How to Write Mathematics

How to Write Mathematics

Proofs, Papers, and Problem Sets
for Students

Gaurav Tiwari

gauravtiwari.org

First edition, 2026.

gauravtiwari.org

To every student who knew the answer
and still lost the marks.

Contents

Preface

Somewhere right now, a student who solved every problem on the sheet is losing a third of the marks. Not for wrong mathematics. For write-ups the grader couldn't follow: a proof that assumes what it promises, a solution with no stated answer, an induction missing its base case. I've watched it happen from both sides of the desk, first as the student, then for years as the teacher and grader, and the pattern never changes. Mathematical writing is a separate skill from mathematics, it's the skill the marks are actually paid through, and almost nobody teaches it.

This book teaches it. It's written for the student in a first proofs course, and for everyone within shouting distance of one: the engineering student writing lab reports full of derivations, the computer science student meeting induction in an algorithms course, the final-year student staring at a blank thesis document. You need no mathematics beyond what your own courses are already teaching. Every example here runs on first-course material, even integers, $\sqrt{2}$, simple limits, because the writing lessons transfer and the mathematics shouldn't get in their way.

A word on how to read it. The four parts build in order: sentences (Part I), proofs (Part II), problem sets and exams (Part III), papers and theses (Part IV). Reading straight through works. But this is a toolbox, not a novel, and I built it for busy people. If a problem set is due Tuesday, start with Chapter 8 and the checklist in Appendix B, then come back for the rest. If a graded proof just came back bleeding, Chapter 7's makeovers are the fastest medicine. The appendices are the working tools: keep them close, use them often, write in the margins.

Two promises about what this book won't do. It won't teach you the mathematics; your courses own that job. And it won't hand you a formula that writes proofs for you, because none exists. What it does instead is show you, rewrite by rewrite, exactly where student write-ups leak marks and how to seal every leak. The skills are small, mechanical in places, and they compound. That's the whole secret, and it fits in a sentence: the mathematics you already do is the hard part; the writing is learnable bookkeeping, and this book is the ledger.

Gaurav Tiwari
gauravtiwari.org

Part I

The Reader Is Always Right

1

Why Good Mathematicians Write Badly

You solved the problem. You saw the trick, you did the algebra, you got the answer. Then the graded sheet came back and half the marks were gone. The comment in the margin said *unclear*, or *justify*, or nothing at all, just a bracket around a paragraph you thought was fine.

Here's the uncomfortable truth behind that bracket: knowing mathematics and communicating mathematics are two different skills. Your courses teach the first one and quietly grade you on both. Nobody sits you down and shows you how a proof is supposed to *read*. This book does.

1.1 The reader's contract

Every piece of mathematical writing is a contract with a reader. The terms are simple. You promise that every claim follows from something the reader already accepts. The reader promises to believe you exactly as far as your justifications reach, and not one step further.

Most lost marks are breaches of this contract. Not wrong mathematics, just claims that appear from nowhere, symbols that were never introduced, steps that live only in your head.

> *The reader owes you nothing. Every symbol you use, you introduce. Every claim you make, you justify. Every step you skip, you lose.*

1.2 Writing is not transcription

Scratch work and finished mathematics are different documents. Scratch work records the order in which you *found* things. A finished proof presents the order in which things are *true*. Confusing the two is the single most common writing mistake students make, and it's the subject of Chapter 5.

Let's make this concrete with the smallest example that shows the disease and the cure.

Definition 1.1. An integer n is *even* if there exists an integer k such that $n = 2k$.

Theorem 1.2. *The sum of two even integers is even.*

A true statement, and an easy one. Watch how easily the write-up still goes wrong. Here is the proof the way it appears on thousands of problem sets every semester.

BEFORE

Let x and y be even. So $x = 2k$ and $y = 2k$. Then $x + y = 4k$, which is even because $4k = 2(2k)$.

The algebra inside is correct. The proof is still broken, because $x = 2k$ and $y = 2k$ says $x = y$. The writer proved that twice an even number is even. The theorem claims more, and the grader will say so. One letter, half the marks.

AFTER

Proof. Let x and y be even integers. By Definition 1.1 there exist integers j and k with $x = 2j$ and $y = 2k$. Then

$$x + y = 2j + 2k = 2(j + k),$$

and $j + k$ is an integer. Hence $x + y$ is even. □

Read the two versions side by side and notice what changed. The rewrite names two independent witnesses, j and k, because the hypothesis gives two independent even numbers. It cites the definition it's using. It ends by checking the exact condition the definition demands, an integer $j+k$ with $x+y = 2(j+k)$. Nothing clever happened. The mathematics was never the problem.

1.3 Your first hostile reader

Students imagine the grader as a mind reader with an answer key. Graders are neither. A grader is a tired person with ninety scripts, twelve minutes for yours, and a rubric that rewards visible justification. They will not reconstruct your intent. They grade what is on the page, and the page is all you control.

This is excellent news. It means the skill that recovers those marks is learnable, mechanical in places, and much easier than the mathematics you've already survived. You'll practice it on proofs in Part II, on problem sets and exams in Part III, and on papers and theses in Part IV.

1.4 What this book will and won't do

This book won't teach you real analysis, and it won't hand you a formula that writes proofs for you. It will show you, with before and after rewrites like the one above, exactly where write-ups leak marks and how to seal them. Small habits, applied every time, compound into work that reads like it came from someone who knows what they're doing. Because by then, you will.

We start where every write-up starts: with the sentence. Mathematical notation has a grammar, and most students were never taught it. That's Chapter 2.

2

The Grammar of Mathematics

Mathematical writing is still writing. Every equation you put on paper lives inside a sentence, and that sentence has a subject, a verb, and a reader trying to parse it. Students who treat symbols as a separate language, exempt from grammar, produce pages that no one can read aloud. That's the test, by the way. If you can't read your solution aloud as English, it isn't finished.

This chapter gives you the grammar rules nobody wrote on the board. There are only a handful. They fix more write-ups than any amount of extra mathematics.

2.1 Symbols are words in a sentence

Read this theorem statement aloud:

$$\text{Let } n \in \mathbb{Z}. \text{ If } 2 \mid n^2, \text{ then } 2 \mid n.$$

You said something like: "Let n be an integer. If 2 divides n squared, then 2 divides n." The symbols expanded into words. That's what symbols are: abbreviations for words, sitting in an ordinary English sentence. A displayed equation is a noun phrase or a clause. The sentence continues through it.

Three rules follow immediately.

Never start a sentence with a symbol. The reader's eye uses capital letters to find sentence boundaries. A sentence that opens with x has no boundary marker, and when the previous sentence also ended in a symbol, the two smear together. Compare:

> Bad: Suppose $x > 1$. $x^2 > x$ follows.
> Good: Suppose $x > 1$. It follows that $x^2 > x$.

The fix costs three words. It's never wrong, so make it a reflex.

Don't let two formulas from different clauses touch. A sentence like "Since $a < b$, c divides d" forces the reader to decide where one formula ends and the next begins. Insert a word: "Since $a < b$, we know c divides d."

Read it aloud before you submit it. Every rule in this chapter is a special case of this one. If the spoken version is gibberish, broken, or ambiguous, the written version is too. Your ear catches what your eye forgives.

> *Every symbol is a word. Every formula lives in a sentence. If you can't read it aloud, you haven't written it yet.*

2.2 The equals-sign epidemic

The equals sign has exactly one meaning: the object on the left is the same as the object on the right. It does not mean "and the next step is," it does not mean "which implies," and it does not mean "I'm still working." Watch it get abused the way graders see every week. The problem: solve $2x + 6 = 10$.

BEFORE

$2x + 6 = 10 = 2x = 4 = x = 2$

Read it aloud: "$2x + 6$ equals 10 equals $2x$ equals 4 equals x equals 2." Taken literally, this asserts $10 = 4 = 2$. The writer meant "and then I did this." The equals sign can't say that. English can:

AFTER

From $2x + 6 = 10$, subtracting 6 from both sides gives $2x = 4$, so $x = 2$.

The same disease has a chronic form: the run-on calculation, where a half page of equalities chains together things that are equal with things that merely follow. When you do a computation, equality chains are fine, and often ideal:

$$(n+1)^2 - n^2 = n^2 + 2n + 1 - n^2 = 2n + 1.$$

Every link here really is an equality of the same object, so the chain is honest. The rule isn't "avoid chains." The rule is: everything joined by = must actually be equal. Steps of reasoning get words: *so*, *hence*, *therefore*, *which gives*.

The same goes for the implication arrow. Writing ⇒ between every line of a solution isn't logic, it's decoration. If you mean "therefore," write *therefore*. Save the arrow for statements about implications, like "we will show that $P \Rightarrow Q$."

2.3 Quantifiers, and the order that changes everything

Two small phrases carry most of mathematics: *for every* and *there exists*. Their order changes what a sentence means, and quantifier-order mistakes are among the most expensive errors a student can make, because they look like typos and grade like misunderstandings.

Compare these two English sentences:

1. Every person has a mother.
2. There is a mother that every person has.

Sentence 1 is true. Sentence 2 claims one shared mother for all of humanity. Formally, the difference is

$$\forall x\, \exists y\, M(y,x) \qquad \text{versus} \qquad \exists y\, \forall x\, M(y,x).$$

Same symbols, opposite order, different world.

Now the mathematical version. "Every nonempty finite set of integers has a largest element" means: for every such set S, there exists an element $m \in S$ with $m \geq s$ for all $s \in S$. The m is allowed to depend on S. Swap the order and you'd be claiming one universal integer that tops every finite set. Nonsense, and yet write-ups claim things like it all the time, silently, by being vague about which variable was chosen first.

Here's the discipline that prevents it:

Introduce every variable before you use it, and say how it enters. Each variable arrives in one of two ways. Either you're handed it ("let $\varepsilon > 0$ be given," it's universally quantified, you control nothing), or you build it ("choose $N = \lceil 1/\varepsilon \rceil$," it's existential, you control everything, and it may depend on anything introduced before it, but nothing after). A proof where every variable has a visible birthplace and a visible dependence order is a proof where quantifiers can't tangle.

That single habit, variables in order of birth, will matter again in the epsilon-delta makeover in Chapter 7.

2.4 Punctuating mathematics

Displayed equations are parts of sentences, so they take punctuation like parts of sentences. If the sentence ends at the equation, the equation ends with a period:

$$x + y = 2(j + k).$$

If the sentence continues, use a comma or nothing, and carry on:

$$x + y = 2(j + k),$$

which is even. Open any well-edited textbook and you'll see periods and commas inside displayed math on nearly every page. Most students never noticed. Now you won't be able to stop noticing.

Two more conventions worth adopting:

- Display an equation when it's important or too big for the line, not as a default. A solution where every trivial step is displayed reads like a slideshow.
- Never stack punctuation after a formula ("$x = 2$.,"). One mark, correctly chosen.

2.5 Symbols or words?

Students err in both directions. Some write pure symbol soup:

$$\forall\varepsilon > 0\ \exists\delta > 0 \text{ s.t. } |x - 2| < \delta \Rightarrow |f(x) - 7| < \varepsilon$$

Others write bureaucratic prose with no symbols at all, spelling out "the quantity obtained by squaring the difference" when $(a - b)^2$ would do.

The working rule: **symbols for objects and relations, words for reasoning.** Objects (n, f, $\mathbb{R}$), operations (n^2, $f(x)$), and relations ($<$, $=$, $\in$) live comfortably in symbols. The connective tissue of an argument, *suppose*, *since*, *therefore*, *it follows that*, *for every*, belongs in words in your prose. Logical symbols like $\forall$, $\exists$, $\Rightarrow$, $\wedge$ are for statements *about* logic and for your scratch work, not for the running text of a proof.

So the symbol soup above becomes:

> For every $\varepsilon > 0$ there is a $\delta > 0$ such that $|f(x) - 7| < \varepsilon$ whenever $|x - 2| < \delta$.

Same content, and now it reads aloud as an English sentence. Appendix A collects the standard translations, symbol by symbol.

2.6 Notation hygiene

Finally, the housekeeping rules. Each one sounds obvious. Each one is violated on half the scripts in any grading pile.

One symbol, one meaning. If n is the size of your set, n can't moonlight as an index in the same solution. When you run out of letters, there are more: m, k, j, subscripts, primes.

One meaning, one symbol. Don't drift from $f(x)$ to y and back for the same function. The reader assumes a change of notation signals a change of object.

Define before use. The first time a non-standard symbol appears, it gets introduced: "let $d(n)$ denote the number of divisors of n." Standard notation ($\mathbb{R}$, π, $|x|$) needs no introduction. When in doubt, introduce.

Respect convention. By long tradition, ε and δ are small positive quantities, n and m are integers, f and g are functions, p is a prime. You can fight convention, and the notation police won't arrest you, but every violation taxes your reader's attention, and you need that attention for the mathematics.

Grammar tells you how to say things. It doesn't tell you what you're saying. Before you can write a proof, you have to be able to take a mathematical statement apart and see exactly what it claims, which hypotheses you're given and which conclusion you owe. That dissection is the next chapter.

3

Definitions, Statements, and the Anatomy of a Theorem

You can't prove what you can't parse. A surprising share of failed proofs die before the first line is written, because the writer never worked out exactly what the statement claims: what's given, what's owed, and which words in it are hiding quantifiers. This chapter is about the dissection you do before any writing starts.

3.1 Hypothesis and conclusion

Every theorem is a deal: *if* you hand me these hypotheses, *then* I owe you this conclusion. The first act of dissection is to separate the two, because they get opposite treatment in your proof. Hypotheses you may use, free of charge, whenever you like. The conclusion you may not use at all. It's the debt you're paying off, and touching it early is the crime called circular reasoning.

Take a statement in theorem form:

> Let a, b, and c be integers. If $a \mid b$ and $b \mid c$, then $a \mid c$.

Given: three integers, and two divisibility facts. Owed: one divisibility fact. Now unpack what the facts *mean*. By the definition of divisibility, $a \mid b$ says there exists an integer j with $b = aj$, and $b \mid c$ says there exists an integer k with $c = bk$. What we owe, unpacked the same way: an integer m with $c = am$. Suddenly the proof is a one-liner waiting to happen: $c = bk = (aj)k = a(jk)$, so take $m = jk$.

Notice what did the work. Not cleverness. Unpacking. We'll return to this move so often that it deserves its name now:

> *When stuck, unpack the definitions. Half of all homework proofs are finished by writing out what the hypothesis means and what the conclusion means, then staring at the gap.*

3.2 The hidden quantifiers

Most theorem statements in English hide their quantifiers. Your job is to drag them into the light before writing, because they dictate the shape of the proof.

> "The square of an odd integer is odd."

No *for all* appears, but it's there: *for every* odd integer n, the number n^2 is odd. A universal claim. To prove it, you must argue for an *arbitrary* odd integer, not a friendly example. The proof's first line is forced: "Let n be an odd integer." One sentence in, and the quantifier already wrote it for you.

Contrast:

> "There is a real number that equals its own square other than zero."

An existential claim. Here one example is a complete proof: "Take $x = 1$; then $x^2 = 1 = x$." Done. No arbitrary anything.

This is the great asymmetry, and it cuts both ways:

- To *prove* "for all": argue for an arbitrary element. To *disprove* it: exhibit one counterexample.
- To *prove* "there exists": exhibit one witness. To *disprove* it: argue that an arbitrary candidate fails.

Students lose marks by proving universal claims with examples, and almost as often, by burying an existence proof under pages of generality when a single witness would do. Read the quantifier first. It tells you which kind of day you're having.

3.3 Converse, contrapositive, negation

Once a statement is in if-then form, three relatives show up, and writing mathematics means never confusing the family members.

Start with a statement $P \Rightarrow Q$: "If n^2 is even, then n is even."

The converse is $Q \Rightarrow P$: "If n is even, then n^2 is even." A different statement with its own truth value. Both happen to be true here, which is exactly why the example is dangerous: it trains bad instincts. ("If n is a multiple of 4, then n is even" has a false converse.) Proving the converse when the problem asked for the statement is a total loss of marks, not partial. Graders see it weekly.

