about the relative order of entries at fixed positions and applies it three times, so that two boundary situations and one interior situation become three instances of one statement rather than three arguments. That is also what protects the boundary case, where an unstated symmetry argument silently breaks.

Small and degenerate cases

> The case $n = 1$ is immediate: [one clause].
> For $n = 1$ the single square contains $1 + 1 - 1 = 1$, consistent with the formula.
> The cases $n = 1, 2$ are checked directly.
> We first rule out [degenerate branch]. In this case [short argument], so the implication holds vacuously. We may therefore assume [the good branch].

Promise: the case is settled, not skipped. *Broken by:* writing "the case $n = 1$ is trivial" when it is the case where your formula divides by zero. If it is immediate, the clause proving it is one line, so write the line.

Case splits

> We distinguish three cases. If $q = 1$, then .... If q has at least two distinct prime factors, then .... Finally suppose $q = p^k$; then ....

Promise: the cases are exhaustive and mutually handled. *Broken by:* a gap between the cases, or by an ordering that puts the hard case first and buries the easy collapses. Order the cases so the collapses come first and the sharp case last; a grader who has seen two cases die in a line each reads the third one knowing that is where the content is.

State the exhaustiveness if it is not visually obvious: "these three cases cover every positive integer q".

Both directions

> *Necessity.* Suppose such a selection exists. Then ..., hence [condition].
> *Sufficiency.* Suppose [condition]. Define ...; then ...
> These work: .... Conversely, suppose f satisfies the condition; we show f is one of these.
> We first show every value lies in S. Conversely, for each $m \in S$ we exhibit a triple with $X = m$. Hence the set of values is exactly S.

Promise: you have proved an equality of sets or an equivalence, not one inclusion. *Broken by:* proving the obstruction and asserting the construction, or the reverse. Each half is roughly half the marks and the verification half is usually the cheap one, which makes omitting it the worst trade on the paper.

Induction

> We prove [claim] by induction on n. For $n = 1$, [base case]. Suppose the claim holds for n; then [step]. This completes the induction.

Promise: base case, inductive step, and a claim whose range you have stated. *Broken by:* a base case that is asserted rather than checked, a step that quietly needs two previous values while the hypothesis supplies one (say "by strong induction" and mean it), or an induction whose claim is not stated as a claim at all.

Closing

> Therefore the smallest possible r is $\pi/2$.
> Combining the two parts, the set of values is exactly S.
> Hence the maximum is 3, attained by [object].

Promise: the question, as asked, is now answered. *Broken by:* finishing in your translated language and leaving the grader to translate back.

1.8 Three kinds of question, three obligations

Read the command term first and let it fix the shape of the write-up before you write a word.

Evaluate, compute, find the value

Obligation: produce the value, and prove it is the value. The answer goes on line one and in a box on the last line, and the two must agree. Between them is a derivation, and a derivation is graded as a chain: every link defensible.

The failure mode here is not incompleteness but an undefended step in the middle, because the presence of a correct final number makes the write-up feel finished. It is not. A correct value with an unjustified division scores 2, the same as no value at all.

Buy yourself a numerical check. In the 2024 B4 model answer the author solves an affine recurrence, writes the closed form, and adds a parenthesis: check $E(1) = 1$, $E(2) = 3/2$. That parenthesis costs six characters and catches the two most common errors, a lost fixed point and an off-by-one in the exponent. Whenever your answer is a formula in n, evaluate it at the smallest legal n and compare with a direct count. Whenever it is a number, estimate it crudely by a different route and confirm the order of magnitude.

Prove, show, deduce

Obligation: a complete chain from the stated hypotheses to the stated conclusion. There is no answer to state, so Slot 1 becomes a one-sentence plan instead: "We prove two statements", or "We show that [reduction], from which the result follows".

The failure mode is the unstated reduction. If your first move changes what you are proving, say so in the sentence that makes the move, and say why the new statement implies the old one.

Watch for hypotheses you never used. At the end, list the hypotheses and tick them. An unused hypothesis means one of three things: you have proved something stronger (say so, in one clause, and enjoy it), you have used it invisibly (find where and name it), or your proof is wrong (much the most likely). The 2011 A6 model answer makes this an explicit step: after establishing the main convergence the author observes that one hypothesis is still unspent and one claim still unproved, and writes "it remains to show $0 < b < 1$", then spends the unused hypothesis on exactly that. The audit is a move, not an afterthought.

Determine all, find all, for which n, what is the largest

This is the one that carries a completeness obligation, and it is where the corpus shows the most marks lost by strong solvers. **A "determine all" question is two theorems.** A grader marks them separately.

The protocol is fixed:

1. State the answer set completely.
2. *Exhibit.* Give the objects, explicitly. For an infinite family, give it as a formula in a parameter with the parameter's range stated.
3. *Verify.* Check that each exhibited object satisfies every condition in the problem. This half is cheap and it is worth roughly half the marks.
4. *Exclude.* Prove nothing else works. This is the other theorem, and it usually needs different tools from the construction.
5. Close by combining: "hence the set of values is exactly S".

Two structural notes from the model answers.

Write the constructions first when they are short, because they calibrate the obstruction. In the 2012 A5 model answer the author lists the working pairs, verifies them, and only then builds the mechanism that excludes everything else. Whatever mechanism you build must not exclude the objects you have already verified, so having them on the page is a live check on the second half as you write it.

Write the bound first when the bound is short and the construction is long. In the 2010 A1 model answer the counting bound is put on record first, isolated and short, because once a grader accepts the bound the construction is mechanical verification.

The threshold variant deserves its own warning. If the answer is "exactly those p with $p \le 1/4$", the counterexample half must be a family indexed by the problem's quantified parameter, not a single configuration. A single configuration can never refute a value just above the threshold. In the 2022 B5 model answer the failure half is a family of weights depending on n, precisely because the statement quantifies over all n; testing $n = 1$ and concluding a threshold of $1/3$ is the trap.

The yes-or-no variant deserves a strategic note. Before you pick a direction, price both. In the 2009 A4 model answer the author notes that answering "no" costs one object plus three verifications, while answering "yes" would cost a full general algorithm. Decide which you believe, then decide which is cheaper to write, and if those disagree spend two minutes deciding which one is true before spending an hour on either. The 2009 A6 model answer does the classification move explicitly: name the true theorem the assertion is generalising, name the upgrade, recall the standard pathology, and only then commit. Two minutes of classification beats an hour of trying to prove a false statement.

1.9 The final read-through

Budget three minutes per problem and treat it as scored, because it is. You are not proofreading for typography. You are running a fixed checklist against the five places marks disappear.

1. *First line and last line.* Does the first line state the answer, does the last line state the answer, and are they the same object? On a "determine all" problem, is the last line an equality of sets rather than one inclusion?
2. *Domains.* Scan for every division bar, every root, every logarithm, every inverse, every arctan, every substitution of the form $a = \cos b$. At each one, is there a clause on the page saying why it is legal? If not, add the clause now; it is one sentence and worth six marks.

3. *Weasel words.* Find every "clearly", "obviously", "it is easy to see", "similarly", "and so forth", "by inspection". For each, either delete the word because the step is instant, or replace it with the clause naming the mechanism. If the honest answer is that you cannot supply the clause, that is a knowledge gap and you have just found it, which is worth more to you than the mark.
4. *Every WLOG has a map.* And the map preserves both hypothesis and conclusion. And nothing downstream uses the labelling you claimed was arbitrary.
5. *Quantifiers and thresholds.* Every "for large n" has a named N. Every constant is defined before it is used. Every claim states the range of n it holds for. Two asymptotic facts combined somewhere have a max of their thresholds written down.
6. *Cases cover everything.* List your cases and check exhaustiveness out loud. Then check the boundary of each case and the smallest instance of the whole problem. $n = 1$, $n = 2$, the empty set, the zero polynomial, the degenerate triangle, the constant function. The corpus is full of solutions that are uniformly correct for $n \geq 3$ and false at $n = 2$, and a grader who checks the smallest case will find it.
7. *Hypothesis audit.* Did you use every hypothesis? Name where each one was consumed. An unspent hypothesis is either a stronger theorem or a hole.
8. *Deliverable audit.* Reread the question sentence itself. Count the things it asks for. Count the things your solution delivers. If the question said "maximum", you owe a bound and an attaining example. If it said "exactly", you owe two inclusions. If it said "unique", you owe existence and uniqueness.
9. *Sanity number.* If the answer is a formula, evaluate it at the smallest legal parameter and check against a direct count.

Do this in the order given, because it runs cheapest first. Items 1 to 3 take forty seconds and catch most of what there is to catch.

One habit to install alongside it: write the checklist items *as you go*, not only at the end. The point of the six-slot skeleton is that most of these checks become impossible to fail. If your degeneracy clearance is Slot 2, item 2 has nothing to find. If your WLOG template forces the map, item 4 has nothing to find. The read-through exists to catch what the templates missed, and as your templates harden it should find less and less. When it finds nothing three papers in a row, you are writing at the level the vow requires.

1.10 Circuit Bank

- **WHEN** the problem asks for a value, a set, a largest k, or a yes or no **DO** write the complete answer as the first line, before any argument **BECAUSE** it converts the proof into a verification with a fixed target, and it makes the halves separately markable
- **WHEN** the problem says "determine all", "find all", or "for which" **DO** plan two theorems, exhibit and verify first, then exclude everything else, each announced **BECAUSE** each half is roughly half the marks and the verification half is the cheap one
- **WHEN** you are about to write "without loss of generality" **DO** name the map, state that it preserves the hypotheses, state that it preserves the conclusion **BECAUSE** an unnamed WLOG is an assumption, and a grader marks it as one
- **WHEN** you are about to write "similarly" or "by symmetry" **DO** run the argument mentally under the relabelling and confirm no step used the specific labels **BECAUSE** if any step did, the cases are not symmetric and the second case does not exist

- **WHEN** a step needs a fact you will use more than once, or that is cleaner with letters than with the problem's numbers **DO** state it as a lemma above the main argument, with a free parameter **BECAUSE** generalising removes the temptation to use irrelevant hypotheses and makes the load-bearing step checkable in isolation
- **WHEN** you rely on a standard result that competition marking may not license **DO** prove it in three or four lines from memory **BECAUSE** a named lemma you cannot cite is a knowledge gap, and four lines removes all doubt
- **WHEN** your route uses a tool whose domain excludes a case **DO** dispose of that case in a short paragraph before introducing the tool **BECAUSE** it makes everything downstream unconditional and removes the need to interrupt the main argument
- **WHEN** you divide, take a root, take a logarithm, or invert **DO** add the clause proving the argument is legal, at the point of use **BECAUSE** this is the single highest frequency gap in twenty years of model answers
- **WHEN** you write "for sufficiently large x" **DO** name the threshold explicitly in terms of a fixed evaluation point **BECAUSE** combining several implicit thresholds without showing they are compatible is where a correct idea drops to two marks
- **WHEN** you construct a map and call it a permutation or a bijection **DO** write the one sentence listing the values each branch takes **BECAUSE** an unverified bijection is an unproved claim, however visible
- **WHEN** your answer is a formula in n **DO** evaluate it at the smallest legal n and compare with a direct count **BECAUSE** it costs six characters and catches the two commonest errors
- **WHEN** you finish and a hypothesis is unused **DO** find where it was consumed, or state the stronger theorem you proved, or find your error **BECAUSE** an unspent hypothesis is almost always a hole
- **WHEN** you translated the problem into another language at the start **DO** translate back in the final sentence, in the problem's own words **BECAUSE** a proof that ends in the wrong language leaves the last step to the grader

Chapter 2

The move catalogue

2.1 How to read this catalogue

A Putnam proof is not a continuous stream of thought. It is a short sequence of named structural moves, each with a cue that fires it, a canonical sentence that executes it, and two or three documented ways of botching it. There are about twenty. Every solution in the 2006 to 2025 corpus is a combination of them, usually three or four. Once you can name them, a solution stops being an act of invention and becomes an act of selection.

This is not a list of theorems. Cauchy and Schwarz is not a move. The pigeonhole principle is not a move either, not as a principle. The move is *pigeonhole with explicitly named boxes*, and the gap between the principle and the move is the gap between a 2 and a 10.

Every entry has the same four parts. **What it is**: the transformation, stated so you could apply it to a problem you have never seen. **The cue**: the guiding light the examiners planted. **The phrasing**: the sentence a full marks script actually contains, not a paraphrase, because in the arena you have no time to invent prose. **How it is botched**: the omissions that turn a correct idea into an 8, or into a 2.

Two facts about the marking. A grader marks the sentence, not the thought; a thought in your head and not on the page scores nothing. And almost every sentence separating an 8 from a 10 is short: eleven words for an interchange, one clause for a degenerate case, one named map for a *without loss of generality*. It is worth repeating: full rigour usually costs under twenty words, and omitting it costs two marks. Be lazy where laziness is free, never where it costs two marks.

2.2 Reduction: shrink the problem before you touch it

Move 1. The lemma cut

What it is. You take the load bearing step out of the main argument, restate it with letters instead of the problem's numbers, prove it once, and run the rest as bookkeeping on top. The lemma should mention as little of the problem as possible: the one extracted in 2006 B4 concerns a real matrix of rank r and a set of r column indices, and mentions neither the cube nor the alphabet the problem was about. That is exactly why it is checkable.

The cue. Four. You are about to do the same computation twice (2008 B4 states $P(a+t) = P(a) + tP'(a) + t^2R$ once and applies it with $t = p$ and with $t = yp^2$). The step is a different kind of mathematics from the rest (2006 B4 separates linear algebra from counting). The problem hands you an ugly constant (2019 B5 replaces 1008 by N in the first line; 2020 B1 proves its lemma for a general block, then instantiates). Or you are about to cite a named theorem.

The phrasing. General data, conclusion, instantiation, in that order.

> Lemma. Let p, q be nonzero polynomials of degrees m, n with $m + n \geq 1$. Then the

> coefficient of x^{m+n-1} in $D(p,q)$ is $(n-m)p_m q_n$. Proof. Only the top two terms of each expansion reach that degree, so … Now apply the lemma with …

The hypothesis $m+n \geq 1$ is not decoration: without it the degenerate pair of constants is silently swallowed. Every hypothesis the lemma needs goes in the lemma, where a grader cannot miss it. In 2019 A5 the order dropping lemma carries its range condition inside its own statement, because outside that range the claim is false.

Cite or prove? Cite what every finalist has met and what would cost a page: 2008 A2 writes that 2008 vectors in a 2007 dimensional subspace are dependent and does not reprove it. Prove, in four lines, anything whose hypotheses a grader might want checked: 2012 B5 proves the supporting line lemma, because a convex function can have a corner. If the honest answer is that the fact is a standard lemma you must know, say exactly that and state the lemma.

How it is botched. Baking the numbers in, so that at the inductive step the lemma cannot be applied to itself: 2017 B6 states its lemma for arbitrary positive integers with sum below p, because deletion preserves both hypotheses and a lemma about one specific list never could. Interleaving machinery with arithmetic, so a later slip takes the machinery marks with it (2006 A3). Dropping a hypothesis from the statement and using it inside the proof.

Where it appeared. 2006 A3, 2006 B4, 2006 B6, 2008 A5.

Move 2. Without loss of generality

What it is. You use a symmetry to fix a coordinate, an ordering, a sign or a scale, so the rest of the proof never mentions the case you disposed of. Every parameter eliminated here is algebra you never write. The four flavours, symmetry, scaling, translation and relabelling, are one move: exhibit a transformation, show it preserves everything the problem cares about, state what it lets you assume.

The cue. A hypothesis or conclusion visibly unchanged by some operation. Two indices appearing symmetrically. An unconstrained scale. Also, negatively: a case split about to double your write up for no mathematical content. 2022 A2 refuses two sign cases in one sentence: replacing p by $-p$ changes neither $p(x)^2$ nor the degree, so we may assume $a_0 > 0$.

The phrasing. Four parts: the map, what it preserves, what it achieves, where it is used.

> *The sloppy version.* Without loss of generality the polygon is the set of nth roots of unity.

> *The full marks version.* Set $Q(z) = (P(z) - b)/a$. Since $a \neq 0$ we have $\deg Q = \deg P$, and $Q(k) = \omega^k$ for $k = 1, \ldots, n$. The statement concerns only the degree, which is unchanged, so it suffices to prove it for Q.

That is 2008 A5: the grader wants the map exhibited and the invariance stated, since an unnamed transformation, or a scaling by $a = 0$, breaks the argument. 2015 A1 does the same for its normalising map $T(x,y) = (\lambda x, y/\lambda)$: $\det T = 1$ so areas are preserved, $(\lambda x)(y/\lambda) = xy$ so the curve is fixed, and T is linear so chords go to chords. Three clauses, and two integrals vanish from the solution. For relabelling the justification clause is still compulsory: 2012 A1 sorts its twelve values and says the claim concerns which three of them work, not their positions.

The discipline, which is the whole move. Three traps. First, the map must preserve the *hypothesis* too: 2009 B6 reduces to odd n by doubling an instruction sequence and checks that the modular instruction survives doubling, which needs a size bound and not merely the congruence. Second, the group may be smaller than you think: 2010 B2 concerns lattice points, where an arbitrary rotation is unavailable, so only the quarter turns may be spent. Third, and this separates the trained writer from the fluent one, some normalisations are illegal. 2013 B4 is the model: the reflex is to centre f, and the write up says outright that this is not available, because one quantity in the

conclusion is not translation invariant; the means are carried through and removed at the end. Before writing *we may assume*, ask which quantity in the conclusion the map moves.

How it is botched. Naked *WLOG* with no map, which 2015 A1 records as a lost mark for an ungeneralised proof. Claiming more than the group supports: 2022 B4 needs reversal, not merely rotation, to control which of two terms is smaller. Forgetting that the transformed object must still satisfy every hypothesis, as in the descent of 2019 B1. Applying an averaging to the integrand when the legality came from the domain, which 2021 A4 states explicitly.

Where it appeared. 2008 A5, 2009 B6, 2010 B2, 2012 A1.

2.3 Selection: name the object the proof will be about

Move 3. The extremal principle

What it is. You manufacture a canonical object out of an unstructured family by taking the largest, the smallest or the first, then use the fact that nothing bigger, smaller or earlier exists. Extremality is not scenery; it is the hypothesis doing the work. In 2021 B5 a directed cycle of minimal length is chosen, and minimality immediately buys chordlessness, since a chord would close a shorter cycle. It is the chordless structure, not the existence of a cycle, that makes the argument run.

The cue. A hypothesis quantified over a family with no distinguished member; a finite set you must say something about; a property that, if it fails, fails somewhere first. 2016 B3: since S is finite, choose A, B, C maximising the area of the triangle, and test every other point against that one fixed object.

The phrasing. Existence of the extremum is part of the sentence.

> The zero set of f in $[0, r]$ is nonempty, closed and bounded, so it has a least element ρ, and $0 < \rho \leq r$. Replacing r by ρ, we may assume f is nonzero on $[0, r)$ and $f(r) = 0$.

That is 2023 A3, and the point is that the auxiliary function about to be defined only exists where f is nonzero; shrinking the interval buys that hypothesis free. Compare 2022 A2, which uses the first failure in place of an induction: let k be the least index with $a_k \geq 0$, so $a_i < 0$ for all $i < k$, and the displayed identity forces a contradiction. Written that way it needs no base case at all.

How it is botched. Taking an extremum that need not exist, on an infinite set or an open interval; if you cannot justify it, you want Move 7. Ignoring the degenerate extremiser: 2016 B3 records that the area maximising triple may be collinear when S contains no genuine triangle. Building a competitor and not checking it is admissible: 2013 A2 must verify the rival satisfies *every* part of the defining condition, not only the one it was built for.

