How to Read Mathematics

How to Read Mathematics

Definitions, Theorems, Proofs, and Papers for Students

Gaurav Tiwari

gauravtiwari.org

Copyright © 2026 Gaurav Tiwari.

All rights reserved. No part of this book may be reproduced in any form without written permission from the author, except for brief quotations in reviews.

First edition, 2026.

gauravtiwari.org

*To everyone who has read the same paragraph four times
and is about to read it a fifth.*

Contents

Preface

You're on page 41.

You've read the same paragraph four times. Every word in it is a word you know. The sentence is grammatical. And nothing's happening. No picture forms, nothing clicks, and the clock says you've been at this desk for fifty minutes and covered a page and a half.

So you draw the obvious conclusion. You're not smart enough for this.

You're wrong, and I can tell you exactly where. Mathematical prose isn't written to be read the way you read everything else. It's compressed, on purpose. A definition is a folded thing that means nothing until you unfold it. A theorem hides half its content in the order of its quantifiers. A proof is the tidy final draft of a messy discovery, with the scratch work deleted, and that deleted scratch work is exactly what would have made it obvious.

Nobody tells you this. You get handed the compressed object, you get graded on whether you can decompress it, and the decompression procedure is left as an exercise for the reader.

This book is the procedure.

It's the companion to *How to Write Mathematics*, and the two halves fit the way you'd expect. That one is about compressing your thinking so a reader can follow it. This one is about unfolding somebody else's. You don't need to have read it first, and you don't need any mathematics past what your own courses are teaching you right now. Every worked passage here runs on first-course material: even integers, $\sqrt{2}$, sequences, continuity. The reading lessons are what transfer, so I've kept the mathematics out of their way.

How to use it. Part I is why the page resists you and what a mathematical text is actually made of. Parts II, III, and IV take the three objects you'll meet ten thousand times, the definition, the statement, and the proof, and hand you a procedure for each. Part V scales up to chapters, whole books, and research papers. Part VI is about the pen in your hand and the people in the room.

Read it straight through if you like. But it's a toolbox, not a novel. Stuck on a definition right now? Go to Chapter 4 and run the drill on the page in front of you. Proof refusing to land? Chapter 10 tells you which of the three readings you're attempting and which one you actually need. Stuck and cross about it? Chapter 13 is a decision tree that ends with you doing something useful instead of reading the paragraph a fifth time. The appendices are working tools. Keep them close, write on them.

Two warnings, both honest.

This book won't make mathematics fast. Nothing makes mathematics fast. What it does is make your slowness productive, which is a different thing and a better one, and it stops you burning forty minutes in a rereading loop that was never going to work.

And it's going to ask you to write while you read. Every technique in here needs a pen. Reading mathematics with your hands in your pockets isn't reading mathematics. It's watching mathematics, and watching doesn't transfer.

That panic on page 41 isn't a verdict on you. It's a signal that the page is folded up and nobody ever taught you how to open it.

Let's open it.

Gaurav Tiwari
gauravtiwari.org

Part I

Why the Page Resists You

1

The Page That Won't Let You In

Here's an experiment I run with students who tell me they're bad at mathematics. I hand them a page from a real analysis textbook, set a timer for fifteen minutes, and ask them to read it. Not solve anything. Just read.

Fifteen minutes later, almost every one of them has done the same thing.

They read the definition at the top and got roughly nothing. Read it again, got roughly nothing again. Moved down to the theorem, because the definition wasn't going anywhere. Recognised none of the notation. Moved to the proof, because it looked like it might explain things. Followed two lines, hit a step that said *and therefore*, stopped, went back to the definition, and started the loop over.

They're not lazy and they're not slow. They're reading mathematics the way they read everything else, and that's the one thing that can't work.

1.1 The compression problem

Ordinary prose has slack in it. A paragraph in a history book says its main idea two or three different ways, gives you an example, tells you what's coming next. You can skim it, miss a third of the words, and still walk out with the argument. That redundancy is a gift from the author, and you've been trained on it since you were five.

Mathematical prose has the redundancy stripped out. It's written to be *correct* and *complete*, not to be absorbed at speed. One printed line can carry six separate ideas, each of which the author expects you to unfold on your own time with your own pen.

When a textbook says

> Since f is continuous on the compact set K, it attains its bounds.

that's fourteen words, and nobody on earth reads it in fourteen words' worth of effort. Figure 1.1 shows what's actually sitting inside it.

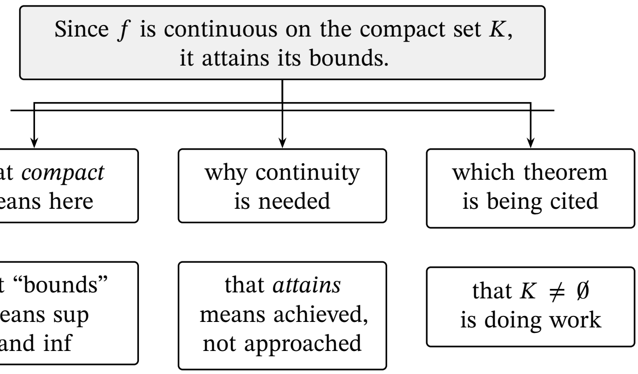

Figure 1.1: One printed line, unfolded. Everything below the rule is content the author expects you to supply.

Six unfoldings. And that's a friendly sentence, written by somebody being kind to you.

It's not badly written. It's written for a reader who unfolds, and unfolding is a learned skill that nobody sat you down and taught.

> *Mathematical text is compressed on purpose. Reading it means decompressing it, on paper, with your own hand. The compression isn't the author being unkind. It's the price of being exact.*

1.2 Four ways students read, and why three of them fail

I've watched the same four reading styles show up over and over. Three are dead ends. Work out which one is your default, because you can't stop doing something you can't see yourself doing.

The passive read

Eyes move, pages turn, nothing gets written down. This feels like studying and produces nothing, because comprehension in mathematics gets tested by production, not recognition. You'll finish a passive read of a chapter feeling fine about it and be unable to state a single definition from it an hour later.

The tell: you've been reading for thirty minutes and your pen hasn't moved.

The linear read

You start at the first word and refuse to move on until you've understood it, so you never reach the paragraph three pages later where the author explains what they were doing. Mathematics texts aren't written linearly. They're written in layers, and the payoff for the definition on page 41 is often the example on page 44.

The tell: twenty minutes on the same object, and you haven't looked ahead once.

The symbol-skip

Your eye slides over $\forall \varepsilon > 0$ the way it slides over a long foreign name in a novel. It registers as *some symbols* and you move on. Since the whole meaning of most mathematical sentences lives in the symbols and their order, you've just skipped the sentence.

The tell: you can't read the line out loud, in English, without stumbling.

The active read

You have a pen. You're writing while you read: examples in the margin, the definition restated in your own words, a small picture, a question you can't answer yet. You look ahead when you're stuck, then come back. You say symbols out loud. You're slower per page and much faster per concept.

Everything in this book is a technique for making that fourth one automatic.

1.3 The rereading loop

The most expensive habit in student mathematics is the rereading loop, and it's expensive precisely because it feels like work.

You read a paragraph, don't get it, and read it again with a bit more determination. The second pass adds nothing. It adds nothing because the reason you missed it the first time wasn't insufficient attention. It was a missing piece somewhere else: a definition you never unfolded, a symbol you can't say out loud, an example you never built.

Determination can't supply a missing example. Only construction can.

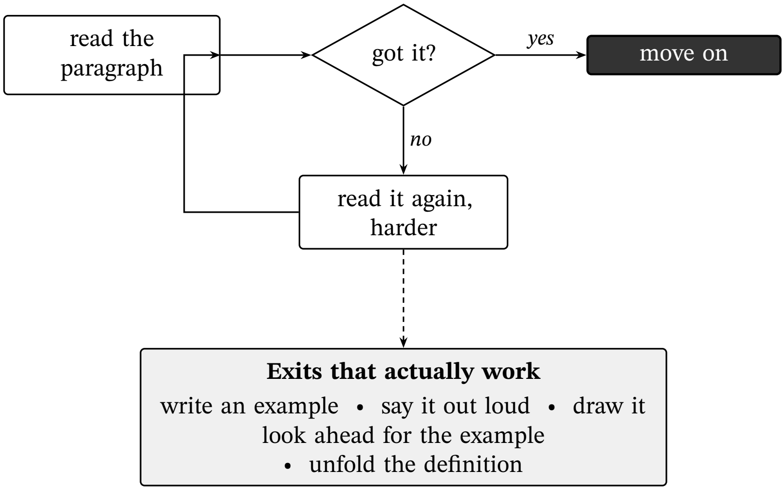

Figure 1.2: The rereading loop has no exit inside itself. Every way out is an action with a pen.

Look at what the working exits have in common. Every one of them changes the page in front of you or the sound in the room. Not one of them is *concentrate harder*.

So when you catch yourself on that second pass with nothing written down, that's the moment. Stop reading. Start producing.

1.4 What reading mathematics actually is

Here's the definition that runs through the rest of this book.

> *To read a piece of mathematics is to be able to reconstruct it: to state the definition in your own words, produce an example and a non-example, say why each hypothesis is there, and rebuild the proof's skeleton without looking.*

High bar, and deliberately behavioural. Notice it never mentions understanding, insight, or feeling comfortable. Those are unmeasurable and they lie to you. Students feel comfortable with things they can't reproduce all the time (it's the most common self-deception in the subject, and it usually surfaces about six minutes into an exam).

Every technique in this book converts the vague question *do I understand this?* into something with a testable answer, like *can I write down a function that satisfies the first two hypotheses and fails the third?*

The bar's lower than it looks, by the way. Reconstruction is mostly bookkeeping once you know what to record, and the recording happens while you read, not in a panicked pass the night before.

1.5 What this book will and won't do

It won't teach you analysis, algebra, or topology. Your courses own that.

It won't make you fast. Five pages an hour is a respectable rate for a first pass through a serious text, and it always will be.

And it won't save you from being stuck, because being stuck is the normal working state of everyone who does mathematics, including whoever wrote the book you're stuck on.

What it will do is give you a procedure for each of the three objects you'll meet ten thousand times. Definitions get unfolded (Part II). Statements get interrogated before you believe them. Proofs get read in a specific order, at one of three depths, on purpose instead of by accident. Then Part V scales all of it up to chapters, whole books read alone, and research papers, which are a different animal with their own rules.

Start where your pain is. Definition refusing to open? That's Chapter 4. Proof? Chapter 10. Just stuck and cross? Go straight to the protocol in Chapter 13 and come back here after.

Next, though, the map. Before you can read a mathematical text efficiently, you need to know what its pieces are and which of them hold the weight, because they don't all deserve the same attention and most students spend theirs in exactly the wrong places.

READ SOMETHING TODAY

Take the page you're stuck on right now. Ten minutes, pen in hand, and write down every symbol you can't say out loud in English. Don't try to understand anything. Just build the list. That list, not the page, is your actual work for today.

2

What a Mathematical Text Is Made Of

A page of mathematics looks uniform. Same font, same margins, one grey block of text after another.

It isn't uniform at all. It's built from about six kinds of object, each wanting a different kind of attention, and students lose enormous amounts of time giving all six the same kind.

The usual version of the mistake: reading an example with the same slow, suspicious care you'd give a definition. Or skipping a remark that contains the only sentence in the chapter explaining what's going on. Learn the parts and you can spend your attention where it pays.

2.1 The six objects

Definitions

A definition creates a word. Before it, the word means nothing. After it, the word means exactly what the definition says and not one thing more.

Definitions are the load-bearing walls. Everything else in the chapter rests on them, and every hour you don't spend unfolding one gets charged back to you with interest when you reach the proofs.

The tell that you haven't really read a definition: you can recite it but you can't produce an object that satisfies it.

Statements

Theorems, propositions, lemmas, and corollaries are all the same kind of object, a claim with hypotheses and a conclusion. The four different words are the author telling you how much it matters and why it's here.

- **Theorem.** A main result. The chapter probably exists for this. Learn it.
- **Proposition.** A real result, one rung down. Useful, quotable, not the headline.
- **Lemma.** A tool. It's there to make a later proof shorter, and it usually contains the actual clever step. Don't skip lemmas because they look small. In plenty of chapters the lemma is where the mathematics happens and the theorem is bookkeeping.
- **Corollary.** Something that follows quickly from what was just proved. Read the statement, then check whether you can see why it follows before you read the proof.

That last one is worth keeping. If you can see it, you understood the theorem. If you can't, you didn't, and the corollary just told you so for free. A built-in comprehension test that costs thirty seconds.

Proofs

A proof is a verification, written in the order things are *true*, not the order they were *found*. That's why proofs so often feel like a magic trick. The rabbit appears in line three and nobody explains where the hat came from.

Part IV is entirely about reading them. For now: a proof is the most compressed object on the page, and it's the one you should attempt last.

Examples

Examples are where meaning lives. The definition tells you the rule; the example tells you what the rule feels like.

Authors put examples immediately after definitions for a reason. A student who bounces off a definition and quits before reading the example directly below it has walked out of the room ten seconds before the explanation.

Remarks

The remark is the author talking in their own voice, off the record. Motivation. Warnings about a common error. Why a hypothesis is there. The honest admission that a definition looks strange but will pay off in Chapter 9.

Per word, remarks are the highest-value text in most books. They're also the most skipped, because they aren't numbered and don't look important.

Exercises

Not decoration, and not homework padding. Textbook exercises regularly contain results the book uses later.

Read the exercise statements even when you're not going to do them. They're the author telling you what you should now be able to do, which is the chapter's learning objective in disguise.

2.2 Attention is a budget, so spend it deliberately

Figure 2.1 maps a typical textbook section with the attention each part deserves on a first working read. The proportions surprise most students. The proof of the main theorem, where they spend most of their time, isn't where most of the value is on a first pass.

The rule behind the picture: attention flows to whatever everything else depends on. Definitions and examples sit upstream. The proof of a theorem is downstream of all of it, which is exactly why attacking it first feels like drowning.

> *Read upstream to downstream: definitions, then examples, then the statement, then the proof. Most students read top to bottom and hit the proof while still carrying an unfolded definition.*

2.3 The dependency graph

Every chapter has a hidden structure: which results need which. Drawing it takes ten minutes and turns a wall of text into a small number of things to learn plus a lot of consequences.

Build it as you read, in the margin or on a separate sheet. The payoff lands at revision time. Once you can see that Corollary 3.8 hangs off Theorem 3.6, which hangs off one lemma and two definitions, you stop trying to memorise eight independent facts and start learning three and a derivation.

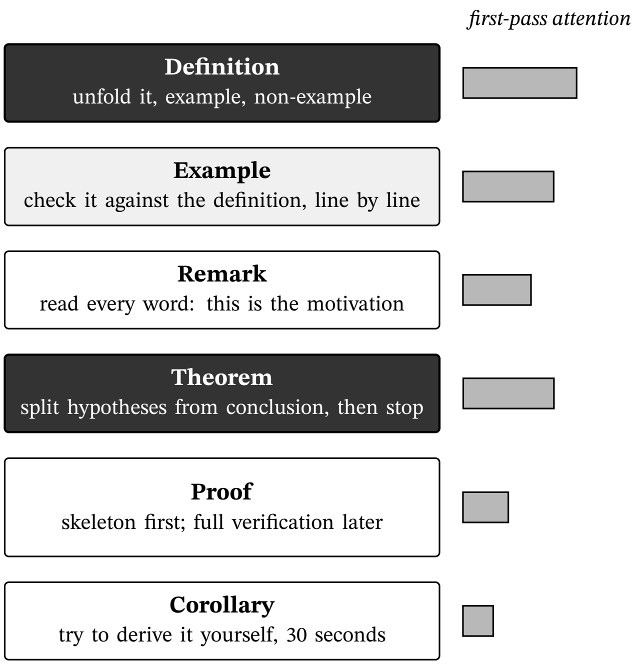

Figure 2.1: A typical section, with first-pass attention. The proof is the biggest object on the page and rarely the best first investment.

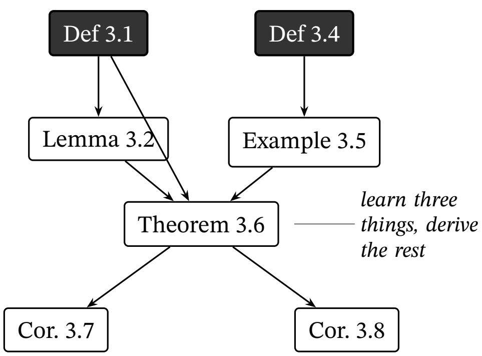

Figure 2.2: The hidden shape of a chapter. Two definitions and one lemma carry everything; the rest is consequence.

There's a second payoff, quieter and more useful. If you can't draw an arrow into a result, you haven't understood why it's in the chapter at all. Better to find that out now than in an exam.

2.4 Where the author hid the intuition

The conventions of mathematical writing push intuition out of the formal objects. A definition can't contain the phrase *think of this as a rubber sheet*; it would stop being a definition.

So the intuition gets displaced, and it lands in four predictable places. Knowing them is worth several IQ points on any text.

1. **The paragraph immediately before a definition.** Often starts *We would like to capture the idea that…* That sentence is the definition in English, and it's frequently the only English version you'll get.
2. **The first example after it.** Chosen to be the smallest object that shows the point.
3. **Remarks.** The author's editorial voice.
4. **The chapter introduction.** Two paragraphs almost everyone skips because they contain no symbols, and which usually say exactly what the chapter is for.