The contrapositive is $\neg Q \Rightarrow \neg P$: "If n is odd, then n^2 is odd." This one is logically *identical* to the original. Proving it is proving the original, and often it's easier, which is why it gets its own template in Chapter 4. When you use it, say so: "We prove the contrapositive." Three words, and the grader relaxes.

The negation is what must hold for the statement to fail: P true and Q false. Not "P false," not "Q false," but the combination: an even n^2 with an odd n. You need negations twice in life: to hunt counterexamples, and to open a proof by contradiction. Negating quantified statements mechanically saves you there: *not (for all x, A)* becomes *there is an x with not-A*, and *not (there is an x with A)* becomes *for all x, not-A*. The quantifier flips, the negation moves inward. Practice until it's a reflex, because a proof by contradiction that starts from the wrong negation proves nothing at all.

3.4 Reading a hard statement

Serious statements stack these features. Here's one from real analysis, the kind that opens a second-year course:

> "Every bounded monotone sequence of real numbers converges."

Dissect before touching a pen. Universal quantifier: every sequence with two properties. Hypothesis 1: bounded, which unpacks to *there exists* M with $|a_n| \leq M$ *for all* n. Hypothesis 2: monotone, say increasing, $a_n \leq a_{n+1}$ for all n. Conclusion: converges, which unpacks to the full ε-N sentence: there exists L such that for every $\varepsilon > 0$ there exists N such that $|a_n - L| < \varepsilon$ for all $n > N$.

Now count what the dissection bought you before any idea arrived. The proof must work for an arbitrary such sequence. It may use M and monotonicity freely. And it owes, first of all, a candidate for L, because the conclusion starts with *there exists*. The proof's skeleton is already visible, and you haven't had a single insight yet. That's the point. Dissection is free, and it's where you should spend your first five minutes on any problem.

3.5 Writing your own definitions

Sooner or later, in a project or thesis, you'll define something yourself. The craft has three rules.

A definition is a contract, not a description. It must give an exact test for membership. "A number is *tidy* if it is fairly close to a perfect square" fails: how close is fairly? "An integer n is *tidy* if $|n - k^2| \leq 1$ for some integer k" works. Every term on the right side must already be defined, or standard.

No circularity. The word being defined cannot appear in its own defining condition, openly or in disguise. You'd be surprised how often a disguise slips in through a synonym.

Convention: definitions use "if," and it means "if and only if." Written mathematics states definitions with a bare *if*: "n is even if $n = 2k$ for some integer k." By convention this *if* is understood as *if and only if*, and the definition cuts both ways: it's your tool for proving a number is even, and equally your tool for using the evenness you're handed. Theorems don't get this courtesy. In a theorem, *if* means *if*, and the difference between *if* and *if and only if* is the difference between one proof and two.

That last point generalizes. An *if and only if* theorem is two theorems sharing a sentence, and its proof is almost always two proofs, labeled visibly: ($\Rightarrow$) and ($\Leftarrow$). Prove one direction and you've done half the job, however long your page looks.

You now know what a statement is: hypotheses you can spend, a conclusion you owe, quantifiers that fix the shape of the argument. The next chapter is about those shapes. There are only six you'll meet this decade, and each comes with opening and closing sentences you can learn like chess openings.

Part II

Proofs

4

The Shape of a Proof

Chess players learn openings. Musicians learn scales. Proof writers learn shapes: a small set of standard argument structures, each with its own first sentence, its own last sentence, and its own obligations in between. There are six that carry an undergraduate degree. This chapter gives you each one as a writing template, because the fastest way to write a clear proof is to know, before you start, what its skeleton must look like.

A warning before the catalog. The shape is the *write-up's* structure, not a recipe for finding the idea. You'll still have to think. But knowing the shape tells you what the first and last sentences must be, and that turns a blank page into a fill-in-the-gap exercise.

4.1 Direct proof

The default shape. Assume the hypothesis, march to the conclusion.

> **Skeleton.** "Let [objects] satisfy [hypothesis]. [Unpack the hypothesis. Compute, deduce, chain facts.] Therefore [conclusion]."

Theorem 4.1. *The product of two odd integers is odd.*

Proof. Let x and y be odd integers. Then $x = 2j + 1$ and $y = 2k + 1$ for some integers j and k. Multiplying,

$$xy = (2j + 1)(2k + 1) = 4jk + 2j + 2k + 1 = 2(2jk + j + k) + 1,$$

and $2jk + j + k$ is an integer. Hence xy is odd. □

Note the two witnesses j and k, one per hypothesis, the lesson of Chapter 1. Note also the shape of the final line: it verifies the exact condition in the definition of odd. A direct proof ends by presenting the conclusion's definition, satisfied.

4.2 Proof by contrapositive

When the conclusion is easier to negate than the hypothesis is to use, prove $\neg Q \Rightarrow \neg P$ instead. It's logically the same statement, and announcing the switch costs one sentence.

> **Skeleton.** "We prove the contrapositive: if [not-Q], then [not-P]. Suppose [not-Q]. [Direct proof from here.] Therefore [not-P]."

Theorem 4.2. *Let n be an integer. If n^2 is even, then n is even.*

Proof. We prove the contrapositive: if n is odd, then n^2 is odd. Suppose n is odd, so $n = 2k + 1$ for some integer k. Then

$$n^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1,$$

which is odd. This proves the contrapositive, and hence the theorem. □

Try proving this directly, starting from "$n^2 = 2m$," and you'll find yourself needing square roots of even numbers and going nowhere. That experience, direct route blocked, negated conclusion friendly, is the signal to switch shapes. The switch itself must be announced. An unannounced contrapositive reads like a proof of the converse, and Chapter 3 told you how that gets graded.

4.3 Proof by contradiction

Assume the whole statement fails, then break something. The shape is powerful and overused. Use it when the negation hands you something concrete to work with, as it does in the most famous example there is.

> **Skeleton.** "Suppose, for contradiction, that [negation of the statement]. [Deduce consequences.] But this contradicts [known fact / earlier line]. Therefore [statement]."

Theorem 4.3. *$\sqrt{2}$ is irrational.*

Proof. Suppose, for contradiction, that $\sqrt{2}$ is rational. Then we may write $\sqrt{2} = a/b$ where a and b are integers with no common factor greater than 1. Squaring both sides and clearing denominators gives

$$a^2 = 2b^2.$$

Thus a^2 is even, so a is even by Theorem 4.2. Write $a = 2c$ for an integer c. Substituting, $4c^2 = 2b^2$, so $b^2 = 2c^2$. Thus b^2 is even, so b is even, again by Theorem 4.2. Now a and b are both even, so they share the factor 2. This contradicts our choice of a and b with no common factor. Therefore $\sqrt{2}$ is irrational. □

Three writing points hide in there. First, "suppose, for contradiction" announces the shape immediately; the reader knows every following line lives inside a doomed assumption. Second, the lowest-terms choice is stated *when the fraction is introduced*, because that's the fact the contradiction will strike. Third, the proof cites the theorem it uses by name and number. Your homework can cite results the same way: "by the theorem from lecture 7," or "by Exercise 3(a)," as long as the course allows it.

One stylistic caution. If your "contradiction" proof assumes $\neg Q$, never touches P's failure, and ends by proving Q outright, it was a direct proof (or a contrapositive) wearing a trench coat. Rewrite it in its honest shape. Graders notice, and the honest version is always shorter.

4.4 Induction

The shape for statements indexed by natural numbers: prove it for the first value, prove each value hands the property to the next. Induction write-ups fail on discipline, not on ideas, so the template matters more here than anywhere.

> **Skeleton.** "We proceed by induction on n. *Base case* ($n = 1$): [verify directly]. *Inductive step*: suppose the statement holds for some $n \geq 1$; that is, [write out the hypothesis]. We show it holds for $n + 1$. [Argument that uses the hypothesis.] This completes the induction."

Theorem 4.4. *For every positive integer n,*

$$1 + 2 + \cdots + n = \frac{n(n+1)}{2}.$$

Proof. We proceed by induction on n.

Base case. For $n = 1$, the left side is 1 and the right side is $1 \cdot 2/2 = 1$. They agree.

Inductive step. Suppose that for some $n \geq 1$,

$$1 + 2 + \cdots + n = \frac{n(n+1)}{2}.$$

Adding $n + 1$ to both sides,

$$1 + 2 + \cdots + n + (n+1) = \frac{n(n+1)}{2} + (n+1) = \frac{(n+1)(n+2)}{2},$$

which is the statement for $n + 1$. This completes the induction. □

The two non-negotiables: the base case gets checked, visibly, even when it's trivial, and the inductive hypothesis is *written out*, not just named. "Assume it's true for n" followed by algebra that never uses the assumption is the most common broken induction in any grading pile. If your inductive step never cites the hypothesis, either the write-up hides the moment it was used (fix the write-up) or the argument never needed it (then it isn't induction; write it directly). Chapter 7 performs surgery on a broken induction at full length.

4.5 Proof by cases

When the hypothesis splits into finitely many situations, treat each one, and prove the same conclusion in all of them.

> **Skeleton.** "We consider two cases. *Case 1:* [first situation]. [Argument.] *Case 2:* [second situation]. [Argument.] In either case, [conclusion]."

Theorem 4.5. *For every integer n, the number $n^2 + n$ is even.*

Proof. Note that $n^2 + n = n(n + 1)$. We consider two cases.

Case 1: n is even. Then $n = 2k$ for an integer k, so $n(n + 1) = 2k(n + 1)$, which is even.

Case 2: n is odd. Then $n + 1$ is even, say $n + 1 = 2k$, so $n(n + 1) = 2kn$, which is even.

In either case $n^2 + n$ is even. □

The obligations: the cases must *cover* every possibility allowed by the hypothesis, and you should say why when it isn't obvious. Label the cases, keep them parallel in structure, and close with the "in either case" sentence that reunites them. A case analysis that silently forgets a case isn't a partial proof. It's no proof, at partial-credit prices.

4.6 Existence and uniqueness

"There is exactly one" is two claims stapled together, and the write-up must visibly do two jobs: produce (existence) and pin down (uniqueness).

> **Skeleton.** "*Existence.* [Exhibit a witness, verify it works.] *Uniqueness.* Suppose u and v both satisfy [the property]. [Deduce $u = v$.]"

Theorem 4.6. *Let a and b be real numbers with $a \neq 0$. The equation $ax = b$ has exactly one real solution.*

Proof. Existence. Take $x = b/a$, which is defined since $a \neq 0$. Then $ax = a(b/a) = b$, so x is a solution.

Uniqueness. Suppose u and v are both solutions, so $au = b$ and $av = b$. Then $au = av$, and since $a \neq 0$ we may divide by a to get $u = v$. □

The uniqueness half has a fixed grammar: assume *two* objects with the property, derive that they're equal. Students regularly prove uniqueness by "deriving" the formula for the solution, which only shows every solution must look like b/a; fine, that is a valid uniqueness argument, but then existence still needs its own verification that b/a actually works. Keep the two labeled halves and nothing goes missing.

Six shapes: direct, contrapositive, contradiction, induction, cases, existence-uniqueness. Between them they cover nearly every proof you'll write before graduate school. What they don't cover is the mess of actually finding the argument, the pages of failed attempts and backwards reasoning that precede every clean proof. That mess has a workflow, and turning it into a final draft is the next chapter.

5

From Scratch Work to Final Draft

Here's a secret that would have saved me a year of confusion as a student: the proofs in your textbook were not discovered in the order they're printed. Nobody sits down and writes "Let $\varepsilon > 0$ be given" as their first thought. The polished proof is a reconstruction, arranged for the reader's benefit, of a discovery that happened backwards, sideways, and with three dead ends the author isn't telling you about.

This has a practical consequence. You are always writing two documents: the one where you *find* the argument, and the one where you *present* it. They have different audiences, different rules, and different directions. Most bad mathematical writing comes from submitting document one with document two's deadline.

5.1 Exploration runs backward

Take the simplest inequality worth proving, a two-variable version of the AM-GM inequality.

Theorem 5.1. *For all real numbers a and b,*

$$\frac{a^2 + b^2}{2} \geq ab.$$

How do you *find* the proof? You start from the thing you want, which you're not allowed to assume, and you push it around:

> Want: $\frac{a^2+b^2}{2} \geq ab$. Multiply by 2: want $a^2 + b^2 \geq 2ab$. Move everything left: want $a^2 - 2ab + b^2 \geq 0$. Hey, that's $(a-b)^2 \geq 0$. Squares are always ≥ 0. Done!

That's honest scratch work, and every step of it is illegal in a final proof, because each line *assumes* the thing being proved and derives something true from it. Deriving truth from your goal proves nothing: false statements imply true ones all day. (From $1 = 2$, add the flipped copy $2 = 1$ to get $3 = 3$. True conclusion, worthless argument.)

But the scratch work found the key: $(a-b)^2 \geq 0$. And here's the move that separates writers from strugglers. **The final proof is the scratch work, reversed.** Start from the unconditional truth you landed on and walk the chain back to the goal, checking each step survives the reversal:

Proof. Let a and b be real numbers. Since every square is nonnegative,

$$(a-b)^2 \geq 0.$$

Expanding, $a^2 - 2ab + b^2 \geq 0$, so $a^2 + b^2 \geq 2ab$. Dividing by 2 gives the result. □

Same steps, opposite direction, completely different logical status. The scratch version begs the question. The reversed version proves it.

> *Explore backward from the goal. Write forward from the truth. The final draft is your scratch work read in reverse, after you've checked every arrow actually runs both ways.*

That last clause earns its place. Steps like "multiply by 2" and "move terms across" reverse cleanly. Steps like squaring both sides, multiplying by a quantity that might be negative or zero, or dropping information do not. The reversal check, done consciously at each link, is where you catch the step that only worked in one direction, before the grader catches it for you.

5.2 Two documents, two standards

Because the two documents do different jobs, hold them to different standards, and keep them physically separate: different page, different file, different notebook. Mixing them is how "want" statements leak into submitted work and read as claims.

Scratch work is for you. It can be messy, wrong, abbreviated, full of arrows and question marks and *want*s. Its one obligation is honesty with yourself: mark what you *know* versus what you *want*, or the leak starts here. Write *want:* in front of every unproved goal, every time. The three-second habit saves entire arguments from circularity.

The final draft is for a stranger. It runs forward, from known things to the goal. Every variable has a birthplace (Chapter 2), every claim a justification (Chapter 1), and the shape announces itself in the first sentence (Chapter 4). Nothing in it refers to the journey: no "after trying several approaches," no "we first attempted induction." The reader is paying for the building, not the scaffolding.

5.3 Finding the key idea

When you reread scratch work before drafting, you're hunting for one thing: the *key idea*, the step that made the problem fall. In the AM-GM exploration it was recognizing $(a-b)^2$. In the $\sqrt{2}$ proof of Chapter 4 it's the lowest-terms choice that the conclusion later violates. Almost every homework-size proof has exactly one key idea surrounded by bookkeeping.

Name it in one sentence before you draft. Two benefits. First, it organizes the write-up: the draft becomes setup, key idea, consequences, conclusion, in that order, and paragraph breaks fall out naturally. Second, if you *can't* name the key idea in a sentence, you don't understand your own argument yet, and drafting now will produce fog. That diagnosis is worth marks: fog in, fog out.

For longer arguments, the key-idea sentence scales up to an outline. Before drafting anything over a page, list the three to five waypoints the argument passes through, in one line each. If a waypoint is a self-contained fact with its own hypotheses and conclusion, pull it out as a lemma: state it before the main proof, prove it separately, cite it by number when needed. One lemma per genuinely separable idea. A proof that pauses for half a page to establish a side fact inline is a proof the reader will lose twice, once going in and once coming back. The lemma also becomes reusable: state it cleanly and you can cite it again in problem 4 without re-proving anything.

5.4 When to start writing

The honest answer: later than you want to. Drafting a proof you haven't finished finding produces the worst documents in mathematics, half exploration and half exposition, wearing neither's virtues. You know you're ready to draft when three things are true:

1. You can state the key idea in one sentence.
2. You know the shape (direct, contradiction, induction, . . .), so you know the first and last sentences already.
3. You've checked the chain runs forward: every step, in final draft order, follows from what precedes it.