Where it appeared. 2009 B3, 2012 A1, 2012 B2, 2013 A2.

Move 4. Infinite descent and well ordering

What it is. You attach a nonnegative integer to every instance, show a bad instance produces a strictly smaller bad instance, and conclude none exists. The well ordering form, take the least bad instance and contradict it, is the same move and shorter to write because it has no base case.

The cue. A self similar family. A valuation you can divide out. A process that visibly reduces something.

The phrasing.

> If every vertex has positive level, divide all four by $1 + i$; the result is again a configuration of the same type and the minimum level has strictly decreased. Hence we may assume some vertex is a unit. The general case follows by multiplying back.

Two clauses there are the entire move. Without *again of the same type* the descent leaves the set it is descending in; without *strictly decreased* it need not terminate. That is 2019 B1, whose named pitfall is the third: reaching the base case and forgetting to multiply back.

The multiplicative form is shorter and more quantitative. 2025 A1 traps a product of the step ratios inside a fixed integer, so only finitely many steps can be nontrivial. One inequality valid for every k replaces *the quantity is nonincreasing, so it stabilises*, which is not enough on its own: stabilisation becomes a conclusion only once you add that the decrease is strict.

How it is botched. Descending out of the set. Stopping at the base case. A measure that is not a nonnegative integer or does not strictly decrease: 2014 B6 front loads a lemma whose only job is to supply that integer, because *keep splitting at mediants* is exactly where a write up silently assumes termination. Assuming a base case you were not given: 2013 B3 descends one element at a time from an arbitrary member to reach the empty set, because only then is the hypothesis about it meaningful.

Where it appeared. 2013 B3, 2014 B6, 2019 B1, 2022 A2.

Move 5. Pigeonhole with explicitly named boxes

What it is. Write down the objects, write down the boxes, count both, and say which count exceeds which. The principle is trivial; the move is the bookkeeping, and the marks live entirely in the bookkeeping.

The cue. A conclusion that two of something coincide, or that some quantity is small, with no indication which. Also any *eventually periodic* conclusion in a finite state space.

> *The sloppy version.* There are more values than intervals, so two must lie in the same interval.
>
> *The full marks version.* The $n+1$ numbers $G(x_1), \dots, G(x_{n+1})$ lie in (A, B), which is covered by the n half open intervals $[A + (t-1)\ell/n,\ A + t\ell/n)$ for $t = 1, \dots, n$. Hence two of them lie in the same interval.

That is 2011 A5: half open endpoints so the cover is exact and nothing is counted twice, both counts written, and downstream a check that the resulting difference stays inside the compact set where the derivative bound was proved. You may not quote a bound on an interval the pigeonhole does not guarantee you stayed inside.

The averaging form is a weak pigeonhole and should be written as one: since the $n+1$ nonnegative gaps sum to 1, at least one is at most $1/(n+1)$. That is 2006 B2, and the discipline is not to identify which gap, nor to assume the minimiser is unique; mere existence is all the argument needs.

One upgrade worth banking: pigeonhole gives a repeat, and a repeat gives *eventual* periodicity only. If the problem wants genuine periodicity, invert the step map, as 2018 B4 does in one line.

How it is botched. Off by one in the box count. Boxes that overlap or fail to cover. Losing strictness: 2013 A4 keeps an inequality strict on purpose and says why, because the final pigeonhole dies at equality and writing ≤ 1 out of habit breaks the last step.

Where it appeared. 2006 B2, 2010 A1, 2011 A5, 2013 A4.

Move 6. The probabilistic and counting existence argument

What it is. You prove something exists by showing an average has a value some member must attain or beat, and you never name the member. Two dialects: expectation, and double counting.

The cue. An existential conclusion with arbitrary data, so no explicit witness could be named; 2021 B3 reads exactly that off its statement and concludes at once that the proof must be an averaging argument. Also a count of locally defined events, which is a first moment question in a combinatorial costume.

> Let I_k be the indicator of the event that π has a local maximum at k, so the number of local maxima is $X = \sum_{k=1}^{n} I_k$. By linearity of expectation, $\mathrm{E}[X] = \sum_{k=1}^{n} \mathrm{P}(I_k = 1)$, and no independence is required.

That is 2006 A4. The final clause does real work: the indicators genuinely interact, since adjacent peaks are impossible, and a candidate who thinks the dependence blocks the step detours into the distribution of the count and loses the problem. Write the clause, not a paragraph.

Double counting: 2013 A1 assumes the conclusion fails, derives a local bound, sums it globally, then evaluates the same sum the other way to get an exact value. The two numbers collide, and every earlier step exists only to manufacture the second, comparable number. The certificate dialect converts an optimisation into two lines: 2014 A4 exhibits a polynomial dominating an indicator on the whole support, takes expectations, then exhibits the distribution making every inequality tight.

How it is botched. Asserting a combinatorial bound by picture: 2021 B1 proves its bound on a lattice point count rather than drawing it, because a grader wants the integrality argument. Assuming all orbits are full: 2020 B4 double counts over the set crossed with the group using only bijectivity, precisely because some orbits are genuinely short. Passing expectation through an infinite sum with no hypothesis, which Move 17 fixes in one clause.

Where it appeared. 2006 A4, 2006 B2, 2013 A1, 2013 A5.

Move 7. Continuity and compactness for existence of extrema

What it is. You obtain an object from a soft topological fact rather than by construction: a continuous function on a compact set attains its bounds; one that changes sign has a zero; one with no zero has constant sign.

The cue. The problem writes *minimum*, *maximum*, *the smallest r*, or asks you to prove a point exists. Treat *minimum* as a claim you must earn.

The phrasing. On a non compact domain, monotonicity does the job and you say so.

> Since $\Delta'(s) = \ln(s+1) - \ln s > 0$, Δ is strictly increasing on $[1, \infty)$, so the minimum is attained at $s = 1$ and equals …

That is 2019 B4, and the named pitfall is the temptation: computing the value at a candidate point and calling it the minimum, on an interval that is not compact. The intermediate value pattern carries its own compulsory clause, continuity of the function you apply it to. 2021 B3: let $I(p)$ be the integral over the disc indexed by p; show I is continuous and that its average over p vanishes; since a continuous function with no zero has constant sign, I vanishes somewhere.

Two refinements recur. A single point has measure zero, so a strict pointwise inequality needs continuity before it can move an integral: 2025 B2 picks an interior point where every factor is visibly positive and lets continuity spread the positivity over a whole disc. The reverse direction settles equality cases: 2006 B5 turns a vanishing integral of a continuous nonnegative function into pointwise vanishing, and adds one clause for the endpoint where the weight vanishes.

How it is botched. Using the extreme value theorem on a set that is not compact, or the intermediate value theorem on a function never said to be continuous. Finding the roots and forgetting the degree count: 2014 B4 and 2024 B5 exhibit an explicit list of evaluation points with a proved sign pattern, get one root per interval, then say that since the degree is n these are all the roots.

Where it appeared. 2006 B5, 2009 B2, 2012 A3, 2014 B4.

2.4 Ladders: induction and its relatives

Move 8. Strong induction, and strengthening the hypothesis

What it is. When the natural statement will not close, you prove a stronger one, chosen so that it is closed under the reduction step. Not generalisation for its own sake: the smallest enlargement that makes the recursion land inside its own hypothesis.

The cue. You reach the inductive step and the hypothesis is one notch too weak. Do not patch. Rewrite the statement, trying these four enlargements in order.

1. **Shift the comparison.** 2012 B4 proves $a_n > \log(n+1)$, not $a_n > \log n$, because the downstream step needs the bound in the form $e^{-a_n} < 1/(n+1)$; the natural looking version leaves a difference of undetermined sign. This is the most useful sentence in the section: write the inductive hypothesis in the shape the step will consume it.
2. **Buy margin.** 2020 B6 adds the $k = 0$ term so the target is $F(n) \geq 1$ rather than $F(n) \geq 0$. In a self similar induction the hypothesis is applied at a smaller index and then added to something, so starting from ≥ 1 leaves room where ≥ 0 gives only ≥ -1 and cannot be repaired.
3. **Free a constant into a parameter.** 2013 B5 promotes its fixed target set to a variable, because the peel produces the same problem with a different sink.
4. **Move to a closed class.** 2022 B5 abandons the literal statement, which has no reason to survive a convolution, and proves positive definiteness instead: the smallest class containing what was wanted and closed under the operation.

The phrasing. Announce the strengthening and its relation to the problem together.

> We prove by induction on $a + b$ the following statement, of which the given hypothesis is the case $a + b = 1$: for all nonnegative integers a, b and all s, t, if $t - s \in I_{a,b}$ then $h(t) - h(s) \in I_{a,b}$.

The induction variable need not be the size of the objects: 2009 B1 inducts on the largest prime factor, because the objects grow arbitrarily large while that invariant strictly decreases. A nonnegative integer that strictly decreases is the only requirement.

How it is botched. Proving the natural statement, finding the gap, and patching with an unproved claim; this is the commonest induction failure in the corpus. An unanchored base case, which 2009 B1 checks as a genuine floor. Asserting the conclusion outside the range the induction supports: 2019 B5 states its identity only on an explicit triangle of indices, because using it one step further is the false step that produces an answer with no correction term.

Where it appeared. 2009 B1, 2012 B4, 2013 B5, 2016 A5.

Move 9. Two variable and structural induction

What it is. Induction whose index is not one obvious integer: two parameters moving together, or objects built by a grammar (sets ordered by inclusion, polygons cut by chords, words extended by a digit).

The cue. A reduction that changes two things at once, or that changes the shape of the object rather than its size.

The technique. Collapse the two parameters into one nonnegative integer and induct on that: 2016 B5 inducts on $a + b$; 2014 B6 inducts on an explicit integer measure and quantifies over all admissible pairs simultaneously, so the step may change the interval as well as the point. Fixing the pair before the induction begins makes the hypothesis unavailable for the smaller pairs the step produces, which is the documented failure.

For structural induction, take the parameter the grammar actually reduces. 2014 B1 inducts on the number of digits, not on N itself, because N and $(N - d_0)/b$ differ multiplicatively and *smaller* would then need its own argument. 2007 A6 inducts on the number of sides and splits the step into a chord case, where cutting gives two smaller admissible polygons, and a no chord case, where the absence of a chord forces every boundary degree up; two cases each strengthening the other is cheaper than one case with exceptions. 2013 B3 makes the point in one line: recognise a graded poset and induct on rank.

Sometimes the right answer is to delete the induction. 2010 B6 replaces an induction over infinitely many cases by a structural decomposition of the index set, proved with the division algorithm, and prefers it because the decomposition makes the sharpness of the bound visible.

How it is botched. Fixing a parameter before the induction starts. Checking one base case when the recursion is k term: 2018 B4 writes three base cases for a three term recursion and states the range explicitly, because the identity it uses would otherwise need an index that does not exist. Inducting on a quantity that does not decrease, which fails silently until checked.

Where it appeared. 2007 A6, 2010 B6, 2013 B3, 2014 B1.

Move 10. Invariants and monovariants, written as a proof

What it is. You attach a quantity to a state and prove a statement about how one move changes it. An invariant does not change; a monovariant changes one way only. The move is not noticing the quantity. The move is stating it as a labelled claim with an explicit quantifier and verifying it on a single arbitrary move.

The cue. Any process, game or rewriting rule. Any *is it possible to reach* or *does it terminate* question. Any *for which n* question about a process, whose answer will be an invariant obstruction on one side and a construction on the other.

> *The sloppy version.* Each move improves the configuration, so the process must eventually stop.

> *The full marks version.* Define $Q = \prod_i a_i^i$. A single move multiplies Q by u^{k-j} with $j < k$, hence by a factor at least 2; and Q is bounded above throughout by [explicit bound]. Hence at most [explicit number] moves occur.

That is 2008 A3, which also carries the invariant half in canonical form: *Claim. For every prime p and every integer $m \geq 0$, the number of indices i with $v_p(a_i) \geq m$ is unchanged by a single move.* Then a verification exhausting the three cases a move can produce, and only afterwards a separate paragraph applying the claim. Labelling it lets the rest of the proof cite it by name instead of re deriving the case analysis each time.

For games there is a second obligation write ups forget: legality. 2020 B2 gives the pattern: name the family, verify the start belongs to it, and prove that from any member, whatever the opponent plays, you have a reply restoring membership. The omitted sentence is always the one guaranteeing that reply is legal.

How it is botched. Proving the invariant on a sub range and applying it outside: 2008 B4 extracts its invariant on a short range and adds the periodicity remark promoting it to every integer, which is named as the standard silent gap. Weights that are not strictly increasing, or a strict inequality never displayed; if the grader cannot see the ratio, the grader cannot see the induction. Forgetting the other half: in 2023 B1 the invariant gives only an injection, and surjectivity needs a separate descent, since an upper bound is not an answer.

Where it appeared. 2008 A3, 2008 B4, 2008 B6, 2011 A4.

Move 11. The algorithm with a termination proof

What it is. You describe a procedure and prove three things: every step is legal, it achieves the goal, and it stops. A strategy is a function on all reachable positions, not a line of play.

The cue. A game. A greedy. A process the problem describes and asks you to control. Any construction built in stages.

> Barbara's strategy is: [one sentence rule covering every cell the opponent could choose]. We check this is always a legal move, by induction on the number of rounds. We check it forces the target condition, by direct computation.

That is the 2008 A2 shape, and describing a strategy through example rounds is marked incomplete however correct the idea. Two refinements. Restate the win condition before inventing any strategy: 2008 A2 opens by saying the player wins if and only if the rows of the final matrix are dependent, so it suffices to force a dependence. And design the greedy to fail: 2008 A6 doubles a set until it cannot continue, and uses the failure of the continuation condition as the key hypothesis, so the process can only halt with a structural identity strong enough to finish. Stating the stopping condition as an equation, rather than lamenting the stall, is the move.

Termination is usually one clause and must not be omitted. 2025 A3: every move consumes a vertex, so the game ends. 2020 A4 records that effective moves occur almost surely, so the process terminates and the expectation is defined. 2013 B6 goes further and quarantines an awkward repetition rule first, replacing a game on histories by an equivalent game on positions.

How it is botched. Presenting one line of play as a strategy, or omitting a branch: 2023 A6 enumerates every branch, including the ones that look like edge cases, with the governing invariant stated once and cited in each. Proving the reply exists but not that it is legal. Omitting termination, so that the count or expectation you then compute is undefined.

Where it appeared. 2006 A3, 2008 A2, 2008 A3, 2008 A6.

2.5 Logical shape: how the proof sits on the page

Move 12. Contradiction versus contrapositive

What it is. A decision, taken before you write, about which of three frames the proof uses: direct, contrapositive, or contradiction. They are not interchangeable in cost.

The decision rule. Use contradiction when the negation hands you something you would not otherwise have. The corpus is unanimous on what that means: an infinite supply of inequalities (2018 A5 negates a statement whose conclusion offers no candidate); a single culprit object (2012 B3 supposes a neighbourhood bound fails, producing one specific team that lost throughout, far easier than reasoning positively about an arbitrary set); a rigid global formula to fight (2014 B3 supposes the rank is at most one and gets an explicit form, whereas *rank at least two* has no canonical starting point); or a finite object where the problem gave an infinite one (2006 A2 supposes a set is finite and lists it).

Use a direct proof when a construction exists. 2019 B3 is the clean example: the structure leaves exactly one candidate line, so the write up defines $w = (Q - I)^{-1}u$, notes $Q - I$ is invertible and $w \neq 0$, and verifies the identity. No determinants, no contradiction.

The genuine contrapositive has its own shape: assume the negation of the conclusion, extract the strongest structural consequence, and push it somewhere a concrete computed value contradicts it. 2017 B3 assumes a value is rational, extracts eventual periodicity of a digit sequence, and transports that periodicity to a *different* evaluation point where the problem already stated the answer. Fighting at the original point proves nothing, because there the hypothesis is unconstrained: a contradiction needs an external anchor to bite against.

The phrasing. Contradiction proofs benefit from being a bootstrap with numbered steps, not a single blow. 2010 B5:

> Suppose, for contradiction, that such an f exists. Step 1: $f' > 0$ everywhere. Step 2: $f'' > 0$. Step 3: $f(x) \geq \alpha x^2$ for large x. Step 4: hence $f'(x) \geq \alpha f(x)^2$ on a ray. Step 5: no positive function satisfies that on an unbounded interval.

There is no decisive idea there. The decisive act is the ordering: each step is the cheapest available upgrade of the last.

How it is botched. Stopping at a qualitative absurdity that is not absurd; the 2010 B5 pitfall is exact, since convex increasing functions with positive derivative exist in abundance, and the marks are lost between the picture and the rate. Negating carelessly: 2018 A5 records that weakening the negation loses the negative axis and makes it false. Wrapping a direct proof in *suppose not*, which costs ink and hides a construction you may need again.

Where it appeared. 2006 A2, 2010 B1, 2010 B5, 2011 A4.

Move 13. Case analysis with an explicit exhaustiveness statement

What it is. You split into cases and prove the split is a split: that the cases cover everything, and, if you are counting, that they do not overlap. The case list is part of the proof.

The cue. Any classification, any count over an inhomogeneous set, any sentence beginning *either ... or*.

> *The sloppy version.* There are three cases, according to which vertex the line is nearest.

> *The full marks version.* Let L be a line dividing T into two regions of positive area. Then L meets the boundary in exactly two points P and Q, and there is a vertex V outside $\{P, Q\}$ such that P and Q lie on the two closed sides through V. We take the three vertices in turn. Finally, no line is counted twice, because ...

That is 2017 B5, which spends four sentences proving the list complete and disjoint before computing anything, then reuses one computation three times. Both named pitfalls are fatal to a count: assuming the line crosses two sides in their interiors, and counting a cevian twice because it is visible from two vertices. The strongest formulation is 2015 B5's: *exactly one of the following holds*, with exhaustiveness proved by a case tree and exclusivity given its own paragraph.

Three disciplines. Let the hypothesis produce the split, as 2022 B3 does by writing out an alternation on a window long enough to contain both configurations it needs and no longer; hand chosen cases are the ones that fail to exhaust. Choose a classifying invariant with provably no third case: 2014 A1 classifies by the smallest prime factor p, giving $m/p \geq p$ with equality exactly when $m = p^2$, where parity or size would give an open ended list. And quarantine the degenerate branch first: 2022 B2 splits off its linearly dependent case before the main argument, because a later sentence is only legitimate outside it.

How it is botched. Overlapping cases in a counting problem, which destroys the constant. A hand chosen list that does not cover. Presenting only the generic configuration: 2016 B3 names both loose ends its construction leaves open and closes each with the smallest separate device available.

Where it appeared. 2006 A6, 2010 A3, 2014 A1, 2014 B4.

Move 14. Construction plus verification

What it is. For any problem whose command term is *determine all*, *find the largest*, *for which n* or *does there exist*, the answer is two theorems, marked separately. One half exhibits the object or family; the other proves nothing else works. A missing half scores near zero however good the present half.

The cue. The command term, and nothing else.

The phrasing. State the answer before either proof, so the reader knows what each half is for, and label the halves.

> The answer is $r = \pi/2$. *Attainment*: take $f = \cos$ and $g = \sin$, which satisfy (a) to (e) with $r = \pi/2$. *No smaller r works*: suppose for contradiction that f, g satisfy (a) to (e) with $0 < r < \pi/2$.

That is 2023 A3, and the accompanying note is the important one: the construction is worth almost nothing and gets one line, and the page goes on the bound. Announcing the split stops the grader hunting. The other standard labelling is necessity and sufficiency, as in 2024 B1, where the sum forces the parameter and an explicit permutation is then displayed and proved to be a bijection.