> *The intuition is never in the definition. It's in the sentence before it, the example after it, and the remark three lines down. Read those first, then come back and read the definition as the precise version of what you already half know.*

2.5 A worked pass over one section

Say you open a book at a section called *Uniform continuity*. Here's the order I'd read it in, with the clock running.

1. **The section's opening paragraph** (1 minute). It'll say something like *continuity is a local condition; sometimes we need one δ that works everywhere at once*. You now know the whole point of the section.
2. **The definition, slowly, out loud** (3 minutes). Watch where the quantifiers sit compared with ordinary continuity. The entire content is in that reordering, and it's invisible unless you go looking.
3. **The first example, with a pen** (5 minutes). Verify it against the definition instead of trusting it.
4. **The non-example** (5 minutes). If the author gives one, work through it. If not, build your own. $f(x) = 1/x$ on $(0, 1)$ is waiting for you, and so is $f(x) = x^2$ on $\mathbb{R}$.

5. **The theorem statement, not its proof** (2 minutes). Split hypotheses from conclusion. Ask why compactness showed up in the hypothesis list.
6. **Now the proof** (15 to 40 minutes), knowing what it's for.

That's half an hour before you touch the proof, and then the proof takes half as long and stays put.

The instinct to jump straight to step 6 is the instinct that costs students their evenings.

READ SOMETHING TODAY

Open the chapter you're working through. Don't read it. Just list every numbered object in it by type: D for definition, T for theorem, L for lemma, E for example, R for remark. Then draw arrows between the ones that depend on each other. Ten minutes, one sheet. You'll know more about that chapter than most of your class will after two hours of reading it front to back.

3

Four Gears

How long should this take? is the question students are most afraid to ask out loud, because they suspect the honest answer will confirm something bad about them.

So here's the honest answer first, and then I'll make it useful.

A serious mathematics text, read properly for the first time, goes at two to six pages an hour. Not a chapter an hour. Pages. A research paper in your own area takes a day. A research paper outside it takes a week, or never. Nobody reads mathematics at reading speed, including the people who wrote it, and anyone telling you otherwise is skimming or lying.

That number isn't a verdict on you. It's a property of the material. And once you accept it, two things become possible: you can plan your week around real numbers, and you can stop reading your own slowness as failure.

3.1 Reading has gears, and the wrong gear wastes the hour

The thing underneath most wasted study time isn't slowness. It's using one gear for everything.

Here are the four, fastest to slowest.

Gear 1: the survey

Ten minutes for a whole chapter. You're not reading, you're building a map. Section headings, the names of the definitions, the statements of the theorems, the last paragraph of the introduction. Don't stop for anything.

At the end you should be able to say what the chapter is about in one sentence and name the three objects it introduces.

Use it before every first reading, before a lecture, and when you're deciding whether a chapter is relevant to you at all.

Gear 2: orientation

Thirty to forty minutes for a section. You read the definitions and statements properly, work the examples, and *skip every proof*.

The goal is to know what's true and what depends on what, without paying the verification cost yet. Most students have never once read mathematics this way, and they're startled by how much it gives them for the time.

Use it the day before the lecture that covers the material, or when you need a result now and will verify it later.

Gear 3: the working read

Two to six pages an hour, pen moving the whole time. Definitions unfolded, examples built, proofs read for their skeleton, questions in the margin.

This is the gear the rest of this book is mostly about, and it's the one that makes material stick. Use it for your current courses and for anything you'll be examined on.

Gear 4: verification

One page an hour. Sometimes less. You check every step of every proof, supply the missing algebra in the margin, and test each hypothesis by trying to break the theorem without it.

Brutal, slow, and the only gear that produces real mastery. Use it on the four or five results a whole course rests on, on anything you'll build research on, and on the proof technique you keep failing to reproduce.

3.2 Choosing the gear on purpose

The gear isn't a mood. It follows from what you need the text for, and you should be able to say the reason out loud before you open the book. Four questions decide it, and Figure 3.2 runs them in order.

Two consequences, both of which cut against student instinct.

Most of what you read shouldn't be read in gear 4. Verification is expensive and your hours are finite. Spending them all on the first three pages of a chapter is why so many students know Chapter 1 of every book and Chapter 6 of none.

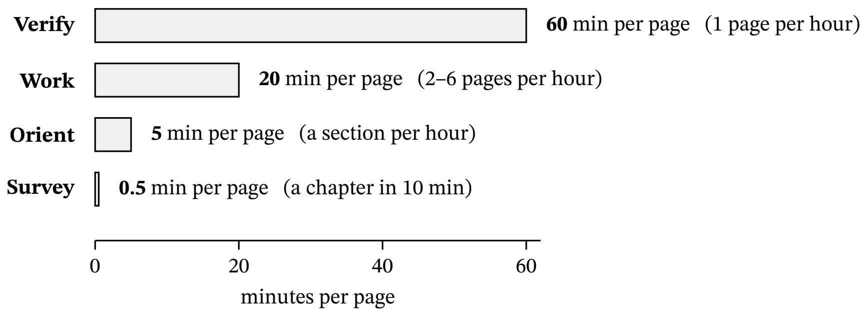

Figure 3.1: Four gears, by how long a page takes. Verification costs about a hundred times what a survey costs per page, which is why choosing the gear matters more than reading faster.

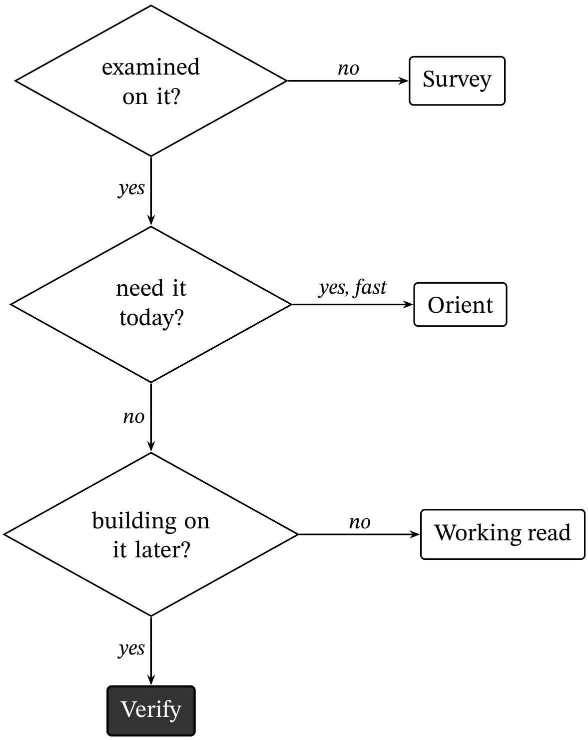

Figure 3.2: Pick the gear before you start, not after you have already lost an hour in the wrong one.

And most of what you read shouldn't be read in gear 1 either. Surveying feels wonderfully productive and teaches you almost nothing by itself. It's a map, not a journey.

3.3 Time budget arithmetic

Do this once a term and it'll change how you plan.

Take a course. Count the pages of assigned reading. Multiply by four minutes a page for a working read. That's your reading budget, and it's usually a shock.

Twelve weeks, fifteen pages of dense text a week, is 180 pages. At four minutes a page that's twelve hours of reading on top of lectures and problem sets, and you don't have twelve spare hours.

So you triage. Deliberately, rather than by accident:

- Everything gets a survey. Ten minutes a chapter, non-negotiable. Reading a chapter you haven't surveyed is walking into a city without a map.
- The core results, the four or five everything else uses, get gear 4. You'll spot them from the dependency graph in Chapter 2, or by noticing which theorems the lecturer keeps citing.
- Everything the problem sets touch gets gear 3.
- Everything else gets gear 2, and you accept knowing what's true there without knowing why.

That last acceptance is the professional move, and students fight it hardest. Working mathematicians hold hundreds of results as *true, proof not personally verified*. It isn't laziness. It's how anyone gets anything done inside a finite life.

What matters is knowing which category each result is in, and never mistaking gear 2 knowledge for gear 4 knowledge the night before an exam.

3.4 How to structure two hours

Two-hour block? Don't spend it in a single mode. This is the pattern that works, over and over:

- **0–10 min.** Survey what you're about to read, or resurvey what you read last time.
- **10–55 min.** Working read. Pen down constantly. Stop at the first proof that resists, and mark it instead of fighting it.

- **55–65 min.** Break. A real one, away from the desk.
- **65–105 min.** Back to the marked proof in gear 4, or exercises on what you just read. Exercises after reading, always. The retrieval is where the learning actually happens.
- **105–120 min.** Close the book and write, from memory, the definitions and statements you met. Half a page.

That last fifteen minutes is the most valuable block on the list and the first thing everyone cuts.

If you keep one habit from this chapter, keep that one. End every session by writing down, without looking, what you just read. It turns a session that would have evaporated by Thursday into something you still own in March.

READ SOMETHING TODAY

Before your next reading session, write two lines at the top of the page: what you're reading, and which gear. Then set a timer. At the end, note whether the gear was right. Do that for a week and you'll start choosing gears automatically, which is the whole point.

Part II

Definitions

4

Unpacking a Definition

Definitions are where reading mathematics is won or lost.

Everything downstream is an elaboration of objects a definition created. Every theorem, every proof, every exercise. And a definition is the most compressed thing on the page: one sentence, often four or five quantifiers, meaningless until you unfold it.

This chapter is a fixed procedure for that unfolding. Six steps, ten to twenty minutes the first few times, about five once it's a habit. Run it on every definition that matters.

Ten minutes here saves you two hours in the proofs. I've never seen that trade fail.

4.1 The drill

Step 1: say it out loud

Not read it. Say it, in English, without the symbols.

Can't? Then you've found your problem, and it isn't the concept, it's the notation. Go to Chapter 6, build the symbol list, come back.

Saying it out loud does something silent reading can't. It forces you to straighten the sentence out and commit to a reading of the quantifiers. You can't mumble past $\forall \varepsilon > 0\ \exists N$ when you're speaking. Either you say "for every positive epsilon there exists an N" or you stall, and the stall is information.

Step 2: name every object and its type

Two columns. Left, the symbol. Right, what kind of thing it is: a real number, a positive real, an integer, a set, a function from where to where, an element of what.

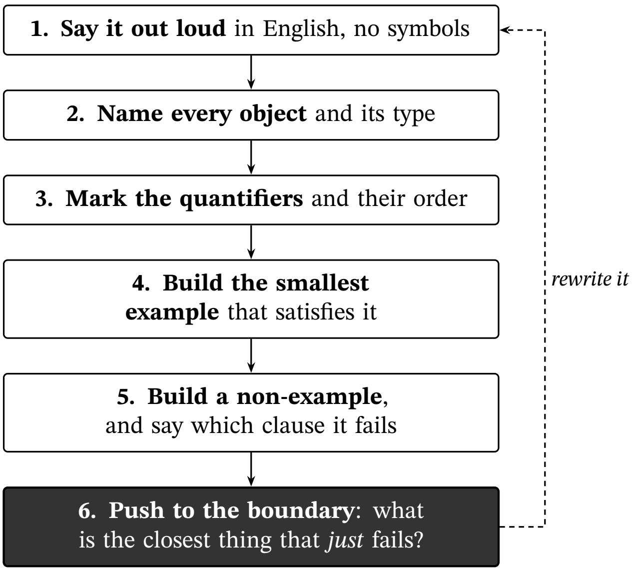

Figure 4.1: The unpacking drill. You are not finished at step 4; the boundary case is where the understanding is.

Half of all confusion with a definition is type confusion. Reading ε as an index. Reading N as a bound on the sequence rather than on the subscript. The list takes ninety seconds and kills the whole category.

Step 3: mark the quantifiers and their order

Nobody teaches this and everybody needs it, so it gets its own section below.

Step 4: build the smallest example

Not any example. The smallest one you can build with your own hands, where you can check every clause explicitly.

If the definition is of a bounded sequence, don't reach for something clever. Write $a_n = 1$ for all n and verify it against the definition symbol by symbol.

Trivial examples aren't beneath you. They're the only ones where you can see the whole machine at once.

Step 5: build a non-example

An object that fails, plus the clause it fails on.

"$a_n = n$ isn't bounded" doesn't count. "$a_n = n$ isn't bounded because for any proposed M we can take $n > M$, so no single M works for all n" is the answer. And notice that writing it made you use the quantifier structure.

Step 6: push to the boundary

The best question you can ask about any definition: what's the closest thing to satisfying it that still fails?

A sequence bounded above but not below. A function continuous everywhere except one point. A set that's closed but not bounded. This is where the definition's edges become visible, and edges are what proofs live on.

> *You haven't understood a definition until you can produce something that only just fails it. Anyone can find an object a mile from the boundary. The boundary is the definition.*

4.2 Quantifier order is the whole game

Two sentences. Identical symbols. They differ only in the order of two quantifiers, and one is a far stronger condition than the other.

> (a) For every x there exists $\delta > 0$ such that ...
>
> (b) There exists $\delta > 0$ such that for every x, ...

In (a), δ gets chosen after x is known, so it can depend on x. A different δ at every point. In (b), δ is chosen first and has to work for every x at once.

That single swap is the whole difference between continuity and uniform continuity. Between pointwise and uniform convergence. Between a dozen other pairs of concepts students spend entire terms confusing.

The mechanical rule is simple and worth memorising.

> *A quantified variable may depend on everything quantified* before *it and nothing quantified* after *it. Read left to right and the dependencies are exactly the reading order.*

So when a definition has more than one quantifier, draw the dependency arrows. Figure 4.2 does it for a convergent sequence, which is the definition first-year students most often never quite pin down.

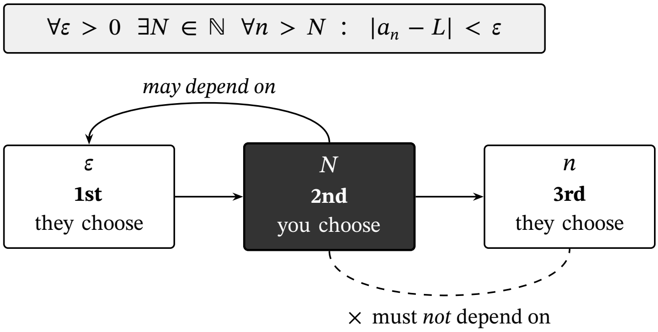

Read it as a game. An opponent hands you ε, as small as they like. You must answer with a single N. Then they may pick any n beyond your N, and you must still be within ε of L.

Figure 4.2: The ε–N definition, with dependency arrows. N may depend on ε. Nothing may depend on n except the conclusion.

Adopt that game reading permanently. Every $\forall \ldots \exists \ldots$ definition is a two-player game: whatever's universally quantified gets handed to you by an opponent, whatever's existentially quantified is your move, and reading order is turn order.

Once you see definitions as games, quantifier order stops being an abstract rule and becomes the obvious question of who moves first.

4.3 Worked: unpacking "bounded"

Let's run the whole drill on something small enough to see all of.

Definition 4.1. A sequence (a_n) of real numbers is *bounded* if there exists $M > 0$ such that $|a_n| \le M$ for all $n \in \mathbb{N}$.

Step 1, out loud. "A sequence is bounded if some single positive number is bigger than the size of every term."

Step 2, types. (a_n) is a sequence of reals, so a function $\mathbb{N} \to \mathbb{R}$. M is a real number, positive. n is a natural number, the index. Notice M isn't a term of the sequence and isn't indexed by anything.

Step 3, quantifiers. $\exists M\ \forall n$. So M comes first, and one M has to work for the whole sequence.

Flip it to $\forall n\ \exists M$ and every sequence in the world becomes bounded, because for a single term you can always find something bigger. That's a useful check, by the way: a definition that's satisfied by every object is almost certainly being read with the quantifiers backwards.

Step 4, smallest example. $a_n = 1$ for all n. Take $M = 1$. Then $|a_n| = 1 \leq 1$ for all n. Done. Second one, still small: $a_n = (-1)^n$, with $M = 1$ again.

Step 5, non-example. $a_n = n$. Suppose some M worked. Take any natural number n bigger than M; then $|a_n| = n > M$, so M fails. No M works, so the sequence is unbounded.

Notice how the proof wrote itself once the quantifier order was clear. To break an $\exists M$, you take an arbitrary candidate M and defeat it.

Step 6, boundary. What's the closest thing to bounded that fails? A sequence bounded above but not below, like $a_n = -n$. Or one that's bounded on any finite stretch and grows without limit, like $a_n = \sqrt{n}$, which grows arbitrarily slowly and is still unbounded.

That second one is the useful boundary object. It kills the intuition "if it grows slowly enough it stays bounded," which a surprising number of students believe without ever having said it out loud.

Total time: about eight minutes. You own that definition now in a way no amount of rereading would have produced.

4.4 The two failure modes

Nearly every misread definition lands in one of two buckets, and they need different fixes.

Type confusion. You're wrong about what kind of object a symbol is. Fix: step 2, the type list. Cheap, and completely effective.

Scope confusion. You're wrong about what depends on what. Fix: step 3, the dependency arrows, plus the game reading.

Scope confusion causes most of the lost marks in analysis, and it never announces itself. A student with scope confusion feels perfectly comfortable and produces proofs where N quietly depends on n, which is exactly the error that makes a proof of a false statement look correct.