Once those hold, drafting is fast. It's transcription plus grammar, usually fifteen minutes for a homework proof, and revision (Chapter 14) will be light because the structure was right before the first word landed.

If the deadline arrives while item 3 still fails, you have a gap. The professional move is to submit the argument with the gap *labeled*, not smoothed over. Exactly how to phrase that, and why an honest gap outscores a fake bridge every single time, is covered with the rest of exam craft in Chapter 8.

You now have the full production line: dissect the statement, explore backward, find the key idea, choose the shape, write forward. What's left is the finish: the sentence-level choices, the words that build trust and the words that burn it. That's style, and it's next.

6

Proof Style

Structure gets you a correct proof. Style gets you a *believed* one. Two write-ups of the same argument, both technically complete, can leave a grader in two different moods, and the difference is a few dozen small word choices. This chapter is about those choices. None of them is hard. All of them compound.

6.1 The banned list

Some phrases appear only in student proofs, never in good ones, and every grader reads them the same way.

"Clearly" / "obviously" / "trivially." These words don't make a step clear. They announce that you've decided not to justify it, and they dare the reader to disagree. Here's the grader's private test, worth internalizing: if the step really is immediate, the word is unnecessary. If it isn't, the word is a confession. Either way the word does nothing for you. When you catch yourself typing "clearly," one of two things is true. Either the justification is one short phrase, in which case write the phrase ("since squares are nonnegative"), or it's longer, in which case the step deserves a real sentence. There's a third possibility, the expensive one: you typed "clearly" because you *couldn't* justify the step. Graders have seen ten thousand clearlys, and they know exactly where the bodies are buried.

"It is easy to see." Same disease, more syllables. If it's easy, show it fast. If seeing it took you twenty minutes, it isn't easy, and the sentence is false advertising.

"We must show" used as "we have shown." Stating the goal is good practice. Stating it and then never returning to discharge it is the fog machine running through Chapter 7's makeovers. Keep *want* and *have* vocabulary rigidly separate: "we must show X" opens a task; only "therefore X" closes it.

"By simple algebra" hiding six lines of algebra. Skipping routine manipulation is fine and often right, but name what was done: "expanding and collecting terms gives." The reader can then redo it. "Simple algebra" that the reader can't reproduce in the margin wasn't simple.

6.2 Who is speaking? Pronouns and tense

Mathematical English has settled, over a century, on a few defaults. You can deviate, but you should know what the defaults are and what deviations signal.

Use "we," and mean it. The convention "we" means *the writer and the reader, walking together*: "we now show," "we conclude." It reads as collaboration, not royalty. "I" is reserved for genuinely personal statements ("I learned this proof from N. Sharma"), which are rare in homework and papers.

Or use the imperative. "Let x be even. Consider x^2. Note that..." Commands to the reader are standard, compact, and fine. Most good proofs mix "we" sentences and imperative sentences freely.

Avoid the passive fog. "It can be shown that" is the worst of all worlds: no actor, no proof, no citation. Who shows it? If you: show it. If a textbook: cite it. If nobody: you've written a rumor.

Present tense, throughout. Mathematics happens in an eternal present: "we see," not "we saw"; "it follows," not "it followed." The one exception is referring to your own document's past: "as we proved in Lemma 2."

6.3 Calibrating detail

The hardest style question, because it has no fixed answer: how much do you justify? Write every arithmetic step and a two-line proof bloats to two pages of noise. Skip too much and the grader can't distinguish insight from luck.

The rule that resolves most cases: **write for a classmate one week behind you.** Not for the professor, who needs nothing, and not for a stranger off the street, who needs everything. A peer who attends your course but missed the current week. Steps that were routine *last* month get a phrase or nothing. Steps that use *this* month's ideas get shown. The new idea in your proof, the key idea of Chapter 5, always gets shown in full, no matter how quick it looks, because it's the part the grade is actually for. Nobody ever lost a mark showing the key step too clearly.

Three refinements:

- *The course sets the floor.* In week one of a proofs course, "the sum of two integers is an integer" may itself be the point. By week ten it's beneath mention. Same fact, different course moment, different treatment. When the problem statement says "prove carefully," the floor drops: unpack everything.
- *Cited results must be citable.* You may use anything the course has established, by name or number: "by the Mean Value Theorem." You may not use next semester's machinery to vaporize this semester's exercise. Proving a limit exercise "by L'Hôpital" in the week the course is defining limits earns zero, correctly: the exercise existed to test the definition.
- *When genuinely unsure, err long, in the right place.* Extra detail costs the reader seconds. A missing step at the crux can cost you the problem. But put the generosity at the crux, not in the arithmetic.

6.4 The connectives, and what they promise

Proofs move by connective words, and the good ones form a precise little vocabulary. Each promises the reader something different:

- *Therefore, hence, thus, so*: what follows is a consequence of what you just read. (Interchangeable; vary them for rhythm.)
- *Since, because, as*: justification incoming before the claim.
- *That is, in other words*: same content, restated. No new claim.
- *Moreover, furthermore*: a second fact, parallel to the last, not derived from it.
- *But, however*: tension; usually feeding a contradiction or an exception.
- *In particular*: a special case of what was just established, usually the one we actually need.

Two abuses to avoid. Don't write *therefore* when you mean *moreover*: claiming consequence where there's only sequence is a logic error dressed as a style choice, and graders read it as the former. And don't let every sentence start with *so*: monotone connectives make even correct proofs feel like a shrug. Read the proof aloud (the Chapter 2 test) and your ear will flag the drone.

6.5 Paragraphs, display, and the shape on the page

A proof longer than four sentences needs paragraphs, and the paragraph breaks should track the argument's phases: setup, key step, consequences, conclusion. In a case analysis, one paragraph per case. In an induction, base case and inductive step each get their own. The white space is doing signposting work, telling the reader "phase complete, breathe here."

Display an equation (centered, own line) when it's the destination of a computation, when it will be referenced later, or when inline would force a line break mid-formula. Displayed equations are emphasis; emphasize everything and nothing is emphasized.

If a displayed equation will be used again, number it and cite the number, using your document's numbering, and refer to it as "by (3)." Never "by the equation above": after the grader's eye travels, above is ambiguous, and in two-column exam booklets it's meaningless.

End every proof with a visible full stop for the argument. The boxed $\square$ does it in print (this book's proofs end with one). In handwritten work, the classical *QED*, a box, or the sentence "this completes the proof" all work. What you're buying is the reader's certainty that nothing is missing, that the page didn't just run out.

> *Style is trust management. Every honest "since," every discharged "we must show," every visible ending deposits trust. Every "clearly" at a shaky step withdraws it. Graders mark the balance.*

You now hold all the components: grammar, dissection, shapes, workflow, style. Time to watch them work together on real, broken proofs. The next chapter takes five write-ups of the kind graders see every week, diagnoses them line by line, and rebuilds each one.

7

Proof Makeovers

Theory is over. This chapter is an operating room. Five write-ups, each the kind of thing graders see every single week, each one containing correct mathematical instincts wrapped in a broken presentation. For each: the submitted version, the diagnosis, the rebuild, and what changed. Read them slowly. These five patients cover, between them, most of the marks you'll ever lose to writing.

7.1 Makeover 1: an induction that assumes too much

The problem. Prove that for every positive integer n,

$$1 + 3 + 5 + \cdots + (2n - 1) = n^2.$$

BEFORE

> Assume the formula is true for all n. Then adding the next odd number gives $n^2 + (2n + 1) = (n + 1)^2$, so it's true for $n + 1$ too. Also it works for $n = 1$ since $1 = 1$. QED.

Diagnosis. Three separate injuries.

1. "Assume the formula is true for all n" assumes the entire theorem. If it's true for all n, there's nothing to prove. The inductive hypothesis is the statement for *one* value, n, and the step transfers it to $n + 1$. This single word, *all*, is the most common fatal wound in student inductions, and many graders stop reading right there.
2. The base case arrives as an afterthought ("also"), unlabeled, at the end. Order matters for readability: foundation first, then the ladder.
3. "Adding the next odd number" adds it to one side only, and never states what equation the addition happens inside. The algebra is right; the sentence containing it is unbuilt.

AFTER

> *Proof.* We proceed by induction on n.
> *Base case.* For $n = 1$, the left side is 1 and the right side is $1^2 = 1$. They agree.
> *Inductive step.* Suppose that for some $n \geq 1$,
>
> $$1 + 3 + 5 + \cdots + (2n - 1) = n^2.$$
>
> The next odd number is $2n + 1$. Adding it to both sides,
>
> $$1 + 3 + 5 + \cdots + (2n - 1) + (2n + 1) = n^2 + 2n + 1 = (n + 1)^2,$$
>
> which is the statement for $n + 1$. By induction, the formula holds for every positive integer n. $\square$

What changed. "For all n" became "for some $n \geq 1$," the one-word repair that turns circular into valid. The hypothesis is displayed in full, so the reader sees exactly what's assumed and watches it get used. The base case moved to the front and got its label. Closing sentence states what the induction delivered.

7.2 Makeover 2: an epsilon-delta proof written backwards

The problem. Prove from the definition that $\lim_{x\to 2}(3x+1) = 7$.

BEFORE

$|(3x+1)-7| < \varepsilon$, so $|3x-6| < \varepsilon$, so $3|x-2| < \varepsilon$, so $|x-2| < \varepsilon/3$. Let $\delta = \varepsilon/3$. QED.

Diagnosis. This is scratch work, submitted (Chapter 5 predicted it). The first line assumes the inequality the proof exists to establish, then works backward to discover δ. The discovery is genuinely correct, and every grader can see the student understood the mechanics. But as written, the logic flows from goal to given: it proves "if the limit works, then $\delta = \varepsilon/3$," which is backwards. Worse, no variable is introduced: ε falls from the sky, δ arrives after being needed, and the quantifier order of Chapter 2 ("for every ε there exists δ") appears nowhere.

AFTER

Proof. Let $\varepsilon > 0$ be given. Choose $\delta = \varepsilon/3$. Suppose x satisfies $0 < |x-2| < \delta$. Then

$$|(3x+1)-7| = |3x-6| = 3|x-2| < 3\delta = \varepsilon.$$

Thus for every $\varepsilon > 0$ there is a $\delta > 0$ such that $|(3x+1)-7| < \varepsilon$ whenever $0 < |x-2| < \delta$, which is the definition of $\lim_{x\to 2}(3x+1) = 7$. □

What changed. The variables now enter in quantifier order: ε is *given* (universal, we control nothing), δ is *chosen* (existential, built from ε), x is then arbitrary within δ of 2. The computation runs forward from $|x-2| < \delta$ to the ε bound. And notice: the backward scratch work didn't die. It's exactly how $\delta = \varepsilon/3$ was found. It just doesn't appear on the final page, only its conclusion does. Keep doing the backward work. Stop submitting it.

7.3 Makeover 3: the Venn diagram that wasn't a proof

The problem. Prove that for all sets A, B, C,

$$A \setminus (B \cup C) = (A \setminus B) \cap (A \setminus C).$$

BEFORE

Draw the Venn diagram for the left side and shade it. Draw the diagram for the right side and shade it. The shaded regions are the same, so the sets are equal.

Diagnosis. The diagram is true and the proof is absent. A three-circle Venn diagram is one picture of one configuration; it illustrates the identity but doesn't establish it for *all* sets, and "the shadings look the same" is an observation about ink. Diagrams are magnificent scratch work and honest illustrations (Chapter 12 defends them properly), but a set equality is proved by showing each side contains the other, or by a chain of equivalences on membership. The grader's margin note here writes itself: *a picture is not a proof.*

AFTER

Proof. We show each side contains the other by tracking an arbitrary element. Let x be arbitrary. Then $x \in A \setminus (B \cup C)$ means $x \in A$ and $x \notin B \cup C$. Since $x \notin B \cup C$ holds exactly when $x \notin B$ and $x \notin C$, this is equivalent to

$$(x \in A \text{ and } x \notin B) \text{ and } (x \in A \text{ and } x \notin C),$$

that is, to $x \in (A \setminus B) \cap (A \setminus C)$. Since the memberships are equivalent for every x, the two sets are equal. □

What changed. Everything: this was a replacement, not a repair. The element-chasing template ("let x be arbitrary," unpack both memberships, connect them) is *the* standard shape for set identities, and it's worth memorizing as a skeleton. One honest subtlety: the middle step used De Morgan's law for *not in a union*. If your course proved it, cite it; if not, its one-line proof belongs in your write-up too. Every equivalence in a chain like this must be individually defensible; the chain format doesn't launder weak links.

7.4 Makeover 4: the prime that wasn't

The problem. Prove that there are infinitely many primes.

BEFORE

Suppose there are finitely many primes $p_1, p_2, \dots, p_n$. Let $N = p_1 p_2 \cdots p_n + 1$. Then N is prime, because dividing N by any p_i leaves remainder 1. But N is bigger than every p_i, so it's a new prime. Contradiction. QED.

Diagnosis. The most famous wrong sentence in student mathematics: "then N is prime." It doesn't follow, and it isn't needed. What follows from the remainders is only that no p_i divides N. The conclusion should be that N's prime factors (it has at least one, since $N \geq 2$) are missing from the alleged complete list. The distinction is not pedantry: with real numbers in place of the story, $2 \cdot 3 \cdot 5 \cdot 7 \cdot 11 \cdot 13 + 1 = 30031 = 59 \times 509$, which is not prime. A grader who writes "30031?" in your margin has made exactly this point. The rest of the structure (assume finite, build N, contradict completeness) is sound and stays.

AFTER

Proof. Suppose, for contradiction, that there are only finitely many primes, say $p_1, p_2, \dots, p_n$, a complete list. Let

$$N = p_1 p_2 \cdots p_n + 1.$$

Since $N \geq 2$, it has at least one prime factor q. For each i, dividing N by p_i leaves remainder 1, so $p_i \nmid N$. Hence q, which does divide N, is not any of the p_i. This contradicts the assumption that $p_1, \dots, p_n$ was a complete list of primes. Therefore there are infinitely many primes. □

What changed. One claim was replaced by a weaker, true one: not "N is prime" but "N has a prime factor outside the list." The proof got *shorter* and stronger at once, a pattern worth noticing: the repair for an overclaim is usually the modest claim you can actually defend, and the modest claim usually suffices. Also note "a complete list" stated up front:

that's the exact clause the contradiction strikes, planted where the reader can see it, the same craft point as the lowest-terms clause in Chapter 4's $\sqrt{2}$ proof.

7.5 Makeover 5: uniqueness with tangled epsilons

The problem. Prove that a convergent sequence of real numbers has exactly one limit.

BEFORE

Suppose $a_n \to L$ and $a_n \to M$. For all ε, $|a_n - L| < \varepsilon$ and $|a_n - M| < \varepsilon$. So $|L - M| < 2\varepsilon$, which is as small as we want, so $L = M$. QED.

Diagnosis. The idea is fully present, and every individual inequality is nearly right. What's missing is the quantifier scaffolding that makes it a proof.

1. "For all ε, $|a_n - L| < \varepsilon$" is false as written: it holds only for n beyond some threshold, and the thresholds for L and for M may differ. The two convergences produce two N's, and the write-up must reconcile them.
2. Which n is "$|L - M| < 2\varepsilon$" about? The triangle inequality that produces it needs an actual a_n, so some specific n must be exhibited, one beyond both thresholds.
3. "As small as we want, so $L = M$" is the right instinct, waving at a real lemma. Write the lemma's use properly: a nonnegative number smaller than every positive number is zero.