Why the verification half is where the marks are. A construction is a construction only once every stated hypothesis has been checked against it, not merely the one it was designed for: 2016 A6 verifies its global constraint on the whole interval in a separate paragraph, because the bound it came from used only two points. The verification is also where the design is tested: 2025 A3 verifies four properties of a constructed involution, one per clause of the pairing lemma, in the same order, and a non canonical choice would fail one silently. And on a *no* answer the verification is the whole proof: 2009 A4 answers a *must* question with one object and three checks, two of them a sentence each and the third carrying the entire argument. Its author's first decision was which of *yes* and *no* to buy, on the grounds that *no* costs one object and three checks while *yes* costs a full algorithm.

How it is botched. Submitting the easy half and asserting the other; 2012 A5, 2011 A4, 2022 B4 and 2025 B6 all carry this pitfall, in both directions. Verifying only at the smallest parameter value, or only generically: 2018 B1 records a construction that silently needs a parity condition. Presenting the discovery process as the proof: 2014 B2 uses an optimisation heuristic privately and writes an elementary verification, because justifying the heuristic costs more ink than the check it found.

Where it appeared. 2007 A2, 2008 B5, 2009 A4, 2010 A1.

Move 15. Counterexample construction

What it is. Disproving a universally quantified claim by exhibiting one instance and verifying it. The write up order is fixed: exhibit, verify every hypothesis, compute, conclude.

The cue. *Prove or disprove. Must it be true.* Any hypothesis of the form *for all m*, which you read as *I get to choose m*.

> We disprove the assertion with $f(x, y) = y \sin(2\pi x)$. This function is smooth on the plane, so it satisfies the hypotheses. All four boundary integrals vanish, so the assertion would demand an interior point where both partial derivatives vanish. The second partial vanishes only at $x = 1/2$, and there the first equals $-2\pi y$, which is nonzero on the open interval.

That is 2009 A6, written in the order a grader checks it, and it never explains where the example came from; the design work stays out of the script entirely. Choose the instance costing least algebra: 2017 B1 takes the parameter value that collapses its transformation to the identity for free.

Two auxiliary uses are worth installing. Kill your own first instinct before committing: 2020 B2 disposes of a plausible but false invariant in one line. In a finished script that is optional; in the arena it is what stops you spending twenty minutes proving a wrong answer. And audit the hypotheses by deleting them: 2018 A6 states where each hypothesis is used, exactly once each, and constructs the configuration showing one cannot be dropped. A proof that never uses a stated hypothesis is a proof of a different, probably false, statement.

How it is botched. An example failing one hypothesis, most often a regularity condition on a closed region. Verifying part of the requirement, such as computing three of the four boundary

integrals. Over solving: 2015 B2 restricts to the residue class where the auxiliary quantity is unconditionally determined, because one witness is enough and resolving the general case is a harder problem nobody is grading.

Where it appeared. 2009 A6, 2012 A6, 2015 B2, 2016 A5.

Move 16. Bounding both ways to force equality

What it is. You never compute the quantity; you trap it. Three sub shapes, and you should know which you are in before writing a line. *Bound plus attainment*: prove $X \leq c$ for every admissible configuration, then exhibit one with $X = c$; this is Move 14 in numerical clothing. *The squeeze*: trap the quantity between two expressions with the same leading behaviour. *Forcing equality*: you have $X \leq c$ plus a rigidity that upgrades it.

The cue. *Largest*, *smallest*, *best possible*, or a limit you cannot evaluate. Also a hypothesis giving you a number and an inequality giving you a competing number, which is a contradiction waiting to be read off, as in 2010 B1 where Cauchy and Schwarz produces a bound below the value the hypothesis asserts.

The examples. For the squeeze, 2016 A2 traps a discrete extremal quantity between two bounds whose ratios to n share a limit, and declines to pin the quantity exactly, since the question asked only for a limiting ratio. 2011 A3 pins a constant exactly using $\sin x < 1$ above and $\sin x \geq 2x/\pi$ on $[0, \pi/2]$ below, both integrated against the same power. For forcing equality, 2010 B2 converts a bound on a difference of distances into an equality using integrality, and 2008 B3 records that equality in the total requires equality in every coordinate at once.

The free mark. Read the equality case off the inequality's own proof rather than searching for it. In 2016 A6 equality forces a double root at an interior maximum, which with the endpoint conditions determines the extremiser completely.

> *The bound.* ... Hence $S \leq 2/3$ for every admissible sequence. *Attainment.* Take $a_k = 3 \cdot 4^{-k}$. Then $\sum a_k = 1$ and $4^k a_k = 3$ for every k, so every application of the inequality above is an equality; a direct computation gives $S = 2/3$. Therefore the maximum is $2/3$.

How it is botched. Proving one direction cleanly and reaching for a weaker unrelated technique for the other, ending with a construction that only works past the true threshold; 2010 B3 avoids this by reading one counting inequality in both directions, which also certifies the threshold is exact. An equality case that is not admissible, because the bound used only some of the constraints. Proving the bound for a convenient sub family and extending illegitimately: 2006 B4 names this, since the change of basis in question does not preserve the relevant set.

Where it appeared. 2006 B3, 2006 B4, 2007 A2, 2007 B2.

2.6 Analysis discipline: where the marks actually leak

Move 17. Interchange of limits, sums and integrals

What it is. Swapping the order of two limiting operations, where that includes an infinite sum, an integral, a derivative and a limit. This is the single most common place a Putnam write up leaks points, and the cheapest to fix. A grader looks for exactly one sentence; eleven words is a typical price.

The cue. Any time you move a $\sum$ past an $\int$, a lim past a $\sum$, a derivative inside an integral, or reorder a double sum. If two symbols change places, you owe a clause.

The ladder of licences. Use the cheapest that applies. A heavier theorem is not safer; it is more surface area to be wrong on.

1. **The sum is finite.** Nothing to justify, and you say so in five words. 2013 A6: *all sums are finite, so the interchange is trivial*. That clause pre empts the only convergence question available.
2. **Everything is nonnegative.** Tonelli applies with no further work. 2021 B2 uses the clause on a double sum and refuses to prove finiteness first: nonnegativity makes both iterated sums equal in the extended reals, and finiteness is a consequence of the bound about to be proved, not a prerequisite.
3. **Absolute convergence, with a deliberately crude bound.** 2016 B6 opens with the bound as the licence for everything after it: the family is dominated termwise by a convergent double series, hence absolutely summable, hence summable by any partition of the index set into blocks. The bound is wasteful on purpose, since only finiteness is needed. 2014 A4 buys every rearrangement in a probability computation with one sentence about a finite third moment.
4. **Uniform convergence, with the tail exhibited.** 2015 B6 writes the uniform tail bound explicitly and integrates term by term. The instructive part is the trap it names: the Weierstrass M test fails there, because the natural majorant series diverges, so *by the M test* would be a false citation; the uniformity came from alternation. Equally, invoking dominated convergence without exhibiting a dominating function is marked down even when the conclusion is true.
5. **No theorem at all: never leave finite sums.** The laziest and therefore the best. 2017 B4 concerns a conditionally convergent series and removes the question entirely by manipulating the Kth partial sum throughout, writing *every step so far is an identity between finite sums; nothing has been rearranged*, then taking limits once at the end with one justified sentence per limit. Splitting such a series into two divergent pieces and cancelling formally is worth close to nothing.

The same discipline covers derivatives. 2012 A6 expects to differentiate a hypothesis three times in three parameters and refuses to justify each separately: it repackages the integral as a double antiderivative and states, once, that continuity makes it continuously differentiable. One regularity sentence covers all three differentiations and removes any need to differentiate under an integral sign.

> *The sloppy version.* Interchanging the order of summation and integration, we obtain ...
>
> *The full marks version.* Every term is nonnegative, so by Tonelli's theorem the interchange below is legitimate: $\sum_{n\geq 1}\sum_{k=1}^{n} = \sum_{k\geq 1}\sum_{n\geq k}$.

How it is botched. Swapping silently; this is the most expensive omission in the corpus relative to its cost, and it is what a marker on an analysis flavoured problem is actively hunting. Citing a theorem whose hypothesis fails; a false citation is worse than none, because it shows the check was not made. Rearranging first and justifying afterwards on a conditionally convergent series, where a grader is entitled to assume the worst. And certifying convergence of the object after computing it: 2014 A3 certifies that an infinite product converges to a positive limit before the telescoping evaluation; 2016 B1 and 2012 B4 prove monotone and bounded before substituting into a recursion, since solving for a limit not yet shown to exist is circular.

Where it appeared. 2006 A4, 2011 B5, 2012 A6, 2012 B4.

Move 18. Telescoping and summation by parts

What it is. You convert a sum into a difference of two evaluations by finding a sequence whose consecutive difference reproduces the summand. Summation by parts is the same move on a product, the discrete analogue of integration by parts, and its commonest disguise is multiplying a series by $1 - z$ to annihilate a constant step in the coefficients.

The cue. A summand built from the same raw pieces as some natural auxiliary quantity. Coefficients decreasing by a constant step. A product over an iterated map. A recursion with computable increments.

The phrasing, in four compulsory parts. Build the auxiliary sequence from the algebraic shape of the summand and say so: 2011 A2 constructs it from the same data rather than guessing a closed form from small cases; 2020 A4 substitutes $w(N) = (N+1)u_N$ and states in the same breath that the factor is chosen so the awkward coefficient cancels, which turns a rabbit from a hat into a choice. Prove the termwise identity as a displayed equation, with the range on which no denominator vanishes. Sum the finite range and use the word *telescopes*: 2019 B2 writes no ellipses at all, because a proved termwise identity plus that word is complete. Check the boundary index: 2011 A2 names the pitfall, since its identity is derived only for $n \geq 2$.

If the sum is infinite, sum to N first. 2016 B1 telescopes the finite sum to an explicit difference and then lets $N \to \infty$. A finite telescoping sum needs no convergence justification; passing to the limit termwise inside an infinite sum does.

For summation by parts, 2018 B2 sees coefficients decreasing by a constant step and reaches for the device that kills a constant step, subtracting a shifted copy, that is, multiplying by $z - 1$; the compulsory extra clause is to rule out the root $z = 1$ that the multiplication introduced. One further use worth banking: name a constant you cannot evaluate, prove only that it exists, and arrange the algebra so it cancels, as 2017 B4 does, taking the limit before the cancellation.

How it is botched. Writing the collapse with ellipses and never proving the termwise identity. Dividing by a quantity that vanishes somewhere in the range. Cancelling two versions of the same constant before taking limits. Presenting a guessed closed form with no derivation and inducting on it, which is correct but fragile, since a wrong guess leaves you with nothing.

Where it appeared. 2010 A6, 2011 A2, 2014 A3, 2014 B3.

2.7 Algebraic leverage

Move 19. The identity theorem and generic parameters

What it is. Two moves sharing one mechanism. The identity theorem: polynomials of bounded degree agreeing on an infinite set are equal, so a condition holding for all values of a continuous parameter collapses to finitely many coefficient conditions. The generic parameter: introduce a parameter to make an object invertible, prove the statement for all but finitely many values, and close by degree counting.

The cue. A condition quantified over a continuum. An expression visibly polynomial in some parameter. A step needing an inverse that might not exist. An asymptotic statement about a quantity you know is discrete.

The phrasing. The compulsory clause is the degree bound, stated before the vanishing is used.

> For $r > 0$ we have $M_P(r) = \sum_i r^i \lambda_i(P_i)$. The right hand side is a polynomial in r of degree at most …; since it vanishes for all $r > 0$, an infinite set, every coefficient vanishes.

That is 2009 B4, and its named pitfall is precise: you must first justify that homogeneity lets you pull r^i out of the integral, and you must say the expression is a polynomial of bounded degree, not merely that it has infinitely many roots. For the generic parameter, 2015 A6 is the model. Write $A_t = A - tI$ and $B_t = B - tI$; both sides are polynomials in t of degree at most n, so it suffices to prove equality for all but finitely many t, and since $\det(B - tI)$ is a nonzero polynomial, B_t is invertible for all but at most n values. The degree bound precedes the appeal to genericity, because that sentence is what a grader checks, and the shift is by tI rather than a perturbation of the other matrix, which would destroy the hypothesis.

One further application worth installing. When an analytic estimate meets a discrete quantity, upgrade the limit to an exact identity and leave analysis behind: 2007 A4 cages a convergent quotient

inside an interval containing at most one admissible value, deduces it is exactly constant beyond a threshold, and then applies the identity theorem. Limits alone never determine a polynomial. One explicit interval is a proof; a fog of epsilons is not.

How it is botched. Invoking *infinitely many roots* without saying the object is a polynomial of bounded degree. Confusing a congruence with an equality of remainders: 2009 B6 proves the identity in the polynomial ring first and then specialises with an explicit size bound, since the congruence alone is not the remainder. Perturbing the object that carries the hypothesis rather than the one needing invertibility.

Where it appeared. 2007 A4, 2007 B5, 2009 B4, 2009 B6.

Move 20. Degrees of freedom and dimension counting

What it is. You replace a combinatorial or geometric question by a statement about the dimension, rank or kernel of a linear map and let linear algebra count. Permutation counting has almost no leverage; rank over a field has a great deal.

The cue. More objects than the dimension allows. A condition that is linear once you name the right field. A count you would like to be the size of a fibre.

The phrasing. Name the translation, state the equivalence, and announce that the original objects are gone.

> Identify the three colours with the elements of the field with three elements, so a colouring is a vector. A face is admissible if and only if the sum of its three edge values is nonzero. The colourings to be counted are exactly the preimage of the set of vectors with no zero coordinate. Since the map is surjective, rank and nullity give every fibre size 3^{10}; there are 2^{20} admissible targets; hence the answer is $2^{20} \cdot 3^{10}$.

That is 2017 A6, which gives the answer in structured form because the structure is the proof. The load bearing claim is the equivalence, checked in all three cases rather than asserted, because the whole count is wrong if the rank is not full. The companion move is announcing the toolkit exchange out loud: 2008 B6 translates a permutation count into a determinant over the field with two elements, notes the signature may be ignored there, and writes *it therefore suffices to determine when the matrix is invertible*. The price of the exchange, an answer known only modulo two, is exactly what was asked for.

Two precision points the corpus flags as expensive. Rank identities valid over the reals are not valid in general: 2011 A4 uses only the inequality $\operatorname{rank}(MM^{\mathsf{T}}) \leq \operatorname{rank}(M)$, arranging both of its bounds to concern the same quantity so the contradiction sits in one display. And taking a real part can double a rank bound: 2009 A3 concludes rank at most *two*, not one, since the real and imaginary parts of a rank one complex matrix are two separate real rank one pieces.

How it is botched. Quoting a dimension bound without stating the dimension; 2008 A2 writes the number rather than leaving the grader to subtract. Multiplying determinants of rectangular factors, which 2023 B6 avoids by matching its tool to the shapes of its factors. Building a basis where a witness would do: 2009 B4 proves a linear functional nonzero by naming one witness and quoting rank and nullity.

Where it appeared. 2006 B4, 2008 A2, 2008 B6, 2009 A3.

2.8 The catalogue as a Circuit Bank

Twenty moves, and the point of naming them is that in the arena you do not search, you match. Read these until the left hand side alone brings the right hand side to mind, then write them under time pressure, because recognition is not production and only production is marked.

- **WHEN** a step is needed twice, carries a different kind of mathematics, or is stated with the problem's ugly constants **DO** cut it out as a lemma with letters and hypotheses inside the statement **BECAUSE** only a general lemma can be applied to itself inside an induction, and machinery marks then survive a later slip
- **WHEN** you are about to write *without loss of generality* **DO** name the map, what it preserves, what it achieves, and the step that uses it **BECAUSE** a normalisation is legal only if the map fixes hypothesis and conclusion, and that clause is what the grader checks
- **WHEN** an unstructured family with no distinguished member **DO** take the largest, the smallest, or the first failure, and say why it exists **BECAUSE** extremality is the hypothesis doing the work, and a minimal object forbids the configurations you cannot handle
- **WHEN** a self similar family closed under a size reducing operation **DO** descend on an explicit nonnegative integer, checking the reduced object is admissible and the decrease strict **BECAUSE** those two clauses are the whole termination proof, and you must then multiply back
- **WHEN** a conclusion that two things coincide, with no indication which two **DO** write the objects, write the boxes as an exact half open cover, count both **BECAUSE** off by one and overlapping boxes are the only failures, and both are visible in the bookkeeping
- **WHEN** arbitrary data and an existential conclusion **DO** average, and take the member at or beyond the average **BECAUSE** no explicit witness exists when the data are arbitrary, so the proof must be averaging or counting
- **WHEN** the problem writes *minimum* or *smallest* **DO** earn the word: monotonicity on a non compact set, compactness on a compact one **BECAUSE** attainment is a claim, and evaluating at a candidate point proves nothing about it
- **WHEN** the natural inductive statement will not close **DO** strengthen it in the shape the later step consumes, or move to a class closed under the reduction **BECAUSE** patching a stalled induction costs more than restating it
- **WHEN** a reduction that changes two parameters or the shape of the object **DO** collapse to one nonnegative integer and quantify over all instances at once **BECAUSE** fixing an instance first makes the hypothesis unavailable for the smaller ones the step produces
- **WHEN** a process, a game, or a rewriting rule **DO** state the invariant as a labelled claim, verify it on one arbitrary move, and prove the restoring reply is legal **BECAUSE** legality is the sentence write ups omit
- **WHEN** you are describing a strategy or a greedy **DO** write the rule in one sentence covering every position, then legality, correctness, termination **BECAUSE** a strategy is a function on all reachable positions, and an example round is not a rule
- **WHEN** the conclusion offers no object to construct **DO** use contradiction, since the negation supplies inequalities, a culprit object or a rigid formula **BECAUSE** contradiction earns its cost only when it hands you something
- **WHEN** you split into cases **DO** write *exactly one of the following holds* and prove exhaustiveness and exclusivity before computing **BECAUSE** overlap destroys a count and a gap destroys the theorem
- **WHEN** the command term is *determine all* or *find the largest* **DO** state the answer first, then two labelled halves, spending the ink on the half that is not obvious **BECAUSE** it is two theorems marked separately, and a missing half scores near zero

- **WHEN** you have constructed an object **DO** verify every stated hypothesis against it, not only the one it was designed for **BECAUSE** a construction motivated by one condition can silently violate another
- **WHEN** a universally quantified claim you doubt **DO** read *for all m* as *I choose m*, then exhibit, verify, compute, conclude **BECAUSE** one instance disproves a universal, and the design work never enters the script
- **WHEN** *largest*, *best possible*, or a limit you cannot evaluate **DO** trap it: bound plus attainment, or a two sided squeeze with a common leading term **BECAUSE** a bound alone never proves sharpness
- **WHEN** two limiting symbols change places **DO** name the cheapest licence in one clause at the point of use **BECAUSE** finite beats nonnegative beats absolutely convergent beats uniform, and a heavier theorem is more surface area to be wrong on
- **WHEN** the series is conditionally convergent **DO** never leave finite sums; manipulate the Nth partial sum and take limits once at the end **BECAUSE** staying finite removes the question that decides the problem
- **WHEN** a condition holds for every value of a free parameter **DO** say the expression is a polynomial of bounded degree in it, then collapse to coefficient conditions **BECAUSE** the degree bound is the clause a grader checks
- **WHEN** a combinatorial condition that is secretly linear **DO** name the field, restate the condition linearly, and count by rank and nullity **BECAUSE** permutation counting has no leverage and rank has a great deal

Two closing instructions, and they matter more than any single entry.