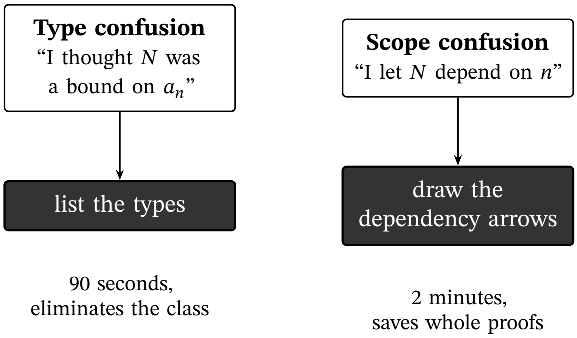

Figure 4.3: Two failure modes, two fixes. Neither is cured by rereading.

4.5 When a definition still won't open

Sometimes you run the drill and step 4 beats you. You can't build a single example.

That isn't failure, it's information, and it means one of three things.

The definition depends on an earlier definition you never unpacked. Most common by far. Go back one level and run the drill there. Keep recursing until you hit something you can build an example of, then come forward. Reading mathematics is often a depth-first search, and students treat getting sent backwards as a defeat when it's the algorithm working correctly.

The examples are genuinely hard. Some definitions have no small examples and the author knows it. That's what the paragraph after the definition is for. Read ahead. Your example is almost certainly sitting right there.

You're missing a piece of notation. You can't build an example of an object described in symbols you can't say. Back to Chapter 6.

None of the three gets fixed by rereading, which brings us back to Chapter 1: when a page resists, the way through is always an action, never more determination.

READ SOMETHING TODAY

Take the most recent definition from your course. Run all six steps on one sheet, writing each step's output. Time yourself. If step 6 defeats you, that's the most useful thing you'll learn this week, because "what's the closest object that just fails this?" is a far better question to bring to a lecturer than "I don't understand this."

5

Examples, Non-Examples, and the Boundary

A definition is a fence. It splits the world into objects that satisfy it and objects that don't.

Reading a definition well means knowing where the fence runs, and you can't get that from the text alone. You have to walk it, which means building objects on both sides and, mostly, objects that sit right up against it.

Most students build one example, feel reassured, and move on. That's like establishing that Paris is in France and concluding you know where the border is.

5.1 The example bank

For every definition that matters, you want four objects. Write them in the margin next to the definition and keep them for the rest of the course.

1. **The trivial example.** The smallest, dullest thing that works. Constant sequence, identity function, one-element set. Its job is to let you check every clause explicitly.
2. **The typical example.** What the author had in mind. Usually the one printed right below the definition.
3. **The near miss.** Something that fails, but only just, and you can name the single clause it fails.
4. **The pathological example.** The famous monster that proves your intuition isn't a theorem. Every serious definition has one, and they're worth collecting on sight.

Four objects, ten minutes. Then when a theorem later says "let f be continuous but not uniformly continuous," you don't have to manufacture

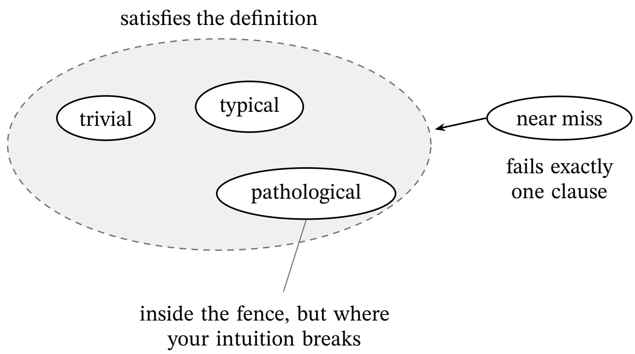

Figure 5.1: The four objects you want for every definition. The near miss and the monster do the teaching; the trivial example only does the checking.

anything, because $1/x$ on $(0, 1)$ has been sitting in your margin since week three.

5.2 Verifying an example is an exercise, not a courtesy

When a textbook says "for example, $f(x) = x^2$ is continuous," the temptation is to nod and move on.

Don't nod. Check it, against the definition, with the definition physically in view.

This feels like busywork on the easy cases. It isn't, and the reason becomes obvious the third time it happens: the act of checking teaches you the shape of the verification, and that shape is exactly what proofs of theorems look like.

A student who has verified fifteen examples against the ε–δ definition has, without noticing, learned the standard architecture of an ε–δ proof. A student who nodded fifteen times has learned nothing and will meet that architecture for the first time in an exam.

> *Checking an example against a definition is the cheapest proof practice in existence. The examples are pre-selected to work, so you can't get stuck, and the moves are identical to the ones you'll need under exam conditions.*

5.3 Building a non-example on purpose

Non-examples are where the information is, and there's a procedure for building them. Take the definition, find its main quantifier, and negate it deliberately.

Say the definition is $\exists M\ \forall n : |a_n| \leq M$. The negation is $\forall M\ \exists n : |a_n| > M$.

Now read that negation as a recipe, because that's exactly what it is. It says: hand me any M and I'll produce an index n that beats it. So build a sequence where you can always produce that n. Try $a_n = n$. Then $a_n = \sqrt{n}$. Then $a_n = \log n$, each more interesting than the last.

This is the most transferable trick in the chapter.

> *To build a non-example, negate the definition first and read the negation as instructions. Students who work by intuition find the obvious non-examples. Students who negate first find the interesting ones.*

Keep the mechanics of negation at your fingertips, since you'll do this hundreds of times:

- $\neg(\forall x\, P(x))$ becomes $\exists x\, \neg P(x)$.
- $\neg(\exists x\, P(x))$ becomes $\forall x\, \neg P(x)$.
- Quantifiers flip in order, left to right, and the innermost statement gets negated last.

So the negation of "$\forall \varepsilon > 0\ \exists N\ \forall n > N : |a_n - L| < \varepsilon$" is "$\exists \varepsilon > 0\ \forall N\ \exists n > N : |a_n - L| \geq \varepsilon$."

Out loud: there's some fixed tolerance the sequence keeps violating, no matter how far out you go.

That English sentence is what non-convergence *means*, and most students never see it, because they never negate.

5.4 Probing the boundary

Once you have objects on both sides, start walking them toward the fence. The question never changes: what's the least I can change to break this?

That figure is doing double duty. It's a technique for understanding definitions, and it's quietly the technique for understanding theorems, which is Chapter 7's job.

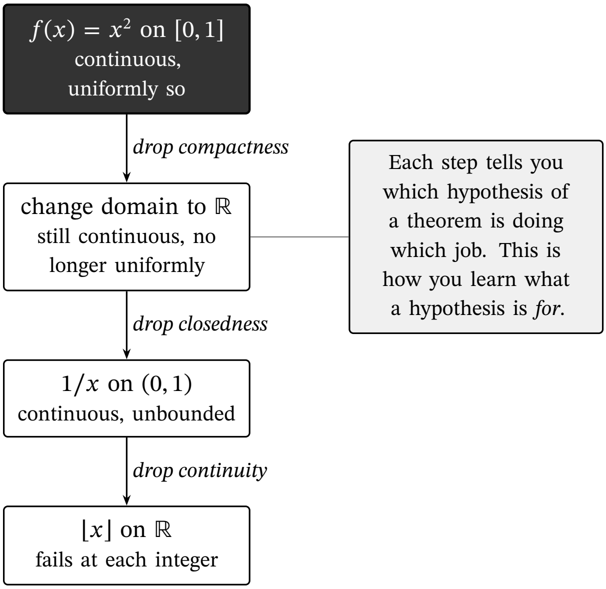

Figure 5.2: Probing. Start from a known example and weaken one hypothesis at a time until the object falls out of the definition.

Every hypothesis in every theorem exists because some counterexample lives just beyond it. Find the counterexample and the hypothesis stops being arbitrary decoration and becomes the obvious thing to require.

5.5 Collect monsters

Every field has a small set of pathological objects that get cited over and over, because each one kills a plausible false belief.

Keep a page at the back of your notebook and add to it all year. In a first analysis course the list grows to include, roughly in order of appearance:

- $a_n = (-1)^n$: bounded, not convergent. Kills "bounded implies convergent".
- $\sum 1/n$: terms go to zero, series diverges. Kills "terms to zero implies convergent".
- $f(x) = \sin(1/x)$ on $(0,1)$: continuous, no limit at 0.
- $f(x) = x\sin(1/x)$: continuous at 0 if you set $f(0) = 0$, not differentiable there.

- The indicator function of $\mathbb{Q}$: nowhere continuous. Kills almost every naive intuition about functions at once.
- $1/x$ on $(0,1)$: continuous, unbounded, not uniformly continuous.

Six objects. Between them they refute most of the false things a first-year student is tempted to believe, and they'll keep showing up in exams for two years, because examiners like them for exactly the same reason.

The habit generalises. When you meet a new area, ask early: what are this subject's monsters? In group theory it's the small non-abelian groups. In topology it's the topologist's sine curve and the long line.

Every experienced reader carries a private zoo, and the zoo is what stops them believing plausible false things.

READ SOMETHING TODAY

Take the definition you unpacked in Chapter 4 and build all four bank objects for it. Then negate the definition on paper, formally, flipping the quantifiers, and read the negation out loud in English. If the English version surprises you, you've just learned something the text never said.

6

Notation: Say It Out Loud

There's a specific, common, entirely fixable condition where a student isn't confused about the mathematics at all.

They're confused about the marks on the page.

They can't pronounce $\bigcup_{n=1}^{\infty} A_n$, so it enters their mind as a shape rather than a sentence, and shapes can't be reasoned with.

The test takes five seconds. Point at any displayed line in your textbook and read it aloud, in English, as a sentence. Stumble, and that line isn't available to you yet, however long you stare at it.

6.1 Notation has parts of speech

Mathematical notation is a language with a small grammar, and almost nobody gets taught it explicitly. Every symbol on the page is doing one of five jobs.

Relation (the verb): $= \; \leq \; \subseteq \; \in \; \equiv \; \rightarrow$ ← find this first

Operation (builds objects): $+ \; \times \; \cup \; \circ \; \sum \; \int$

Quantifier (says how many): $\forall \; \exists \; \exists!$

Connective (joins claims): $\wedge \; \vee \; \neg \; \Rightarrow \; \Longleftrightarrow$

Delimiter (groups): $(\,) \quad \{\,\} \quad |\,| \quad \langle\,\rangle$

Figure 6.1: The five jobs a symbol can do. Finding the relation first is what makes a long formula readable.

So here's the practical move: to read a long formula, find the relation first. A displayed line is a sentence, and the relation is its main verb. Everything left of it is one noun, everything right of it is another, and the line claims those two nouns stand in that relation.

Get the verb and you can unpack the nouns at leisure.

Take a line that looks intimidating:

$$\left|\sum_{k=1}^{n} a_k b_k\right|^2 \le \left(\sum_{k=1}^{n} a_k^2\right)\left(\sum_{k=1}^{n} b_k^2\right).$$

Find the verb. It's the $\le$ in the middle. Everything else is two nouns. The left one is "the square of the size of a sum of products". The right one is "the product of two sums of squares". The sentence says the first is no bigger than the second.

You've just read Cauchy–Schwarz out loud without doing any mathematics, and the formula has stopped being a wall.

> *A displayed formula is a sentence. Find the main relation, name the two sides in English, and only then look inside them. Reading a formula symbol by symbol is reading a sentence by pronouncing each letter.*

6.2 The same symbol means different things

Notation is overloaded, aggressively, and authors assume you'll disambiguate from context. The vertical bar alone does at least six jobs.

One rule resolves all six: look at the type of the thing inside. Which is why the type list from Chapter 4 pays off twice. If you know G is a group, $|G|$ can't be an absolute value, and the ambiguity evaporates.

The same overloading hits $\subset$ (proper subset to some authors, any subset to others), $\mathbb{N}$ (contains zero or doesn't, depending on where the author trained), log (base e in analysis, base 2 in computer science, base 10 in old engineering texts), and (a, b), which is an open interval, an ordered pair, a greatest common divisor, or an inner product, depending on who's writing.

None of this is sloppiness. There just aren't enough symbols. But it makes one habit compulsory.

> *Every book gets its own notation table. Conventions differ between authors, and assuming this book uses the last book's conventions is a reli-*

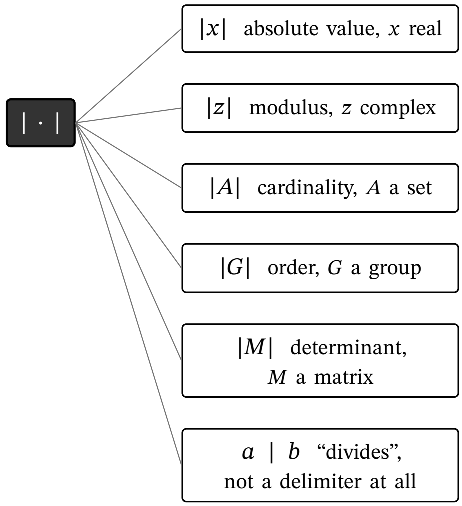

Figure 6.2: One symbol, six meanings. Context disambiguates, and context means *what kind of object is inside.*

> *able way to misread a theorem.*

6.3 Build the notation table

Two columns, front page of your notes for that book. Symbol on the left. On the right: how you say it, what type of object it is, where it was defined.

Symbol	Say it	Type / where
A^c	"A complement"	set; p. 12
$\overline{A}$	"closure of A"	set; Def. 2.7
A'	"A prime", derived set	set; Def. 2.9
$\|\|x\|\|$	"norm of x"	real ≥ 0; p. 31
$f\|_S$	"f restricted to S"	function; p. 44
$\liminf a_n$	"lim inf"	extended real; Def. 3.4

Look at $\overline{A}$ and A' in that table. In a lot of analysis texts one is the closure and the other is the set of limit points, and they aren't the same set. A student without the table will conflate them inside a week and then be quietly wrong in every proof after that, in a way that's very hard to spot from the inside.

Fifteen minutes to build for a whole book. Saves hours.

Add to it every time a symbol makes you pause, including symbols you think you know. Especially those.

6.4 Dummy variables are free, and this matters

In $\sum_{k=1}^{n} a_k$, the letter k means nothing outside the sum. It's a placeholder. Rename it j or i or anything else not already in use and nothing changes. Same for the bound variable in $\forall x\, P(x)$ and the integration variable in $\int_0^1 f(t)\,dt$.

Sounds like a triviality. It's the source of two very common reading errors.

First, students see $\sum_k a_k$ and $\sum_j a_j$ in adjacent lines and think something changed. Nothing changed.

Second, and worse, students fail to rename when they should. When a proof already has an ε in play and you introduce a second one, reusing the letter merges two different quantities in your head, and the proof becomes incomprehensible for reasons that have nothing to do with the mathematics.

Rename freely. Paper is cheap.

> *If two different quantities share a letter in your working, rename one immediately. Most "I can't follow this proof" moments are really "two objects are wearing the same name" moments.*

6.5 Say it out loud: a starter list

Build the habit on the symbols you meet daily. The full table lives in Appendix E. These are the ones students most often can't pronounce, which turns out to be a decent proxy for the ones they can't use.

You see	You say
$\forall\varepsilon > 0$	"for every positive epsilon"
$\exists! x$	"there is exactly one x"
$f : A \to B$	"f maps A into B"
$x \longmapsto x^2$	"x goes to x squared"
$A \subseteq B$	"A is contained in B"
$\bigcup_{n=1}^{\infty} A_n$	"the union over n from one to infinity of A n"
$\{x \in \mathbb{R} : x > 0\}$	"the set of real x such that x is positive"
$a_n \to L$	"a n tends to L"
lim sup	"lim sup", the eventual upper bound
$\sim$	depends: "is asymptotic to", "is equivalent to"

That last row isn't a joke. When a symbol's pronunciation depends on context, write both meanings in your table and note which one this book uses. A symbol you can say two ways and haven't decided between is a symbol you can't use in a proof.

6.6 A five-minute drill that fixes this permanently

Open your textbook at any page with displayed formulas. Read every displayed line out loud as an English sentence. Don't try to understand anything. Just pronounce.

Five minutes a day for a fortnight and the symbol layer disappears. You stop seeing marks and start seeing sentences, and it feels exactly like the moment a foreign language stops being a decoding exercise.

Students who do this tell me the same thing every time. The mathematics didn't get easier. The page stopped fighting them, and that turns out to have been most of the difficulty.

READ SOMETHING TODAY

Start your notation table today, in the front of your notes for one specific book. Seed it with every symbol from the chapter you're reading, including the ones you're sure about. Then read that chapter's displayed formulas aloud and count how many you stumble on. That number, not the page count, tells you how much of the chapter is actually available to you.

Part III

Statements

7

Reading a Theorem Before Its Proof

Here's a habit that changes how much you get out of a textbook, and it costs four minutes.

When you reach a theorem, don't read the proof. Read the statement, work on it, and only then look down.

Almost everyone does the opposite. Eyes skim the statement, and the real attention starts at the word *Proof*. That order guarantees you're verifying an argument for a claim you haven't understood yet, which is like checking somebody's arithmetic without knowing what they were computing.

Four minutes of statement work makes the proof two to five times faster to read. I've watched this in tutorials for years and the effect isn't subtle.

7.1 Split it

Every theorem is hypotheses plus conclusion. Draw the line physically. Write **H** beside each hypothesis and **C** beside the conclusion, and break the compound hypotheses into numbered pieces.