AFTER

Proof. Suppose $a_n \to L$ and $a_n \to M$. Let $\varepsilon > 0$ be given. Since $a_n \to L$, there is an N_1 such that $|a_n - L| < \varepsilon/2$ for all $n > N_1$. Since $a_n \to M$, there is an N_2 such that $|a_n - M| < \varepsilon/2$ for all $n > N_2$. Fix any $n > \max(N_1, N_2)$. By the triangle inequality,

$$|L - M| \le |L - a_n| + |a_n - M| < \frac{\varepsilon}{2} + \frac{\varepsilon}{2} = \varepsilon.$$

Thus $0 \le |L - M| < \varepsilon$ for every $\varepsilon > 0$. A nonnegative real number less than every positive real number is zero, so $|L - M| = 0$, and hence $L = M$. □

What changed. Every quantifier now has its scaffolding: ε given first, the two N's built from it and named separately, the working n fixed beyond both. The $\varepsilon/2$ split is the classic trick for landing exactly on ε, and it's worth stealing; so is $\max(N_1, N_2)$, the standard repair whenever two limit statements must hold simultaneously. And the final step is now a stated fact rather than a gesture. Compare the lengths: the honest version is four sentences longer. That's the entire price of the marks.

7.6 Makeover 6: half a biconditional

The problem. Prove that an integer n is even if and only if n^2 is even.

BEFORE

Suppose n is even. Then $n = 2k$ for some integer k, so

$$n^2 = (2k)^2 = 4k^2 = 2(2k^2).$$

Therefore n^2 is even, and hence n is even if and only if n^2 is even.

Diagnosis. The submitted direction is correct. The word "hence" is not. An if-and-only-if statement contains two debts:

$$n \text{ even} \Rightarrow n^2 \text{ even} \quad \text{and} \quad n^2 \text{ even} \Rightarrow n \text{ even.}$$

The draft pays the first and quietly declares both paid. This is a common failure because the two directions feel nearly identical in English. They are not interchangeable in a proof.

The missing direction also has a trap. Starting from $n^2 = 2m$ and trying to take a square root gives $n = \sqrt{2m}$, which does not exhibit n as twice an integer. The contrapositive is cleaner: show that if n is odd, then n^2 is odd.

AFTER

Proof. First suppose n is even. Then $n = 2k$ for some integer k, and

$$n^2 = 4k^2 = 2(2k^2),$$

so n^2 is even.
For the converse, we prove the contrapositive. Suppose n is odd, so $n = 2k + 1$ for some integer k. Then

$$n^2 = (2k+1)^2 = 4k^2 + 4k + 1 = 2(2k^2 + 2k) + 1,$$

which is odd. Hence if n^2 is even, n cannot be odd and must be even. Both directions hold, so n is even if and only if n^2 is even. □

What changed. The connective was treated as an instruction, not decoration. Each direction received its own paragraph. The harder direction announced its contrapositive before using it, and the final sentence closed the biconditional only after both halves existed.

7.7 Makeover 7: the solution lost by division

The problem. Find all real numbers x such that

$$x^2 = x.$$

BEFORE

Divide both sides by x:

$$x = 1.$$

Therefore the only solution is $x = 1$.

Diagnosis. Dividing by x assumes $x \neq 0$. But $x = 0$ is one of the solutions, so the first instruction deletes an answer before the work begins. This error is especially expensive because the algebra after it is flawless. A neat page can still be wrong at its first move.

There are two honest repairs. Split into the cases $x = 0$ and $x \neq 0$, then divide only in the second. Or factor and use the zero-product property. The second is shorter here because it keeps both possible zero factors visible at once.

AFTER

$$x^2 = x \quad \Longleftrightarrow \quad x^2 - x = 0 \quad \Longleftrightarrow \quad x(x-1) = 0.$$

By the zero-product property, either $x = 0$ or $x - 1 = 0$. Hence the solutions are $x = 0$ and $x = 1$. Substitution verifies that both satisfy the original equation.

What changed. The rewrite used equivalences, so it did not lose or invent solutions. It named the rule that produces the two cases, and it checked the answers in the original equation. That last check takes one line and catches exactly the kind of damage an illegal division causes.

7.8 Makeover 8: cases that were not cases

The problem. Prove that $n(n + 1)$ is even for every integer n.

BEFORE

There are two cases. If n is even, then the product is even. If $n + 1$ is odd, then the product is even because n is even. Therefore $n(n + 1)$ is always even.

Diagnosis. The two claimed cases describe the same situation. If n is even, then $n + 1$ is odd. An odd value of n, such as $n = 3$, belongs to neither case as written. The proof happens to state a true theorem, but its case split is not exhaustive.

A case proof needs a partition: every allowed input belongs to at least one case, and overlapping cases should be avoided unless the overlap is deliberate. For an integer, “even or odd” is the partition waiting to be used.

AFTER

Proof. Let n be an integer. If n is even, then $n = 2k$ for some integer k, and

$$n(n + 1) = 2k(n + 1),$$

which is even.
If n is odd, then $n + 1$ is even, so $n + 1 = 2k$ for some integer k. Then

$$n(n + 1) = 2kn,$$

which is even. Every integer is either even or odd, so the two cases cover all integers. Therefore $n(n + 1)$ is even for every integer n. □

What changed. The cases now cover the domain because the proof says why they do. Each case exhibits the product as twice an integer, which is the definition of evenness. The last sentence closes both the case split and the theorem instead of assuming the reader will infer either.

7.9 The pattern across all eight

Look back at the diagnoses. Different courses, same three failures in rotation: assuming what’s owed (makeovers 1 and 2), claiming more than what follows (3 and 4), leaving quantifiers unscaffolded (2 and 5), paying only half a logical debt (6), and destroying the domain through an illegal move or a broken case split (7 and 8). Now look at the repairs. Not one required a new mathematical idea. Every fix was structural: reverse the direction, weaken the claim to the true one, give each variable its birthplace, protect exceptional values, and make the cases cover what the theorem promised.

> *Almost every broken proof already contains its own repair. The mathematics you found is the hard part, and you did it. The rewrite is bookkeeping, and bookkeeping is learnable.*

That's the closing argument of Part II. Part III takes these skills to the place where most of your grade actually lives: the weekly problem set and the exam hall, where time is short, partial credit is real, and how you present what you *couldn't* finish matters almost as much as what you could.

Part III

Problem Sets and Exams

8

Solutions That Earn Full Marks

A proof and a problem-set solution are cousins, not twins. A solution often contains computation, modeling, and a final answer alongside any reasoning, and it's graded by a person moving fast, against a rubric, at eleven at night. This chapter is about writing for that person. Everything from Part II still applies. What's added is packaging: the structure that lets a tired grader find your marks in thirty seconds, because if they can't find the marks, you don't get the marks.

8.1 The anatomy of a full-credit solution

Strong solutions to computational problems share a five-part skeleton, and the parts take one sentence each:

1. **Setup.** Name the variables and what they measure, including units. "Let w be the width of the field in meters."
2. **Plan.** One sentence saying what you're about to do. "We express the area as a function of w alone, then maximize it."
3. **Execution.** The mathematics, written with Part II's grammar: equations in sentences, steps justified, key transitions worded.
4. **Answer.** Stated in a full sentence, in the problem's own terms, with units, visibly marked (box it or set it on its own line).
5. **Check.** One line verifying the answer against the problem: plug it back, test an endpoint, sanity-check the size.

Watch the skeleton carry a classic:

> **Problem.** A farmer has 100 meters of fencing and wants to enclose the largest possible rectangular area. What dimensions should the rectangle have?
>
> **Solution.** Let w and ℓ be the width and length of the rectangle in meters. The fencing constraint is $2w + 2\ell = 100$, so $\ell = 50 - w$, with $0 < w < 50$.
>
> We write the area as a function of w alone and maximize it. The area is
>
> $$A(w) = w(50 - w) = 50w - w^2.$$
>
> Completing the square, $A(w) = 625 - (w - 25)^2$, which is largest exactly when $w = 25$. Then $\ell = 50 - 25 = 25$.
>
> Therefore the maximum-area rectangle is a 25×25 meter square, with area 625 square meters.

Check: the perimeter is $4 \times 25 = 100$ meters as required, and nearby dimensions do worse: a 20×30 rectangle uses the same fence but encloses only 600 square meters.

Every part is visible: variables born with units, a one-line plan, justified algebra, the answer as a sentence in the farmer's terms, and a check that costs two lines. Total overhead beyond the bare algebra: maybe four sentences. Those four sentences are what full-credit work looks like from the grading side.

Two of the five parts deserve their own attention.

The plan sentence is the highest-value sentence in the solution. It converts a wall of algebra into a story the grader can follow, and if you botch the arithmetic downstream, it's the plan sentence that proves the method was right, which is where most of the partial credit lives.

The check catches disasters at the cheapest possible moment. Negative lengths, probabilities above 1, a "maximum" worse than an endpoint: thirty seconds of checking converts these from lost problems into fixed ones. Write the check down. A visible check also signals a habit graders respect, and when your answer is wrong despite a right method, the written check often documents that you knew something was off, which itself is worth sympathy and sometimes marks.

8.2 How graders actually read

It helps to know the mechanics on the other side of the desk. I've graded piles of scripts, and the routine is close to universal. The grader has a rubric: this many points for the setup, this many for the main step, this many for the conclusion. They do not read your script like a novel. They *scan* for the rubric items, in order, and they've already read forty attempts at the same problem, so they know exactly what right looks like.

Consequences, each one actionable:

- **Rubric items you performed silently score as absent.** If the rubric pays two points for checking the base case and your induction "obviously" starts at $n = 1$ without showing it, that's two points gone. Do the steps visibly.
- **Structure is scanning speed.** Labeled cases, displayed key equations, paragraph breaks between phases, a boxed answer: every one of these shortens the grader's search, and a grader who finds every rubric item in thirty seconds is a grader in a good mood.
- **Neatness is credibility, not decoration.** A script with crossed-out false starts, arrows rerouting the reading order, and answers wedged into margins forces the grader to reconstruct your document before judging your mathematics. Some of that reconstruction will go against you. If the write-up dissolved into chaos, rewrite the problem on a fresh page. Five minutes, repaid.
- **The first line sets the prior.** Graders calibrate fast. A script that opens with clean setup and a plan sentence gets read charitably at the smudged step on line nine. A script that opens with unexplained symbols gets read suspiciously everywhere. Unfair? Slightly. Universal? Completely.

8.3 When you can't finish

Every student meets problems that won't die by the deadline. What you write in that situation separates the ones who understand partial credit from the ones who donate marks.

Never bluff the bridge. The catastrophic move is writing the missing step as if it were done: an "it can be shown" spanning the exact gap you couldn't cross. Graders grade that problem all week; they know precisely where the hard step is, and a bluff at the hard step reads as either dishonesty or blindness to the difficulty. Both are expensive, and the damage leaks into how the rest of your script gets read. The banned list of Chapter 6 exists for style reasons; at a genuine gap, it's not a style issue anymore, it's your credibility.

Label the gap and keep going. The professional format:

> "I can show the result assuming the following claim, which I was not able to prove: *every x in S has property P.* Assuming the claim, the argument proceeds as follows…"

This is not an admission of defeat; it's a map of exactly what you know. It collects the marks for the dissection (you identified the crux), for everything after the gap, and for mathematical maturity. Conditional progress, honestly labeled, is real mathematics; papers get published with that structure.

Show the failed attempt if it's substantial. "I attempted induction; the inductive step fails because the hypothesis for n says nothing about the new element, as follows…" documents understanding of both the technique and the obstacle. Two lines of honest post-mortem routinely outscore a page of confident nonsense.

Answer the sub-questions you can. Multi-part problems often gate later parts on earlier ones. The convention: you may *use* part (a) to solve part (b) even if you couldn't prove part (a). Say "using the result of (a), which I was unable to prove" and take the (b) marks cleanly.

> *Partial credit is not charity for wrong answers. It's full payment for the parts you got right, and it's only payable when the grader can see, labeled and legible, exactly which parts those are.*

8.4 Exams: the same craft at speed

Everything above holds in exams, compressed. The adjustments that matter:

Triage before writing. Read every question first. Grade inflation lives in the easy problems you never reached because problem 2 ate forty minutes. Bank the certain marks first; the hard problem will still be there.

Skeleton first when time is short. If five minutes remain and the proof needs fifteen, write the shape: "By induction on n. Base case: [verified]. Inductive step: the key idea is to apply the hypothesis to the subset missing x…" A correct skeleton with the key idea named collects a surprising fraction of the marks, because it's the rubric's spine.

Don't erase, redirect. Crossing out a wrong start costs one clean line through it. Erasing costs time and often destroys work that was partially creditable. One line, the word "no," move on.

Leave the check for the walk-through. Your last five minutes belong to a units-and-sanity pass across every answer: signs, sizes, domains, the question actually asked. Every exam room in the world contains someone who found the right w and reported it as the area. Don't be the reason that sentence stays true.

One chapter remains in Part III, and it's the one nobody teaches: the system around the writing. When to start, how to move from attempt to write-up, what to do with solved problems so they keep paying rent. Craft is what you do on the page; the system is what makes the page arrive on time.

9

The Homework System

Writing quality collapses under time pressure, and most problem sets are written under maximum time pressure because they were started the night before. So the last chapter of Part III isn't about sentences at all. It's about the system that gets you to the writing desk with hours to spare and solved problems in hand. None of it is complicated. All of it is the difference between the craft of this book being usable and being a nice idea you didn't have time for.

9.1 The two-session rule

The single highest-leverage change: **touch the problem set twice, on different days.**

Session one, the day the sheet appears. Twenty to thirty minutes, no obligation to solve anything. Read every problem. Classify each one: routine (I know the shape already), stretch (I see where it lives but not the key idea), or mystery (no idea). Solve any routine ones on the spot while the classification is warm. Then stop.

This first pass does two jobs. It kills the false comfort of "the set doesn't look too bad" (you now *know*), and it hands the mysteries to the most underrated problem-solving engine you own: background processing. A problem read on Tuesday gets solved in the shower on Thursday shockingly often. A problem first read the night before the deadline gets no shower time at all. The mathematical folklore calls this incubation; every working mathematician relies on it, and the only price is reading the sheet early.

Session two, days later. The real work: stretches and mysteries, scratch paper, the backward exploration of Chapter 5. You arrive already knowing the terrain, with subconscious progress banked.

The write-up is a third, separate pass, and this is the rule students resist until they try it: *don't write final drafts during the solving session*. Solving-you and writing-you have different jobs (that was Chapter 5's whole point), and interleaving them means every dead end gets a beautiful write-up and every real solution gets a tired one. Solve everything solvable, then write everything, in problem order, in one focused block. The write-up block is fast because it's transcription plus grammar: fifteen minutes per proof, on material you fully understand.

9.2 The twenty-minute rule and getting help

Struggle is where the learning is, but struggle has diminishing returns. A workable protocol:

- **Before seeking any help, give the problem twenty minutes of genuine effort**: statement dissected (Chapter 3), definitions unpacked, one shape attempted, one example computed. Not twenty minutes of staring.

- **Then escalate in order of independence**: reread the relevant section, hunt for a similar worked example, ask a classmate for a *hint* (not a solution), go to office hours with your failed attempt in hand. “Here’s where my induction breaks” gets you a targeted nudge; “I don’t get problem 3” gets you a re-lecture you’ll forget.
- **After any hint, walk it home alone.** The hint plus your own completion is learning. The hint plus transcription is a loan against exam day, when no hints attend.

9.3 Collaboration and the honesty line

Working with others is usually allowed and genuinely good for you. The line, at almost every institution, is: **discuss ideas together, write alone.** Whiteboard sessions, shared examples, arguing about whether the contrapositive helps: all fine. What crosses the line is writing from someone else’s write-up, or dictating yours.

The clean protocol: solve together at the whiteboard if you like, then erase the board, go home, and reconstruct the argument yourself from a blank page. If you can’t reconstruct it, you didn’t have it, and it’s better to find out now than in the exam hall.

Credit your collaborators in writing: “I discussed problems 2 and 5 with Priya Nair.” One sentence at the top of the set. It’s required at many places, respected at all of them, and it converts a potential integrity question into a non-event. The same disclosure rule covers any other help: a textbook you consulted beyond the course text, a forum thread, an AI assistant if your course permits one. State what you used and for which problem. The policies vary by course; the honesty convention doesn’t. Undisclosed help is the definition of the problem, and disclosed help almost never is.