When you finish a problem, do not ask whether your solution was correct. Ask which of the twenty moves it used, in what order, and whether the model answer used the same ones. That is the juxtaposition: your chain of moves laid next to theirs. Where the chains diverge is the bifurcation point, and the move you failed to fire is the one to install.

And the phrasings here are to be produced, not recognised. Reading *every term is nonnegative, so by Tonelli the interchange is legitimate* and thinking *yes, that is right* is consumption, and consumption is not analysis. That sentence has to be self generated, in the arena. Copy each phrasing out once. Then close the book, write the same move on a different problem from memory, and check what you left out. What you left out is the mark.

Chapter 3

Inside the author's head

3.1 What the model answer deletes

A published Putnam solution is the last ten per cent of the work, printed: the transcript of a machine with the machine removed. Read one and you get an answer, a definition, a lemma, four paragraphs, a box. Read a hundred and you notice that the solutions do not lie, they simply never speak. They never say why the substitution was the one to try, which routes the author priced and dropped, why the lemma was stated before it was proved, or how the author knew that 2018/2048 was not a number to be estimated but an expression to be recognised. Every one of those decisions happened. None is on the page.

That silence is what makes model answers feel like magic. It is also entirely reconstructible. Across the 2006 to 2025 corpus the same mental operations recur with almost boring regularity; there are about ten of them, they fire in a near fixed order, and they have cues visible in the question text exactly as the moves of the previous chapter do. What follows is that machine, reassembled from the back end of 775 model answers and compressed into a protocol you can run under exam conditions.

Everything here is about deciding, not writing. Chapter 1 gave you the architecture of the page and Chapter 2 the catalogue of moves; this chapter sits before both, and it ends with a verification pass that also happens before you write. If you find yourself doing any of it after the first sentence of the script is on the page, you have run the machine in the wrong order and you will pay for it in ink.

3.2 The first sixty seconds

Read in reverse order

You do not read a Putnam problem from the top. You read it in the order of information density, which is roughly the reverse of the order it is printed in.

First the command term. *Determine all*, *prove that*, *evaluate*, *for which n*, *find the largest*. This phrase fixes how many separate theorems you owe, and Chapter 1 has told you that a missing half of a two part obligation scores near zero however good the other half is. The 2011 A4 author knew before touching the mathematics that a *for which n* question is graded as two proofs, a construction and an obstruction.

Second the conclusion, not the hypotheses. The conclusion tells you the shape of the last line, and that shape constrains every route which can reach it. The 2015 A5 author read the conclusion as a named theorem before starting and went looking for the machine whose output is that kind of statement. The 2013 A6 author read the word *nonempty* and the strict inequality and knew at once that the Fourier argument alone gives only ≥ 0, so the last step was planned before the first.

Third the hypotheses, and you count them. Fourth the constants, which get their own treatment below.

The three things you write down before you attempt anything

Three lines, on scrap, in under a minute. They are not working. They are the specification of the work.

Line one: the obligation count. How many separate deliverables the command term demands, and what each is. *Determine all*: obstruction plus construction. *Find the largest*: bound plus attaining example. *For which n*: the n that work, by construction, and the rest, by an invariant. *There exists*: one object plus a verification of every listed condition. This is a one line contract and the rest of your time is spent discharging it. The 2014 A6 author wrote the answer n^n and then two headings, achievability and bound, before proving either.

Line two: the hypothesis ledger. List the hypotheses and leave a blank next to each. Every hypothesis has a job; examiners do not include decoration. When you finish the solve, every blank should carry the line where that hypothesis is spent. The 2011 A5 author counted four hypotheses and four jobs before writing a line, and knew from the count alone that the proof had four moving parts and no case analysis. That is not tidiness but a live correctness check: a route leaving a strong hypothesis unspent is proving a different theorem, probably a false one.

Look at 2016 A5. You are given a finite group G generated by g and h, and the hypothesis that the order of g is *odd*. Every argument that never mentions that parity is wrong, because the conclusion fails for $G = \mathbb{Z}/2\mathbb{Z}$ with $g = h$ the nonidentity element.

> *The sloppy version.* Since g and h generate G, every element is a word in $g^{\pm 1}$ and $h^{\pm 1}$, and by collecting letters we may write it in the required alternating form; the length bound follows since G is finite.

> *The full marks version.* Put $T = \{gh,\ gh^{-1},\ g^{-1}h,\ g^{-1}h^{-1}\}$; an admissible expression is exactly a product of elements of T. Now $(gh)(g^{-1}h)^{-1} = g^2 \in \langle T\rangle$, and since the order d of g is odd, $g = (g^2)^{(d+1)/2} \in \langle T\rangle$; hence $h = g^{-1}(gh) \in \langle T\rangle$ and $\langle T\rangle = G$. (The hypothesis cannot be dropped: for $G = \mathbb{Z}/2\mathbb{Z}$ with $g = h$ the nonidentity element, no product of elements of T is g.)

The sloppy version is not merely thin. It is a proof of a false statement, and its falsity is detectable in ten seconds by the ledger.

Line three: the shape. One word, from a fixed list, described next.

Classify by shape, not by topic

The instinct after A level and Further Maths is to classify by topic: this is number theory, that is analysis. Topic classification is nearly useless in the arena, because it tells you which chapter of a textbook to open and the Putnam is not testing chapters. What you want is the *shape*: the logical form of the thing you must produce. The corpus separates into about eight, and the shape fixes the skeleton of the write up before you know a single mathematical fact about the problem.

1. **Image.** What values does this expression take. A property every value has, and a family hitting every value with that property. 2019 A1.
2. **Extremal.** Largest, smallest, best constant. A bound for all configurations, and one configuration attaining it. Never one argument doing both. 2014 A6, 2023 B4.
3. **Existence.** One named object plus a verification of every listed condition, in the order listed. Cheapest shape on the paper, because a construction needs no cleverness to verify.

4. **Impossibility.** One invariant, its value on the start state, a proof it does not change, its value on the target. 2008 B6.
5. **Classification.** Find all f, all n, all pairs. Two inclusions, both written, both announced; the verification half is cheap and not optional. 2008 B5, 2022 A1.
6. **Evaluation.** A limit, sum, integral, determinant. One deliverable, exact, with convergence settled once at the top.
7. **Game.** A rule crisp enough that a machine could execute it, a proof it is always executable, a proof it achieves the goal. 2008 A2, 2013 B6.
8. **Asymptotic.** The cheapest deliverable that suffices, and no more. 2017 A5 asks only for *arbitrarily large*, so the author proved recurrence and refused to reach for equidistribution.

Naming the shape also buys a prohibition, worth more than the skeleton. If the shape is extremal you know in advance that you will write two independent arguments and that fusing them is a mistake. The 2014 B1 author is explicit: fabricating a counterexample is a construction task, ruling out all others is a classification task, and forcing them into one argument obscures which tool is doing the work.

The constants are a fingerprint

Putnam problems contain no arbitrary numbers. A constant appearing in the conclusion and nowhere in the hypotheses was manufactured, so it is the output of a specific computation, so it names the method. Reading constants is the highest yield thing you can do in the first sixty seconds, and almost nobody does it. The cleanest instance is 2018 B6: show that the number of sequences of length 2018 with terms in $A = \{1, 2, 3, 4, 5, 6, 10\}$ summing to 3860 is at most $2^{3860}(2018/2048)^{2018}$.

> *The published line.* Note first that $\frac{2018}{2048} = \frac{1009}{1024} = \frac{1}{2} + \frac{1}{4} + \cdots + \frac{1}{64} + \frac{1}{1024} = \sum_{j \in A} 2^{-j}$.
>
> *What produced it.* The bound is (something)$^{\text{sum}}$ times (something)$^{\text{length}}$, and only a product turns into powers like that, so the argument is a product over the 2018 positions. The number $2048 = 2^{11}$ appears nowhere in the problem, so it was manufactured; cancel and it is $2^{10} = 1024$, which is what the largest alphabet element 10 produces in an exponent of 2. So compute the one quantity the alphabet and the base 2 make together, $\sum_{j \in A} 2^{-j}$, and see whether it is $1009/1024$. It is: $\frac{63}{64} + \frac{1}{1024} = \frac{1008+1}{1024}$. The problem is over; what remains is finding the argument in which that sum is raised to the power 2018.

The whole solve happened in the second block, before any mathematics was written, and the argument then writes itself: weight every sequence s in the full product set A^{2018} by $2^{-\sigma(s)}$, note the total factorises as $(\sum_{j \in A} 2^{-j})^{2018}$ by the distributive law, note that every sequence in S carries the same weight 2^{-3860} and that all weights are positive, so $|S|2^{-3860}$ is at most the total.

The candidate who does not read the constant treats 2018/2048 as a number to be bounded, reaches for Stirling estimates of multinomial coefficients, and burns the session. The margin is thin: relaxing to all positive parts gives about 10^{1158} against a target of about 10^{1149}, so a wasteful route is dead rather than painful. That check is worth sixty seconds on its own, because a thin margin says the intended argument is exact and structural rather than an estimate with room in it.

- **WHEN** a constant in the conclusion does not appear in the hypotheses **DO** factor it, and try to build it out of the objects the problem does name **BECAUSE** examiners do not invent numbers, so the constant is the output of the intended computation and therefore names the method

- **WHEN** the requested bound has an exponent equal to one of the problem's parameters **DO** look for a product with one factor per unit of that parameter **BECAUSE** only products become powers, so the exponent tells you the indexing set of the intended product
- **WHEN** the crude bound you can compute in ten seconds overshoots the target by orders of magnitude **DO** abandon every estimate based route immediately **BECAUSE** a thin margin certifies that the intended argument is exact or structural, and no amount of care will rescue a wasteful one

3.3 Small cases: what the data is for

What you are looking for

Compute $n = 1, 2, 3$. Everybody knows to do this. Almost nobody knows what they are looking at, so the data gets used for the one thing it cannot do, pattern guessing, and not for the five things it can.

One: the target. You are choosing a destination, not gathering evidence. The 2019 A5 author computed $p = 3$ and $p = 5$ in scratch, saw 1 and 2, conjectured $(p-1)/2$, and aimed the entire proof at that known destination. Not one numeral of the reconnaissance appears in the script.

Two: whether the answer is uniform in n. The most underused signal in the corpus. If the answer to a game is *Bob wins for all n*, that uniformity says the winning argument is itself uniform, so look for one invariant rather than a recursion over cases. The 2025 A3 author read it exactly that way and refused to tabulate further.

Three: whether the smallest case is false. Check $n = 1$ before anything else. The 2007 B6 author checked, found the claim false there (the left side is 1 while both bounds are $e^{-1/4} < 1$, and no constant repairs it since $1^{Cn} = 1$), and spent one sentence saying so. Naming the failing case costs a clause; hiding it costs the paper.

Four: the boundary index. Where does the general argument stop reaching? An identity derived for $n \geq 2$ leaves the $n = 1$ term of the sum unverified, and that is a real gap. 2011 A2 records exactly this pitfall.

Five: which factor is small. When an expression is a product, the small cases tell you the smallest values each factor takes, and that is what a construction is built from. In 2019 A1 you factorise

$$A^3 + B^3 + C^3 - 3ABC = (A + B + C) \cdot \tfrac{1}{2}\big((A-B)^2 + (B-C)^2 + (C-A)^2\big),$$

and the small cases are there to tell you that the second factor equals 1 for triples like $(n+1, n, n)$ and 3 for consecutive triples $(n+1, n, n-1)$. Those two facts are the entire construction: the value is then s or $3s$ with $s = A + B + C$ under your control.

When not to bother, and why the table never appears

Two situations make small cases a trap. The first is an exponential state space with no mechanism in the data: in 2025 A3 the state space has size 3^n, so tabulating $n = 2$ and $n = 3$ costs ten minutes and yields a table rather than a reason; the author computed the parity of the arena, used it as a signpost, and stopped. The second is underspecified data: in 2011 A2 the sequence (b_j) is not pinned down, so no closed form guessed from numbers can be trusted without the identity behind it, and the author built the auxiliary telescoping sequence from the algebraic shape of the general term instead.

It is worth repeating, because this is where the most marks in the corpus are thrown away. A table plus the words *and similarly in general* scores at most one or two out of ten, however obviously right the pattern is. Exploration is not evidence, and the 2011 A4 author decided in advance not to write a line of it.

The sloppy version. For $p = 3$ we get $n = 1$ and for $p = 5$ we get $n = 2$. The pattern is clear, so $n = (p-1)/2$ in general.

The full marks version. The answer is $n = (p-1)/2$. Write $D = x\frac{d}{dx}$, so that $D(x^k) = kx^k$ and hence $q = D^{(p-1)/2}s$ where $s = \sum_{k=0}^{p-1} x^k$. Lemma: if u is nonzero with $\operatorname{ord}(u) = m$ and $1 \le m \le p-1$, then $\operatorname{ord}(Du) = m-1$. *Proof.* Write $u = (x-1)^m v$ with $v(1) \neq 0$; the product rule gives $Du = (x-1)^{m-1}(\cdots)$ whose bracket evaluates at 1 to $m\,v(1) \neq 0$, using $1 \le m \le p-1$. Applying the lemma along the chain, the orders visited are $p-1$ down to $(p-1)/2$, all inside the permitted range.

Notice where the hypothesis lives. The range condition $1 \le m \le p-1$ sits inside the lemma statement rather than being mentioned in passing, because outside that range the claim is false: D annihilates $(x-1)^p$ in characteristic p. An unchecked induction here is not untidy, it is wrong.

- **WHEN** you have computed three small cases and seen a pattern **DO** write the conjectured answer at the top of the script and delete the table **BECAUSE** the table selects the destination and earns nothing; the marks are in the mechanism that reaches it
- **WHEN** a general derivation was carried out for $n \ge 2$ **DO** verify $n = 1$ separately, in one clause **BECAUSE** the boundary index is where the argument stops reaching, and it is the cheapest place on the paper to lose two marks
- **WHEN** the answer to a *for which n* question turns out to be the same for every n **DO** stop tabulating and look for a single invariant **BECAUSE** a uniform answer is a message that the intended argument is uniform, so a recursion over cases is the wrong shape

3.4 Generating routes, and the cost model that prunes them

You do not find a route, you price three

The largest structural difference between the model answer author and the strong amateur is that the amateur finds a route and starts, while the author generates two or three and prices them before committing. The 2010 A6 record is explicit: the author knew three complete proofs of the same statement and chose one, and the recorded pitfall is not a mathematical error at all. It is sinking twenty minutes into repairing the first route's uniformity failure when a route with no such failure was sitting there.

Route generation is not free association. It is a short fixed sweep over the shape you named in the first minute. For an impossibility: an invariant, a counting collision, a valuation argument, a minimal counterexample. For an evaluation: a closed form then a limit, a telescoping, a symmetry that kills the integral, a known asymptotic. For a bound: a named inequality with exponents matched to the target, a smoothing argument, a dual certificate, an explicit competitor. You are enumerating a list you already own from Chapter 2. Three candidates is enough; four means you have not classified the shape properly.

The five costs

Then you price them, and the pricing is the skill.

Ink. Literal length, which is time, which is the only resource you cannot buy back. The 2015 A3 author noticed the divisor sum was multiplicative and turned eight terms into three factors, not because the eight term route fails but because it is more writing and more places for an arithmetic slip. The 2009 A2 author refused to solve for two auxiliary functions the question never asked about: every extra function solved for is wasted ink on a timed paper.

Theory footprint. How much machinery must be true, cited correctly, and have its hypotheses checked. This decided 2010 A6: a proof needing measure theory risks an error under exam conditions, while the tail argument fits in six lines. Theory is not free even when you know it, because a cited theorem drags its hypotheses onto your page.

Slip risk. How many places can a sign, an index or a factor go wrong invisibly? The 2012 B6 author reduced modulo 2 *before* substituting $m = (p-2)/3$, because substituting first produces $(p-2)(p+1)/6$ and invites a division slip. Same answer, half the exposure.

Argument count. Here the naive instinct is backwards. Two clean arguments beat one contorted argument covering both jobs, and the corpus is unanimous. But two where one would do is waste: the 2014 B2 author collapsed a two piece sign comparison into one pointwise inequality valid everywhere by shifting the kernel by its value at the switching point. The question is not *how many* but *are they doing different jobs*.

Verification burden. Can a grader check this line by line, or must they accept a judgement? The 2010 B3 author chose an explicit frozen configuration over an abstract impossibility argument because a grader verifies in one line that no move is legal, whereas the abstract route must rule out every clever sequence. The 2014 A6 author could have written *tensor product* in one word and wrote coordinates instead.

Ranking, out loud, on the page

Rank them in a line of scrap, in words, then commit: *invariant, six lines, no theory; measure route, correct, needs uniformity check; block aggregation, medium.* You write it so that at minute twenty, when the chosen route bites, you look at the ranking instead of inventing a fourth. One refinement: where the shortest route needs a named theorem you may not remember under pressure, keep a mechanical fallback. The 2012 A6 author wrote the named route and held an elementary mean value theorem argument in reserve. Robustness to memory failure is a cost too.

- **WHEN** you have a route and the impulse to start writing **DO** stop and name two others, then price all three on ink, theory, slip risk, argument count and verification burden **BECAUSE** the amateur repairs the first route that stalls; the cost of generating a second candidate is two minutes and the cost of not generating it is the paper
- **WHEN** two candidate routes reach the same conclusion and one needs a named theorem **DO** take the one with the smaller theory footprint unless it is more than twice the ink **BECAUSE** a cited theorem drags its hypotheses onto your page, and every hypothesis is a place a grader can open a question
- **WHEN** your route would still work if one of the stated hypotheses were deleted **DO** treat that as a defect in the route, not a bonus **BECAUSE** Putnam hypotheses are never decoration, so an unspent hypothesis means you are proving something the question did not ask, and usually something false

3.5 The inversion: what would make this easy

Here is the operation that most reliably converts a hard problem into a bookkeeping problem, and it is almost never visible in the published text. You assume the answer, work backwards from the last line, and ask what structure would have to hold for the last line to be one step away. Then you go and build that structure.

This is not wishful thinking. It is a constraint propagation, and it is legitimate precisely because the answer's shape is usually forced long before the answer is proved.

Version one: read the endgame off a degree or exponent count. In 2011 B6 the author compared the required bound with the degree of the polynomial in play, saw the ratio two, and knew from that

alone that the endgame would be a multiplicity count against a degree p polynomial; everything else was reverse engineered from there. In 2011 B5 the author counted n^2 pairs available for Cauchy and Schwarz against a hypothesis of size n, computed $n^4/n = n^3$, saw it matched the target exactly, and knew the last step before writing the first.

> *The published line.* Applying Cauchy and Schwarz to the n^2 pairs $\sqrt{u}$ and $1/\sqrt{u}$ gives $n^4 \le (\sum u)(\sum 1/u) \le (\sum u) \cdot 6An$, so $\sum u \ge n^3/(6A)$.
>
> *What produced it.* The target is a cubic lower bound. I have n^2 objects and one hypothesis of size n. Cauchy and Schwarz on n^2 pairs produces n^4 on the left. Divide by the hypothesis and I get n^3. That is the target, so this is the inequality, and the only remaining question is whether the constant it produces is positive.

That last clause is not an afterthought. Taking $n = 1$ in the hypothesis shows $A \ge \pi/2$, so $B = 1/(6A)$ is positive; without that check the conclusion could be vacuous. The author knew the edge case in advance because the ledger had flagged it.