Look at what the split turns up. The sentence looked like one hypothesis and one conclusion. It's four and three.

Students who haven't split it will, in an exam, prove C1 and stop, having genuinely believed they answered the question.

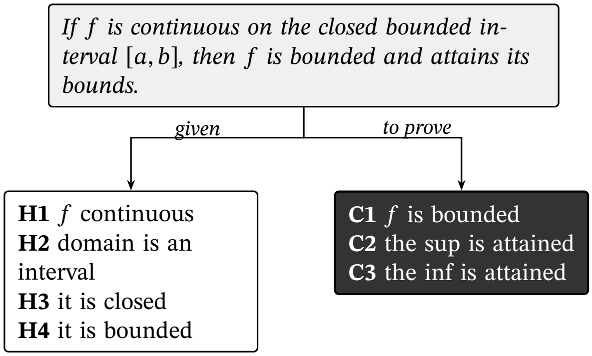

Figure 7.1: Splitting a statement. The conclusion is one claim; the hypotheses are usually three or four claims wearing one sentence.

7.2 Restate it in your own words

Close the book and write the theorem in English, without symbols, the way you'd explain it to someone in your class who missed the lecture. Then open the book and compare.

The gap between your version and the printed one is the most useful diagnostic in this book. Nine times out of ten it's a lost quantifier or a dropped hypothesis, and you've just found a hole in your understanding for two minutes' work instead of finding it in March.

7.3 The hypothesis audit

Now the question that separates reading from processing. Why is each hypothesis there?

Every hypothesis in a published theorem is load-bearing. Authors don't add conditions for decoration; if one could be dropped, a referee would have made them drop it. So each is stopping some specific disaster, and your job is to find the disaster.

The technique is the drop-one test. Take the hypotheses one at a time, delete one, and hunt for a counterexample to what's left.

Do that once for a theorem and you'll never again wonder why the statement says "closed and bounded" instead of just "interval". You'll also pick up three counterexamples for your monster page on the way.

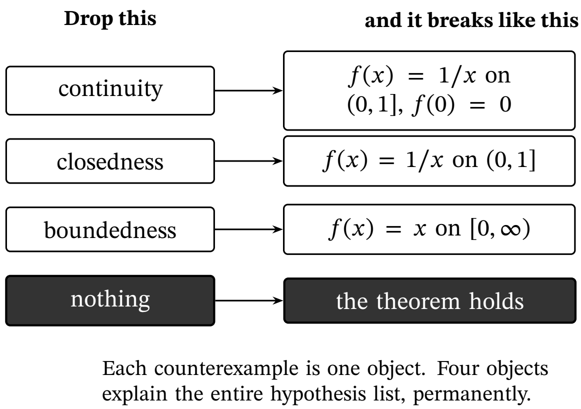

Figure 7.2: The drop-one test on the extreme value theorem. Each deleted hypothesis has a counterexample waiting.

> *A hypothesis you can't break the theorem without is a hypothesis you haven't understood. Find the counterexample that each one is holding back.*

Can't find a counterexample after five minutes? Don't grind. Either the hypothesis is there for a technical reason the proof will reveal, in which case read the proof and watch for the line where it gets used, or it's genuinely redundant, which happens occasionally in textbooks and is a small discovery worth making.

That second habit, *watch for the line where the hypothesis gets used*, is worth carrying into every proof you read. It turns a proof from a wall of steps into a story with a cast, and you know the cast in advance.

7.4 Test it on your example bank

You built four objects for each definition in Chapter 5. Now spend them.

Run the theorem on the trivial example. Run it on the pathological one. Both should satisfy it, since the theorem is true, but watching *how* they satisfy it tells you what the theorem is really saying.

The pathological case is the informative one. If a theorem holds even for the indicator function of the rationals, the theorem isn't about niceness. If it visibly fails for that function, then that function violates a hypothesis, and finding which one is a free comprehension check.

7.5 Predict the proof from the statement's shape

Before you read the proof, predict its opening line.

You can do this reliably, because the logical form of a statement dictates the first move of any proof of it. This is the highest-leverage table in Part III.

If the statement says	**the proof starts**
for all x, $P(x)$	"Let x be arbitrary."
there exists x with $P(x)$	"Take $x = \dots$"
$P \Rightarrow Q$	"Assume P."
$P \iff Q$	two proofs, one each way
exactly one x	existence, then $x_1 = x_2$
$A = B$ for sets	$A \subseteq B$ and $B \subseteq A$
true for all $n \in \mathbb{N}$	base case, then induction step
no x satisfies P	"Suppose such an x existed."

Figure 7.3: Statement form determines the proof's opening move. You can write the first line of most proofs before reading them.

Cover the proof with your hand and write its first line yourself.

Match? Then you understood the statement's logical shape, and the proof will read quickly, because you're following a route you already sketched.

No match? Stop and find out why. The mismatch is telling you that you misread the statement's structure, which is a much bigger problem than not following a step.

This habit also transfers straight to writing your own proofs, which is the companion volume's job. The forms in Figure 7.3 are the skeletons you'll reach for when the theorem is yours to prove and nobody has written the proof yet.

7.6 Four minutes, in order

To put it together:

1. Split hypotheses from conclusion, numbering each. (1 min)
2. Restate the whole thing in English, from memory. (1 min)
3. Ask what each hypothesis is stopping. (1 min, or flag it)
4. Predict the proof's opening move from the statement's form. (1 min)

Then read the proof. Chapter 10 is about doing that at the right depth.

READ SOMETHING TODAY

Find the most recent theorem in your notes. Cover the proof. Do the four steps on paper, then uncover it and check your prediction of the opening line. Right or wrong, you now know something specific about your own reading, which is worth more than having read the proof.

8

Testing a Statement Against the Edges

Textbooks are mostly right, and that breeds a bad habit. Students read a theorem and believe it immediately, because it's printed.

Belief is free and teaches nothing. What teaches is the five minutes you spend trying to break the statement before you accept it.

You'll fail to break it. It's a theorem. But the failure is informative in a way that agreement never is, because every attempt that dies tells you which hypothesis killed it, and after four attempts you understand the statement from the inside.

8.1 The test battery

Run these in order. Most take seconds.

Probe 2 deserves a note, because degenerate cases are where mathematics hides its conventions. Is the empty set compact? Is the empty sum zero or undefined? Is the empty product one? Is $0! = 1$?

Every one of those has an answer that was chosen to make theorems come out clean, and noticing the choice is noticing something real about how the subject was built.

Probe 6 is the least practised and the most interesting. If a theorem concludes "f is bounded", ask whether it could have concluded "f is bounded and attains its bounds". Sometimes it couldn't, and the reason is a counterexample worth knowing. Sometimes it could, and you've just noticed the author stated a weaker version for a reason that shows up two pages later.

Before believing a statement, try:

1. The trivial object. Constant function, one-point set, $n = 1$.

2. The empty case. Empty set, empty sum, $n = 0$.

3. The extremes. Very large, very small, zero, negative.

4. Your monsters. The pathological objects you collected.

5. Drop a hypothesis. Does it still hold? Why not?

6. Strengthen the conclusion. Could it claim more?

Figure 8.1: The test battery. Six probes, most of them instantaneous, applied before you believe anything.

8.2 Worked: testing Bolzano–Weierstrass

Theorem 8.1 (Bolzano–Weierstrass). *Every bounded sequence of real numbers has a convergent subsequence.*

Don't read a proof. Test it.

Trivial object. $a_n = 7$ for all n. Bounded, yes. Convergent subsequence? The whole sequence converges to 7. Passes, and notice the subsequence was allowed to be the entire sequence. Useful: "subsequence" here doesn't mean "proper subsequence".

The oscillating case. $a_n = (-1)^n$. Bounded. Doesn't converge. But the even-indexed terms are all 1, and that subsequence converges. Passes.

And this is the example that shows why the theorem says *subsequence*. Without that word it's false.

Extremes. What about a bounded sequence that hits every rational in $[0, 1]$, in some enumeration? Wildly non-convergent, and the theorem still promises a convergent subsequence.

That should feel surprising. Sit in the surprise, because it's the whole content of the theorem: however badly a bounded sequence behaves, some thread inside it settles down.

Drop the hypothesis. $a_n = n$. No subsequence converges, because every subsequence is unbounded. So boundedness is doing real work, and now you know what it stops.

Strengthen the conclusion. Could it say "and the limit is unique"? No: $(-1)^n$ has subsequences converging to 1 and to -1. So uniqueness is off the table, and you've just discovered for yourself why the subject later needs lim sup and lim inf.

Five probes, six minutes, no proof read. You know what the theorem says, why each word is there, and roughly what shape its proof has to be, since something has to produce the thread that settles.

> *Testing a statement isn't a check on the author. It's how you turn a printed sentence into something you've personally handled. Believe the theorem afterwards, not before.*

8.3 The counterexample instinct

Students are often uneasy about trying to break a printed theorem, as though doubting it were rude.

Get over that quickly. Counterexample hunting is the primary mode of understanding in mathematics, and it's how research actually proceeds.

The instinct also protects you against a specific and very common disaster: the false statement you half remember. Under exam pressure, students routinely write things like "every continuous function is differentiable", or "a bounded sequence converges", or "if $\sum a_n$ converges then $\sum a_n^2$ converges".

None of those survives contact with someone who habitually tests statements against $|x|$, $(-1)^n$, and $a_n = (-1)^n/\sqrt{n}$. The habit is a defensive weapon.

8.4 Sharpness: is this the best possible version?

Once a theorem survives your probes, ask what a researcher would ask. Is it *sharp*? Can the hypotheses be weakened or the conclusion strengthened at all, or is this the exact boundary?

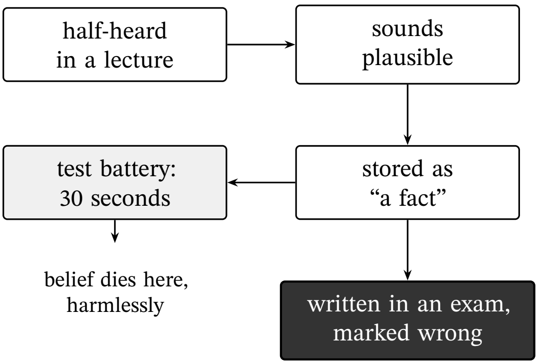

Figure 8.2: Where false beliefs come from, and where the test battery intercepts them.

A theorem is sharp when each hypothesis has a counterexample showing you can't do without it, and the conclusion can't be improved. So the drop-one test from Chapter 7 plus probe 6 above is precisely a sharpness audit.

Run it and a theorem stops being a fact and becomes a landscape. Here's the region where this is true, and here, one step outside, is where it fails.

Textbooks rarely say “and this is sharp” out loud, so most students never notice the question exists. Noticing it is one of the clearest markers of a mature reader.

8.5 When your counterexample seems to work

Every so often you'll build something that appears to break a printed theorem. Don't email the author yet. In order of likelihood:

1. You misread a quantifier. Check the order again.
2. Your object fails a hypothesis you didn't notice, often an implicit one like “non-empty” or “$n \geq 1$”.
3. You're using a different definition of a term than the book is. Extremely common with $\subset$, $\mathbb{N}$, and “increasing” versus “strictly increasing”.
4. The book has a typo. This does happen, maybe once per textbook.
5. The book is wrong. Rare, and worth writing up.

Work down the list in order. Whatever the outcome, the twenty minutes you spend resolving it teaches you more than the twenty pages you'd otherwise have read, because you're doing mathematics now instead of receiving it.

READ SOMETHING TODAY

Take any theorem from this week's lectures and run all six probes. One line per probe. If probe 5 defeats you, if you can't work out what a hypothesis is stopping, write that hypothesis on a card and take it to your next tutorial. It's the best question you can ask a lecturer, and they'll enjoy answering it.

9

The Neighbourhood of a Theorem

No theorem stands alone.

Around every statement sits a small neighbourhood of related ones: its converse, its contrapositive, the version with a hypothesis weakened, the version with the conclusion strengthened, the special case that's famous in its own right.

Reading a theorem well means locating it in that neighbourhood, because the neighbourhood is where the exam questions come from and where the misunderstandings live.

9.1 The four statements

Given "if P then Q", you can build four sentences, and students conflate them constantly.

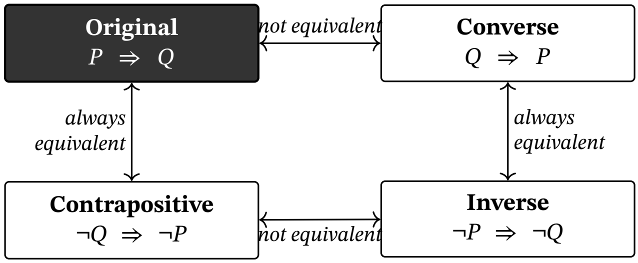

Figure 9.1: The four statements. Only the diagonal pairs are logically equivalent, and the converse is where almost every student error lives.

Two facts from that picture are worth burning in.

A statement and its contrapositive are the same statement. Not similar. The same. That's why a proof by contraposition is a complete proof and not a weaker substitute, and why flipping to the contrapositive when you're stuck costs nothing and sometimes makes the whole thing obvious. "If n^2 is even then n is even" is awkward head-on and immediate as "if n is odd then n^2 is odd".

A statement and its converse are unrelated. Both can be true, either can be true alone, both can be false. Every time a student writes "since all squares are non-negative, every non-negative number is a square" they've assumed a converse, and it's the single most common logical error in undergraduate work.

9.2 Always ask about the converse

Finish a theorem, ask immediately: is the converse true? Three possible answers, each worth having.

Yes, and the book proves it. Then the theorem is really an "if and only if", and the two halves usually get proved separately. Note which half is hard. That asymmetry is often the interesting content.

No, and here's a counterexample. Best case. Write it down. This is the common one, and the counterexample is usually small. Continuity doesn't imply differentiability: $|x|$. Terms going to zero doesn't imply the series converges: the harmonic series. One object each, settling a whole class of confusions.

Unknown, or hard. Occasionally the converse is a real theorem or an open problem. Knowing that a question you can state in one line is hard is itself an education in the shape of the subject.

> *Every theorem you read gets one written line after it: the converse, and whether it holds. Ten seconds of writing prevents the commonest error in student mathematics.*

9.3 The generality ladder

Second axis of the neighbourhood: how general is this? Almost every theorem sits on a ladder, with special cases below it and generalisations above.

Locating a theorem on its ladder does three things. It tells you what the proof will need, since a more general statement usually forces a more abstract argument. It tells you what you can safely quote later. And it's how

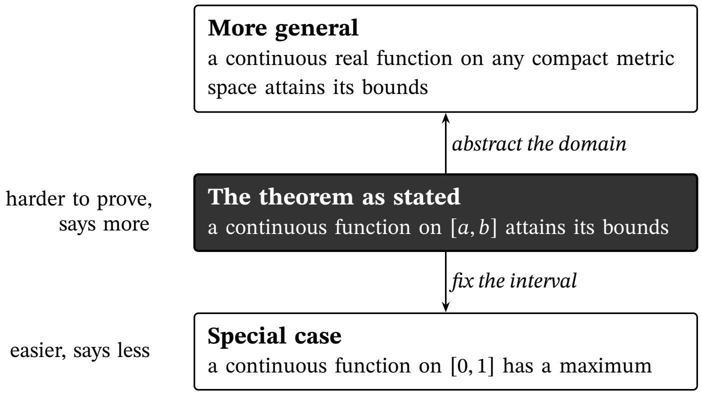

Figure 9.2: One theorem, three rungs. Reading a result means knowing which rung you are on and what the rungs above and below look like.

a subject stops feeling like an unrelated pile of results and starts feeling like a few ideas at different altitudes.

Then when you meet the compact-metric-space version in a later course, recognising it as your old friend from first year, standing one rung higher, is worth a great deal. Students who never built the ladder meet it as a new and frightening object.

9.4 Where does it sit in the chapter?

Third axis, local: what does this theorem depend on, and what depends on it? You built that graph in Chapter 2. Use it.

A theorem with many arrows out of it is a tool. Learn its statement cold even if you never reread its proof. A theorem with no arrows out is either the chapter's destination or an aside, and knowing which changes how much attention it deserves.

Watch for the theorem that gets cited in a later chapter, too. When a proof on page 130 says "by Theorem 3.6", that's the author telling you Theorem 3.6 mattered. Go back and mark it.

Textbooks have no highlighters, so citation frequency is the only emphasis they've got, and it's invisible unless you're watching for it.

9.5 A worked neighbourhood

Take the mean value theorem: if f is continuous on $[a,b]$ and differentiable on (a,b), then there's $c \in (a,b)$ with $f'(c) = \frac{f(b)-f(a)}{b-a}$.

Converse? There's no clean one. The conclusion is an existence claim, so "if such a c exists then f is differentiable" is plainly false, and building the counterexample takes a minute and is a good exercise.

Contrapositive? If no such c exists, then f fails continuity on $[a,b]$ or differentiability on (a,b). Occasionally useful, and a good check that you've got the hypotheses straight.

Down the ladder. Rolle's theorem is the special case $f(a) = f(b)$. Most books prove Rolle first and derive the mean value theorem from it, which is the ladder being climbed in front of you.

Up the ladder. Cauchy's mean value theorem generalises to two functions, and Taylor's theorem with remainder is the next rung again. So when Taylor turns up in a later chapter looking unmotivated, the ladder is where its motivation went.