9.4 Handwriting or type?

Both are fine nearly everywhere; choose deliberately.

Handwriting is faster for scratch and for heavy notation, and it’s the exam-day medium, which means it needs practice. If your handwritten sets come back with “illegible” anywhere, treat that as a failing grade for the writing system, not a personality trait: slow down 15 percent, double-space, and rewrite final drafts on fresh paper.

Typesetting shines for longer sets, revision (editing without recopying), and anything you’ll want to keep. The standard tool in mathematics is LaTeX, and if you go that road, my *LaTeX for Students* covers it end to end; Appendix D of this book ships ready-made templates either way. Two warnings from years of watching students adopt it. First, don’t learn LaTeX the night a set is due; learn it on a weekend, on last week’s already-solved set. Second, typeset polish can camouflage logical holes from *you* (a beautifully set wrong proof still reads wrong to the grader), so the read-aloud test of Chapter 2 stays mandatory either way.

9.5 The error journal

The highest-return fifteen minutes of your mathematical week happen after the graded set comes back. Most students glance at the score and file the pain away. Instead, keep an **error journal**: one notebook (or file) where every lost mark gets one line in three columns.

> *Where* (set 4, problem 2) *What* (assumed the hypothesis for all n; the makeover-1 disease) *Rule* (write “for some n,” display the hypothesis)

Two things happen. Within weeks, the *What* column starts repeating, because you don't have twenty weaknesses, you have three wearing different problems as disguises. And the *Rule* column becomes a personal pre-submission checklist that outperforms any generic one, including Appendix B's, because it's fitted to your actual failure modes. Before each exam, the journal is your entire review of writing craft: ten minutes, every mark you've ever donated, and the rule that gets each one back.

> *You don't have twenty weaknesses. You have about three, recycled. The error journal finds them, and three named weaknesses can be fixed.*

9.6 The submission checklist

Last pass before anything leaves your hands, ninety seconds per set:

1. Every problem attempted, or its gap honestly labeled (Chapter 8).
2. Every solution ends with a stated answer or a closed proof, visibly.
3. Names, set number, collaborator line, page order, pages stapled or one PDF.
4. Spot-check the two rules atop your error journal, on every problem.
5. The read-aloud test on the one solution you're least sure of.

That closes Part III, and with it the mathematics you write for a grade. Part IV changes the audience: papers, project reports, and theses, where no rubric exists, the reader owes you even less patience, and structure does the work that a course's problem numbering used to do for free. The good news: everything transfers. The better news: almost none of your competition knows that.

Part IV

Papers, Reports, and Theses

10

The Architecture of a Mathematical Paper

At some point, usually in your final year, the assignment changes species. Instead of ten problems due Friday, it's "write up your project" or "submit your thesis draft," and suddenly nobody is numbering your tasks for you. The blank document is now yours to architect, and most students have never been shown the blueprint.

Here it is. Mathematical papers have converged, over a century, on a standard structure, and readers navigate by it the way drivers navigate by road signs. Use the standard structure and your reader always knows where they are. Invent your own and every reader pays a toll on every page.

10.1 Why not IMRaD

If you've seen scientific writing guides, you know IMRaD: Introduction, Methods, Results, Discussion. Lab sciences live by it. Mathematics doesn't, and it's worth understanding why, because the reason *is* the architecture of a math paper.

In an experimental paper, the evidence for a claim lives in the world: the data, the procedure, the instruments. Methods and Results report on them. In mathematics, the evidence for a claim is the proof, a self-contained object made of the same stuff as the claim itself. There's no separate "methods" to report. So the mathematical paper organizes around a different spine: **statements and their proofs**, arranged so that every statement is precise when met and every proof arrives after the reader has what they need to follow it.

10.2 The standard blueprint

The parts, in the order the reader meets them:

Title. A noun phrase naming the actual contribution, specific enough that a stranger can tell whether to open it. "A bound for cycle lengths in dense graphs" beats "Some results in graph theory" every time it's played.

Abstract. Three to six sentences: what is proved, in what setting, and, if it fits, why one should care. Self-contained: no citations, no undefined notation, because abstracts travel alone in databases and search results. Written last (Chapter 11).

Introduction. The context, the question, your answer stated informally, and a map of the paper. Big enough to deserve its own chapter; it gets one (Chapter 11).

Preliminaries (or *Notation and background*). The definitions, conventions, and known results your paper needs. Two disciplines govern the section. Include only what's used: every definition here is a promissory note, and readers notice notes that never come due. And cite known results with precise pointers ("[4, Theorem 3.1]"), stated in your own notation, so your reader never has to leave your paper mid-proof to translate someone else's.

Main results. Your theorems, stated in full, with their proofs. Structure options below.

Examples and applications. Worked instances, special cases, comparisons with prior bounds. In student work this section is chronically undersized, which is backwards: examples are where a reader (and an examiner) checks that they, and you, actually understand the theorems. A paper whose main theorem never meets a concrete instance leaves doubt about whether it ever could.

References. Chapter 13's territory.

The proportions surprise students: in a strong short paper, the introduction plus preliminaries often run a third of the length. That's not padding. That's the part that makes the rest readable.

10.3 Ordering the mathematics: statements first

Within the main results, one decision shapes everything: in what order do statements and proofs appear? The professional default is **statement first, proof after, always**, and never a proof before its theorem. Never make the reader wander through an argument wondering what it's for. The derivation-first style ("let's compute and see what we find, and now we've proved…") mirrors how you discovered it; Chapter 5 already made the case that discovery order is for scratch paper.

Beyond that default, two architectural patterns cover nearly all student papers:

Top-down: main theorem first, supported by lemmas after. State the headline result at the top of the section; prove it assuming two or three clearly stated lemmas; then prove the lemmas. The reader gets the destination early and can see why each lemma earns its keep. This is the right shape when the lemmas are technical and the theorem is the story.

Bottom-up: lemmas first, assembled into the theorem. Build the small results in order, each proved as met, and let the main theorem's proof become a short assembly at the end. Right shape when the lemmas are individually meaningful and the reader can enjoy the staircase.

Both work. What doesn't work is drift: three pages of unmotivated lemmas with no announced destination (bottom-up gone wrong), or a main proof interrupted every paragraph to establish side facts inline (top-down gone wrong; the lemma-extraction rule of Chapter 5 applies with double force in papers). If you choose bottom-up, spend one sentence at the section's start announcing the destination, so the staircase visibly climbs toward something. "The goal of this section is Theorem 3.4, which we assemble from three lemmas" costs eleven words and buys three pages of reader patience.

10.4 Numbering, naming, and the rhythm of the page

Number everything referable, in one series. Theorems, lemmas, propositions, definitions, remarks: give them one shared counter per section (Definition 2.1, Theorem 2.2, Lemma 2.3), the scheme this book uses per chapter. A single series means a reader hunting Lemma 2.3 flips pages monotonically. Equations get their own numbers, and only the equations that are actually cited: a page where every displayed line is numbered reads like a spreadsheet.

Choose the right honorific. The conventions carry information. A *theorem* is a main result, the thing the paper exists for; use it sparingly or it stops meaning anything. A *proposition* is a result of standalone interest below headline grade. A *lemma* is a stepping stone whose purpose is serving another proof. A *corollary* follows from a nearby result with little or no extra work. A *remark* is prose that wants a number: a caveat, a sharper version you won't prove, a pointer. Misgrading these confuses readers who steer by them, and examiners notice a thesis whose every observation is a Theorem.

Let statements breathe. The visual rhythm of a good math paper alternates: prose paragraph (motivation, transition), then a set-off statement, then its proof, then prose again. When a proof runs long, signpost its phases inside ("*Step 1: reduction to the bounded case.*"), exactly the paragraph discipline of Chapter 6 scaled up. When a proof would derail the section's flow entirely, state the result, gesture at the idea in one sentence, and move the full proof to a final section or appendix, with a pointer both ways.

10.5 The thesis variant

An undergraduate or master's thesis is a paper wearing a coat: same skeleton, more fabric. The differences that matter:

- **The background chapter is a real chapter.** In a paper, preliminaries compress against a professional reader. A thesis is examined partly on whether *you* command the background, and its stated audience is usually "a fellow student": so definitions come with motivation, examples, and your own prose, not just citations. This is the one place where writing more than the minimum is the professional move.
- **Chapters need their own architecture.** Each chapter opens with a paragraph saying what it does and how it sits in the whole; each closes by handing off to the next. The chapter-opening map paragraph, cheap to write, is the single most examiner-pleasing habit I know.
- **Honest attribution is structural.** A thesis mixes surveyed material with your contribution, and the document must make the boundary unmissable, chapter by chapter: "Sections 3.1 and 3.2 follow [7]; the results of Section 3.3 are original." Vagueness here isn't a style problem. It's the difference between a contribution and a copy, and Chapter 13 returns to it.

> *Architecture is a promise-keeping machine. The title promises a topic, the abstract promises results, the introduction promises a route, and every section either keeps a promise or breaks one. Readers forgive hard mathematics. They don't forgive broken promises.*

The blueprint has one room we walked past quickly, and it's the room where papers are accepted, theses pass, and readers decide in ninety seconds whether you can write: the introduction. It gets the next chapter to itself.

11

Introductions, Abstracts, and Motivation

Here's how your paper or thesis actually gets read. The abstract gets read by everyone who finds the document. The introduction gets read by everyone who opens it. The proofs get read by three people, one of whom is your examiner and one of whom is you. Yet students spend weeks polishing proofs and forty minutes, at the end, on the introduction. This chapter inverts that budget.

11.1 What an introduction owes the reader

A mathematical introduction has four jobs, in order. Think of them as four questions the reader is silently asking, in the order they ask them.

1. What is this about, and why should anyone care? Context first. Not the history of mathematics since Euclid: the specific line of questions your work continues, in a few sentences a fellow student could follow. The oldest teaching trick in mathematics works here: open with the concrete instance before the general machinery. An introduction that starts with a small, sharp example of the phenomenon ("consider what happens to $x^2 + x + 41$ as x runs through $0, 1, 2, \ldots$") earns the abstraction it introduces next paragraph.

2. What exactly is the question? One paragraph that ends with the problem stated plainly. If the precise statement needs machinery the reader doesn't have yet, state an honest informal version now and promise the precise one: "roughly, how fast can such sequences grow? (Theorem 1.2 makes this precise.)"

3. What's the answer? Your main result, stated as early as you can manage it, informally if necessary, with a pointer to the formal statement. This is the proof-first instinct applied to a whole document: the reader gets the verdict up front, then decides how much of the evidence to inspect. Students bury the result on page nine out of modesty or suspense. There is no suspense in mathematical writing; there is only navigation. Nobody reads a theorem faster for having been kept from it. And when your result is partial (a special case, a bound that's probably not sharp), say exactly that, up front. Framing what you have honestly beats implying what you don't; examiners, like graders, know precisely where the hard part was.

4. How does the paper get there? The roadmap paragraph: "Section 2 fixes notation. Section 3 proves the key lemma, the heart of the argument being [one-phrase idea]. Section 4 assembles the main theorem, and Section 5 shows the bound is attained." One sentence per section, and if you can smuggle in the key idea of the proof, do: a reader who knows the trick in advance follows everything else at double speed.

Do those four jobs in that order and your introduction is, by construction, better than most published ones.

11.2 Related work, without the book report

Somewhere in the introduction (or a short section right after), you place your work among its neighbors. The student failure mode is the book report: a paragraph per paper, summarizing each, connected by "also." The fix is to organize by *idea*, not by paper:

> The problem was raised in [3], which settled the case $d = 2$. For general d, the best known bound was $O(n^2)$, obtained independently in [1] and [6] by different methods. Our approach extends the counting argument of [6], replacing the fixed partition with an adaptive one.

Three citations, one narrative, and the reader learns where you stand relative to each: what you use, what you improve, what you merely coexist with. That last sentence, the one naming exactly what your contribution changes, is the sentence examiners hunt for. Write it deliberately.

Two ethics notes that double as style notes. Cite generously at the boundary of your contribution: claiming novelty for a known idea, even innocently, is the most expensive mistake a young writer can make, and one honest "as in [6]" immunizes you. And describe prior work respectfully: today's clumsy-looking bound was yesterday's breakthrough, and your examiner may have proved it.

11.3 The abstract: written last, read first

The abstract is not the introduction's opening paragraph. It's a free-standing summary for a reader deciding whether to fetch the paper at all, and it obeys harsher constraints: no citations, no unexplained notation, no roadmap ("Section 3 then shows…" means nothing to someone who hasn't opened the document).

A four-sentence skeleton that fits most student work:

1. The setting, in standard terms. ("We study colorings of finite graphs in which adjacent vertices receive distinct colors.")
2. The question or gap. ("It is known that $k + 1$ colors always suffice; how few can be guaranteed when triangles are forbidden?")
3. Your result, as precisely as the no-notation rule allows. ("We prove every triangle-free graph of maximum degree k admits a proper coloring with $O(k/\log k)$ colors.")
4. Method or scope, one sentence. ("The proof is probabilistic and yields a polynomial-time algorithm.")

Write it after everything else is frozen, because it summarizes the final document, not the plan you had in March. Then test it the way it will be used: read alone, on a screen, by someone who will see nothing else. If a sentence doesn't survive that isolation, it doesn't belong.

11.4 Motivation that respects the reader

A word about the word "motivation," because students hear it as "add a paragraph pretending this is useful." Real motivation is none of that. It's answering, honestly, the question *what becomes possible or clear if this is settled?* Acceptable honest answers include: it completes a pattern (cases $d \leq 3$ were known); it removes a hypothesis someone needed; it connects two things not previously connected; it's the simplest unknown case of something famous. "It is intrinsically interesting" is also honest, but then the introduction's job is to *transmit* that interest: show the surprising small case, the pattern that begs explanation. Interest asserted is interest lost; interest demonstrated is motivation.

> *The introduction is not the trailer for your paper. It is the paper, at one-tenth scale: question, answer, and route, all visible. The sections after it are the same document with the resolution turned up.*

With architecture and introduction in hand, what remains of Part IV is the finishing trades: the notation you choose and the tables that tame it, the figures and examples that carry understanding, the citations that pay your debts, and the revision passes that turn a draft into a document. Notation and figures first.

12

Notation, Figures, and Examples

Three supporting players decide how a long document feels to read: the notation you choose, the figures you draw, and the examples you work. None of them is "the mathematics," and all of them are why one thesis reads like a guided tour while another, with identical theorems, reads like a customs inspection. This chapter is the craft of all three.

12.1 Choosing notation

Chapter 2 gave the hygiene rules: one symbol one meaning, define before use, respect convention. Long documents add a design problem on top: you're choosing a small language the reader must hold in memory for fifty pages. Principles that pay:

Make it mnemonic. T for a tree, C for a cycle, $d(v)$ for degree, $\mathcal{F}$ for a family of sets. Every mnemonic letter is one the reader never has to look up. Arbitrary letters (z for a tree, Q for degree) each cost a little attention forever.

Match structure to typography. Related objects should look related: vertices u, v, w; their labels $\ell(u), \ell(v), \ell(w)$; sequences indexed as $a_1, \ldots, a_n$. And different kinds of objects should look different: sets in capitals, elements in lower case, families in script. The reader's eye learns the system in a page and reads type from shape ever after.

Don't decorate needlessly. Every subscript, prime, hat, and tilde must earn its place. If there's only one graph in the whole paper, it's G, not G_0. Saving decorations also keeps them meaningful when you genuinely need $\hat{G}$ later.

Follow your field's dialect. Your thesis lives in some literature; use its standard symbols even when you privately prefer others. The reader who comes from that literature (your examiner came from that literature) reads standard notation for free.

And when the document is heavy, give the reader a notation table. A half-page table in the preliminaries, symbol beside meaning beside the page or section where it's defined, converts fifty look-backs into fifty glances. For any thesis, I'd call it mandatory kindness.

12.2 Figures: illustration, not testimony

Makeover 3 in Chapter 7 convicted a Venn diagram of impersonating a proof. That verdict was about *testimony*: a figure can't certify a general claim. But banish figures entirely and you've thrown away the fastest channel into a reader's head. The resolution is a clean division of labor: **the prose and symbols carry the logic; the figure carries the intuition.** Both, together, in the same document, each doing its own job.