Version two: read the construction off the equality case. Having proved an upper bound, do not hunt for an attaining example: ask what the bound's proof would need in order to be tight, and the answer usually describes the example completely. In 2014 A6 equality requires n^n independent vectors, so those vectors must be a standard basis, which forces every row of each matrix to be a standard basis vector and indexes the whole family by tuples. The author does not search, the author reads.

Version three: reverse the optimisation until the feasible set is convex. In 2023 B4 fixing the output gives a curved, non convex constraint surface, while fixing the total resource gives a polytope. So the author swaps them: maximise the output at fixed resource, then take the least resource for which that maximum reaches the target. The two are equivalent because the output is strictly increasing in the final variable, and that equivalence is the one sentence the write up must contain.

Version four: choose the parameters of a named inequality by matching degrees. 2012 B2 picks the conjugate exponents in Hölder by matching degrees against the target rather than trying exponents in turn. There is no search: the target's homogeneity determines them uniquely.

Version five: build the object already in the shape you need. 2007 B5 knows the target carries coefficient 1 on the kth power, so it constructs the object in that shape rather than producing $\prod(kq - n + j) = 0$ and dividing by k^k afterwards. Same mathematics, one fewer chance to lose a factor.

- **WHEN** you have proved an upper bound and now need an attaining example **DO** write down the conditions under which every inequality in your proof is an equality, and read the example off them **BECAUSE** the equality case is a complete specification of the extremiser, so searching for it is wasted work
- **WHEN** an optimisation has a curved constraint set **DO** swap the roles of objective and constraint and check the swap is an equivalence in one sentence **BECAUSE** convexity of the feasible set is worth more than naturalness of the formulation, and the equivalence usually costs one monotonicity remark
- **WHEN** you can compute the degree, order of growth or dimension of both sides **DO** do it first and read off what the last step must be **BECAUSE** matching counts fixes the endgame before any algebra, and every intermediate step then has a known destination to be checked against

3.6 Reformulation: deciding the problem is about a different object

The most common move in the whole corpus, and the one that makes model answers look like they arrived from nowhere, is the decision that the problem is not about the object it names. The first

two lines introduce a new object, the original is never mentioned again until the conclusion, and you are left wondering how anyone knew. There are six triggers, all visible in the question.

Trigger one: the given condition is not closed under the operation you need. The sharpest and easiest to miss. In 2008 B5 the hypothesis is that $f(q)$ has denominator exactly b when q does. Every analytic tool available subtracts, and *denominator exactly b* is not closed under subtraction, while *denominator divides m* is a subgroup and therefore is. So the author weakens the hypothesis, runs the argument on the weak version, and repays the strength once, at $q = 1/b$. Dropping it costs one line; keeping it costs a coprimality lemma repeated at every rational. The recorded pitfall is a candidate who keeps the exact condition and asserts that $q + 1/n$ has denominator bn for all n, which is false: the script then contains a false lemma.

Trigger two: the constraint is awkward in the given letters and standard in others. 2016 A5 again: *exponents in* $\{-1, 1\}$ is unmanageable as written, but renaming the four permitted syllables as an alphabet T makes it a word length question, in which generation, reachability and the length bound are all standard.

Trigger three: the conclusion mentions a quantity the hypotheses do not. Introduce the missing object and restate every hypothesis in its language. 2007 B2 notices that both given conditions are integrals of f from the same point, and refuses to keep working with f once one antiderivative absorbs both as boundary values, which are far easier to reason about than a pair of integral equalities.

Trigger four: an existential over an infinite family. Direct construction would need you to control infinitely many objects at once; contradiction converts it into a claim about one fixed object. 2006 A2 assumes the set is finite, writes it as $\{b_1, \ldots, b_m\}$, and defeats that one list with a gadget parametrised only by $t = \max(B) + 1$, never by the individual b_i, which keeps the proof uniform instead of casework growing with m.

Trigger five: the object named is not the object constrained. In 2013 B5 the condition never reads the values of f on part of the domain, so those coordinates are free and factor out as n^k; the content of the step is certifying that nobody else's orbit consults them, and the certificate is minimality of the first entry time. When the constraint has smaller support than the object, count on the support.

Trigger six: the count is easier on a superset. In 2018 B6 the constrained set does not factorise and the full product set does, so the author enlarges the problem and designs the weight to be constant exactly on the constrained set. Enlarging is legal because the weights are positive, and that positivity is the sentence the grader looks for.

The discipline: equivalence, not implication

Reformulation is the highest yield move and the most dangerous, because the restatement must be an *equivalence* and candidates routinely prove only an implication. The 2007 A4 author checks that the substitution is invertible and degree preserving before using it, so the restatement is genuine, and then never mentions the original objects again until the final line.

> *The sloppy version.* Substituting $x = (10^n - 1)/9$, the problem becomes a statement about polynomials in x, which we now prove.
>
> *The full marks version.* The map $P(x) \mapsto P\big((10x + 1)/9\big)$ sends polynomials of degree d with rational coefficients to polynomials of degree d, and is invertible with inverse $Q(y) \mapsto Q\big((9y - 1)/10\big)$. It therefore suffices to prove the restated claim, and we do so; the original statement follows by applying the inverse map.

There is one further discipline, and it is a scoring rule rather than a mathematical one. Check the conditions the *question* stated, not the reformulated version. The 2019 B6 record is explicit:

having derived the construction from a reformulation, the author verifies against the question's own conditions, because that is what the examiner marks.

- **WHEN** the hypothesis as stated is not closed under the operation your method needs **DO** weaken it to the nearest closed statement, run the argument on the weak version, and repay the strength once at a single chosen point **BECAUSE** every analytic and algebraic tool you own operates on a closed family, and a lemma quietly applied outside its closure is a false lemma
- **WHEN** you introduce a new object in the first two lines **DO** state the correspondence as an equivalence, name the inverse, and say what it preserves **BECAUSE** a one way substitution proves a different theorem, and an unstated invariance is exactly the sentence a grader marks you down for
- **WHEN** a constrained sum will not factorise **DO** enlarge to a superset on which it does, and design the weight so it is constant on the original set **BECAUSE** positivity of the weights licenses discarding the excess, and factorisation is worth far more than tightness

3.7 Commitment, and the discipline of not thrashing

There is a moment, usually between minute five and minute ten, when you stop generating and start executing. That moment is the bifurcation point, and it must be a decision, not a drift.

The commit criterion is not *I think this will work*. It is: **you can state the last line**. If you can write down, in words, the sentence that will end the proof, you have a route. If you can only say *I will try induction*, you have a direction, and a direction is not a route. The 2014 B4 author decides the shape of the write up before computing anything: one lemma, an evaluation identity, plus one claim, a positivity estimate. The 2013 A1 author knows the proof must end in a numerical collision, because the only usable fact in the problem is a global total, and therefore knows every earlier step exists only to produce a second comparable number. Those are routes.

Once committed, the discipline is not to thrash. Thrashing has a signature: the *repair loop*, in which you hit an obstruction, patch it, and the patch creates a second obstruction. Two patches is the limit. At the second, ask whether the obstruction is local or is the route telling you it does not fit. An obstruction that recurs in a different disguise is structural, and structural obstructions are avoided, not repaired.

Set the re-decision point before you start, not when you are panicking. For a three hour paper: eight minutes to classify and commit, then one hard checkpoint at the twenty five minute mark at which you either see the last line from where you stand or go back to the ranked list and take the second entry. That checkpoint is the whole of the anti thrashing machinery, and the 2010 A6 pitfall exists because the counterfactual candidate had none.

One exception: deliberately designing a route to fail is not thrashing. In 2008 A6 the author built a greedy procedure knowing it would stall, because the point was never that the greedy reaches the goal, but that the only way it can stop is with a structural identity strong enough to finish the proof. Stating the stopping condition as an equation, rather than lamenting the stall, is the move. In 2013 B5 the author tried the peel, saw it produced the same problem with a differently sized sink, and rewrote the statement so the sink size is free. That is reading failure as data about the correct statement.

- **WHEN** you are about to start executing **DO** say the last line of the proof out loud; if you cannot, you have a direction and not a route **BECAUSE** a route is a path to a named destination, and everything you write en route can then be checked against it
- **WHEN** you have patched the same obstruction twice in different disguises **DO** stop and return to the ranked route list **BECAUSE** a recurring obstruction is structural, and structural

obstructions are avoided rather than repaired

- **WHEN** a procedure you designed stalls **DO** write the stopping condition as an equation before deciding the route has failed **BECAUSE** a stall is a hypothesis you did not have a minute ago, and in the corpus it is more often the key to the proof than the end of one
- **WHEN** a recursion produces the same problem with different parameters **DO** promote the parameters to free variables in the statement and prove the general version **BECAUSE** the generalisation is not ambition, it is the exact closure the recursion demands, and without it the inductive hypothesis does not apply to its own output

3.8 Knowing a lemma is worth isolating before you have proved it

Chapter 1 gave you the decision procedure for lemma or inline. What it did not explain, because it belongs here, is how the author knows to isolate a lemma *before the proof of it exists*. That looks impossible. It is not, because the lemma is generated by the intended endgame, not by the proof.

Three mechanisms, in decreasing order of frequency.

The lemma is what the endgame consumes. You reverse engineered the last line in the previous section; now look at what it needs as input. That input is your lemma, and you can state it before you can prove it. In 2012 B4 the author writes the difference $d_{n+1} - d_n$ using the recursion, and the two inequalities needed to make that difference negative are then visible on the page. The lemma is not guessed, it is generated.

The lemma is the closure the recursion demands. If your induction produces an instance of the same problem with a different parameter, the lemma is the statement with that parameter free. 2013 B5 promotes the target set to a free parameter; 2012 B4 shifts the index inside the comparison function and proves $a_n > \log(n+1)$ rather than the natural looking $a_n > \log n$, because the downstream step needs the bound as $e^{-a_n} < 1/(n+1)$. Patching later fails: $e^{-a_n} < 1/n$ does not beat $\log \frac{n+1}{n}$ and leaves the difference of undetermined sign. The hypothesis was written in the shifted form because the author knew, before the induction started, what the step would need.

You are about to write the same computation twice. 2008 B4 states $P(a+t) = P(a) + tP'(a) + t^2R$ once and applies it at two values of t; naming the increment as a free parameter is what makes one lemma serve both. If you catch yourself writing a computation already on the page, stop mid sentence and name it. Cheapest lemma trigger in existence, and it fires on a purely clerical observation.

A fourth signal protects marks. Isolate as a lemma anything whose hypotheses you would otherwise check inline in the middle of the main argument: hypotheses stated inside a lemma cannot be missed by the grader and, more importantly, cannot be missed by you at the second application. 2019 A5 puts its range condition inside the lemma statement because outside that range the claim is false.

One last thing the corpus does that nobody teaches: state the lemma in letters, not numbers, and choose the letters so it mentions as little of the problem as possible. The lemma extracted in 2006 B4 concerns a real matrix of rank r and a set of r column indices, and mentions neither the cube nor the alphabet the problem is about. That is why it is checkable. 2006 A3 gives the scoring reason: machinery stated in general survives even if a later arithmetic step slips.

- **WHEN** you have reverse engineered the last line **DO** write down what that line consumes as input and state it as a lemma immediately, before you can prove it **BECAUSE** the endgame specifies the lemma completely, so the lemma is derived rather than guessed and the proof of it has a fixed target
- **WHEN** you are writing a computation you have already written on this page **DO** stop, go back, name the varying quantity as a free parameter and state the general version once **BECAUSE**

one general statement costs less ink than two special ones and removes the second opportunity to slip

- **WHEN** a step needs a hypothesis checked and you are in the middle of the main argument **DO** lift the step into a lemma whose statement carries the hypothesis **BECAUSE** a hypothesis inside a lemma statement cannot be missed by the grader or forgotten by you at the second application
- **WHEN** a lemma you are about to write mentions the problem's specific numbers **DO** restate it with letters and delete every feature of the problem it does not need **BECAUSE** a lemma that mentions the problem is not separable from it, and separability is the entire point of stating it

3.9 Stuck: the perturbation checklist

You will get stuck. Being stuck is not a failure state, it is the normal condition of the middle twenty minutes. What separates the author from the amateur is that the author has a fixed list to run rather than a feeling to sit with. Run it in order. Each item is one minute.

1. Weaken the claim to exactly what you need. You are almost certainly proving something stronger than the question asked. 2012 B4 asks whether a limit is finite, not what it is, so the author reduces to monotone plus bounded and takes the crudest lower bound available, since a decreasing sequence needs any lower bound and not a sharp one. 2017 A5 asks for *arbitrarily large n*, which needs recurrence, not density or a limit. 2012 B2 needs strictness from a tool with an equality case and, rather than analysing that case, weakens the constant slightly and disposes of it in one line.

2. Strengthen the induction hypothesis. If the inductive step nearly closes, the hypothesis is nearly right, and the repair is a shift or a free parameter, as in 2012 B4 and 2013 B5. A step that misses by a bounded amount needs that amount added to its hypothesis.

3. Weaken the hypothesis to make it stable. The rarest reflex on this list. If your hypothesis is not closed under the operations your method uses, replace it by the nearest closed statement and repay the difference once. 2008 B5.

4. Manufacture a second measurement of the same quantity. One number and no way to move: get a second number for the same thing and let them collide. 2013 A1 is the model. Nonnegative integers are placed on the twenty faces of an icosahedron with total 39, and you must show two faces sharing a vertex carry equal numbers. Suppose not. Each of the twelve vertices sees five pairwise distinct nonnegative integers, so its sum is at least $0 + 1 + 2 + 3 + 4 = 10$, giving at least 120 overall. But each face has three vertices, so the same total is exactly $3 \times 39 = 117$.

> *The published line.* Since the five values at a vertex are pairwise distinct nonnegative integers, their sum is at least $0 + 1 + 2 + 3 + 4 = 10$; summing over the twelve vertices gives at least 120, whereas each face is counted three times, so the total is $3 \cdot 39 = 117$, a contradiction.
>
> *What produced it.* The only number in this problem is 39. A qualitative argument cannot use a number, so the proof must end in an arithmetic collision. That means I need a second expression for the same total. Summing over vertices is the only other way to sweep all the faces, and the distinctness hypothesis is exactly what turns each vertex into a lower bound. So the shape is fixed before I know any of the numbers: bound below by distinctness, evaluate exactly by double counting, collide.

Note the pitfall the corpus records: starting the minimum at 1 rather than 0 because the integers are described as nonnegative. Here the weaker bound still closes ($12 \times 15 = 180 > 117$), but the habit breaks the method on any problem with a tighter margin.

5. Go to the extremal case. The smallest counterexample, the minimal witness, the largest element, the first index where something fails. 2013 A2 contradicts a minimum with a smaller

witness; 2009 B3 classifies a closure condition by extremal peeling. The extremal element carries a free hypothesis no generic element carries, and that hypothesis is usually the whole proof.

6. Quotient out a symmetry. Which symmetries does the problem have, and have you spent them? The 2009 B4 author checks the symmetries of the domain before splitting into cases, knowing an anti invariant function integrates to zero over an invariant domain, and records this as the first thing he tries on any symmetric integral. 2013 A5 goes further and chooses notation that absorbs the symmetry: the vector area rather than the unit normal, since the normal is undefined for collinear triples and defined only up to sign otherwise, and both defects vanish if you never normalise. One notational decision, two paragraphs deleted.

7. Add a parameter. Introduce a free variable the problem did not mention, run the argument, then specialise. Highest leverage item on the list, and the one you will never reach for until it is installed as a reflex. 2007 B6 isolates its analytic estimate as a lemma with a free parameter and uses it twice, since the upper and lower products differ only in a constant inside an exponent. 2011 A2 keeps the initial constant unresolved as $(b_1 + 2)/2$ and substitutes once at the end, to see exactly what the initial condition was responsible for. Generalisation frequently makes the proof *shorter*: 2008 A5 replaces the roots of unity by an arbitrary $t \neq 0, 1$ because that removes every temptation to use $\omega^n = 1$, which is irrelevant.

8. Count the same thing twice, in the other order. The double count proper: a sum over rows equals a sum over columns. 2013 A4 gets its mean this way; 2013 A1 as above.

9. Enlarge the domain until it factorises. 2018 B6. A constrained object is hard, a product set is easy, and a positive weight lets you throw away the difference.

10. Return to the raw definition. When an algebraic manipulation fails at a degenerate point, manipulate the definition instead. 2011 B3 cannot make the square root manipulation work at $f(0) = 0$, so rather than forcing it, it returns to the difference quotient of f/g and multiplies back by g, which continuity keeps away from zero near 0. This is also your standing answer to *why is this the right object*: a circle is the set of points equidistant from the centre, so obviously the point of interest is the centre.

11. Change what you are inducting on. Not the objects but a derived invariant. 2009 B1 runs strong induction on a derived quantity rather than on the size of the objects, with the accompanying check that the invariant has a minimum the induction can reach, so the recursion is well founded rather than an unanchored descent.

12. Read the ledger. Which hypothesis have you not yet spent? In the corpus this resolves more stalls than any single technique, because an unspent hypothesis is a strong statement sitting unused. The 2018 B5 author uses the audit as a live correctness check during the solve, not as decoration afterwards.

Run one to twelve in order and you will clear most stalls in under ten minutes. If none of them fires, that is information too: go back to the ranked list and take the second route.

3.10 Verification, which happens before you write

A habit that separates a 9 from a 10 and costs almost nothing: the verification pass happens *before* the write up. Once the script is on the page you will re-read your own algebra and see what you meant rather than what you wrote. Verify while the argument is still on scrap.

Check one: a small case against the final formula. Not against the derivation, and preferably a boundary case. 2013 B5 has exponents like $n - k - 1$ in play, so the author expected an off by one and spent five seconds on $k = n$: there Y is empty, the lemma returns 1, and the count is $n^n = n \cdot n^{n-1}$, as it must be since every function then qualifies. 2023 A5's answer contains constants checkable no other way, so one case costing three lines catches every sign and factorial slip.

Check two: the ledger. Every hypothesis spent, with a line number. If one is unspent, either you have a stronger theorem, which happens, or a broken proof, which is likelier. Ask the falsification

question directly: *would this argument also prove the statement in a case where it is false?* Any argument for 2016 A5 that never mentions the order of g answers yes.

Check three: the seams. If the answer is a classification with cases, check the cases agree at their boundaries. 2022 A1 does this explicitly: as one parameter tends to its critical value from below, both thresholds tend to the same limit, so the excluded band closes and the answer matches the neighbouring case. Boundary continuity is a free proof read, because a classification whose cases do not glue almost always contains an algebra slip. Thirty seconds, and it catches a class of error re-reading never will.

Check four: verify by a different identity, not by re-reading. The 2006 A6 author knew that in a computation of that length the risk is a dropped factor rather than a wrong method, so he built in a check using a different identity. Re-reading confirms your intention; a second independent computation confirms the result. 2012 B6 checks the exact count at $p = 11$ to protect the block boundaries. 2017 B6 writes the count as (expected) minus (bias), so an error of sign or magnitude is visible without redoing the algebra.

Check five, for constructions only, and the most commonly failed check on the paper. Re-verify a constructed example against every hypothesis of the problem, not only the one it was built for. 2013 B2 records exactly this: you make one inequality tight and forget the object must also be monic, or positive, or of the right degree. The construction half of an extremal problem is the cheap half and it is thrown away this way more often than any other.