Hypothesis audit. Why continuity on the closed interval but differentiability only on the open one? Because the conclusion is about an interior point, and demanding differentiability at the endpoints would throw out perfectly good functions like $\sqrt{x}$ on $[0,1]$.

That asymmetry looks fussy right up until you ask the question, and then it's obviously right.

Fifteen minutes of neighbourhood work, and the mean value theorem has stopped being a sentence to memorise. It's a place you know.

READ SOMETHING TODAY

Pick a theorem you'll be examined on. Write its converse and decide whether it's true. Write its contrapositive. Name one rung up and one rung down the generality ladder. Half a page. Do that for the five main theorems of a course and you'll walk into the exam holding the course's structure instead of its contents.

Part IV

Proofs

10

Three Ways to Read a Proof

Students treat reading a proof as one activity with a binary outcome. You got it or you didn't.

It's three activities. They cost wildly different amounts, and most of the frustration in reading mathematics comes from attempting the wrong one without knowing you're doing it.

10.1 The three readings

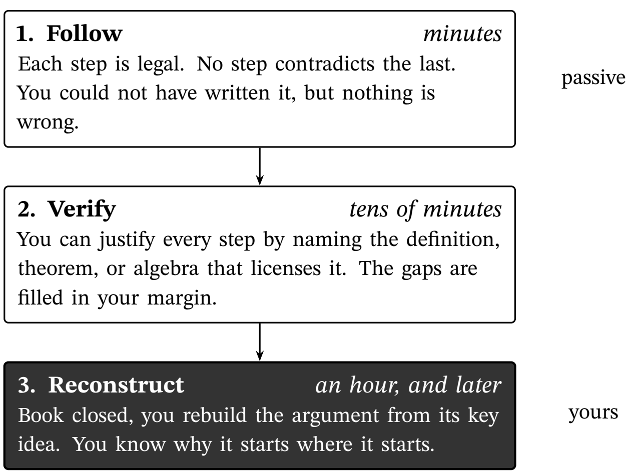

Figure 10.1: Three readings of the same proof. Each is a legitimate stopping point; the mistake is not knowing which one you are attempting.

Following

You read line by line and nothing objects.

This is where most students stop, and it produces that familiar experience: understanding a proof in the library and being unable to reproduce a line of it that evening. Following is real, but shallow. It checks local consistency and nothing else.

It's the right stopping point for results you only need to quote, which is most results in any book.

Verifying

Now you interrogate each step. Why is this equality true? Which theorem licenses this? Where did that $\varepsilon/2$ come from?

You write the missing algebra in the margin, and you find, on average, two or three steps per proof that you can't justify.

Those steps are the proof's actual content.

Reconstructing

You close the book and rebuild it. Not from memory of the words, from memory of the idea: this proof works by contradiction, the key move is to consider the least element, everything else follows.

Manage that and the proof is yours, and it'll still be yours in June. Chapter 12 is entirely about this test.

> *Following is not understanding. The gap between "I followed it" and "I could rebuild it" is where every exam disappointment lives.*

10.2 Choosing a depth on purpose

Same rule as the gears in Chapter 3, applied to proofs. Decide before you start, and write the decision at the top of the page.

- **Follow** the proofs of results you only need to cite.
- **Verify** the proofs of anything the problem sets touch.
- **Reconstruct** the four or five proofs a course is built on, plus every proof whose technique recurs.

That last clause matters more than the count.

Some proofs are worth reconstructing not because the result is important but because the *technique* will reappear twenty times. The $\varepsilon/2$

trick. Nested interval bisection. The pigeonhole set-up. The "consider the smallest counterexample" opening.

Learn the technique once, properly, and every later proof that uses it drops from a verify-level read to a follow-level one. That's the main way experienced readers get fast. Not by reading faster. By having already reconstructed the twenty techniques everything else is assembled from.

10.3 Worked: one proof at three depths

Theorem 10.1. *$\sqrt{2}$ is irrational.*

Proof. Suppose $\sqrt{2} = p/q$ with p, q integers, $q \neq 0$, and the fraction in lowest terms. Then $p^2 = 2q^2$, so p^2 is even, so p is even. Write $p = 2r$. Then $4r^2 = 2q^2$, so $q^2 = 2r^2$, so q is even. But then p and q are both even, contradicting lowest terms. □

Reading 1, follow (ninety seconds). Every sentence follows from the last. No objections. Done.

Reading 2, verify (fifteen minutes). Now the interrogation:

- "The fraction in lowest terms": can we always assume that? Yes, by cancelling common factors. But that's a fact about the integers, and it's worth noticing we used it.
- "p^2 is even, so p is even": not obvious, and not the same as the definition. It needs its own proof, by contraposition: if p is odd, $p = 2k+1$, so $p^2 = 4k^2+4k+1$ is odd. Most textbook presentations quietly assume this. Write it in the margin.
- "$4r^2 = 2q^2$, so $q^2 = 2r^2$": division by 2, legal in the integers here because both sides are even. Fine, but check it.
- The contradiction: we assumed lowest terms and derived a common factor. So the contradiction is with the assumption, not with anything about $\sqrt{2}$. Notice where the pressure got applied.

Three of those four went straight past you at follow depth. The second one in particular is a whole lemma travelling incognito.

Reading 3, reconstruct (thirty minutes, then again next week). Book closed. What's the idea? *Assume it's rational in lowest terms and show both numerator and denominator have to be even.*

One sentence, and it generates the whole proof. From it you can rebuild every line, and you can now answer the question that separates owning a proof from remembering one.

Does the same argument show $\sqrt{3}$ is irrational? Yes, with "even" replaced by "divisible by 3" and the little lemma adjusted. Does it show $\sqrt{4}$ is irrational? No, and finding the exact line where it breaks is the best five minutes you'll spend on this proof.

> *The test of a reconstruction isn't whether you can repeat the proof. It's whether you can say what happens when you change the theorem.*

10.4 Read structurally before sequentially

One more habit, and it's close to a superpower on long proofs.

Before reading a proof line by line, read only its *structure words*: the Suppose, the Let, the It suffices to show, the Conversely, the Hence, the By Theorem 3.2. Skip everything else on the first pass.

A three-page proof usually has six or seven of those, and they carry its whole architecture. Reading them first takes forty seconds and turns the sequential read from an unmapped slog into a walk along a route you've already seen.

That's Chapter 11's subject, and it's the technique that makes long proofs tractable.

READ SOMETHING TODAY

Take a proof you read this week and decide, in writing, which of the three depths you actually reached. Then push it one level deeper. If you followed it, verify it: interrogate every step and count how many you can't justify. That count is the honest measure of what your first read was worth.

11

Finding the Skeleton

Every proof has a skeleton: three to six moves that carry the argument. Everything else is connective tissue, algebra, and bookkeeping.

A reader who sees the skeleton can hold a three-page proof in their head. A reader who doesn't is trying to memorise ninety sentences, which can't be done and shouldn't be attempted.

Extracting the skeleton is mechanical. Five minutes, and it's the most useful thing in Part IV.

11.1 Structure words are the skeleton's joints

Proofs get written with a small, fixed vocabulary of signposts, and each one announces a structural move. Learn the list and proofs stop being opaque.

When you see	**The move is**
Let x be arbitrary	universal claim opening
Suppose, for contradiction	the whole proof is indirect
It suffices to show	the goal has just changed
By Theorem 3.2	an external tool is entering
Conversely	second half of an iff
Without loss of generality	a symmetry is being used
Define $g(x) = \ldots$	the key construction, often the idea
Base case / Inductive step	induction
Case 1, Case 2	exhaustive split
Hence, therefore, so	routine consequence, low information

Two of those repay attention.

It suffices to show is the most important phrase in mathematical writing and the most skimmed. It means the author has swapped the goal for a different goal, and if you miss it you'll spend the rest of the proof wondering why they're proving the wrong thing. Underline it. Every time.

And *Define* $g(x) = \ldots$ is where proofs hide their ideas. When a proof invents an auxiliary object out of nowhere, that invention is almost always the entire content and the rest is verification. Mark it, then ask the question that matters: how would anyone have thought of that?

11.2 Annotate in the margin

Read the proof once for signposts only, writing a structural label beside each. Figure 11.1 shows what a marked-up proof looks like.

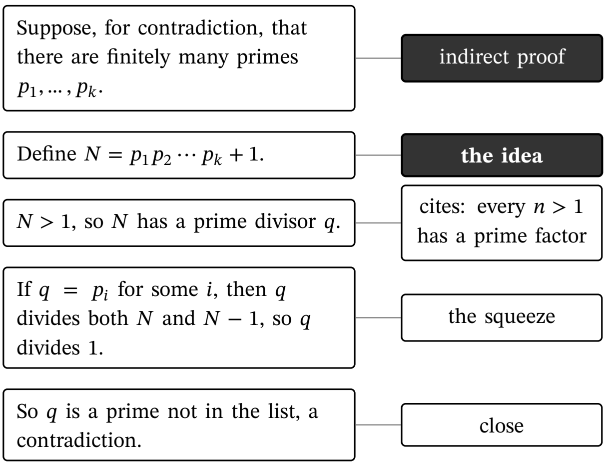

Figure 11.1: A proof with its skeleton in the margin. Six labels, and the three pages become one idea.

Now read only the right-hand column. Indirect proof, build N from the list plus one, use the prime-factor fact, squeeze, contradiction.

That's Euclid's proof, complete, in five phrases, and you'll have it for the rest of your life.

> *If you can't write the skeleton in five phrases, you haven't finished reading the proof. If you can, you can stop.*

11.3 Find the load-bearing step

In every proof there's one step where the work happens. Everything before it is set-up, everything after it is consequence.

Finding it is the difference between a proof you can compress and a proof you have to carry whole. It's usually one of these:

- an object invented from nowhere (the auxiliary function, the clever set, the specific ε);
- a choice with a strange-looking value ($\varepsilon/2$, $\delta = \min(1, \varepsilon/5)$, $N = \lceil 1/\varepsilon \rceil$);
- a re-goaling, announced by *it suffices*;
- the single citation of a big theorem.

Everything else, and I mean the great majority of most proofs, is algebra you could have done yourself given the idea.

Those strange-looking constants deserve their own note, because they cause more distress than anything else in a first analysis course.

When a proof opens with "let $\delta = \min(1, \varepsilon/5)$", nobody derived that before writing it. The author did the proof with an unknown δ, found at the end that they needed $\delta \le 1$ and $5\delta \le \varepsilon$, then went back and wrote the answer at the top.

The strangeness is an artefact of the write-up. It isn't in the mathematics.

> *Odd-looking constants are back-filled. They were computed at the end of the proof and moved to the front. Read them as "some δ small enough", carry on, and check at the end that the choice works.*

That one piece of information removes what is, in my experience, the biggest single source of despair in first-year analysis. Students believe the author saw $\varepsilon/5$ by inspiration and conclude they'll never be able to do this.

Nobody saw it. It was reverse-engineered, and you're allowed to reverse-engineer too.

11.4 The four shapes

Most proofs in the first two years have one of four global shapes, and spotting the shape from the first two lines tells you what to expect and where the proof has to end up.

Here's the practical use. When you get lost in the middle of a proof, ask which shape you're in.

Indirect? Then the destination is an absurdity and every step is heading there, which makes an otherwise baffling line like "we now have $q \mid 1$"

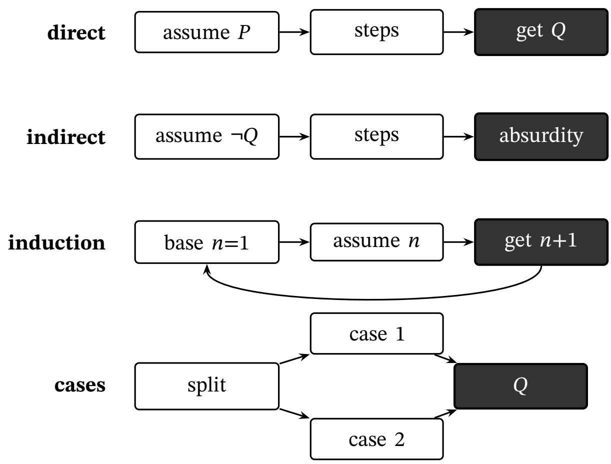

Figure 11.2: Four global shapes. Identify which one you are in from the opening line, and you know what the last line must be.

suddenly legible as the arrival. Induction? Then the inductive hypothesis has to get used somewhere, so go and find where.

Being lost inside a proof is nearly always being lost about the shape, not about a step.

11.5 Compressing a long proof

For a proof that runs pages, do this on a separate sheet:

1. Read only structure words. Write them in a list. (2 min)
2. Name the shape. (30 sec)
3. Find and star the load-bearing step. (2 min)
4. Write the skeleton in five phrases. (2 min)
5. Now read the whole thing, sequentially, with the skeleton beside you.

Seven minutes of preparation, and the sequential read that follows is a different experience. You're checking a route instead of exploring an unmapped one.

READ SOMETHING TODAY

Take the longest proof in your current chapter. Don't read it. Extract the structure words, name the shape, write the five-phrase skeleton.

Then read it properly and see how much of the difficulty was just not knowing where you were going.

12

The Reconstruction Test

Everything in Part IV has been building to one test, and it's the only honest measure of whether you've read a proof.

Close the book. Write the proof. Compare.

That's it. Ten minutes, uncomfortable the first several times, and it's the difference between students who are surprised by their exam results and students who aren't.

12.1 Why recognition lies

When you reread a proof it feels familiar, and your brain files familiarity as understanding.

That isn't a character flaw, it's how memory works, and it's why rereading is the most popular and least effective study method in existence.

Reconstruction removes the cue. With no page in front of you there's nothing to feel familiar about, so what you actually possess becomes visible. The first time a student does this with a proof they were sure they knew, the result is usually a shock. Three lines, then blank.

That shock is the whole value. It arrives in your room in October instead of in the exam hall in May.

> *Familiarity is not knowledge. The only evidence that you own a proof is that you produced it with the book shut.*

12.2 The three levels of reconstruction

You don't have to rebuild everything in full detail. Three levels, each a legitimate target.

Level 1: the statement
Write the theorem exactly, hypotheses and all. Two minutes. Everything you are examined on needs this.

↓

Level 2: the skeleton
The five phrases, the shape, the load-bearing step. Five minutes. This is the working target for most proofs.

aim here by default

↓

Level 3: the full proof
Every line, with justifications. Twenty minutes or more. Reserve it for the course's core results and recurring techniques.

Figure 12.1: Three levels of reconstruction. Most results need level 2; a handful need level 3.

Level 1 pays out far beyond its cost. A student who can state every theorem in a course precisely, and nothing else, will outperform a student who half-knows several proofs and misremembers the hypotheses.

Statements first. Always.

12.3 The gap list

Reconstruction produces something more valuable than the reconstruction itself. A list of exactly where you failed.

Keep it. Every time you stall, write one line describing the stall. Don't look anything up until you've finished the whole attempt. Then open the book and write, beside each stall, what you were missing.

Gap lists have a striking property. They repeat.

Across a whole course, a student's stalls cluster into four or five recurring items. Almost always the same missing lemma, the same unfamiliar technique, the same definition that never got properly unpacked.

Once you see the cluster, a vague feeling of struggling with the course turns into a short, specific list of things to fix. And the list is usually shorter than you feared.

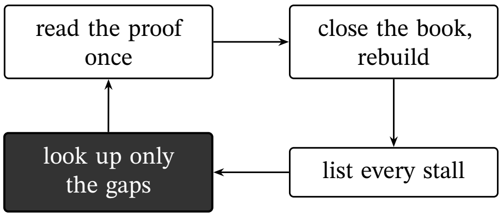

Second time round, start at "close the book". Three passes spread over two weeks beats ten passes in one evening.

Figure 12.2: The reconstruction loop. The gap list is the output that matters; the rebuilt proof is a by-product.

12.4 Spacing the attempts

Do the reconstruction three times. Same day, three days later, two weeks later. Each attempt runs faster than the last, and the third usually takes four minutes and works.

This is retrieval practice, the best-evidenced study technique there is. In mathematics it has a particular advantage: you're rebuilding an argument rather than recalling a fact, so each attempt strengthens the connections between steps, which is exactly what you need when an exam asks you to adapt a technique instead of repeating a proof.

Schedule the attempts in whatever system you already use. No system? Write the date of the next attempt at the top of the sheet. The system matters far less than whether the second attempt happens at all, which for most students it doesn't.

12.5 Reconstruction as a source of questions

There's a quieter benefit. Rebuild a proof, get stuck at the same place twice, and you've found a genuine question. Genuine questions are the currency of tutorials.

Compare two students at office hours. One says "I don't understand Theorem 4.6." The other says "I can rebuild 4.6 up to the point where we choose $\delta = \min(1, \varepsilon/5)$, and I don't see where the 1 comes from."

The second gets a two-minute answer that fixes the problem permanently. The first gets the proof read out to them again, which they won't remember either.

> *A reconstruction that fails produces a precise question. A reread that fails produces a vague complaint. Precise questions get answered.*

12.6 What to reconstruct, and what to let go

You can't reconstruct everything, and trying is how students burn out in week five. Reconstruct:

- every theorem *statement* in the course (level 1, all of them, no exceptions);
- the skeletons of everything the problem sets touch (level 2);
- the four or five results the course is built on, plus every proof whose technique recurs (level 3).