In practice:

- **Draw the generic case, and flag its hidden assumptions.** A figure is one instance pretending to be all instances. Make it as unspecial as you can (no accidental symmetry, no conveniently equal lengths), and when the picture can't help assuming something (the triangle looks acute; the sets intersect), either add the caption note "the argument does not use this" or handle the other configuration in the text. Figures mislead precisely at their unexamined assumptions.
- **Caption for the skimmer.** Readers meet figures out of order, while flipping. A caption should say what the figure shows and what to notice: "Figure 2: the adaptive partition after two rounds; note the isolated block on the right, which drives the bound in Lemma 3.2." Never just "Figure 2."
- **Reference every figure from the text.** "As Figure 2 illustrates…" If no sentence needs to say that, the figure is decoration; cut it. (The same test this book applies to displayed equations, Chapter 6, and for the same reason: emphasis channels clog when overused.)
- **In proofs, figures accompany, never replace.** The honest formula appears in Makeover 3's rebuild: element-chasing prose that stands alone, plus, if you like, the diagram beside it labeled *illustration*. An examiner who distrusts your figure should be able to cover it with a hand and lose nothing but speed.

12.3 Examples as load-bearing structure

Student drafts treat examples as garnish: something sprinkled on after the theorems, if pages remain. Professionals build with them. Three structural uses, in ascending order of importance:

The running example. Choose one small, rich instance early (one graph, one function, one equation) and revisit it as each new definition and theorem lands: here's the definition, here's what it says for our example; here's the theorem, here's the bound it gives for our example. The running example gives the reader a home base, and it gives *you* a continuous correctness check. Half the subtle errors I've caught in my own drafts surfaced the moment a new theorem met the old example and produced nonsense.

The boundary examples. For every definition you introduce, the reader silently wonders: what qualifies, what barely qualifies, what barely fails? Answer with a matched pair: one object satisfying the definition non-trivially, one just missing it. The pair does more than any paraphrase to fix the concept, and choosing the pair forces you to understand your own definition's edges. The same service, one level up: after a theorem, an example showing the hypothesis can't simply be dropped (the conclusion fails without it). That example certifies your theorem is telling the truth about its own terms.

The sharpness example. The professional finishing move on any inequality or bound: an instance showing the bound can't be improved, or a remark honestly conceding you don't know. "The constant 4 is attained by the following family, so Theorem 2 is sharp" upgrades a result from *true* to *finished*, and examiners read the sharpness discussion as a direct measurement of mathematical maturity. If you can't settle sharpness, the honest sentence ("we do not know whether 4 can be replaced by 2") is itself professional: it hands the next student a problem and shows you saw the question.

> *Definitions tell the reader what's true. Examples tell the reader what it means. A document with theorems and no examples has announced its results and communicated nothing.*

12.4 Tables, and when prose beats all three

A last calibration. Tables earn their keep exactly where prose wastes it: parallel data varying in one or two dimensions (the bound for each small case; the comparison of three algorithms' costs). If you find yourself writing the same sentence four times with different numbers, that's a table. Conversely, a table with one row, or a figure of a straight line, or a worked example of a definition that was already transparent: those are apparatus without payload, and each one dilutes the reader's trust that your apparatus signals substance. The test is uniform across this whole chapter: every figure, table, and example exists to make some specific sentence's job easier. Name the sentence, or cut the prop.

Notation, figures, examples: the reader-facing surfaces of your document. Two chores remain before a draft can be called a paper. The first is settling your debts, saying precisely what you took from where. That's citation, and it's shorter and higher-stakes than students think.

13

References and Citation

Citation in mathematics is a precision instrument, not a bibliography-padding ritual. A well-cited paper tells the reader, at every moment, which bricks are yours and exactly where the borrowed ones came from, down to the theorem number. This chapter covers what to cite, how to point, the mechanics of a reference list, and where the plagiarism line actually runs in a subject where everyone proves the same classics.

13.1 What needs a citation

The working test: **cite whatever you did not prove and is not common knowledge at your document's level.** Both halves need calibrating.

Results you use but don't prove: cited, with precision (next section). This includes the contents of your background chapter, even when you re-prove them in your own words: the reader must know the result isn't yours.

Common knowledge: not cited. The quadratic formula, the triangle inequality, the definition of a group: citing these reads as insecurity. The boundary moves with level; in a first proofs course, the well-ordering principle might deserve a pointer, while in a graduate thesis it's furniture. When genuinely unsure, a pointer to a standard textbook costs one bracket and never harms.

Two cases students miss:

- **Definitions that vary.** "Graph" (loops allowed?), "ring" (unit required?), "natural numbers" (is 0 in?). When your document's claims depend on the choice, cite the convention you follow: "we follow [2] in requiring rings to have a unit." This is cheap insurance against a reader testing your theorem with the other convention.
- **Provenance of the problem.** If the question came from a course project list, a supervisor, or an open-problems survey, say so in the introduction. Questions have authors too.

13.2 Pointing precisely

The unit of mathematical citation is not the paper. It's the *statement*. The professional format names both:

> By a theorem of Dirac [4, Theorem 3], every graph with minimum degree at least $n/2$ contains a Hamiltonian cycle.

The bracket-pair format "[4, Theorem 3]" (or "[4, p. 72]" when numbering fails you) is the difference between a pointer and a homework assignment: bare "[4]" aimed at a

300-page book asks your examiner to go find your evidence themselves. They notice being asked.

Three habits complete the craft:

- **State borrowed results in your own notation**, in full, before using them (Chapter 10 said this; citation is where it's enforced). The citation then certifies your restatement, so double-check that your restatement says what the source says: no strengthening in translation, and hypotheses carried over whole. Misquoting a theorem, even downward, is an error in *your* paper.
- **Attribute in the statement, not just the bracket**, when the result has a name: "(Dirac, 1952)." Named results carry their history as useful metadata, and using the names correctly marks you as someone who's read the literature's map.
- **Cite what you actually read.** If you learned Dirac's theorem from a textbook, cite the textbook ("see, e.g., [2, Theorem 10.1.1]"), or cite both original and textbook. Citing a 1952 paper you've never opened propagates whatever errors live in the chain of paraphrases you actually used, under a false badge of verification.

13.3 The reference list, mechanically

The reference list itself has one virtue, consistency, and one tool worth adopting: let your typesetting system build the list (in LaTeX, BibTeX or biblatex; *LaTeX for Students* walks through it, and Appendix D's paper template arrives wired up). You maintain a small database of entries; the system formats, alphabetizes, and, crucially, includes only what you actually cited, so the list can't drift out of sync with the text as drafts evolve.

Whatever tool, the quality bar per entry: every author with initials, full title, journal or publisher, year, and pages or article number; for anything online-only, a DOI or stable URL. Harvest entries from the source itself, not from auto-generated citation exports, which are wrong often enough to embarrass you. Before submission, click or look up every entry once: a reference list with a broken pointer or a mangled name is the first page an examiner finds errors on, and it primes them to hunt for more.

13.4 Where the plagiarism line runs

Mathematics has a genuine subtlety here, so it deserves plain words. Every generation proves the infinitude of primes; ideas circulate freely; your background chapter *is* other people's mathematics. None of that is plagiarism. The line runs at **presenting someone else's exposition or contribution as yours**, and it's crossed in three recognizable ways:

Copied prose. Taking sentences or their lightly-edited skeletons from a source into your write-up without quotation or attribution. In mathematics the temptation is strongest in background sections, where you feel there's "only one way to say it." There isn't. Close the book, write the section from your own understanding, reopen the book to check correctness, and cite it as the source of the material. If a formulation is so good you want it verbatim, quote it, marked as a quotation, with a pointer. (Rare in math prose, but legitimate.)

Copied proofs presented as original. Reproducing a proof from a paper, a textbook, a forum answer, or an AI assistant, in a context that implies the argument is your contribution. The repair is one honest clause: "we follow the proof of [6]" or "this argument was suggested by an answer on Math StackExchange [9]." With the clause, it's scholarship. Without it, it's the real thing, and modern examiners have the same search engines you do.

Unattributed help. The thesis-scale version of Chapter 9's collaboration line. Supervisors expect to be thanked, not cited, for guidance; but a key idea contributed by another student, a forum, or a tool belongs in the acknowledgments or in the text, matching your institution's rules. The instinct that something "doesn't count if I don't mention it" is exactly backwards: mentioned, almost everything is fine; unmentioned, almost nothing is.

> *Attribution is never the confession that weakens your work. It's the frame that makes your own contribution visible: every honest "following [6]" sharpens the edge around the part that has no citation, because that part is yours.*

Your document now has architecture, an introduction, notation, figures, examples, and settled debts. It's a complete draft, which is to say: it's ready to be taken apart. The last chapter of Part IV is about revision, the pass where drafts become papers, and it starts with the hardest instruction in this book: put it down and walk away.

14

Revision: Where Writing Actually Happens

Nobody writes well. Some people revise well, and from the outside the two are indistinguishable. Everything this book has taught, grammar, shapes, style, architecture, gets its second and real chance during revision, because the first draft was written by someone too close to the mathematics to see the page. Revision is the process of becoming a stranger to your own work, then serving that stranger.

This chapter is a system: four passes, each with one job, plus the craft of using other people's eyes. It scales down to a homework proof (where the passes take minutes) and up to a thesis (where they take weeks).

14.1 First, become a stranger

The prerequisite for all of it: **cooling time**. Reread a proof the minute you finish it and you'll read what you meant. Reread it tomorrow and you'll read what you wrote. The gap between those two documents is where every unjustified step and every undefined symbol hides.

Practical dosage: overnight for a problem set (the two-session system of Chapter 9 builds this in), two or three days for a paper section, a full week for a thesis chapter before its final revision. Deadlines compress the dose, but even twenty minutes and a walk beats zero. The draft is not losing value while it cools. It's developing.

14.2 Pass 1: the logic audit

Coldest read, biggest question: *is it true, and is it proved?* You're not polishing sentences yet; polished wrongness is worse than rough wrongness because it survives longer. In this pass:

- **Check every "therefore."** For each consequence claimed, name to yourself the exact justification: a hypothesis, an earlier line, a cited result. The connective vocabulary of Chapter 6 makes these checkpoints findable; makeover logic (Chapter 7) tells you what failure looks like at each.
- **Audit every variable's birth** (introduced before use? right quantifier? right dependence order?), the Chapter 2 discipline, now applied wholesale.
- **Interrogate boundary behavior.** Empty set, $n = 0$ or 1, equality cases in inequalities, division by things that might be zero. First drafts are written for the generic case; graders and examiners read for the edges.

- **Re-dissect the statement** (Chapter 3) and confirm the proof proves *it*: all of it, not a special case, not its converse, and using every hypothesis. An unused hypothesis is either a sign your theorem is stronger than stated, which is good news worth having, or a sign you've mis-proved it, which is better news now than later.

14.3 Pass 2: the stranger's read

Now read for the reader: start to finish, at reading speed, ideally **aloud**. This is the full-document version of Chapter 2's sentence test, and it catches what line-by-line auditing can't: the lemma used before it's believable, the notation defined too far from its heavy use, the paragraph where three ideas fight for one topic sentence, the "rhythm" failures where six displayed equations arrive with no prose oxygen between them.

Mark problems; don't fix them mid-read (fixing re-immerses you, and you lose the stranger's eyes you paid cooling time for). One symbol in the margin per problem type: a question mark for "didn't follow," a wiggle for "reads badly," a circle for "define/cite this." Then fix in batches afterward, worst first.

14.4 Pass 3: the compression pass

First drafts overexplain in the wrong places, always. This pass has one instruction: **every sentence defends its existence, or dies.** The usual suspects, in order of yield:

- Scaffolding sentences that narrate the writing rather than the mathematics ("we now proceed to consider the next case"; the case label already says so).
- Restatements wearing "in other words" that add no other words worth having.
- Detail spent on routine steps while the key idea got one line (rebalance: compress the routine, expand the crux, the Chapter 6 calibration applied retroactively).
- Hedges ("it seems that," "basically") that survive from the draft's uncertain hours. You've since done the logic audit: now commit, or state the genuine limitation plainly.

Cutting is where drafts most visibly become papers. A tight page signals an author in control, and examiners read control as mastery. If a cut hurts, save the fragment in a scraps file; it's amazing how rarely you ever go back for one, and knowing they're saved makes the knife easier to use.

14.5 Pass 4: the mechanical sweep

Last, the checklist pass, done with the document treated as an artifact rather than an argument: numbering all references resolve (no "Theorem ?? "), notation table matches usage, figures referenced and captioned (Chapter 12), citations pointing at statements (Chapter 13), spelling, punctuation inside displays (Chapter 2), consistent fonts for consistent objects. Appendix B packages this as a run-through list. It's the least intellectual pass and the one readers judge fastest, because mechanical errors are the only kind every reader can see.

14.6 Borrowing eyes

After your own four passes, one stranger's single read is worth three more of yours. Craft points for using readers well:

Ask for the right failure. Hand a peer your draft with one question: "mark the first place you stop following." That's the data you can't generate yourself. General impressions ("looks good!") are worthless socially pleasant noise.

Different readers, different instruments. A classmate in your course finds the steps that were only clear in your head. A student one course behind (this book's calibration reader from Chapter 6) finds what's genuinely unclear versus merely compressed. Your supervisor finds the mathematical issues, so don't spend supervisor reads on typo-hunting; that's what Pass 4 was for.

Receive like a professional. Every confusion report is correct by definition (the reader is the measurement), even when their proposed fix is wrong. "The proof is fine, you just misread" may be true and is still a defect report: something on that page taught the misreading. Fix the page, not the reader. Getting good at receiving critique now pays compound interest later; the referee reports and review cycles of professional life are this exact game at higher stakes.

> *The first draft is you explaining to yourself. The final draft is you explaining to someone else. Revision is the scheduled, teachable process that converts one into the other, and no mathematical talent substitutes for it.*

14.7 Declaring it done

Perfectionism has a failure mode too: the thesis polished past its deadline, the problem set rewritten a fourth time while tomorrow's set goes unstarted. Done, for a student document, means: the four passes are complete, one outside reader has stopped finding confusions, and the submission checklist (Appendix B) comes up clean. After that, ship it. Writing is a craft, and a craft improves by finishing things, not by orbiting them. Your next document inherits everything this one taught you; the error journal (Chapter 9) is the ledger where the inheritance is recorded.

That completes the tour: from the sentence, through the proof and the problem set, to the paper and the thesis. What remains are the tools you'll reach for while working: the phrase glossary, the checklists, the error catalog, and the templates. They're the appendices, and unlike most appendices, they're meant to get dirty. Keep them within arm's reach; that's where the writing happens.

15

A Small Paper, Built in Public

A chapter-sized explanation can hide behind momentum. A paper cannot. Its title makes a promise, its introduction creates a debt, every definition changes what later sentences are allowed to mean, and its main theorem has to survive being read by someone who did not watch the work happen.

This chapter follows one miniature expository paper from rough notes to a finished version. The mathematics is accessible: counting binary strings that do not contain two consecutive ones. The point is not the result. The point is watching the document become trustworthy.

15.1 The question before the paper

Suppose a binary string uses only 0 and 1. How many strings of length n contain no occurrence of 11?

For small n, the count can be found by hand:

n	allowed strings	count
0	ε	1
1	$0, 1$	2
2	$00, 01, 10$	3
3	$000, 001, 010, 100, 101$	5

Here ε denotes the empty string. The counts $1, 2, 3, 5, \ldots$ suggest Fibonacci numbers. That observation is scratch work, not yet a theorem. Three things remain unpaid:

1. define the indexing convention for the Fibonacci numbers;
2. explain why the recurrence continues for every n, not only the four rows inspected;
3. connect the recurrence and the initial values to the claimed closed identification.

The scratch page already contains the paper's spine. It does not yet contain a paper.

15.2 Finding the argument

Let a_n be the number of valid binary strings of length n. A valid string of length n ends in either 0 or 1.