- **WHEN** you have a final formula with exponents or factorials in it **DO** substitute the smallest or most degenerate admissible value and check the formula by hand **BECAUSE** a boundary case costs three lines and catches every off by one, and off by one is the single most common surviving error in the corpus
- **WHEN** you are about to start writing the script **DO** run the hypothesis ledger and confirm each hypothesis is spent at a named line **BECAUSE** an unspent hypothesis is either a stronger theorem or a broken proof, and you must know which before you commit ink
- **WHEN** your answer is a classification with two or more cases **DO** take limits at each case boundary and check the cases agree **BECAUSE** cases that do not glue at their seams almost always indicate an algebra slip inside one of them, and the check costs thirty seconds
- **WHEN** you want to check a long computation **DO** recompute the answer by a different identity rather than re-reading the same lines **BECAUSE** re-reading confirms your intention rather than your result, so it cannot detect a dropped factor
- **WHEN** you have constructed an example to attain a bound **DO** check it against every hypothesis of the problem separately, including the ones it was not designed for **BECAUSE** an example that satisfies the tightness condition but violates a stated constraint scores nothing, and this is the most common way the easy half of an extremal problem is lost

3.11 What you delete, and why

The author, having done everything above, sits down and deliberately deletes most of it. Understanding exactly what gets deleted is what stops the model answer feeling as though it was written by an alien.

Seven categories go, each for a reason.

One: the exploration that produced the conjecture. Exploration is not evidence and consumes lines that earn nothing. 2011 A4, 2019 A5.

Two: the design equation that produced a constant. In 2021 B2 the scaling weight 4^k appears with no explanation, because the search for it happened offstage: the author already knew the answer and the extremiser from the equality condition and wrote the proof backwards. A full marks

solution does not need the design equation. You do, which is why it belongs in your notes and not in the script.

Three: the heuristic that located the extremiser. The 2014 B2 author used a bang bang and bathtub heuristic privately, but citing either by name would require justifying it from scratch, which costs more ink than a direct elementary comparison. The discovery tool and the presentation tool are different tools.

Four: machinery that found the route but cannot be made rigorous cheaply. The 2019 A4 author found the example by testing a translation invariant hypothesis on characters, but making that rigorous for merely continuous f needs distributions, so the finished proof verifies the example by hand.

Five: commentary on unused hypotheses. The 2011 A5 author noticed the second derivatives are never used and decided not to say so, since it is irrelevant to the grader. Note the tension with the ledger: you audit for yourself, then stay silent. The exception is an unused hypothesis conspicuous enough that a grader will wonder whether you missed a condition, in which case one clause naming the redundancy converts a worry into a demonstration of understanding (2012 A3).

Six: generalisation, unless quarantined after the proof is closed. 2018 B6 exposes the free parameter and shows the problem's constant is optimal for the method, but in a remark, after the required proof is complete, so the grader is not invited to demand justification of a minimisation the problem never required. 2013 A6 separates its free strengthening the same way.

Seven: pictures. The 2022 A1 author draws none, because a picture cannot certify the boundary case $|a| = 1$. Draw all you like on scrap; a diagram in the script is decoration unless it is the object itself.

Against those seven, the things *never* deleted however obvious they feel: the hypothesis check before a citation, the statement that both directions have been done, the degenerate case, the sentence naming the map behind a *without loss of generality*, and the nonvanishing of anything you divided by. Each costs under twenty words and each is worth two marks.

The rule that generates the whole list, and the one sentence to carry into the arena:

> Delete anything whose absence leaves no question a grader can ask. Keep anything whose absence leaves a question a grader can ask, however obvious the answer.

That is the entire detail calculus. Exploration leaves no question, because the conjecture is proved on the page regardless of where it came from. A missing hypothesis check leaves a question, because nothing on the page settles it. Apply the test to each line as you write and the length of your solution becomes an output rather than a decision.

3.12 The arena protocol

Everything above, compressed to something you can run under time pressure. Numbered, because there is a protocol.

1. **Minute 0 to 1. Read backwards.** Command term, conclusion, hypotheses, constants. Write the obligation count. Write the shape in one word.
2. **Minute 1 to 2. The ledger.** List every hypothesis with a blank next to it. Count them. The count is the number of moving parts.
3. **Minute 2 to 4. Constants and small cases.** Factor any manufactured constant and try to build it from the problem's objects. Compute $n = 1, 2, 3$ unless the state space is exponential or the data underspecified. Check whether the smallest case is false. Write the conjectured answer.
4. **Minute 4 to 6. Invert.** Assume the answer. Ask what the last line must be. Count degrees, dimensions, orders of growth on both sides and see what has to match.

5. **Minute 6 to 8. Generate and price.** Name three routes. Price each on ink, theory footprint, slip risk, argument count, verification burden. Rank them in words on scrap. Commit to one, and confirm you can say its last line.
6. **Minute 8 onwards. Execute on scrap, not on the script.** State the lemma the endgame consumes before proving it. If you write a computation twice, stop and parametrise it.
7. **At any stall, run the twelve.** Weaken the claim; strengthen the induction hypothesis; weaken the hypothesis for stability; get a second measurement; go extremal; quotient a symmetry; add a parameter; double count; enlarge until it factorises; return to the definition; change the induction variable; read the ledger.
8. **At minute 25, the checkpoint.** Can you see the last line from where you stand? If not, return to the ranked list and take the second route. Two patches on the same obstruction is the limit.
9. **Before writing, verify.** Boundary case against the final formula. Ledger complete. Seams glue. Recompute by a second identity. Constructions checked against every hypothesis, not only the one they were built for.
10. **Write.** Answer first. Setup. Lemmas in letters. Main argument. Degenerate cases. Closing statement. Delete exploration, design equations, heuristics, unusable machinery, unused hypothesis commentary, pictures. Keep every hypothesis check, both directions, every degenerate case, every named map, every nonvanishing statement.
11. **Final read.** One pass, reading only the sentences a grader could question. Not the algebra: you have already verified that, and re-reading algebra you wrote fifteen minutes ago confirms your intention rather than your result.

Two remarks. Steps 1 to 5 take eight minutes and feel, the first fifty times, like eight minutes of not doing mathematics. That feeling is the thing to overcome: those eight minutes are the difference between a solution and a session. Speed is a byproduct; writing faster will never cut time, and starting sooner will actively cost it.

And the protocol is not linear in practice. Small cases feed the inversion, the inversion changes the route ranking, the ranking sometimes sends you back to reformulate. What is fixed is the order of the first pass and the position of the checkpoints. Everything else circulates.

3.13 Circuit Bank

The chapter, as reaction circuits. These are the ones to install first, because they fire before any mathematics is chosen and therefore determine everything downstream.

- **WHEN** you finish reading a problem **DO** write the obligation count, the hypothesis ledger and the one word shape before touching the mathematics **BECAUSE** the shape fixes the skeleton of the page and the obligation count is what a missing half of the marks is measured against
- **WHEN** a constant appears in the conclusion that is absent from the hypotheses **DO** factor it and try to construct it from the problem's own objects **BECAUSE** a manufactured constant is the output of the intended computation and therefore names the method
- **WHEN** you compute small cases **DO** use them to select the destination and to find the smallest values of each factor, then delete them **BECAUSE** small cases choose a target and are never evidence; a table plus *and similarly in general* scores at most two
- **WHEN** you have one candidate route **DO** generate two more and price all three on ink, theory footprint, slip risk, argument count and verification burden **BECAUSE** choosing between routes is cheaper than repairing one, and the recorded cost of not choosing is twenty minutes

- **WHEN** you can say the last line of the proof out loud **DO** commit and stop generating **BECAUSE** a route is a path to a named destination; anything less is a direction, and directions are what thrashing is made of
- **WHEN** you are stuck for more than a minute **DO** run the twelve perturbations in order rather than staring **BECAUSE** being stuck is the normal state of the middle twenty minutes, and the difference is having a list instead of a feeling
- **WHEN** you are about to write the script **DO** verify first: boundary case, ledger, seams, second identity, construction against all hypotheses **BECAUSE** once the script exists you will read your intention rather than your result
- **WHEN** you are deciding whether a line stays **DO** ask whether its absence leaves a question a grader can ask **BECAUSE** that single test generates the entire detail calculus and makes solution length an output rather than a decision
- **WHEN** you finish the required proof and see a generalisation **DO** close the proof first, then put the generalisation in a clearly separated remark **BECAUSE** a generalisation inside the proof invites a demand to justify work the problem never required

Chapter 4

Where the marks leak, and the six hours in the arena

4.1 The leak table

Chapter 1 told you the mark scheme is a cliff and not a slope. Chapter 2 gave you the twenty moves. This chapter tells you where, precisely, the fall happens, and it tells you with counts rather than with folklore, because folklore is what everybody else revises from.

The raw material is 775 harvested items covering all 240 problems set from 2006 to 2025. Each carries a *pitfall* field, one sentence naming how a candidate who has the mathematics still loses the marks, and a *priority* from 1 to 5 recording how load bearing the move is. I read all 775 pitfall sentences and tagged them by leak. An item can carry two tags, so the items column does not sum to 775 and is not meant to. The problems column is how many distinct problems, out of 240, document the leak at all. Rows are in frequency order except row 13, where I have merged two clusters the corpus separates only by which symbol you wrote. The last column is the share of that leak's items at priority 4 or 5; the corpus baseline is 80%, so anything above that is a leak which, when it appears, sits on a step the proof cannot survive losing.

	Leak	Items	Problems	Pri 4 or 5
1	The missing edge or degenerate case	236	160	81%
2	The half you did not write	89	81	84%
3	The cited theorem whose hypothesis you never checked	48	44	81%
4	Arithmetic and sign slips in the final computation	45	43	62%
5	Case analysis that is not exhaustive	44	39	82%
6	Induction with a base or a step weaker than it needs	35	32	80%
7	The unproved "clearly" and "similarly"	30	27	77%
8	The construction given without the verification	27	25	78%
9	Unjustified interchange of limits, sums and integrals	25	23	92%
10	Ink nobody is grading	25	21	56%
11	Notation collisions and unbound variables	22	22	77%
12	The pattern submitted as a proof	18	17	61%
13	Existence assumed: the limit, and the maximum	30	29	80%
14	A "without loss of generality" that is not	14	13	71%
15	Circularity	8	7	88%
16	Answering a different question	6	6	67%

The shape of the table is the finding

Read the shape before you read the entries.

Across the whole corpus, 348 of the 775 pitfall sentences, on 203 of the 240 problems, have one grammatical form: *X without Y*. Dividing without showing the divisor is nonzero. Citing without checking the hypothesis. Substituting without stating the range where the substitution is defined. Interchanging without naming the licence. Constructing without verifying. That is 45% of every documented way to lose marks in twenty years, and it is a single sentence pattern.

Understand what that means, because it reorganises your revision. The corpus is not mainly a record of candidates who could not do the mathematics. It is a record of candidates who did the mathematics and did not write the clause. A leak of the form *X without Y* closes by writing Y, Y is almost always one clause, and you already know Y: you performed the check in your head and judged it not worth the ink. It was worth six marks. You are not revising to learn new mathematics here. You are installing a writing reflex.

Two orderings, and you need both

It is tempting to name unjustified interchange as *the* leak, because it is the one everybody warns about and because Chapter 2 gave it a whole move. The count says otherwise and the count is worth having exactly right. Interchange sits low on the frequency list, tenth of sixteen: 25 items on 23 of the 240 problems, about 3% of the corpus. The missing edge case is documented on 160 of the 240. If you drill in frequency order, you drill edge cases first, by a factor of nine.

But frequency is not the only ordering, and here is where interchange earns its reputation. Of its 25 items, 23 sit at priority 4 or 5, a share of 92%, the highest of any cluster in the table with more than ten items, against a corpus baseline of 80%. Interchange is rare and nearly always fatal. It shows up on a specific kind of problem, the analysis flavoured one, and when it shows up it is sitting under the load bearing step, so the deduction is not a shaved mark, it is the difference between 8 and 2. Compare the arithmetic slip at 62%, or ink nobody is grading at 56%: those are frequent and usually survivable.

So run two lists. The frequency list tells you what to sweep for on every single problem you write, twelve times a day, because it will be there. The lethality list tells you which problems demand a slow, deliberate licence check before you leave the page. In practice: sweep for edge cases always; sweep for interchange whenever an analysis symbol has moved.

4.2 The catalogue

Each entry has the same four parts, because you will be applying them under time pressure and a fixed shape is faster to apply. What the leak is. The exact sentence in a write up where it happens. How a grader sees it. The one clause that closes it.

Leak 1. The missing edge or degenerate case

What it is. Your argument is correct on the generic instance and undefined, false, or unproved on one instance at the boundary. The boundary is one of a short list: $n = 0$ or $n = 1$; the empty set, the empty sum, the empty product; equality in an inequality you treated as strict; a denominator, a leading coefficient or a scaling factor that could vanish; the endpoint of an interval; the coincident, collinear or tied configuration; a hypothesis stated for positive integers when the problem said any integer.

The exact sentence. It is almost never a sentence you wrote. It is the sentence you did not write, and it hides in an innocuous connective. Dividing both sides by $a_1 a_2 \cdots a_n$. Writing $a = \cos b$. Applying the formula for $n \geq 3$ when the problem said $n \geq 1$.

Three from the corpus, all of them problems where the mathematics was done and the clause was not.

The 2006 A5 solution divides by $a_1 \cdots a_n$ and by $1 - \omega$. The hypothesis given is that θ/π is irrational, and it exists in the problem for exactly one reason: it is what makes those two quantities nonzero. The model answer opens with a numbered Step 0 discharging all of it at once, since θ/π irrational means no $\theta + k\pi/n$ is congruent to 0 or $\pi/2$ modulo π, so every a_k is finite and nonzero and $\omega \neq \pm 1$. One sentence, placed before anything else, paying for every division in the proof.

The 2018 B4 solution rests entirely on the substitution $a = \cos b$, defined only for $|a| \leq 1$, which the problem does not hand you. Its first paragraph rules out $|a| > 1$ separately and then writes "we may therefore assume $|a| \leq 1$". Write $a = \cos b$ with no comment and your proof covers a proper subset of the hypotheses.

The 2007 B5 problem says "for any integer n". Its model answer chooses the floor characterisation $q \leq n/k < q + 1$ precisely because it makes the sign of n irrelevant, and remarks that no sign assumption was used. A write up that tacitly assumes $n \geq 0$ has proved half the theorem.

How a grader sees it. Not as an oversight. As a case, and a case is either present or absent. This is the single most important sentence in the chapter: *a grader treats a missing case as a missing case, not as something you obviously knew*. An undefended division in the middle of an otherwise complete proof is the reason a solved problem scores 2.

> *The sloppy version.* Dividing through by $a_1 a_2 \cdots a_n$ and simplifying, the ratio equals ni^{n-1}.
>
> *The full marks version.* Since θ/π is irrational, every a_k is finite and nonzero, so the division below is legitimate. Dividing through by $a_1 a_2 \cdots a_n$ and simplifying, the ratio equals ni^{n-1}.

The closing clause. At the point where the tool enters, one clause naming the hypothesis that supplies the nondegeneracy: *since [hypothesis], [quantity] is nonzero, so [operation] is legitimate.* And when the route uses a tool whose domain excludes an instance, the case goes *before* the tool, not at the end. Handled first, everything downstream is unconditional. Handled last, the grader reads the whole proof unsure whether it was ever legal.

Leak 2. The half you did not write

What it is. The problem asked for two theorems and you delivered one. This is the whole family of completeness failures: "determine all" answered with the family but not the exclusion; "find the largest k" answered with a bound but no attaining construction, or with a construction but no bound; an "if and only if" proved in one direction; a claimed minimum with no proof it is a minimum; a set characterised without proving the characterisation is exhaustive.

The exact sentence. The tell is a full stop where a paragraph should be. You write "…hence f' is constant, so $f(x) = \pm x + n$" and put down your pen. You have proved that any solution must have that shape; you have not proved that anything of that shape is a solution, and on 2008 B5 that verification is the cheap half. The model answer states the answer first and labels both halves: "The answer is $f(x) = \pm x + n$ with $n \in \mathbb{Z}$. These work: [three lines]. Conversely, suppose f satisfies the condition; we show it is one of these."

Same shape in the other direction. The 2010 A1 solution proves the bound $k \leq (n+1)/2$ and then says "equality is attained by the construction below", because an upper bound alone is not evidence of achievability and "largest k such that" is a claim about achievability. The 2025 A2 solution, asked for the best constants, splits the work into necessity, evaluating $\rho(x) = \sin x/(x(\pi - x))$ at the contact points including the limit as $x \to 0^+$, and sufficiency, proving the two inequalities with those constants in hand. Proving only that $1/\pi$ and $4/\pi^2$ work answers a different and easier question.

How a grader sees it. As two separately marked deliverables, because that is what they are. One half well written is not 5 out of 10; the bands do not have a 5. It is typically 2. Meanwhile the

missing half is often the four minute one: on 2008 B5 it is a substitution, on 2010 A1 two lines of parity bookkeeping.

The closing clause. State the answer in the first line, then announce both halves before you write either: *We prove two things: first that every such f has this form, and second that every function of this form works.* Announcing them makes them visible as two deliverables, and visible deliverables get marked separately. Then spend the ink on the half that is not obvious and dispatch the other in three lines. Do not skip it because it is obvious. Obvious is not a mark band.

Leak 3. The cited theorem whose hypothesis you never checked

What it is. You invoke a named result by name and the grader cannot tell whether you verified its hypotheses, because you did not write that you had. Worse, you invoke one whose hypotheses fail.

The exact sentence. "By Pappus's theorem, the volume is the area times the distance travelled by the centroid." Pappus fails when the region crosses the axis of revolution, and the 2006 A1 model answer states the hypothesis check inside the same sentence as the citation: since the disc lies entirely in $r > 0$, the volume equals the area times the distance travelled by the centroid. One relative clause, and the shortcut reads as justified rather than as a formula from nowhere.

The 2015 B6 model answer goes further and explicitly declines to cite the Weierstrass M test, because the natural majorant $\sum 1/(2\ell - 1)$ diverges; the uniformity there comes from alternation, not from size. Naming an inapplicable theorem reads worse than naming none, because it proves the check was not made.

How a grader sees it. A named theorem with no hypothesis clause is a claim, and they will check whether the hypothesis holds. If it holds, you may keep the marks; you have merely made them work. If it does not, you lose the step and everything after it. A false citation is strictly worse than an unnamed step, because it converts "did not write it" into "did not know".

The closing clause. Put the check and the citation in one breath: *Since [hypothesis holds here], [theorem] gives [conclusion].* And when you cannot remember whether the theorem needs the hypothesis, do not cite it: prove the three line version instead, as 2012 B5 does with the supporting line lemma rather than quoting the supporting hyperplane theorem. Four lines removes all doubt and costs less than a paragraph of hedging.

Leak 4. Arithmetic and sign slips in the final computation

What it is. The method is right, the structure is right, and the number is wrong. A dropped factor, a flipped sign, an index off by one, a double counted intersection.

The exact sentence. It is the last displayed line. The 2006 A6 pitfall is exact: a factor of two lost in the radial density or in the folding of a double integral leaves an answer of exactly the right shape, and it is invisible to a re-read, because a re-read checks the shape. The 2007 A1 slip is dropping the sign in $-B/(2A)$, which silently swaps which discriminant condition you end up solving. The 2011 A1 slip is double counting the four points where the excluded column meets the excluded rows.

How a grader sees it. Here is the one piece of good news in the catalogue. A slip that does not touch the argument is a 9, not a 2, which is why this leak sits at 62% priority 4 or 5, well below the 80% baseline. It costs one mark, not six. But it costs one mark twelve times if you let it, and one mark twelve times is the difference between 120 and 108. Your vow does not have room for it.