Everything else you're allowed to follow once and let go. Write "followed, not reconstructed" in the margin and move on with a clear conscience.

The point of the label is that in May you'll know which category each proof is in, instead of discovering it under exam conditions.

READ SOMETHING TODAY

Pick the most important proof you've met this term. Read it once, carefully. Close the book. Rebuild it on paper without looking, writing down every point where you stall. Then open the book and fill the gaps. Put the sheet somewhere you'll find it in three days, and do it again. That second attempt is the one that does the work.

13

Stuck: A Protocol

Being stuck is the normal working state of everyone who does mathematics.

That sentence is true, well known among professionals, and almost never said to students, who therefore experience being stuck as evidence of unsuitability instead of evidence that they're doing the thing correctly.

What separates people isn't how often they get stuck. It's what they do in the next twenty minutes. And most students do the worst available thing, which is to reread with more determination, because nobody ever handed them a list of alternatives.

Here's the list.

13.1 Name the stuck first

Six kinds of stuck. Different cures. Applying the wrong cure wastes the evening, so spend thirty seconds working out which one you're in before you do anything.

The last one needs saying loudly.

Where did this idea come from? isn't a failure of comprehension. It's a research-level question, it often has no good answer, and a proof can be completely understood without it.

Students routinely refuse to move past a proof because they can't see how anyone would have thought of the auxiliary function, treating a question no textbook answers as a personal deficiency. Write it in the margin, note that it's unresolved, and go on. The answer, when it comes, usually arrives months later from having seen the same trick three more times.

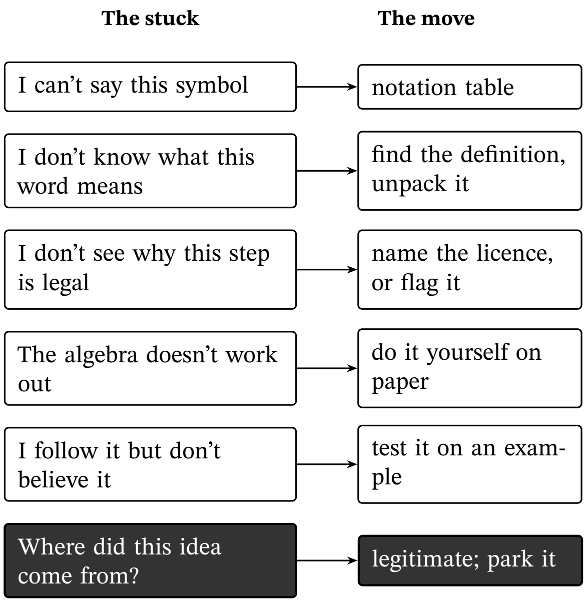

Figure 13.1: Six kinds of stuck, six different actions. The wrong cure is worse than no cure, because it feels like progress.

13.2 The protocol

When stuck, run this. It's built so that every branch ends in an action, and no branch ends in rereading.

The twenty-minute rule

Twenty minutes on one obstacle is the limit. After that the returns go negative. You're not thinking any more, you're anxious, and anxiety reads the same paragraph forever.

At twenty minutes, mark the spot with a symbol you use consistently, and read on.

That feels like giving up. It isn't. Mathematics texts are full of forward references and delayed explanations, and a startling proportion of obstacles dissolve two pages later when the author finally supplies the example they were saving. Reading on is a technique, not a retreat.

> *Twenty minutes, then mark it and move. The obstacle you skip is often explained by the text you haven't reached yet.*

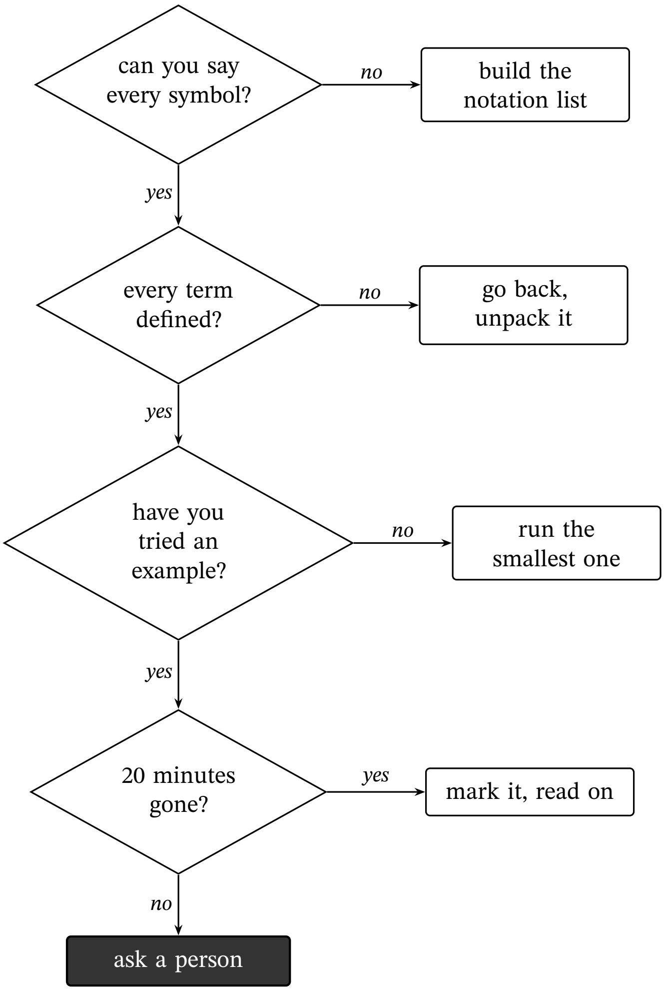

Figure 13.2: The stuck protocol. Every path terminates in something you do, and one path terminates in stopping, which is also something you do.

Marking, and coming back

Use one consistent mark. Mine is a small box in the margin, filled in when it's resolved.

At the end of a session, list the unresolved boxes on a single page. Your open questions sheet. Look at it at the start of the next session.

Two things happen to that list, both good. Some items resolve themselves after sleep or after a later section, and you cross them off having done nothing. The rest pile up into exactly the set of questions worth taking to a tutorial, and they're specific questions, which get useful answers.

13.3 Escalation, in order

When the protocol runs out, escalate. Order matters, because each step costs more of somebody's time than the last.

1. **A different book.** Enormously underrated. Every subject has three or four standard texts, and the treatment that defeats you in one is often transparent in another. Rudin's terseness and Abbott's patience cover much of the same first-year analysis; if one isn't landing, the other is in the library.
2. **The examples again.** Not the same example, a different one. Nine times out of ten the obstacle is that you've been reasoning about an object you can't picture.
3. **A peer.** Explaining your stuck to another student forces you to make it precise, and that alone resolves it surprisingly often, before they've said a word.
4. **The lecturer.** Bring the precise question from your open list, not "I don't understand section 4". Show what you tried.
5. **Leave it.** Genuinely. Some things only become clear after the next chapter, or the next course. Flag it, move on, and don't carry it as a debt.

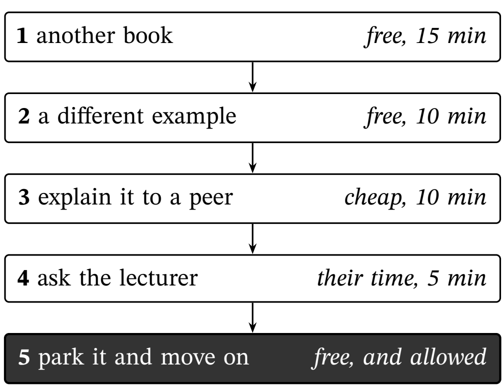

Figure 13.3: Escalation ladder. Cost rises as you climb, so climb only when the rung below has been used.

13.4 The stuck that is really tiredness

One diagnostic before you escalate anything. How long have you been working?

Mathematical reading degrades sharply after about ninety minutes, and the degradation is invisible from the inside. It presents as a hard passage.

So if you've been at it for two hours and the page has become impenetrable, the page is probably fine. Stop, walk somewhere, come back tomorrow. Students routinely spend a miserable extra hour proving to themselves that they're stupid when they were only tired, and the same paragraph opens in four minutes the next morning.

READ SOMETHING TODAY

Start an open questions sheet today. One page, front of your notes. Every time you hit the twenty-minute limit, write the question on it with the page number, mark the margin, and read on. Review it at the start of your next session and cross off everything that quietly resolved itself. Most weeks that's a third of the list.

14

When the Proof Uses a Technique You Have Never Met

Some proofs are hard because the definitions are dense. Others are hard because the proof makes a move you did not know was legal.

Rereading does not fix the second kind. You can understand every symbol and still miss the argument because the author is using a technique as a single unit while you are seeing six unrelated steps. The repair is to pause the theorem, learn the move on a smaller object, and then return.

14.1 The example: a board that cannot be tiled

Remove two opposite corner squares from an ordinary 8×8 chessboard. Can the remaining 62 squares be covered by 31 dominoes, each domino covering two squares that share an edge?

A first reaction is usually geometric. Try arrangements. Start in one corner. Rotate a domino. Get stuck, erase, and try another arrangement. That search can continue for hours because the board has too many local choices.

Then the proof does this:

> Colour the board in the usual black-and-white pattern. Opposite corners have the same colour, so removing them leaves 30 squares of one colour and 32 of the other. Every domino placed along an edge covers one black and one white square. Hence 31 dominoes would cover equal numbers of the two colours, which is impossible. Therefore the mutilated board cannot be tiled.

The arithmetic is easy. The move is not. Why was colouring the board allowed, and how would you think of it on another problem?

14.2 Name the move, then translate it

The technique is an *invariant argument*. An invariant is a quantity or property that stays unchanged while the allowed moves happen.

Here the allowed move is placing one domino. No matter where the domino goes, it uses one black and one white square. So the difference

$$\#(\text{uncovered black}) - \#(\text{uncovered white})$$

does not change when a domino is placed: both counts fall by one.

At the start, after the corners have been removed, the difference is $30 - 32 = -2$. A fully tiled board would have difference $0 - 0 = 0$. The allowed move cannot change the difference, so the start can never reach the finish.

That is the proof in technique-language:

1. define a state;
2. identify the allowed move;
3. find a quantity the move cannot change;
4. compute it at the start and at the desired finish;
5. if the values differ, the finish is unreachable.

The colouring is not decoration. It manufactures the quantity that refuses to change.

14.3 Read the proof in three layers

When a technique is new, do not try to remember the polished proof at once. Read it at three depths.

Layer 1: local validity. Check every sentence. Opposite corners really do share a colour. Every edge joins opposite colours. A domino really does cover adjacent squares. This tells you the proof is valid, but not yet why it was found.

Layer 2: the reusable skeleton. Remove the chessboard nouns: allowed moves preserve a quantity; the desired state has the wrong value; therefore it is unreachable. That sentence is what transfers.

Layer 3: the discovery question. Ask what the allowed move does predictably. A domino always covers one square of each colour. That regularity suggests counting colours. On another problem, the predictable thing may be parity, total weight, a remainder modulo 3, or orientation.

Most readers stop after layer 1 and say they followed the proof. They did. But following is not owning. The second and third layers are what let the technique appear again next month.

14.4 Build a toy problem

Before returning to the theorem, use the same technique somewhere too small to hide it.

Place five coins on a table, all showing heads. One move flips exactly two coins. Can you reach a position with exactly four tails?

The answer is yes: flip two heads, then the remaining two heads. But can you reach exactly three tails?

Each move changes the number of tails by -2, 0, or 2. Its parity never changes. The initial number of tails is 0, which is even. Three is odd, so that state is unreachable.

The coins remove the geometric distraction. The invariant is simply parity. Once that smaller argument feels natural, the chessboard proof stops looking like a trick and starts looking like an instance.

14.5 The technique card

Keep one half-page per unfamiliar technique. The card is not a definition copied from the book. It is a retrieval device.

> **INVARIANT ARGUMENT**
>
> **Signal:** repeated allowed moves; question asks whether one state can reach another.
> **Skeleton:** find a quantity preserved by every move; compare start and target.
> **Small example:** flipping two coins preserves the parity of the number of tails.
> **Full example:** each domino covers one black and one white square.
> **Failure mode:** checking that the quantity survives some moves, not every allowed move.
> **Search question:** what does one move always change by an even amount, leave fixed, or move between fixed classes?

Five lines are enough. The card should let you recognize the move, rebuild its logic, and avoid its most common misuse.

14.6 Return to the original proof

Now close this chapter and reconstruct the chessboard argument without looking.

If you remember only "black and white squares," you have remembered the example. If you remember "find what every allowed move preserves and compare start with finish," you have learned the technique.

Then change the problem. Remove two corners of opposite colours. The colour count is now 31 and 31, so the invariant no longer rules out a tiling. That does not prove a tiling exists. It proves only that this particular obstruction has disappeared.

This distinction matters whenever a technique is new:

> *An obstruction can prove impossibility. The absence of that obstruction does not prove possibility. Every proof technique has a direction, and learning the direction is part of learning the technique.*

14.7 When to look up the technique properly

Do the one-level lookup when any of these happens:

- the technique carries the main theorem rather than one lemma;
- the same move appears twice in the chapter;
- you can follow the example but cannot state the reusable skeleton;
- changing a hypothesis leaves you unable to predict whether the argument survives.

Look for an expository treatment or a textbook section named after the technique. Read one definition, two worked examples, and one failure case. Then return to the original proof the same day.

Do not begin a full course in combinatorics because one proof uses an invariant. You need a working handle first. Depth can come after the technique has earned a second appearance.

READ SOMETHING TODAY

Find a proof in your current course that uses a named technique you have not learned before. Make a six-line technique card: signal, skeleton, small example, course example, failure mode, search question. Then reconstruct the proof with the book closed. If the skeleton survives after the symbols disappear, the lookup worked.

Part V

Longer Texts

15

Reading a Chapter

Everything so far has worked at the scale of one object. One definition, one statement, one proof.

A chapter is thirty of those. Reading it front to back at working depth is both exhausting and ineffective, because the first thing you meet is the thing you're least equipped to understand.

Chapters want three passes. A serious one takes four to six hours spread over several days, and the three passes cost roughly 1:3:2.

15.1 The three passes

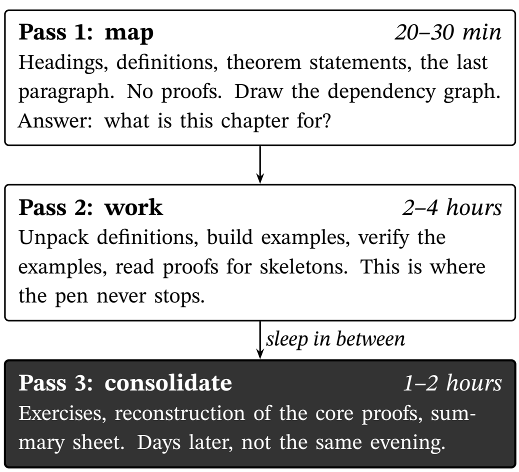

Figure 15.1: Three passes over a chapter. Nobody understands a chapter on the first pass, and nobody is supposed to.

Pass 1: the map

Twenty minutes, no proofs, no stopping. You're building the table of contents the chapter should have had.

On one sheet: every definition's name, every theorem's statement in abbreviated form, and the dependency arrows between them.

At the end you should be able to answer three questions. What's this chapter for? What are the new objects? Which result is the destination?

That sheet is worth more than it looks. It's what you'll revise from in May, and it took twenty minutes in October.

Pass 2: the work

The slow pass, and the one to break into sessions. Two hours at a time, following the block structure from Chapter 3.

Order inside the pass matters. Don't go strictly top to bottom. Go *definitions first, all of them*, then examples, then statements, then proofs.

A chapter typically has three or four definitions, and front-loading them means that when you reach the proofs you aren't carrying an unfolded object.

That reordering is the single biggest change most students can make to how they read a chapter, and it costs nothing.

Pass 3: consolidation

Days later. Deliberately.

Exercises first, because retrieval beats review. Then reconstruct the two or three proofs that matter. Then the summary sheet, which is Chapter 19's subject.

Skip pass 3 and everything in pass 2 was rented rather than bought. It's the pass students cut when time is short, and that's exactly backwards. If time is short, shorten pass 2 and keep pass 3 intact.

15.2 Where the chapter's difficulty actually is

Chapters aren't uniformly hard. In almost every mathematics text the difficulty concentrates in two places: the new definition that generalises something familiar, and the one proof that introduces a new technique.

Everything else is application.

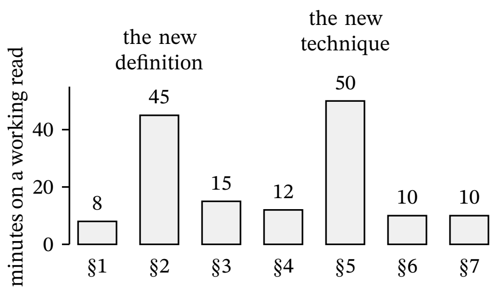

Figure 15.2: One analysis chapter, timed section by section on a working read. Two sections took 95 of the 150 minutes; the other five took 55 between them.

So triage. Identify the two spikes in pass 1 and budget for them: half your working time goes to those two objects, and the rest of the chapter gets what's left.

Students who spread their time evenly across a chapter run out of energy before the second spike, which is usually the important one.

15.3 Reading around a lecture

If the chapter comes with a lecture course, the timing changes and gets better.