If it ends in 0, delete that final 0. What remains is any valid string of length $n - 1$. This gives a_{n-1} possibilities.

If it ends in 1, the preceding symbol must be 0. Delete the final 01. What remains is any valid string of length $n - 2$. This gives a_{n-2} possibilities.

The two classes are disjoint and exhaustive, so

$$a_n = a_{n-1} + a_{n-2} \qquad (n \geq 2).$$

Together with $a_0 = 1$ and $a_1 = 2$, this identifies $a_n = F_{n+2}$ under the convention $F_0 = 0$, $F_1 = 1$.

Notice where the proof actually lives. It is not in recognizing the sequence. It is in the reversible classification:

$$\text{valid strings ending in } 0 \longleftrightarrow \text{valid strings of length } n-1,$$
$$\text{valid strings ending in } 01 \longleftrightarrow \text{valid strings of length } n-2.$$

The table guessed. The bijections proved.

15.3 The first architecture

The smallest paper that carries this result honestly needs five pieces:

1. a title that says which objects are counted;
2. a short abstract stating the result and method;
3. an introduction that poses the counting question and explains why the obvious small cases are not the proof;
4. definitions followed by the recurrence theorem and corollary;
5. one example that checks the indexing against a concrete value.

It does not need a literature review, three pages of general material about Fibonacci numbers, or a section called "Methodology." Those add the appearance of a paper while hiding the argument. Structure serves the debt on the page. Nothing more.

15.4 A plausible first draft

BEFORE

Fibonacci Numbers and Binary Strings
Abstract. Fibonacci numbers occur in many surprising places. In this paper we discuss one such occurrence involving binary strings and show some interesting results.
Introduction. Binary strings are important in computer science. The Fibonacci sequence is also famous. Consider binary strings without consecutive ones. The numbers for the first few lengths are 1, 2, 3, 5 and 8, so the answer is the Fibonacci sequence.
Theorem. The number of binary strings of length n without consecutive ones is F_{n+2}.
Proof. A string can end in 0 or 1. If it ends in 0 there are F_{n+1} possibilities before it. If it ends in 1, it has to end in 01, so there are F_n possibilities before it. Adding gives $F_{n+1} + F_n = F_{n+2}$. The first few values also work. Therefore the result is proved.

The draft is not nonsense. It has the right answer and the right combinatorial idea. But it asks the theorem to prove itself.

15.5 The referee-style diagnosis

A useful reader report separates mathematical problems from writing problems. Mixing them produces the familiar margin note "unclear," which helps nobody.

Mathematical problems

1. The indexing of F_n is never defined. Under one common convention the Fibonacci sequence starts $1, 1, 2, \ldots$; under another it starts $0, 1, 1, \ldots$. The subscript $n + 2$ cannot be checked until the convention is fixed.
2. The proof says there are already F_{n+1} and F_n possibilities in the two classes. Those are the claims being proved. The counting recurrence should first be stated in terms of a_n. Only after its initial values are known may it be identified with the Fibonacci recurrence.
3. "A string can end in 0 or 1" does not by itself establish the counts. The draft must show the deletion maps in both directions. Otherwise it has not ruled out loss or duplication.
4. "The first few values also work" is too vague to supply the initial conditions. Name a_0 and a_1, including the empty string.

Writing problems

1. The title is broad. The paper is not about every relationship between Fibonacci numbers and strings.
2. The abstract uses "surprising," "interesting," and "discuss." None tells the reader what was proved.
3. The introduction spends two sentences announcing importance without giving a reason the reader needs.
4. The theorem arrives before the objects and notation have been defined.
5. The proof ends by gesturing at examples instead of closing the recurrence argument.

This is a repairable report. Every item names a place, a debt, and an action.

15.6 The revised miniature paper

AFTER

Counting Binary Strings with No Consecutive Ones

Abstract. Let a_n count binary strings of length n that contain no adjacent ones. Splitting the strings by their final symbol gives $a_n = a_{n-1} + a_{n-2}$. With $a_0 = 1$ and $a_1 = 2$, it follows that $a_n = F_{n+2}$ for the Fibonacci sequence defined by $F_0 = 0, F_1 = 1$.

1. The counting question. A binary string is a finite sequence whose entries are 0 or 1. Call such a string admissible if it contains no occurrence of 11. Let a_n be the number of admissible strings of length n. We count the empty string as the unique string of length 0, so $a_0 = 1$; the strings 0 and 1 give $a_1 = 2$.
Listing small cases suggests the Fibonacci sequence, but the list does not explain why the pattern continues. The recurrence below does.

2. The recurrence.
Proposition. For every integer $n \geq 2$,

$$a_n = a_{n-1} + a_{n-2}.$$

Proof. Partition the admissible strings of length n according to their last symbol. If a string ends in 0, deleting that symbol gives an admissible string of length $n - 1$. Conversely, appending 0 to any admissible string of length $n - 1$ creates an admissible string of length n. Hence this class has a_{n-1} members.
If an admissible string ends in 1, its preceding symbol must be 0. Deleting the final 01 gives an admissible string of length $n - 2$. Conversely, appending 01 to any admissible

string of length $n-2$ creates an admissible string ending in 1. Hence this class has a_{n-2} members.
The classes are disjoint and contain every admissible string of length n, so their sizes add. This proves the recurrence. □

3. Identification with Fibonacci numbers. Define $F_0 = 0, F_1 = 1$, and $F_n = F_{n-1} + F_{n-2}$ for $n \geq 2$.

Corollary. For every integer $n \geq 0$,

$$a_n = F_{n+2}.$$

Proof. The sequences (a_n) and (F_{n+2}) satisfy the same recurrence. Their first two values agree:

$$a_0 = 1 = F_2, \qquad a_1 = 2 = F_3.$$

Induction on n now gives $a_n = F_{n+2}$ for all $n \geq 0$. □

4. A check. For $n = 4$, the admissible strings are

$$0000,\ 0001,\ 0010,\ 0100,\ 0101,\ 1000,\ 1001,\ 1010.$$

There are eight, and $F_6 = 8$, as the formula predicts.

15.7 What the revision actually did

The revised paper is longer, but length is not the improvement. Four structural decisions did the work.

1. It separated discovery from proof. The small table motivates the question; the partition proves the recurrence.
2. It delayed the Fibonacci notation until the counting sequence had been defined and understood on its own terms.
3. It made both deletion arguments reversible. "Delete" alone can hide loss; "delete, and conversely append" establishes the correspondence.
4. It used the final example as a check, not as evidence for the general claim.

> *A paper becomes readable when every section pays one visible debt. The introduction poses the question. Definitions fix the language. The theorem states the debt. The proof pays it. The example checks the machinery without pretending to replace it.*

Take the same sequence to your own draft. Write one sentence beside each section saying what debt it pays. If two sections pay the same debt, combine them. If a section pays none, cut it. If a debt has no section, the outline has shown you exactly what to write next.

Part V

Toolkit

A

Say It in Words

Chapter 2's rule: symbols for objects and relations, words for reasoning, and everything readable aloud. This glossary is the working reference for that rule. For each symbol: how to read it aloud, and when the words beat the symbol in written prose. The standing advice for running text in proofs: if the symbol is doing *logical* work (quantifying, implying), prefer the words; if it's naming an object or a relation, the symbol is usually fine.

Logic and quantifiers

Symbol	Read it as, and when to prefer words
$\forall x$	"for every x" (or "for all x"). In prose, always write the words.
$\exists x$	"there exists an x such that." In prose, always words: "there is an integer k with…"
$P \Rightarrow Q$	"if P, then Q"; "P implies Q." Words in running text; symbol fine when discussing the implication itself.
$P \Leftrightarrow Q$	"P if and only if Q." The abbreviation "iff" is fine in notes, spelled out in formal work.
$\neg P$	"not P"; "P fails." Words, nearly always.
$P \wedge Q$, $P \vee Q$	"P and Q," "P or Q" (or is inclusive). Words in prose; symbols belong in logic courses.
s.t.	"such that." Fine on a whiteboard; write the words in anything submitted.

Sets

Symbol	Read it as
$x \in A$	"x is an element of A"; "x belongs to A."
$A \subseteq B$	"A is a subset of B"; every element of A lies in B (equality allowed).
$A \subsetneq B$	"A is a proper subset of B." If you write plain $\subset$, say which convention you follow.
$A \cup B$, $A \cap B$	"the union / intersection of A and B."
$A \setminus B$	"A minus B"; the elements of A not in B.
$\emptyset$	"the empty set."
$\{x : P(x)\}$	"the set of all x such that $P(x)$ holds."
$\|A\|$	"the number of elements of A" (finite case); "the cardinality of A." Same bars mean absolute value for numbers; context must make it unambiguous, and if it doesn't, say it in words.
$A \times B$	"the Cartesian product of A and B": ordered pairs (a, b).

Numbers, relations, operations

Symbol	Read it as
$\mathbb{N}, \mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$	"the natural numbers / integers / rationals / reals / complex numbers." State whether $0 \in \mathbb{N}$ if it matters.
$a \mid b$	"a divides b." Note the direction; "b/a" is a number, "$a \mid b$" is a statement.
$a \equiv b \pmod{n}$	"a is congruent to b modulo n."
$\lfloor x \rfloor$, $\lceil x \rceil$	"the floor / ceiling of x."
$n!$	"n factorial."
$\binom{n}{k}$	"n choose k."
$\sum_{i=1}^{n} a_i$	"the sum of the a_i for i from 1 to n."
$\prod_{i=1}^{n} a_i$	"the product of the $a_i \ldots$"
$\approx$, $\sim$	"approximately equal"; "asymptotic to." Never inside a proof step that needs equality; say what the error term is or use $\leq$.

Functions and analysis

Symbol	Read it as
$f : A \to B$	"f is a function from A to B."
$x \longmapsto x^2$	"x maps to x^2." The arrow with the bar describes the rule; the plain arrow describes domain and codomain.
$f \circ g$	"f composed with g": first g, then f.
$f^{-1}(B)$	"the preimage of B under f." Exists whether or not f is invertible; don't read it as "f inverse of B" unless it is.
$a_n \to L$	"the sequence a_n converges to L."
$\lim_{x\to a} f(x)$	"the limit of $f(x)$ as x approaches a."
max, min, sup, inf	"maximum, minimum, supremum (least upper bound), infimum." Write max only when it's attained.
$O(n^2)$, $o(n)$	"big-O of n squared" (at most a constant times n^2 eventually); "little-o of n" (negligible compared to n). State the variable that's growing.

Abbreviations that shouldn't leave your scratch paper

WLOG ("without loss of generality"), TFAE ("the following are equivalent"), resp. ("respectively"), w.r.t. ("with respect to"), iff. Each has a legitimate, spelled-out or carefully deployed professional use, and each, abbreviated in a first course, reads as haste. Special mention for WLOG, which is not decoration: it carries a proof obligation. "Without loss of generality $a \leq b$" is legal only when the general case really does follow from the assumed one (here: by swapping names), and your sentence should make that visible. A WLOG that quietly discards cases is makeover material, and Chapter 4's case-coverage rule is the law it breaks.

B

The Proof-Writing Checklist

Run every proof through this before it leaves your hands. It's sequenced to match how the chapters built it: statement, structure, line-by-line logic, style, mechanics. With practice the whole list takes under two minutes per proof, and it's exactly the list I'd run against your script if you handed it to me.

The statement (Chapter 3)

1. I can point at the hypothesis and the conclusion, and the proof uses the first and never assumes the second.
2. Hidden quantifiers are identified: universal claims get an arbitrary element, existential claims get an explicit witness.
3. If the statement is "if and only if," both directions are present and labeled.
4. Every hypothesis gets used somewhere (if one doesn't, I've checked whether the proof is wrong or the statement is understated).

The structure (Chapters 4–5)

5. The shape is announced: "we prove the contrapositive," "suppose for contradiction," "by induction on n."
6. Induction only: base case verified visibly; hypothesis stated for *some* n, written out in full, and actually used in the step.
7. Cases only: the cases cover all possibilities (said why, if unobvious), and each case ends at the same conclusion.
8. The argument runs forward: no line assumes the goal, and any backward exploration was reversed and reverse-checked before drafting.

The lines (Chapters 1–2)

9. Every variable has a birthplace ("let," "choose," "there exists. . . so fix"), in dependence order.
10. Distinct quantities got distinct letters; no symbol works two jobs.
11. Every "therefore" has a visible reason: hypothesis, earlier line, or cited result.
12. Equals signs join only equal things; implication arrows aren't decorating a calculation.
13. Edge cases are covered: empty set, $n = 0$ or 1, equality in inequalities, zero denominators.

The style (Chapter 6)

14. No “clearly,” “obviously,” “it is easy to see” standing where a reason should be.
15. Every “we must show” is later discharged by a “therefore.”
16. Detail is calibrated: key idea in full, routine steps compressed but named.

The mechanics (Chapter 2, Chapter 8)

17. Read aloud, start to finish: every sentence parses, no sentence starts with a symbol, displayed equations carry their punctuation.
18. The proof ends visibly: $\square$, QED, or a closing sentence.
19. (Problem sets) The five-part skeleton where it applies: setup with units, plan sentence, execution, stated answer, written check.
20. (Problem sets) Gaps honestly labeled, collaborators credited, submission formalities done.

Twenty items is the full sweep. For daily use, your error journal (Chapter 9) tells you which three of these you personally fail; check those three on everything, and run the full list before anything is submitted for real marks.

C

Forty Common Errors

The grading pile's greatest hits. Each entry: the error, why it costs you, and the repair. Skim the bold lines now; return when a graded script sends you here. Cross-references point to the chapter that treats the disease at length.

Logic

1. **Assuming the conclusion.** The proof's first line is the goal in disguise. Costs everything; it's circular. Fix: explore backward on scratch, write forward from knowns (Chapter 5).
2. **Proving the converse.** Asked $P \Rightarrow Q$, delivered $Q \Rightarrow P$. Total loss, not partial. Fix: dissect before drafting (Chapter 3).
3. **Unannounced contrapositive.** Valid argument, reads as error 2. Fix: one sentence, "we prove the contrapositive."
4. **Example as universal proof.** One instance offered for a "for all" claim. Fix: arbitrary element, or recognize the claim is existential.
5. **Generality wasted on existence.** Pages of arbitrary argument where one witness settles "there exists." Fix: exhibit and verify.
6. **Wrong negation launching a contradiction.** Negating "all x have P" as "no x has P." The contradiction proves nothing. Fix: quantifier flips, negation moves inward (Chapter 3).
7. **Fake contradiction.** Assumed $\neg Q$, never used it, proved Q directly. Logically fine, structurally dishonest. Fix: rewrite in the shape actually used (Chapter 4).
8. **Deriving truth from the goal and stopping.** "Want X; X gives true statement; done." False implies true. Fix: check every step reverses; write the reversed chain.
9. **Case analysis with a missing case.** Covered $x > 0$ and $x < 0$; the theorem died at $x = 0$. Fix: name the partition and why it's exhaustive.
10. **Unjustified WLOG.** "Without loss of generality" discarding genuinely different cases. Fix: state the symmetry that restores generality, or do the cases (Appendix A).
11. **Using the theorem to prove the theorem.** Citing the result itself, or a result whose proof depends on it. Fix: know your course's dependency order; cite only what's already established.
12. **"It works for $n = 1, 2, 3$, so always."** Pattern observation as proof. Fix: induction (Chapter 4), or a general argument.

Quantifiers and variables

13. **Inductive hypothesis "for all n."** Assumes the theorem, the makeover-1 disease. Fix: "for some $n \geq 1$," written out (Chapter 7).
14. **Missing base case.** The ladder with no ground. Fix: verify it, visibly, even when trivial.
15. **One witness for two hypotheses.** "$x = 2k$ and $y = 2k$": the Chapter 1 classic. Fix: fresh letter per witness.
16. **Variable used before birth.** A δ or N appears, then gets defined. Fix: introduce in dependence order (Chapter 2).
17. **δ depending on x** (or N on n): the chosen thing depending on what comes after it. Fix: respect quantifier order; choose before the arbitrary element arrives.
18. **Fixed constant treated as arbitrary** (or vice versa). "Let $\varepsilon = 0.1$" proving a for-all-ε claim. Fix: "let $\varepsilon > 0$ be given."
19. **Symbol reuse.** n is the set's size and the running index. Fix: one symbol, one meaning, per document.
20. **Silent domain change.** x was real, becomes an integer mid-proof. Fix: state the domain at birth; restate if it genuinely narrows, with justification.