The closing clause. There is no clause. There is a check, and it is mechanical: evaluate your final formula at the smallest legal instance and compare with a direct count. The 2019 B5 model answer does exactly this, verifying $F_{2N+3} - F_{N+2}$ at $N = 1$ ($p(5) = 3 = F_5 - F_3$) and $N = 2$ ($p(7) = 10 = F_7 - F_4$). It costs six characters and it catches the index shift that would otherwise have produced $F_{2019} - F_{1009}$ with total confidence. Likewise 2013 B1: never trust a freshly derived sign identity without one numeric substitution.

Leak 5. Case analysis that is not exhaustive

What it is. You split into cases and the cases do not cover, or they overlap and you counted the overlap twice. A gap destroys the theorem. An overlap destroys the count.

The exact sentence. "If n is odd … if n is even …" where the even branch quietly assumed n is not divisible by 4, which is the 2010 A4 pitfall. Or "squaring both sides, we obtain", which discards a branch: the 2006 A1 model answer names both branches of $|r^2 + z^2 + 8| \leq 6r$ and shows the left one holds automatically because $r^2 + z^2 + 8 \geq 8 > 0 \geq -6r$, so only the right one is a genuine constraint. One extra clause, and the discarded branch is visibly checked rather than assumed away.

How a grader sees it. They read your case list and ask whether it covers. If it does not, the theorem is unproved, however good the covered cases are. And they read it fast, so a case list that is hard to check is a case list that will be checked hostilely.

The closing clause. One sentence before the split: *Exactly one of the following holds.* Then, where it is not immediate, one clause of justification for the word "exactly". Better still, remove the split: the 2010 A4 model answer writes $n = 2^m k$ with $m \geq 0$ and k odd, a reparametrisation that covers every n with no cases at all. Reparametrising is lazier than exhausting, which is why it is right. Ink costs time and invites error.

Leak 6. Induction with a base or a step weaker than it needs

What it is. Two distinct failures that a grader marks the same way. Either your base case does not anchor the induction, or your inductive step quietly consumes a hypothesis stronger than the one you proved.

The exact sentence. On the base case: "The case $n = 1$ is immediate." The 2018 B4 model answer checks $n = 0, 1, 2$, and it has to, because the recursion is second order; checking one base case, or applying the identity $F_m - F_{m-1} = F_{m-2}$ at $m = 1$ where it silently needs F_{-1}, are both wrong and both are caught by writing the index range down. The 2009 B1 pitfall is the mirror image: citing a base case without checking that no instance exists below it, which leaves the induction unanchored at the bottom.

On the step: the 2020 B6 model answer strengthens the target from $F(n) \geq 0$ to $F(n) \geq 1$ before inducting, because the raw statement dies in the negative run case, where the step yields only ≥ -1 and the induction cannot be repaired without restarting. It then names the exact inequality licensing the hypothesis, $\lfloor (m+1)\alpha \rfloor < m < C_{m-1} \leq n$, because the renormalised index has to be strictly smaller than n or the hypothesis does not apply.

How a grader sees it. An induction is a machine with two moving parts and they check both. An unanchored base is a hole at the bottom, a step using more than the hypothesis gives is a hole in the middle, and neither is repairable by the quality of the rest.

The closing clause. State the induction with its range in the same sentence, and state the strengthened form before you start rather than after it fails: *We claim [strengthened statement] for all $n \geq n_0$, and prove it by induction on n. The cases $n = n_0, \ldots, n_1$ are [checked]. Assume the claim for all indices up to m, where $m \geq n_1$.* Then one clause inside the step certifying that the index you apply the hypothesis to lies in range. Writing the range down is what catches the F_{-1}.

Leak 7. The unproved "clearly" and "similarly"

What it is. Two different sins wearing the same coat. "Clearly" claims a fact without proof. "Similarly" claims a whole argument without writing it.

The exact sentence. On "clearly": the 2008 B1 pitfall is writing that the solution is clearly rational instead of exhibiting the closed form. Cramer's rule gives $g = (C_1B_2 - C_2B_1)/(A_1B_2 - A_2B_1)$, manifestly a ratio of rationals with nonzero rational denominator, and that display is shorter than the sentence defending "clearly".

On "similarly": the 2010 A1 pitfall is handling one parity by hand and asserting the other works similarly, which is exactly where the odd case singleton gets forgotten. And 2008 A4, the purest example in the corpus, writes "and so forth" after two blocks instead of stating and proving the general block formula. The whole problem lives in that formula, and a grader reading "and so forth" at the point where the problem lives does not award the marks.

How a grader sees it. "Clearly" reads as "I did not check this". "Similarly" reads as a claim of symmetry, and they test it by running your first case under the relabelling; if any step used the specific labels, the symmetry is false and your second case does not exist.

The closing clause. For "clearly", delete the word and write the reason; if the reason exceeds a clause, it was a lemma. For "similarly", earn it: *The map exchanging x and y preserves the hypothesis and the conclusion, so the second case follows from the first.* Name the map. The 2008 A1 pitfall shows the cost even on trivial algebra: an unlabelled "similarly" jumping between substitutions, without stating which values were used, is marked as an unjustified leap even when the algebra is nothing.

Leak 8. The construction given without the verification

What it is. You build the object the problem asked for and never check it against the conditions it was supposed to satisfy. Usually you check the one condition it was designed for and forget the other two.

The exact sentence. "Define $g(x) := f(x, 0)$." Full stop. The 2008 A1 model answer does not stop there: it substitutes the definition back into $f(x, y) = g(x) - g(y)$ and confirms the match, and full marks require that closing line, not just the definition. The 2013 B2 pitfall is sharper: the constructed f satisfies the condition it was built for, hypothesis (ii) is never checked, and (ii) happens to hold only because $N = 2$ was chosen small enough to avoid every multiple of 3.

A second species: the construction described by example instead of by rule. The 2008 A2 pitfall is describing a game strategy through an example round, "if Alan plays here, Barbara plays there", instead of a rule covering every cell Alan could choose. A strategy is a function on all reachable positions, and an example round is marked as an incomplete strategy.

How a grader sees it. A construction is an existence claim, proved by exhibition plus verification; exhibition alone is a guess. They check your object against every stated hypothesis, including the ones you did not have in mind when you built it.

The closing clause. After the construction, one short paragraph that walks the hypotheses in the order the problem listed them: *This f satisfies (i) since ..., (ii) since ..., and (iii) since ...* List them all, including the trivial ones. The trivial ones are where the design work quietly broke something.

Leak 9. Unjustified interchange of limits, sums and integrals

What it is. Two limiting operations change places and you did not say why they may. Infinite sums, integrals, limits and derivatives all count, and so does reordering a double sum. This is the rarest of the big leaks and the deadliest per appearance, 25 items with 23 at priority 4 or 5, because it appears on the problems where it is the whole difficulty.

The exact sentence. "Interchanging the order of summation and integration, we obtain ..." The 2015 B6 pitfall records exactly what happens next: that phrase, written with no comment, converts a full solution into a fragment. Chapter 2 gave you the licence ladder, so use the cheapest rung and say which rung you are on.

Four sentences from the corpus, in increasing cost.

All sums are finite, so the interchange is trivial (2013 A6). Eleven words, and it pre-empts the only convergence question a grader could raise.

Every term is positive, so the sum is a well defined element of the extended reals and may be reordered and regrouped freely; the computation below yields a finite value, which justifies the manipulations retrospectively (2015 B4). Note the structure: nonnegativity buys the reordering now, finiteness arrives later as a

consequence, and the sentence says so rather than pretending finiteness was known first. Cheaper still, in 2014 A4: *since* $\mathrm{E}[X^3]$ *is finite, all four series converge absolutely*, one sentence that buys every rearrangement in the proof.

And the expensive rung, when nothing cheaper applies: 2015 B6 opens with "the series is not absolutely convergent, so we must show that the partial sums S_N ordered by k converge to the rearranged value S" and then spends four lines bounding $S - S_N$ explicitly. Four lines, not four pages, because the estimate was arranged to be easy.

How a grader sees it. On an analysis flavoured problem they hunt for this, because it is the step the setter built the problem around, so a silent swap reads as not having seen the problem. A wrong citation is worse: dominated convergence with no dominating function exhibited, or the M test where the majorant diverges, is marked down even when the conclusion is true.

> *The sloppy version.* Interchanging the order of summation, $\sum_{n\ge1}\sum_{k=1}^{n} a_{n,k} = \sum_{k\ge1}\sum_{n\ge k} a_{n,k}$.
>
> *The full marks version.* Every $a_{n,k}$ is nonnegative, so by Tonelli's theorem the two iterated sums are equal in $[0,\infty]$: $\sum_{n\ge1}\sum_{k=1}^{n} a_{n,k} = \sum_{k\ge1}\sum_{n\ge k} a_{n,k}$.

The closing clause. Name the cheapest applicable licence at the point of use, in one clause: finite beats nonnegative beats absolutely convergent beats uniform. And the laziest move, which is the best one, is never to leave finite sums. Manipulate the Nth partial sum throughout, write *every step so far is an identity between finite sums; nothing has been rearranged*, and take the limit once at the end. That is what 2017 B4 does with a conditionally convergent series, and it removes the question rather than answering it.

Leak 10. Ink nobody is grading

What it is. Correct mathematics that earns nothing. It is in the table because the corpus documents it 25 times, and it is the leak that steals the minutes you needed for Leaks 1 to 9.

The exact sentence. There is not one; there is a paragraph, and it should not be there. The 2009 A2 pitfall is solving the full system for f, g and h when the problem asked for one of them. The 2015 B2 pitfall is chasing a closed form for every residue class when the problem only asked to prove or disprove existence. Two variants worth naming: the wrong route pursued past its expiry (2010 A6, twenty minutes sunk into repairing a uniformity failure when a route without that failure was available), and the false statement (2009 A6, a whole session spent proving something false, because each coordinate condition is separately achievable and the claim therefore feels nearly true).

How a grader sees it. They do not see it at all, which is the point: they mark the requested statement.

The closing clause. Before writing a paragraph, ask which of the problem's clauses it discharges. If none, it does not go on the page. If you proved more than was asked, put it in a separated closing remark so the proof visibly ended before the commentary began.

Leak 11. Notation collisions and unbound variables

What it is. A symbol means two things, or means nothing. Includes the off by one, which is a special case: an index whose range you never wrote down.

The exact sentence. The 2024 A4 pitfall is reusing the symbol r after a scaling, so that the r in the conclusion is not the r in the hypothesis. The 2011 A5 pitfall is n boxes for n points rather than $n+1$. The 2017 A5 pitfall is using arg without checking $z_n \neq 0$ and without choosing a continuous branch, so the symbol is not even well defined.

And the purest unbound variable in the corpus, 2009 A1: proving $f(P) = 0$ and never writing that P was arbitrary. Without that sentence the proof has shown the result for one unspecified point, and a strict grader can dock it.

How a grader sees it. They track your symbols, and one that changes meaning halfway is a broken chain they cannot repair, because they do not know which meaning you intended at which line. An unquantified variable at the end is a theorem about nothing.

The closing clause. Three habits, no clause required. Never reuse a letter after a transformation; write r'. Write every index range on the line where the index first appears. And restore the quantifier at the close: *Since P was an arbitrary point of the plane, f vanishes identically.*

Leak 12. The pattern submitted as a proof

What it is. You computed small cases, saw the shape, and wrote the shape down as though computing it had established it.

The exact sentence. "and similarly in general". The 2019 A5 pitfall says what that scores: at most one or two out of ten. The 2006 A2 pitfall is asserting a set is infinite from an observed pattern in small cases, when the underlying move set is not periodic. The 2017 B4 pitfall is verifying a decomposition on the first eight terms, when the residue case check is what turns a pattern into a proof.

How a grader sees it. As reconnaissance, not evidence, and they are right. Small cases tell you what lemma to look for; they do not prove it.

The closing clause. The pattern never appears in the write up as a justification, only as a check at the end. What appears instead is "We claim [closed form] for all $n \geq n_0$", followed by a proof. Reconnaissance stays offstage.

Leak 13. Existence assumed: the limit, and the maximum

What it is. Two faces of one leak, 30 items across 29 problems, separated in the corpus only by which symbol you wrote. Either you compute the value of a limit, a sum or an integral before establishing that it exists, or you treat a supremum as a maximum without compactness or monotonicity.

The exact sentence. For the limit: "Letting $n \to \infty$ in the recursion, $L = f(L)$, so $L = \dots$" The 2016 B1 pitfall is exactly this, asserting a limit exists because a recursion looks like it settles down numerically, with no monotone and bounded justification. The 2012 B4 pitfall names the circularity directly: computing a candidate limit and then arguing in a circle from it. The 2014 A3 pitfall is presenting a telescoping limit as the answer without stating why the infinite product converges.

For the maximum: "Therefore the minimum is $\Delta(1)$." The 2019 B4 pitfall is computing $\Delta(1)$ and asserting minimality with no argument, on an interval that is not compact; its model answer earns the word in one clause, since $\Delta'(s) = \ln(s+1) - \ln s > 0$, so Δ is strictly increasing on $[1, \infty)$ and the minimum is attained at $s = 1$. The sharper version is 2012 B5, introducing a maximising index $s(x)$ and differentiating through it when the supremum need not be attained at any s. That does not lose a mark, it voids the argument.

How a grader sees it. Solving for a limit not yet shown to exist proves nothing at all; it is not a small gap, it is the absence of the theorem. And the word "minimum" in the problem is a specification, so attainment is graded. Evaluating at a candidate point proves nothing about minimality.

The closing clause. Separate existence from value and buy existence with the crudest tool available. *The sequence is decreasing and bounded below by* 0*, hence convergent; let L be its limit*, then and only then substitute into the recursion; 2012 B4 picks the bound 0 deliberately, because a decreasing sequence needs any lower bound and not a sharp one. For attainment: compactness plus continuity, or monotonicity, or a construction that attains by design, as 2012 B5 does by building the witness family so that equality holds at $t = x$. And check the family is nonempty, which is why 2019 B4 closes by exhibiting $f_0(x, y) = \frac{1}{2}(xy \ln(xy) - xy)$.

Leak 14. A "without loss of generality" that is not

What it is. You normalise, and the normalisation either does not preserve the hypotheses, does not preserve the conclusion, or is contradicted by a later step.

The exact sentence. "Without loss of generality, $a_0 > 0$." The 2022 A2 model answer supplies the reason in the same sentence: replacing p by $-p$ changes neither $p(x)^2$ nor the degree, so we may assume $a_0 > 0$. Written without the reason, a grader may simply not accept it. The 2008 A5 pitfall is "without loss of generality the polygon is the nth roots of unity" with no map exhibited and no degree invariance stated. The 2015 A1 model answer names the map and everything it preserves: $T(x, y) = (\lambda x, y/\lambda)$ is area preserving, $\det T = 1$, it fixes the curve, and it sends arcs to arcs and chords to chords.

The nastiest species is fine when written and false three paragraphs later: the 2025 B1 pitfall is invoking a WLOG when a later step uses the specific colour assignment the WLOG was supposed to make irrelevant.

How a grader sees it. An unnamed WLOG is an extra assumption, and they mark it as one. A named one they check against the hypotheses and the conclusion, in that order.

The closing clause. Four components, always, in one sentence: the map, what it preserves, what it achieves, where it is used. *The map T preserves [hypotheses] and [the quantity in the conclusion], and carries any configuration to one with [property], so we may assume [property].* Then, at the later step, confirm you are still inside the normalisation. Omitting the WLOG loses a mark for an ungeneralised proof; stating it without naming the map loses the same mark for a different reason.

Leak 15. Circularity

What it is. A step uses the thing being proved. Rare, at 8 items, and 88% of them load bearing, because when it happens it is not a gap, it is the absence of a proof.

The exact sentence. It is never visible in one line, which is what makes it dangerous. The 2010 A5 pitfall is using $x * x = 0$ before the lemma that establishes it, on the grounds that $x \times x = 0$; hypothesis (ii) is vacuous exactly there. The 2012 B4 pitfall is the ansatz $a_n \approx \log n + c + o(1)$, manipulated before convergence is known.

How a grader sees it. They read your lemma order. If Lemma 3 is used inside the proof of Lemma 2, the chain has no ground floor.

The closing clause. No clause, a structural discipline: number your lemmas and read the proof backwards, asking of each step which earlier number it consumed. Anything consuming a later number is circular. This also catches the 2025 B6 species, a threshold constant depending on the very n you are quantifying over.

Leak 16. Answering a different question

What it is. You proved something true, and it is not what was asked.

The exact sentence. "Hence the sequence is eventually periodic, and therefore periodic." The 2018 B4 pitfall names it flatly: those are not the same statement, and the problem asked for the stronger one. The model answer earns periodicity properly by inverting the step map: $T(u, v) = (v, u + v)$ is a bijection of $(\mathbb{Z}/d\mathbb{Z})^2$ with inverse $(u, v) \mapsto (v - u, u)$, so a repeat propagates backwards to index 0.

Two more species. Proving a supremum bound when an identity was asked (2012 B5, producing $xh_1(t) + h_2(t) \le F(x)$ and never exhibiting a t achieving equality). And the cheapest of all: boxing $34/42$, or boxing $17/21$ without a word, on a paper that explicitly said lowest terms (2015 B4). That model answer's last line is "Finally 17 is prime and does not divide 21, so $\gcd(17, 21) = 1$ and $17/21$ is in lowest terms."

How a grader sees it. They hold your final sentence next to the question and match the words. If the words do not match, the marks do not come, however correct your page.

The closing clause. Restate the conclusion in the problem's own words as your final sentence, with every qualifier it used: lowest terms, smallest, all such, for every integer n, periodic and not merely eventually periodic. Copy the qualifiers off the question paper. Do not paraphrase them.

4.3 The five second sweep

Sixteen leaks is too many to hold at the end of a problem. Compress them into five questions and ask them of every write up before you leave the page. Ninety seconds, and they cover every entry in the table.

1. Did I divide, root, take a logarithm, invert, substitute, or cite? (Leaks 1, 3, 14.)
2. Did the problem ask for two things? (Leaks 2, 8, 13, 16.)
3. Did two limiting symbols change places (Leak 9), or did I name a limit or a maximum? (Leak 13.)
4. Is every case covered, every index in range, every symbol still meaning what it meant? (Leaks 5, 6, 11, 15.)
5. Does my final line use the problem's own words, and does my final number survive one small case? (Leaks 4, 7, 12, 16.)

4.4 The arena: two three hour sessions

Everything above is worth nothing unless it fires under time pressure. So convert it now into an operating protocol, and rehearse the protocol until it is as automatic as the circuits themselves. What follows assumes the standing format: Saturday 5 December 2026, a morning session of three hours holding A1 to A6, a break of about two hours, and an afternoon session of three hours holding B1 to B6. Twelve problems, ten marks each, 120 available. Confirm the exact clock times your site publishes and rehearse against those, not against these.

Three facts about that format decide everything.

Fact one: A1 pays what A6 pays. There is no difficulty bonus. A6 is not worth twenty. The numbering is the setters' estimate of difficulty, and it is a good estimate, but it is an estimate of the field and not a measurement of you, which is what the first ten minutes are for.

Fact two: the bands are $0, 1, 2, 8, 9, 10$**.** Your unit of production is a *closed* problem, and half a proof of A4 is worth about a fifth of a closed A1.

Fact three: 180 minutes does not fit six problems. It fits between two and four. Which two, three or four is the single highest leverage decision you make all day, and you make it in the first ten minutes on bad information.

Minutes 0 to 10: the read through and triage

You do not start solving. You read all six problems and you write nothing except triage notes. Ten minutes, and use all ten.