Before the lecture. Pass 1 only, twenty minutes. You walk in knowing the names of the objects.

The effect on how much of the lecture you absorb is dramatic and out of all proportion to the effort, because a lecture spent decoding what a word means is a lecture spent not following the argument.

During the lecture. Write down what's said, not what's on the board. The board is in the book. The speech isn't. When the lecturer says "the reason we need compactness here is…", that sentence is in no text and it's the most valuable thing in the room.

After the lecture. Pass 2 on the sections covered, within forty-eight hours. After that the marginal value of your lecture notes falls off fast, because you'll no longer remember what your own abbreviations meant.

> *Twenty minutes of pass 1 before a lecture is worth more than an hour of pass 2 after it. Going in cold and hoping to catch up later is the default strategy and the expensive one.*

15.4 When to skip a chapter

Sometimes the right call is not to read it. Legitimate reasons:

- It's marked optional and nothing later cites it. Check the dependency structure of the book, not just the chapter.
- It's an application chapter in a field you don't need. Read the statements, skip the development.
- You're missing a prerequisite from an earlier chapter. Go back. Reading this one now won't work, and it'll cost you an evening to discover that.

Illegitimate reasons: it looks hard, the notation is unfamiliar, the first page defeated you. All three are cured by pass 1, which you haven't done yet.

15.5 A worked chapter transcript

The procedure is easier to trust when you can see the decisions. Here is a constructed but realistic chapter on compactness, timed from the moment the book opens.

00:00–00:04. Headings. The chapter contains: open covers, compact sets, the Heine–Borel theorem, continuous images of compact sets, and the Extreme Value Theorem. The last two sections look like the payoff. I write them at the right edge of the page.

00:04–00:11. Definitions and statements. I record:

> **D1** open cover; **D2** compact = every open cover has a finite subcover.
> **T1** $K \subseteq \mathbb{R}^n$ compact $\Longleftrightarrow$ closed and bounded.
> **T2** continuous image of compact is compact.
> **T3** continuous $f : K \to \mathbb{R}$ on compact K attains max and min.

I do not read the proofs. I draw

$$\text{D1} \longrightarrow \text{D2} \longrightarrow \text{T1} \longrightarrow \text{T2} \longrightarrow \text{T3}.$$

The printed order puts T1 before T2, but the logical destination is T3: compactness is being built so continuous functions must attain extreme values.

00:11–00:18. Examples. The chapter gives $[0,1]$ as compact and $(0,1)$ as noncompact. I add the missing reason:

$$\{(1/n, 1) : n \geq 2\}$$

is an open cover of $(0,1)$ with no finite subcover. The notation looks wrong at first because $(1/n, 1)$ is an interval, not a point. I circle it, say it aloud, and the type error disappears.

00:18–00:23. Difficulty spikes. D2 is the first spike because "every cover has a finite subcover" is a quantifier statement with families of sets. T2 is the second because its proof must move a cover back through a function. I assign 45 minutes to D2 and 60 to T2 in the working pass. T1 gets no time today; its proof is long and the course notes mark it optional.

00:23–00:26. Decision. The chapter is not seven equally important sections. It is one difficult definition, one transport theorem, and an application. My next session starts with D2's examples, not page one.

That is pass 1. Nothing has been mastered, but the chapter now has a shape, two named obstacles, and a budget. Starting the working pass with those three facts is different from starting cold.

READ SOMETHING TODAY

Take the chapter you're in the middle of. Stop reading it. Go back and do pass 1 properly on one sheet: all definitions, all statements, the dependency arrows, the two difficulty spikes. Twenty minutes, and you'll finish the chapter faster than if you'd carried on from where you were.

16

Working Through a Book Alone

Reading a textbook with a lecture course around it is a supported activity. Someone else sets the pace, marks the important results, and answers questions.

Reading one alone is a different problem, and it's the one most people face after graduating, or over a summer, or when the course they need isn't offered.

Self-study fails for predictable reasons, and almost none of them are about ability. It fails on book choice, on pace, on exercises, and on the absence of any test of whether it's working.

All four are fixable.

16.1 Choose the right book, and it isn't the famous one

The most common self-study error is picking the book with the best reputation instead of the one at your level.

Reputation among researchers tracks elegance and economy, which is exactly what a first-time reader doesn't want.

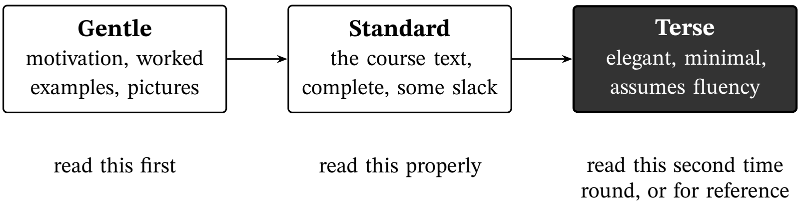

Figure 16.1: The book ladder. Reputation rises to the right; suitability for a first pass falls.

The strategy that works: read a gentle book quickly and a standard book slowly, in that order.

The gentle one hands you the shape and the vocabulary in a fortnight. Then the standard one goes three times faster than it would have, because you're no longer meeting every object for the first time.

Reading two books isn't twice the work. It's usually less total work than reading one hard book once, and far more likely to finish.

16.2 Pace: the fifteen-page week

Set a page target and make it embarrassingly small.

Ten to fifteen pages a week of a serious text, done properly, is a real pace that finishes a 250-page book in about five months. Most people set out to do forty pages a week, manage it for eleven days, and stop.

Write the target down, track it, and treat the streak as the thing you're protecting. Self-study dies from stopping, not from slowness, and nothing dies faster than a plan that requires a good week.

> *The only failure mode that matters in self-study is stopping. Set a pace you can hit in a bad week, because the bad weeks are the ones that decide whether you finish.*

16.3 Exercises: five, chosen, not forty

You have to do exercises. Reading mathematics without them produces the familiar illusion of understanding and nothing else.

But you can't do all of them alone, and you don't need to. Do five per section, chosen like this:

- **One routine.** The first exercise, the one that just applies the definition. It confirms you've got the definition right.
- **Two typical.** The middle of the set, where most exam questions live.
- **One that looks hard.** You may not finish it. Twenty minutes, then park it. This is where the actual learning happens.
- **One the text later uses.** Textbooks quietly delegate results to exercises, and a later proof will assume you did it.

Spotting that last category is a skill worth building. Scan the exercise set for anything that reads like a proposition rather than a computation. Those are the delegated results.

1 routine *5 min*
confirms you read the definition correctly

2 typical *15 min each*
the standard moves, where exam questions come from

1 hard *20 min, then park*
you may not finish; the attempt is the point

1 delegated *as long as it takes*
a result the book will cite later, so it is not optional

Figure 16.2: Choosing five exercises from a set of forty. The two you must not skip are the hard one and the delegated one.

16.4 The absence of a marker

The hardest part of self-study is that nobody tells you whether you're right. You need substitutes, and there are four good ones.

Solutions, used correctly. Plenty of books have solutions to alternate exercises. Use them, but only after a genuine attempt and a written one. Reading a solution to a problem you haven't tried teaches you nothing while feeling like it taught you something, which is the worst combination available.

Reconstruction. The test from Chapter 12 needs no marker. Either you rebuilt the proof or you didn't, and you can see which.

Consistency checks. Test your answers against the special cases from Chapter 8. Does your formula give the right thing for $n = 1$? For the empty set? An answer that fails its own degenerate case is wrong, and you found that out without a marker.

People. Mathematics forums, a study partner, a former lecturer. One person who'll read a proof of yours once a month is worth a great deal, and precise questions get answered by strangers far more often than vague ones do.

16.5 When to abandon a book

Sometimes stopping is right. Abandon a book when:

- you've gone three sessions without writing anything you understood, and the gentle-book move hasn't fixed it;
- you find a prerequisite you genuinely lack, in which case go and get it instead of pushing on;
- the book is a reference work you mistook for a textbook, which happens often and is nobody's fault.

Don't abandon it because it's slow. Fifteen pages a week isn't slow. It's the pace.

16.6 A realistic six-month plan

For a serious 250-page text, self-taught, from a standing start:

- **Weeks 1–2.** Read a gentle treatment of the same subject end to end, fast, no exercises. Build the vocabulary.
- **Weeks 3–20.** The main text at twelve to fifteen pages a week, three passes per chapter, five exercises per section.
- **Every fourth week.** No new reading. Reconstruct the core proofs of the previous three weeks and rebuild the summary sheets.
- **Weeks 21–24.** The last chapters, plus a full pass over all the summary sheets.

That consolidation week is the part people cut, and it's the part that makes the other twenty weeks stick. Put it in the plan as a scheduled item, not as something you'll do if there's time.

READ SOMETHING TODAY

Pick the book you've been meaning to work through. Today, write down three things: the weekly page target (make it smaller than you want), which five exercises per section you'll do, and the date of your first consolidation week. Then read the first fifteen pages this week and nothing more.

17

Reading a Research Paper

A research paper isn't a textbook chapter, and reading it as though it were is why the first encounter is so demoralising.

Textbooks are written to teach. Papers are written to convince a referee and establish priority, by an author writing for maybe two hundred people worldwide who already know the field.

You're not the intended reader. Accept that and papers become tractable, because you can stop trying to understand them the way you understand a chapter and start extracting what you need.

17.1 What a paper is made of

Two entries in that figure are odd enough to justify.

The introduction is the paper. In a well-written mathematics paper the introduction states the main theorem, positions it against what was known, and sketches the idea of the proof in a paragraph. Everything after it is verification.

Read the introduction twice, carefully, and you have most of the value. For most papers you'll read as a student, that's where you should stop.

References, third. Scanning the bibliography early tells you what field you're in, whose work this builds on, and, crucially, whether there's a survey article or a textbook treatment you should be reading instead of this paper.

There often is, and finding it saves days.

17.2 The twenty-minute triage

Before committing to a paper, spend twenty minutes deciding whether to, and at what depth. Read in this order and stop as soon as you have your

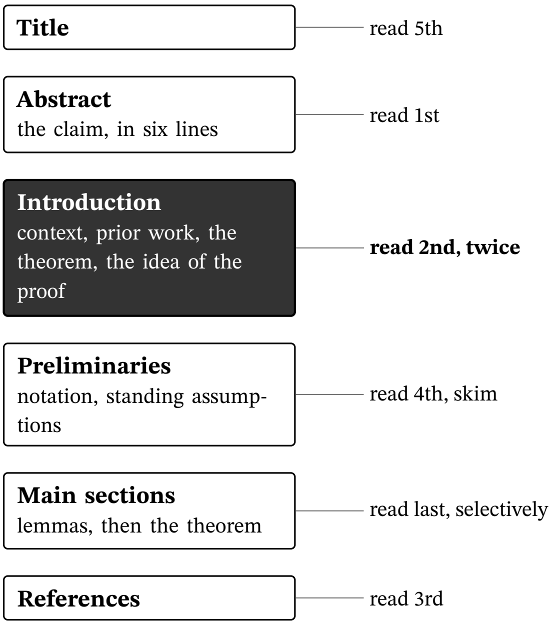

Figure 17.1: The anatomy of a mathematics paper, with the order you should actually read it in.

answer.

1. **Abstract** (2 min). What's claimed?
2. **Introduction** (10 min). What's the theorem, and what's the idea?
3. **Section headings** (1 min). How is it organised?
4. **The main theorem statement** in the body (3 min). Usually stronger and more technical than the version in the introduction.
5. **References** (2 min). Is there a friendlier route to this material?
6. **Decide** (2 min). One of: not for me; introduction only; statement of the main theorem only; full read.

Write the decision down. The default in the absence of a decision is to start reading section 2 and grind to a halt in section 3, having learned nothing and lost an afternoon.

> *Most papers you open should be closed again after twenty minutes with a recorded decision. Reading a paper you didn't choose to read is the main way postgraduate time disappears.*

17.3 Reading for one theorem

The most common legitimate reason a student reads a paper is to extract one thing. A definition, a statement they need to cite, a technique.

That's a different activity from reading the paper, and it's perfectly respectable.

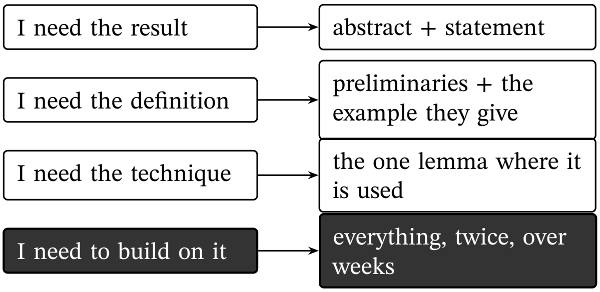

Figure 17.2: Depth by purpose. Only the bottom row requires reading the proofs.

17.4 What makes papers hard, specifically

Four things, and naming them helps, because three have direct fixes.

Compression. Papers are more compressed than any textbook. One line can hide a page of routine argument, marked "clearly" or "a straightforward computation shows". Sometimes it genuinely is straightforward and takes you an hour. Normal, and not a comment on you.

Assumed background. The paper assumes a graduate course you may not have taken. Fix: find the survey or textbook chapter in the references and read that first. Highest-value move in this chapter.

Local notation. Authors invent notation for a paper, define it once in section 2, then use it for thirty pages. Fix: build the notation table from Chapter 6 on a separate sheet before you touch section 3. Non-negotiable for any paper you'll read more than once.

No worked examples. Papers rarely include examples, because the intended reader supplies their own. Fix: supply your own, exactly as in Chapter 5. Take the main theorem and apply it to the simplest object you can think of. If nothing in the paper is concrete enough for that, that's itself a signal about how far the material is from you today.

17.5 The three-pass read, for papers you commit to

If you decide to read a paper properly, three passes over three sittings.

Pass 1, one hour. Introduction, all statements, no proofs. Build the notation table. Write the main theorem in your own words.

Pass 2, two to four hours. The proofs, at follow depth. Mark every step you can't justify but don't stop for them. You're learning the architecture of the argument.

Pass 3, as long as it takes. The two or three lemmas that carry the paper, at verify or reconstruct depth. You won't do this for the whole paper, and you shouldn't.

> *Nobody reads a research paper once. The question is never whether you understood it on the first pass, but what you extracted on this pass and what you'll come back for.*

17.6 Where to find readable ones

Reading papers for the first time? Start with survey articles, expository journals, and lecture notes.

The Monthly, the Notices, and the various *Bulletin* survey pieces are written to be read by people outside the specialism, and they're the right entry point. Some subjects also have long, careful lecture notes posted by their authors, and those are frequently better than anything published.

Three good surveys before your first research paper will make the paper twice as accessible, and it's a far more pleasant way to spend a fortnight.

17.7 A worked paper transcript

This transcript uses a constructed paper so the reading decisions can be shown without borrowing a specialist argument. The paper's title is "A Generating-Function Proof for Restricted Binary Strings." Its main result counts length-n binary strings with no consecutive ones.

00:00–00:02. Abstract.

> Let a_n count binary strings of length n with no adjacent ones. We

> derive the generating function
>
> $$A(x) = \frac{1+x}{1-x-x^2}$$
>
> and recover the recurrence $a_n = a_{n-1} + a_{n-2}$.

My first note is not “understand generating functions.” It is:

> **Claim:** a rational function encodes the counts.
> **Known result:** the counts satisfy a Fibonacci recurrence.
> **Unknown object:** how coefficient sequences become functions.

That is enough to read the introduction.

00:02–00:09. Introduction. The paper says the recurrence has a direct last-symbol proof, then offers generating functions because the method extends when more patterns are forbidden. That sentence changes my reason for reading. The recurrence is not the destination; the transferable method is.

I write one question: *where does the denominator* $1-x-x^2$ *come from?*

00:09–00:12. Headings. Section 2 defines ordinary generating functions. Section 3 converts the recurrence. Section 4 handles the forbidden patterns 11 and 101 together. I decide to read sections 2 and 3 fully and skim section 4 for shape.

00:12–00:15. Main proposition. The proposition assumes $a_0 = 1$, $a_1 = 2$, and the recurrence for $n \geq 2$. It multiplies the recurrence by x^n and sums. I can state every hypothesis. The only unfamiliar move is treating an infinite power series algebraically. That goes on the borrowed-results list.

00:15–00:18. References. One reference is an introductory chapter on generating functions. I mark it as the one-level lookup. I do not open a 600-page combinatorics textbook.

00:18–00:20. Decision.

> **Depth:** technique read.
> **Read now:** sections 2 and 3, follow depth.
> **Borrow:** formal power-series algebra.
> **Leave with:** how a recurrence turns into a denominator.

Twenty minutes produced a route and one exact question. If I had started at the first displayed series and demanded full understanding, I would still be on page two.

17.8 The second sitting: follow the technique

The next day, the key calculation is

$$\sum_{n\geq 2} a_n x^n = \sum_{n\geq 2} a_{n-1} x^n + \sum_{n\geq 2} a_{n-2} x^n.$$

I label the three sums before manipulating them. The left side is $A(x) - a_0 - a_1 x$. The first sum on the right is $x(A(x) - a_0)$; the second is $x^2 A(x)$. Substitution gives

$$A(x) - 1 - 2x = x(A(x) - 1) + x^2 A(x),$$

so

$$A(x)(1 - x - x^2) = 1 + x.$$

Now the denominator is explained. It is the recurrence, with the lag-one and lag-two terms recorded as x and x^2. That single sentence is my technique card.