Algebra and computation

21. **Dividing by something that might be zero.** Fix: case it, or note why it's nonzero, in writing.
22. **Multiplying an inequality by something that might be negative.** Direction flips unremarked. Fix: sign check, stated.
23. **Squaring both sides as an equivalence.** $x = 2$ and $x^2 = 4$ aren't the same information. Fix: track direction; squaring is one-way without sign control.
24. **Square roots forgetting $\pm$** (and $\sqrt{x^2} = x$ instead of $|x|$). Fix: absolute values until signs are known.
25. **Equals-chain containing a non-equality.** The Chapter 2 epidemic. Fix: words for reasoning steps.
26. **Term-by-term operations that aren't legal.** Distributing limits, roots, or logs over sums at will. Fix: name the rule being used; if you can't name it, it probably isn't one.
27. **Boundary-blind formulas.** Sum formulas at $n = 0$, $\binom{n}{k}$ with $k > n$, log of something possibly zero. Fix: the edge-case sweep (Chapter 14).
28. **Calculation replacing the question.** Computed the derivative, never answered "where is it increasing?" Fix: the answer sentence, in the problem's terms (Chapter 8).

Style and presentation

29. **"Clearly" at the crux.** The banned list's flagship (Chapter 6). Fix: the reason-phrase, or the honest sentence.
30. **Sentence starting with a symbol.** Fix: three words of runway (Chapter 2).
31. **Undefined nonstandard notation.** Reader meets $d(n)$ cold. Fix: define before use.
32. **Unreadable-aloud symbol soup.** Logic symbols doing prose's job. Fix: symbols for objects, words for reasoning.
33. **Naked displayed equations.** No punctuation, no sentence membership. Fix: displays are clauses (Chapter 2).

34. **“We must show” never discharged.** Goal stated, never closed. Fix: every want matched by a therefore (Chapter 6).
35. **Undeclared proof ending.** The argument just stops. Fix: □ or a closing sentence.
36. **Scratch work submitted.** Backwards exploration on the final page, makeover 2’s disease. Fix: two documents, reversal check, transcription (Chapter 5).

Mechanics and conduct

37. **The bluffed bridge.** “It can be shown” across the exact gap. Costs credibility beyond the problem. Fix: label the gap, argue conditionally (Chapter 8).
38. **Illegible or chaotic script.** The grader reconstructs, resentfully. Fix: fresh-page rewrite; structure as scanning speed.
39. **Uncredited collaboration or sources.** An integrity question where a sentence would have prevented one. Fix: the disclosure line (Chapters 9 and 13).
40. **Imprecise citation.** “[4]” pointing at a book. Fix: theorem-level pointers, “[4, Theorem 3]” (Chapter 13).

Forty errors, and not one requires more mathematical talent to fix. That’s the whole thesis of this book, stated one last time as a list.

D

Templates

Three skeletons you can copy onto a blank page (or into a blank file) so that no document in your student career starts from nothing. Each line in brackets is an instruction to replace, and each template is the structural summary of the chapters it points to. If you typeset your work, my *LaTeX for Students* builds these same skeletons as working files, cover to cover; the structure below is identical either way, on paper or in type.

Template 1: the problem-set solution

For each problem (Chapters 8 and 9):

> **Problem k.** [Restate the problem in one line if the course expects it; skip if the sheet travels with your script.]
>
> *Setup.* Let [variables, with types and units]. [Known facts translated into notation.]
>
> *Plan.* We [method, one sentence].
>
> [Execution: the mathematics, in sentences. Shape announced if it's a proof. Key equation displayed. Steps justified.]
>
> *Answer.* Therefore [answer in the problem's own terms, with units, boxed or set on its own line].
>
> *Check.* [Plug-back, endpoint, or sanity test, one line.]

Top of the set, once: name, course, set number, and the collaboration line: "Discussed problems [list] with [names]."

Template 2: the proof

The all-shapes skeleton (Chapters 3 through 6); delete what the shape doesn't need:

> **Claim.** [Full statement, hypotheses and conclusion explicit.]
>
> *Proof.* [Shape announcement if not direct: "We prove the contrapositive: …" / "Suppose, for contradiction, that …" / "We proceed by induction on n."]
>
> [Objects introduced, in dependence order: "Let … be given. Choose …"]
>
> [Hypotheses unpacked via definitions.]
>
> [The argument. Induction: labeled base case, then hypothesis displayed and used. Cases: labeled, exhaustive, reunited at the end. Existence-uniqueness: two labeled halves.]

[Closing line verifying the exact conclusion, or naming the contradiction.] □

Template 3: the paper or thesis chapter

Structure per Chapters 10 through 14:

Title: [noun phrase naming the contribution]

Abstract: [4 sentences: setting; question; result; method or scope. No citations, no notation. Written last.]

1. Introduction. [Concrete opening instance. The question, stated plainly. Main result, early and informal, pointing to its formal number. Related work organized by idea, with the "what we change" sentence. Roadmap, one line per section.]

2. Preliminaries. [Only what's used. Conventions declared. Known results stated in your notation with theorem-level citations. Notation table if heavy.]

3–k. Main results. [Statement-first. Top-down (theorem, then lemmas) or bottom-up (lemmas building, destination announced). One numbering series. Proofs signposted by steps; long side-proofs deferred with pointers.]

k+1**. Examples / applications.** [Running example revisited. Boundary examples for the definitions. Sharpness discussion, or the honest open question.]

References. [Built by your bibliography tool; every entry verified against the source.]

[Thesis only: chapter-opening map paragraphs, expanded background chapter, and the attribution sentences marking surveyed versus original work, per chapter.]

The revision pass, attached to all three

Whatever the template produced, it ships only after (Chapter 14): cooling time, the logic audit, the stranger's read-aloud, the compression pass, the mechanical sweep with Appendix B, and, for anything that matters, one borrowed reader asked the one right question: "where did you first stop following?"

E

Further Reading

Short and honest. Everything here shaped this book; each entry says who it's actually for.

Paul Halmos, "How to Write Mathematics" (1970). A long essay, not a book, and still the best thing ever written on the subject. Opinionated, quotable, aimed at working mathematicians but readable the summer after your first proofs course. Read it once now and again before your thesis; it grows as you do. Free copies circulate through university libraries and the AMS.

Steven G. Krantz, *A Primer of Mathematical Writing* (AMS). The professional's handbook: papers, referee reports, grant prose, talks. Overkill during coursework, exactly right the month you start writing a thesis or your first paper. Where my Part IV compresses, Krantz expands.

Nicholas J. Higham, *Handbook of Writing for the Mathematical Sciences* (SIAM). The reference shelf: comprehensive, meticulous, from English usage through publishing workflow. Not a cover-to-cover read; a superb place to settle arguments. If your department library owns one writing book, it's this.

Franco Vivaldi, *Mathematical Writing* (Springer). The closest cousin to this book: an undergraduate course text with exercises, more formal in style. If you want problem sets that drill the skills of my Parts I and II, Vivaldi supplies them.

Kevin Houston, *How to Think Like a Mathematician* (Cambridge). Half study skills, half proof technique, warm in tone. Strongest on the thinking that precedes the writing, so it pairs naturally with my Chapters 3 and 5. Best read in your first proofs semester.

G. Pólya, *How to Solve It* (Princeton). Not a writing book at all, and on the shortlist anyway: the classic on the exploration phase, the scratch-paper thinking this book keeps telling you to do and then conceal. Timeless, cheap, everywhere.

Francis Su, "Guidelines for Good Mathematical Writing." A few pages of advice for exactly your level, free online, written by one of the great mathematical teachers. Hand it to a classmate who won't read a whole book; it agrees with this one on everything that matters.

And when you're ready to typeset any of it, from problem sets to the thesis: my *LaTeX for Students* picks up precisely where this book's templates leave off.

F

Rewrite Workshop and Solutions

Reading a diagnosis feels easier than producing one. These exercises make you mark the debt before you see the repair. Cover the solution, rewrite the passage, and then compare decisions rather than wording. More than one sentence can be good. The mathematical obligations are less negotiable.

F.1 Exercise 1: repair an equals chain

BEFORE

$$x^2 - 5x + 6 = 0 = (x-2)(x-3) = x = 2 \text{ or } x = 3.$$

Your task. Keep the calculation compact, but make every relation true and state the rule that creates the two answers.

Solution.

$$x^2 - 5x + 6 = 0 \quad \Longleftrightarrow \quad (x-2)(x-3) = 0.$$

By the zero-product property, $x - 2 = 0$ or $x - 3 = 0$. Hence $x = 2$ or $x = 3$. Substitution verifies both solutions.

The original chain used an equals sign between an equation, a factorization, and a conclusion. The repair uses equivalence only where the solution set is preserved, then changes to prose when a rule is applied.

F.2 Exercise 2: expose the quantifiers

BEFORE

A function is continuous when close inputs give close outputs.

Your task. Rewrite this as a usable definition at a point a, then translate the order of control into one sentence.

Solution. A function f is continuous at a if for every $\varepsilon > 0$ there exists a $\delta > 0$ such that

$$|f(x) - f(a)| < \varepsilon \quad \text{whenever} \quad |x - a| < \delta.$$

The reader chooses the required output tolerance ε first; only then may we choose an input tolerance δ that works for every eligible x.

The informal line contains the right intuition, but it does not say who chooses "close," in which order, or whether one δ must work for every ε.

F.3 Exercise 3: turn discovery around

BEFORE

We want $|x^2 - 4| < \varepsilon$. This is $|x-2||x+2| < \varepsilon$. Near 2, $|x+2| < 5$, so it is enough that $|x-2| < \varepsilon/5$. Choose $\delta = \min(1, \varepsilon/5)$.

Your task. Keep the discovery, but write the final argument from the hypothesis toward the goal.

Solution. Let $\varepsilon > 0$ be given and choose

$$\delta = \min\{1, \varepsilon/5\}.$$

If $|x-2| < \delta$, then $|x-2| < 1$, so $1 < x < 3$ and therefore $|x+2| < 5$. Hence

$$|x^2 - 4| = |x-2||x+2| < 5|x-2| < 5\delta \le \varepsilon.$$

Thus $\lim_{x\to 2} x^2 = 4$.

The scratch version correctly found the bound. The rewrite introduces the choice before using it and shows why the auxiliary restriction $\delta \le 1$ controls $|x+2|$.

F.4 Exercise 4: finish the question asked

BEFORE

$$f'(x) = 3x^2 - 12x = 3x(x-4).$$

The problem was: determine where $f(x) = x^3 - 6x^2 + 5$ is increasing.

Your task. Add only what the answer needs.

Solution.

$$f'(x) = 3x(x-4).$$

This derivative is positive when $x < 0$ or $x > 4$, zero at 0 and 4, and negative on $(0,4)$. Therefore f is increasing on $(-\infty, 0)$ and $(4, \infty)$.

The derivative was intermediate work. A complete solution performs the sign analysis and answers in intervals, the language of the question.

F.5 Exercise 5: untangle overloaded notation

BEFORE

Let n be the number of vertices. For each vertex n, let n be its number of neighbours. Then $\sum n = 2m$.

Your task. Give each object one job and make the summation index visible.

Solution. Let $G = (V, E)$ be a finite graph with $n = |V|$ vertices and $m = |E|$ edges. For each vertex $v \in V$, let $\deg(v)$ denote its degree. Then

$$\sum_{v \in V} \deg(v) = 2m.$$

The rewrite separates the graph, its sets, their sizes, a generic vertex, and its degree. The equation can now be read aloud without guessing which n is meant.

F.6 Exercise 6: replace an introduction that advertises

BEFORE

Graph theory is a vast and important area with many applications. In this paper we study trees and obtain some interesting results.

Your task. Open with the exact question and result for a short paper proving that every tree on n vertices has $n-1$ edges.

Solution. A tree is a connected graph with no cycles. We prove that every finite tree with n vertices has exactly $n-1$ edges. The proof removes a leaf and uses induction, making precise why connectedness can be maintained while one vertex and one edge disappear together.

The repair names the object, result, and method. It does not need to declare the field important or promise "interesting" material.

F.7 Exercise 7: write the abstract after the result exists

BEFORE

This paper will investigate sequences and will attempt to find a formula for a particular recurrence. Several methods will be considered.

The paper proved: if $a_0 = 2$ and $a_{n+1} = 3a_n$, then $a_n = 2 \cdot 3^n$.

Solution. For the recurrence $a_0 = 2$ and $a_{n+1} = 3a_n$, we prove $a_n = 2 \cdot 3^n$ for every $n \geq 0$. The proof is a direct induction, and the closed form shows that each step multiplies the initial value by one further factor of 3.

An abstract reports the finished result. Future tense and "attempt" belong in a proposal, not in the document after the theorem has been proved.

F.8 Exercise 8: point a citation at the claim

BEFORE

It is known that every finite integral domain is a field [4].

Your task. Assume source [4] is a textbook and the result is Theorem 7.3 on page 142. Make the dependency and pointer visible.

Solution. Every finite integral domain is a field [4, Theorem 7.3, p. 142]. We use this result below to conclude that the nonzero residue classes modulo a prime have multiplicative inverses.

The citation now lets the reader locate the exact result and explains why it enters this paper. A bare bracket proves only that a book exists.

F.9 Exercise 9: separate picture from proof

BEFORE

The graph clearly crosses the axis once, so the equation has exactly one solution.

Your task. Suppose f is continuous and strictly increasing, with $f(0) < 0 < f(1)$. Use the figure only as illustration.

Solution. Because f is continuous and changes sign between 0 and 1, the Intermediate Value Theorem gives at least one $c \in (0,1)$ with $f(c) = 0$. Because f is strictly increasing, it cannot take the value 0 at two distinct points. Hence the solution is unique. The figure illustrates this crossing but is not used as evidence.

Existence and uniqueness are separate debts. The first comes from continuity and the sign change; the second comes from strict monotonicity.

F.10 Exercise 10: perform the stranger's read

Take one page of your own current work and mark:

1. every symbol used before it is defined;
2. every connective whose promised direction is unpaid;
3. every line that a reader cannot justify from the previous one;
4. every example that is being asked to act as a proof;
5. the sentence that answers the actual question.

Then rewrite the page without adding new mathematics. If the new page is longer, check whether the added words expose reasoning. If they only announce that a step is "obvious," cut them and write the reason.

There is no model solution for this exercise because the page must be yours. The test is concrete: a classmate should be able to point to the birthplace of every symbol and the reason for every transition without asking what you meant.

A Note from the Author

Thank you for reading. Books like this one live or die on word of mouth from students, so if it saved you marks or headaches, two minutes of your help matter more than you'd guess: leave a short, honest review wherever you bought it, or hand your copy to the classmate whose proofs need it most.

The tools in the appendices, the proof-writing checklist, the error catalog, and the templates, are meant to live on your desk, not inside a closed book. You'll find printable companions and more of my writing for students at **gauravtiwari.org**. Questions and corrections are welcome there too; future editions will owe their sharpest fixes to readers who wrote in.

Also by Gaurav Tiwari

LaTeX for Students
A Practical Guide to Beautiful Documents and Mathematical Typesetting

The natural companion to this book. Everything you just learned to write, from Tuesday's problem set to the thesis, and how to make it look as good as it reads: documents, equations, figures, references, and ready-to-use templates, built step by step for students with no prior experience.

Other titles

- *Indian Polity: Complete Notes for UPSC Prelims & Mains*
- *Website and Blog Copywriting*
- *Marketing Essentials for Bloggers*
- *The Affiliate Bloggers System*
- *Ultimate Guide to Backlink Building Strategies*

All titles are available on Amazon. Find the full, current catalog at **gauravtiwari.org**.