For each problem, in order, spend roughly ninety seconds and produce four things on your scratch sheet:

1. **The command term.** Determine all, evaluate, prove, show that there exists, find the largest. This fixes your obligations before you know anything else, and it is what tells you whether the problem is one theorem or two.
2. **The field.** Number theory, real analysis, linear algebra, combinatorics, polynomials, probability, geometry. Book 1 is organised this way for exactly this moment.

3. **The cues, and which circuit they fire.** Write the circuit name if one fires. "Positive terms plus a double sum: Tonelli." "Determine all: two halves." "An unstructured family with no distinguished member: extremal." If nothing fires, write "no circuit", and that is information, not failure.

4. **A time estimate and a probability.** How many minutes to close, and how likely you are to close it. Be brutal. Your first estimate is the one to trust; the optimism arrives later.

Ranking by expected marks per hour

Now rank. Not by difficulty, not by taste, and not by which one you find interesting. By rate.

For each problem let p be your honest probability of closing it and t the hours you think it will take. The rate is $10p/t$ marks per hour. Some representative numbers from a real triage:

Problem	Read	p	t (min)	Rate
B1	circuit fired, route visible	0.85	30	17.0
B2	circuit fired, one gap to fill	0.60	45	8.0
B4	field known, no circuit yet	0.40	60	4.0
B3	field known, route unclear	0.20	75	1.6
B5	no circuit, no field	0.10	120	0.5
B6	no circuit, no field	0.08	150	0.3

Read the rate column and then read it again. B1 pays fifty times what B6 pays per minute of your life, not because B1 is easier in some abstract sense but because both carry ten marks and B1 costs a fifth of the time. Fact one is doing all the work.

The consequence, as an order you follow without negotiation: **you write the two highest rate problems to completion before you start anything else.** Not in parallel. Not "I will just see whether B4 opens up first." Closed and verified, one after the other. Two closed problems is 20; four opened problems is 8. Same mathematics, two and a half times the marks, and the difference is a scheduling decision.

The second consequence is the one candidates resist. If B5 and B6 are both at rate 0.5 or below, they are not going to be attempted as full solutions in this session. That is not defeat, it is arithmetic; what you do with them instead is the partial credit section below, worth roughly 2 to 4 marks for 25 minutes, a rate that beats attempting either of them properly.

A caveat worth clarifying: the triage is made on bad information, so it is wrong, not in the top two but in the middle. That is why the protocol re-triages twice, and why the estimate is written down. An estimate you did not write down cannot be revised, only forgotten in favour of whatever you happen to be doing.

The clock: how the 180 minutes are spent

1. **0 to 10. Read and triage.** No solving. Write the four notes and the rate ranking.

2. **10 to 60. The top two.** Close them, writing the final version directly onto the answer paper rather than a draft to be copied later.

3. **60 to 65. Re-triage.** Re-read the remaining four with fifty minutes of warmed up mathematics behind you. Rates change: a problem you rated 0.2 often jumps because something you wrote in B1 fired a circuit. Rewrite the ranking.

4. **65 to 140. The next one or two.** In rate order, with the abandonment rules below running the whole time, and a second thirty second re-triage at minute 100 to confirm you are still on the highest rate problem left.

5. **140 to 180. The last forty.** No new problem is opened after minute 140. None. This block is described separately below because it is where the marks that were nearly earned become marks that were earned.

When to abandon, and when to come back

Three abandonment tests. Any one of them fires, you park.

The ten minute no new object test. Set a mental checkpoint every ten minutes and ask whether you produced a new object: an equation, a reformulation, a named lemma, a reduction, a counterexample, a case eliminated. Rewriting the same expression in different notation is not a new object, and neither is the same substitution with different constants. Two consecutive checkpoints with no new object and you park, immediately, mid sentence if necessary.

The missing fact test. If your route requires a theorem you cannot state precisely, park. You will not derive it in the arena and you cannot cite what you cannot state. Write the reduction, name the fact, move. That note is worth 1 or 2 marks; chasing the fact is worth nothing.

The rate test. If you have spent forty minutes on a problem you estimated at forty five and you are not two thirds through, both p and t were wrong, so recompute them and compare against the best unstarted problem.

Parking, not quitting. When you park, you spend sixty seconds writing the park note on the scratch sheet for that problem:

> *The park note.* Parked at minute 78. Have: reduced to showing g has at most two zeros on $[0, 1]$. Need: a bound on g'', or a convexity argument. Tried and dropped: direct differentiation (gets a cubic, no sign information); IVT on g alone (not enough). Next thing to try: Rolle twice, using the second derivative hypothesis I have not yet used.

That note costs a minute and it does two things. It makes the return cheap, because you re-enter where you left rather than re-deriving for six minutes. And it is nine tenths of a bankable partial write up, which is what it becomes at minute 145 if you never return.

When to come back. Two triggers only. At a scheduled re-triage, or when something you wrote on a different problem fires the circuit the parked one needed, which is why park notes stay face up. Never return because a problem is nagging at you. Nagging is not information.

One check before you park. List which hypotheses your attempt consumed. If a hypothesis is untouched, that is where the route is, and it is worth five more minutes rather than a park. The 2011 A5 pitfall is an argument that never uses the decay bound on g; any such argument is wrong, because without that bound the conclusion is false. An unspent hypothesis is a free signpost.

Banking partial credit honestly

This section is worth more marks than any other in the book, because it is the one that applies to the four problems a session you will not close.

First, what actually pays. A 2 is significant progress: the main idea, a substantial reduction, the crux lemma proved. A 1 is some progress: a real reduction, a correct nontrivial lemma, one case of several handled completely. A 0 is everything else, and everything else includes restating the problem in your own notation, computing small cases, writing "the answer appears to be 4^k", and describing what you would do with more time. Reconnaissance is not evidence, and that rule does not relax because the clock is running out.

Second, the structure of an honest partial. Four parts, in this order:

1. **The target.** One sentence naming what you set out to prove, in the problem's words.
2. **The claim you have proved.** Stated precisely, as a lemma, with its own label.

3. **The proof of that claim.** Complete. Full rigour. Every leak in Part One swept. This is the part being marked, so it is marked to the same standard as a full solution.
4. **The declaration.** One or two sentences naming exactly what remains and what you were unable to do.

Now the part that feels wrong and is not. *Writing "I have not proved Lemma 2" costs nothing and gains everything.*

> *The hidden gap version.* ... so it suffices to show that g has at most two zeros on $[0, 1]$. Since g'' is continuous and g vanishes at the endpoints, g has at most two zeros. Therefore F is injective, as required.

> *The declared gap version.* ... so it suffices to show that g has at most two zeros on $[0, 1]$ (Lemma 2). Given Lemma 2, the conclusion follows: [three lines, complete].
>
> I have not proved Lemma 2. I can show that g has at most two zeros whenever g'' has constant sign, by Rolle applied twice; the remaining case, where g'' changes sign on $[0, 1]$, I was unable to close.

Take the two apart, because the accounting is exact.

The hidden gap version asserts a false implication in its second sentence: continuity of g'' and vanishing at the endpoints do not bound the zero count. A grader finds it does not follow and now has to decide what you knew. They cannot award the reduction cleanly, because you presented it as complete rather than as a reduction, so there is no labelled object to give marks to. And they have been told, by you, that you believe an untrue thing. Realistic mark: 1, possibly 0.

The declared gap version claims exactly what it has. The reduction is labelled, so it is markable. The conditional argument from Lemma 2 is complete, so it is markable. The partial progress on Lemma 2 itself, one case of two, is markable. And the grader spends no effort deciding what you meant. Realistic mark: 2.

One mark a time when the alternative was a 1, two when it was a 0, and it happens three or four times a session. That is the difference between a mid table score and a serious one. There is no risk on the other side, and this is the point candidates get wrong: **you cannot lose marks for honesty, because a grader who spots the gap deducts for it either way.** They read for gaps. That is the job. The only thing hiding a gap changes is whether they also conclude you did not see it, and that conclusion moves you from 2 to 1.

Two more rules for banking.

Do not write a solution you know is wrong. If you have found a counterexample to your own lemma, say so and stop. A page of consequences drawn from a false lemma scores 0 and consumes twelve minutes.

Never write "it is easy to see" in a partial. In a full solution it is a leak. In a partial it is a lie about the very thing you are being marked on, and it is the fastest way to convert a 2 into a 0.

The last forty minutes

Minutes 140 to 180. Three blocks, in this order, and the order is not negotiable because each block pays more per minute than the next.

140 to 155. Convert the parks. Take every parked problem and turn its park note into an honest partial by the four part structure above. Fifteen minutes, three or four problems, four minutes each. You are not solving anything here; you are transcribing what you already have into a markable form. This block routinely produces 3 to 6 marks and it is the highest paid quarter hour of the session.

155 to 172. The verification pass. On the problems you closed. Detailed below.

172 to 180. The physical sweep. Headers, page order, cross out scratch, confirm every sheet you intend to hand in is a final version and every sheet you do not is not in the pile.

What you do not do between 140 and 180: open a new problem, chase one more idea on a parked problem, or improve the exposition of a solution that is already complete. All three feel productive and all three pay less than the block they displace.

The verification pass

Seventeen minutes, roughly five per closed problem. Do not re-read for pleasure; read hostilely, as though the page were somebody else's and you were being paid to find the hole.

1. **The first line.** Does the solution open with the complete answer, in the problem's own words? If the problem said "determine all", is the answer a set and not an example?
2. **The operations sweep.** Find every division, root, logarithm, inverse, substitution and citation on the page. Physically point at each one. Each wants a clause. This is Leak 1 and Leak 3, and it is the highest yield ninety seconds you will spend.
3. **The obligations sweep.** Count the deliverables in the question and count the announced blocks in your solution. If the question has two and your page has one, you know what to do with the next four minutes.
4. **The symbol sweep.** Scan for any letter used twice, any index without a stated range, any variable that appears in the conclusion but was introduced inside a proof.
5. **The small case.** Evaluate your final formula at the smallest legal instance and compare against a direct count. Six characters.
6. **The unused hypothesis.** Find where each was consumed. An unspent hypothesis is almost always a hole; occasionally it means you proved something stronger, in which case say so in one closing remark.
7. **The last line.** Does it restate the conclusion in the problem's words, with every qualifier the problem used? Lowest terms. Smallest. For every integer n. Periodic, not eventually periodic.

The physical discipline of the write up

Two piles of paper, and they never mix. The scratch pile is where you think. The answer pile is where you write. Nothing moves from the first pile to the second.

Write the final version. Do not hand in scratch. Absolute. Scratch is a record of your search; a solution is a certificate of verification; they are different documents and the second is shorter. Handing in scratch means handing in the small cases, the abandoned substitution and the exploratory algebra containing three sign errors you never fixed because you did not need to. A grader reading that has to find the proof inside it, and the proof they find is the weakest reading available.

The rest of the discipline, every item there because the corpus or the mark scheme puts it there.

1. One problem per sheet. Start every problem on a new sheet, however short the solution, because sheets get separated and marked independently.
2. Header every sheet: problem number first, then whatever identifier your site requires, then "page k of m". Write the problem number before you write anything else, always.
3. Write the answer statement, then leave a blank line, then the proof. The blank line is not decoration; it is what makes the answer findable in two seconds.
4. Number your lemmas and your steps, with a claim in each header. A grader who wants to check Step 3 can then find it, which is worth more than any amount of polish.
5. Do not copy from scratch. Copying is where the sign slips enter, because copying is not reading. Write from the structure in your head, consulting scratch only for computed constants, and check those against the small case test.

6. Cross out abandoned work with a single diagonal line and write "superseded". Do not scribble it into illegibility and do not leave it uncrossed. Uncrossed wrong work on an answer sheet is read as part of your solution.
7. Legible ink, normal size. Do not micro write to fit a page; ask for more paper, and if you run out of room mark the break with "continued on page 3" and "continued from page 2". A proof whose middle cannot be found is a proof with a gap.

The break, and the second session

Two hours between sessions, three rules. Do not discuss the A problems with anyone: the information arrives too late to use, and the one thing it reliably does is tell you A3 was easy. Eat and drink, which is arithmetic rather than wellness advice, since the B session is 60 marks scored by a brain that has already worked three hours. And reset the triage completely: B1 is not harder because A6 was hard, and the commonest way to lose the afternoon is to carry the morning's mood into the first ten minutes and mis-triage because of it.

4.5 Pre-flight checklist: the morning of Saturday 5 December 2026

Read this once on the morning and read nothing else. It is not revision; it is the loading of a protocol you have already installed.

1. Ten minutes of reading before any solving. All six problems. Four notes each: command term, field, circuit, rate. Use the whole ten minutes.
2. Rank by $10p/t$, not by taste, not by difficulty, not by which one is interesting.
3. Close the top two before starting anything else. Closed and verified. Not in parallel.
4. A1 pays what A6 pays. Whenever the arena tempts you towards a hard problem, say that sentence to yourself and go back to the ranking.
5. State the answer in the first line. Every time.
6. Two deliverables means two announced blocks. Count the question's clauses before you start writing, not after.
7. Every division, root, logarithm, inverse, substitution and citation gets its clause, at the point of use.
8. Dispose of the degenerate case before you introduce the tool whose domain excludes it, not at the end.
9. Two limiting symbols changing places is one clause, and the cheapest licence wins. Finite beats nonnegative beats absolutely convergent beats uniform. If it is conditionally convergent, never leave finite sums.
10. Ten minutes with no new object, twice in a row: park. Write the park note. Move.
11. Re-triage at minute 60 and again at minute 100.
12. Nothing new after minute 140. Convert the parks at 140, verify at 155, sweep the paper at 172.
13. Declare every gap you know about, in one sentence, naming exactly what remains. It cannot cost you a mark and it routinely earns two.
14. Write the final version. The scratch pile does not get handed in.
15. Last line in the problem's own words, with every qualifier copied off the paper.
16. The afternoon is a fresh 60 marks. Do not discuss the morning.

4.6 Circuit Bank: the leak closing clauses

These are the sentences, and they are to be produced, not recognised. Copy each one out by hand once. Then close the book, write the same clause into a solution of your own from memory, and see what you left out. What you left out is the mark.

- **WHEN** you divide, take a root, take a logarithm, invert, or substitute **DO** write *since [hypothesis], [quantity] is nonzero, so the [operation] below is legitimate*, at the point of use **BECAUSE** the missing edge case is documented on 160 of the 240 problems in the archive, and a grader treats a missing case as a missing case
- **WHEN** your route uses a tool whose domain excludes an instance **DO** clear that instance in a short paragraph before the tool appears **BECAUSE** handled first everything downstream is unconditional; handled last the grader reads the whole proof unsure it was ever legal
- **WHEN** the command term is *determine all*, *find the largest*, or *if and only if* **DO** announce both halves by name before writing either, then spend the ink on the hard half and three lines on the cheap one **BECAUSE** one half well written is not five marks, it is two, and the missing half is usually the four minute one
- **WHEN** you name a theorem **DO** put the hypothesis check in the same sentence, *since [hypothesis holds here], [theorem] gives [conclusion]*, and if you cannot recall the hypotheses, prove the three line version instead of citing it **BECAUSE** a false citation is worse than silence because it proves the check was not made
- **WHEN** your answer is a formula in n **DO** evaluate it at the smallest legal n and compare with a direct count **BECAUSE** six characters, and it catches the index shift an unaided re-read cannot see, because a re-read checks the shape
- **WHEN** you split into cases **DO** write *exactly one of the following holds*, and prefer a reparametrisation that removes the split entirely **BECAUSE** a gap destroys the theorem and an overlap destroys the count, while $n = 2^m k$ with k odd has no cases at all
- **WHEN** you begin an induction **DO** state the range in the same sentence as the claim, strengthen it into the shape the later step consumes, and certify inside the step that the index used lies in range **BECAUSE** writing the range down is what catches the base case that silently needs F_{-1}
- **WHEN** you are about to write *clearly* or *obviously* **DO** delete the word and write the reason; if the reason exceeds a clause it was a lemma **BECAUSE** *clearly* is read as *I did not check this*
- **WHEN** you are about to write *similarly* or *by symmetry* **DO** name the map and confirm no step used the specific labels **BECAUSE** if a step did, the cases are not symmetric and your second case does not exist
- **WHEN** you have constructed an object, a strategy or an algorithm **DO** walk the problem's hypotheses in the order it listed them, trivial ones included, and write a strategy as a rule covering every reachable position **BECAUSE** the design work satisfies the condition you had in mind and quietly breaks another, and an example round is not a rule
- **WHEN** two limiting symbols change places **DO** name the cheapest applicable licence in one clause: finite, then nonnegative, then absolutely convergent, then uniform with the tail exhibited **BECAUSE** 25 items in the corpus and 23 at priority 4 or 5; it is rare and it is nearly always fatal
- **WHEN** the series is conditionally convergent **DO** manipulate the Nth partial sum throughout and take the limit once at the end **BECAUSE** staying finite removes the question instead of answering it

- **WHEN** you are about to name a limit L **DO** prove existence first with the crudest available tool, then substitute **BECAUSE** solving for a limit not yet shown to exist proves nothing, and reasoning back from a candidate is circular
- **WHEN** the problem writes *minimum* or *maximum* **DO** earn the word: compactness plus continuity, or monotonicity, or a construction that attains by design **BECAUSE** evaluating at a candidate point proves nothing about minimality, and a supremum need not be attained anywhere
- **WHEN** you write *without loss of generality* **DO** name the map, what it preserves, what it achieves and the step that uses it, then confirm at that step you are still inside the normalisation **BECAUSE** an unnamed WLOG is an extra assumption and marked as one; the nastiest species is contradicted three paragraphs later
- **WHEN** you transformed or rescaled **DO** use a new letter, never the old one **BECAUSE** the r in your conclusion must be the r in your hypothesis, and a symbol that changes meaning is a chain a grader cannot repair
- **WHEN** you finish a universally quantified proof **DO** restore the quantifier in the last sentence: *since P was arbitrary, ...* **BECAUSE** without it you proved the result for one unspecified point
- **WHEN** you found the answer by computing small cases **DO** delete the small cases from the write up and state *we claim [closed form] for all $n \geq n_0$*, then prove it **BECAUSE** *and similarly in general* scores at most one or two out of ten, and reconnaissance is not evidence
- **WHEN** ten minutes have produced no new object, twice **DO** park, write the park note, move to the next problem in the rate ranking; if the route needs a theorem you cannot state, record the reduction, name the fact and stop **BECAUSE** a new object is an equation, a reformulation, a lemma, a reduction or a case eliminated, and chasing a fact you cannot state is worth nothing
- **WHEN** you are about to finish and a hypothesis is unused **DO** find where it was consumed, or find your error **BECAUSE** an argument that never uses a given hypothesis is usually wrong, because without that hypothesis the conclusion often fails
- **WHEN** you know there is a gap and the clock is at minute 145 **DO** write *I have not proved Lemma 2*, then name precisely what remains **BECAUSE** it cannot cost a mark, because the grader finds the gap either way, and it converts a 1 into a 2 by making the reduction a markable object
- **WHEN** a paragraph does not discharge one of the problem's clauses **DO** do not write it **BECAUSE** correct mathematics answering a question nobody asked is documented 25 times, and it is the ink that steals the minutes the licence clauses needed
- **WHEN** you have a complete solution and forty minutes left **DO** convert the parks first, verify second, tidy the paper third **BECAUSE** converting parks pays three to six marks a session and improving prose that is already complete pays zero

One closing instruction, and it outranks every entry above.

At minute 130, with two problems closed and B4 half open, you will want to keep pushing B4 because it feels close. Feeling close is not information; it is what an unsolved problem feels like from the inside at minute 130 every single time. The protocol was written by a version of you who is calm, who has the whole clock in view, and who is not currently inside the problem. Trust it over the feeling. That is what it is for.