I still have not justified every operation on an infinite series. For this read, I borrowed that algebra and recorded the debt. The paper's method is now mine at follow depth, which is exactly what the decision asked for.

READ SOMETHING TODAY

Find a paper related to something you're studying, ideally one cited in your course notes. Do the twenty-minute triage and write down your decision. If the decision is "introduction only", do that, and write the main theorem in your own words. You've now read a research paper, which is the milestone. Depth comes later.

18

Reading Above Your Level

Sooner or later you'll need something that sits above where you are. A paper using a theory you haven't studied. A chapter assuming a course you haven't taken. A book everyone in your field has read and you haven't.

The instinct is to go back and learn the prerequisites properly, in order, from the beginning.

Sometimes that's right. Usually it's a trap, because the prerequisite has prerequisites, and three weeks later you're reading about vector spaces instead of the thing you needed.

There's a better strategy, and it's what working mathematicians actually do.

18.1 Depth-first, with a budget

Read the hard thing. When you hit something you don't know, go back exactly one level, learn only what you need, and return immediately.

Never go back two levels on the same trip.

The move that makes it work is borrowing a statement without its proof.

If the paper uses a theorem you've never seen, you don't need to be able to prove it. You need to know what it says, what its hypotheses are, and roughly what it's for. Fifteen minutes, not three weeks, and it's enough to carry on reading.

> *You may use any theorem you can state correctly, whether or not you can prove it. Insisting on proving everything from the ground up before using anything isn't rigour. It's a way of never arriving.*

Keep a page called *borrowed results*. Every time you take a statement on trust, write it there with its source.

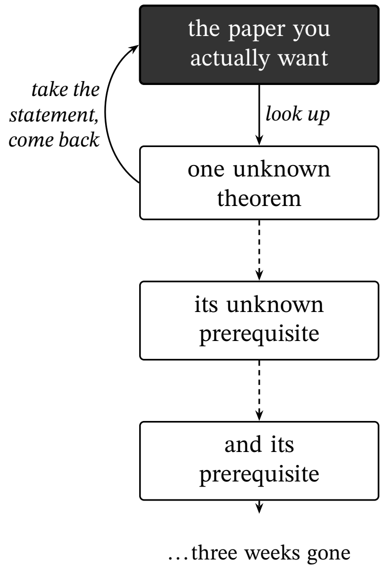

Figure 18.1: Depth-first with a budget. Go back one level, take what you need, come straight back. The dashed path is how people lose months.

Two things follow. You have an honest record of what your understanding is standing on, which matters if you're going to write anything. And the list becomes a reading plan: when a borrowed result gets used for the third time, that's the signal to go and learn it properly.

18.2 Breadth-first, when depth-first fails

Depth-first has a failure mode. If the unknowns arrive faster than you can absorb them, four per page, you're not close enough to the material and no amount of looking things up will fix it.

The signal is quantitative. If more than about one in four sentences contains an object you can't state the meaning of, stop. At that density comprehension doesn't happen, and you'll spend six hours producing nothing.

Then switch to breadth-first: put the hard text down, find a treatment one level below it, read that first.

Not defeat. Routing. A fortnight on the level below is almost always faster than a month of grinding at the level above.

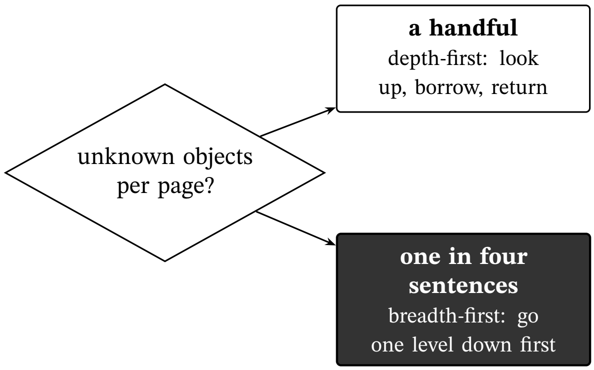

Figure 18.2: The density test. Count the unknowns on one page and let the count choose your strategy.

18.3 Reading for the shape, not the content

There's a third mode, and it's the one that makes attending seminars worthwhile even when you understand a quarter of them.

Read the hard text for its shape only. What kind of object is this about? What form does the main theorem take: existence, uniqueness, classification, bound? What machinery does it lean on? What does the author think is hard, judging by where the pages go?

You can answer all of those without following a single proof, and they're real knowledge. They're how you build a map of a field before you can walk around in it, and they make the second encounter, six months later, dramatically easier.

It's also the honest answer to "how do people sit through seminars they don't understand?" They're not following the argument. They're collecting shape, names, and the two or three moments where the speaker says something surprising.

18.4 The three questions to leave with

Whenever you read something above your level, extract these three, in writing. Five minutes, and they're what makes the attempt worth having made.

1. **What's the main claim?** One sentence, your own words, even if crude.

2. **What's the one object I most need to learn?** Not five. One. That's your next study target.
3. **What would I need to read this properly?** Name the course or the chapter. Now you have a route.

A student who reads a hard paper and comes away with those three lines has done well. A student who reads the same paper for six hours, follows none of it, and writes nothing down has done the same reading and kept none of it.

18.5 On the feeling of being out of your depth

Being out of your depth is uncomfortable in a specific way. It feels like a verdict rather than a position.

Worth knowing: it doesn't go away with seniority. Professionals feel it at every seminar outside their specialism, at every conference, and in most of the papers they open.

What changes isn't the feeling. It's the response, which becomes procedural. Triage, borrow, extract the three questions, decide whether to come back.

That's all this chapter is. A procedure to run instead of a judgement to accept.

READ SOMETHING TODAY

Find something one clear level above you: a paper cited in your course notes, or a chapter from a graduate text on the same subject. Give it forty minutes. Run the density test to choose depth-first or breadth-first, keep a borrowed-results list, and finish by writing the three questions. Forty minutes, three lines, and a route.

Part VI

With a Pen, and With People

19

The Artifacts

This book has quietly asked you to build six things.

Each showed up where it was needed. Here they are together, because collectively they're your reading system, and the system is what survives the term.

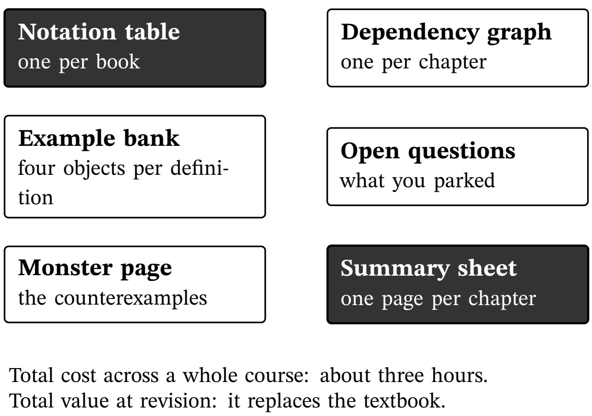

Figure 19.1: Six artifacts. Three live in the book, three live in your notes, and all six get used at revision time.

19.1 Marking the book itself

Use a small, consistent set of marks. Consistency is the whole point. You want to flip through a chapter in June and read your own annotations without decoding them.

Here's the set I use. Borrow it until you build your own.

Mark	Means
box in the margin	unresolved; on the open questions sheet
filled box	resolved later
star	the load-bearing step of a proof
double line	a definition I have unpacked
→ p. 41	this depends on that
“why?”	I follow it but do not believe it yet

Write in the book. If it’s a library copy, use sticky notes or keep a page of annotations keyed by page number.

The rule about not writing in books is a rule for novels.

19.2 The summary sheet

One page per chapter, written from memory after pass 3:

- every definition, in your own words, with one example and one non-example;
- every theorem statement, exactly, with hypotheses split out;
- the skeleton of each proof you reconstructed, in five phrases;
- the dependency graph;
- your monsters from this chapter.

And the rule that makes it work: write it from memory first, then open the book and fill the gaps in a different pen.

The second colour shows you exactly what you didn’t have, which is more useful than the sheet itself.

> *Anything you can only write with the book open is something you don’t yet know. The two-pen summary sheet makes that visible instead of comfortable.*

By the end of a course you’ve got ten to twelve of these pages. That’s what you revise from. You won’t reread the textbook in May, and knowing that in October is what makes the sheets worth writing.

19.3 Reading with your hands: the four moves

Underneath every artifact is one idea, and it’s worth stating on its own, because it’s the thesis of this book.

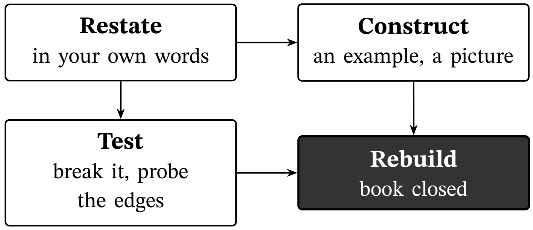

Figure 19.2: The four hand movements that constitute reading mathematics. Eyes alone are not reading.

Everything in this book is one of those four, applied to a definition, a statement, a proof, a chapter, or a paper.

Forget every technique here and remember only that reading mathematics means restating, constructing, testing, and rebuilding, and you'll still have most of the value.

19.4 A worked summary sheet

Here's what one looks like for a first-analysis chapter on sequences, squeezed to fit this page.

> **Ch. 3 — Sequences**
>
> **Defs.** *Bounded*: one M works for all n. Ex: $(-1)^n$. Non-ex: $\sqrt{n}$. *Convergent*: $\forall\varepsilon\ \exists N\ \forall n > N$. Game: they pick ε, I pick N.
>
> **Thms.** 3.4 convergent $\Rightarrow$ bounded (converse false: $(-1)^n$). 3.6 Bolzano–Weierstrass: bounded $\Rightarrow$ convergent subsequence. 3.9 monotone + bounded $\Rightarrow$ convergent.
>
> **Skeletons.** 3.4: take $\varepsilon = 1$, finitely many terms outside, max. 3.6: bisection, nested intervals, pick one term per interval.
>
> **Graph.** Def bounded → 3.4 → 3.6; Def sup → 3.9.
>
> **Monsters.** $(-1)^n$, $\sqrt{n}$, $1/n$, harmonic partial sums.

A chapter, on one page, and it took twenty-five minutes to write after the reading was done. Nothing in it was copied. Every line was recalled and then checked.

READ SOMETHING TODAY

Write a summary sheet for the last chapter you finished, from memory, in one pen. Then open the book and fill the gaps in a second pen. Look at how much is in the second colour. That proportion, not your feeling about the chapter, is where you actually stand.

20

Lectures, Groups, and Good Questions

Reading mathematics is mostly solitary, and then it isn't.

Lectures, study groups, office hours, and seminars are all reading with other people, and they follow different rules from reading alone. Most students get poor value from all four, for reasons that are entirely mechanical and entirely fixable.

20.1 A lecture is not a text

The commonest lecture mistake is trying to copy down everything on the board.

The board is a compressed version of the textbook, and the textbook is already in your bag. What isn't in your bag is what the lecturer *says*.

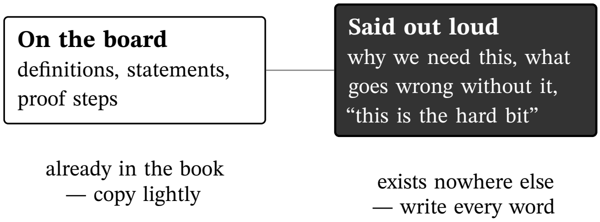

Figure 20.1: What to capture in a lecture. The board is already written down somewhere; the speech is not.

Sentences worth catching word for word: "the reason this hypothesis is here is…", "this is the step people get wrong", "you should think of this as…", "we'll use this again in chapter 7", and anything starting "intuitively".

Those are the sentences that got edited out of the textbook, and they're why attending beats reading the notes.

The practical scheme: divide the page. Board content on the left, abbreviated. Spoken content on the right, in full sentences. Afterwards, the right-hand column is what you reread.

20.2 Prepare twenty minutes, gain an hour

From Chapter 15, and it earns the repetition, because it's the highest-return habit in this book. Do pass 1 on the material before the lecture. Twenty minutes, no proofs, just the names of the objects.

The mechanism is simple. In an unprepared lecture your attention goes on decoding notation and vocabulary, and the argument passes you completely. In a prepared one the vocabulary is familiar and your attention is free for the argument, which is the only thing worth being in the room for.

20.3 Study groups that work

Most study groups are two hours of social time with a textbook open.

The ones that work have a structure, and the structure is the whole difference. Three formats worth using:

The teaching round. Each person gets assigned one result before the meeting and teaches it at the board for ten minutes. Teaching is reconstruction under pressure, and it exposes every gap you have. Easily the most effective format, and the most avoided.

The gap swap. Everyone brings their open questions sheet. Round the table, five minutes per question. Most get answered by somebody in the room, and the ones that don't are the questions to take to the lecturer, now with the authority of four people being stuck on the same thing.

The counterexample game. Someone states a plausible false claim and everyone races to break it. Fifteen minutes of this beats an hour of rereading, and it builds exactly the instinct from Chapter 8.

> *If nobody in the group is writing on a board, it isn't a study group. Explaining out loud, to people who will interrupt, is the fastest diagnostic tool a student has, and it's free.*

20.4 How to ask a good question

A good question has three parts and takes thirty seconds to build. The difference in the answer you get is enormous.

1. Where you are
"In the proof of Theorem 4.6, at the line where we set $\delta = \min(1, \varepsilon/5)$..."

2. What you tried
"I checked that $\varepsilon/5$ makes the second term work..."

3. The precise gap
"...but I can't see what the 1 is protecting against."

Figure 20.2: The anatomy of a question that gets a useful answer. All three parts are needed; the third is the one students omit.

Compare that with "I don't understand Theorem 4.6". The second version makes the lecturer guess where you are, and their best guess is to explain the whole thing again, which is what you already failed to absorb once.

There's a bonus, and it isn't small. Building the three-part question forces you to locate your own confusion precisely, and a decent fraction of the time the question answers itself before you ask it.

Not wasted work. That's the technique working.

20.5 Office hours, and the thing nobody tells you

Lecturers like specific questions.

A student who turns up with a marked-up page, an open questions sheet, and a precise gap is the best part of most academics' week, and they'll give you far more than the answer. The context, the history, the reason the definition is the way it is, the warning about what breaks in the next course.

What they can't do much with is "can you go over section 4". Not because they're unwilling. There's just nothing to grip.

Go early in the term, not the week before the exam. Go with the sheet. Two visits with a precise question beat a term of attending and saying nothing.

20.6 Seminars you don’t understand

You’ll go to talks where you follow the first ten minutes and then lose the thread entirely. That’s universal, including among the professors in the room.

The move is the one from Chapter 18. Stop trying to follow the argument and start collecting shape. What kind of object, what kind of theorem, what machinery, what did the speaker seem to think was hard, what were the two moments the audience reacted to?

Write those five things down. Leave with a page.

Then look up one thing afterwards. One definition, one name, one theorem. Do that after every seminar for a year and you’ll have built a map of a field almost by accident.

READ SOMETHING TODAY

Before your next lecture, spend twenty minutes on pass 1 of the material. In the lecture, split your page and write the spoken column in full sentences. Afterwards, read only that column. Then take one precise three-part question to office hours this week, even if you have to manufacture the occasion.

21

What Changes When You Get Good

Nobody tells you what the destination looks like, so students assume it's fluency. Open a text, understand it at reading speed.

That never arrives. Not for anyone. And expecting it is why so many capable people decide they aren't cut out for this.

Here's what actually changes.

21.1 You get stuck just as often

The frequency doesn't drop. It might rise, because you're reading harder things now.

What changes is the response time. Where you used to spend forty minutes rereading and feeling worse, you now spend ninety seconds working out which kind of stuck this is and taking the matching action.

The obstacle is the same size. The cost of hitting it has collapsed.

21.2 You know what you don't know

The clearest marker of a mature reader isn't confidence. It's precision about their own ignorance.

They can tell you exactly which result they're taking on trust, which definition they've unpacked and which they haven't, which proof they could rebuild and which they followed once in 2019.

Students tend to have a single undifferentiated feeling about a subject instead. It's fine, or it's a disaster.

Replacing that feeling with a list is most of what this book has been trying to do.

> *The goal was never to understand everything. It was to know, at all times, which category each thing is in.*

21.3 You read for the shape first

At some point you stop starting at the top of the page. You go to the statements, then the examples, then the structure of the proof, and only then the lines.

It stops being a technique you're applying from a book and becomes what reading feels like.

The reordering happens gradually, and then one day you notice you've been doing it for months.

21.4 Your example bank does the work

The most underrated change.

After a year of collecting examples, non-examples, and monsters, you meet a new theorem and the objects arrive on their own. Someone states a plausible claim and $(-1)^n$ turns up in your head uninvited. You read "continuous but not differentiable" and $|x|$ is already there.

That's what mathematical maturity mostly is, in practice. Not raw intelligence, and not memory for proofs. A large, well-organised collection of objects you've personally handled, arriving when needed.

21.5 You get comfortable with partial understanding

The last change is the hardest to accept and the most freeing.

You'll read a paper, extract the main theorem and one technique, take four lemmas entirely on trust, close it, and be satisfied. You'll go to a seminar, follow a third, leave with five lines and one thing to look up, and count it a good hour.

The undergraduate instinct says that's cheating. That real understanding means proving everything from the axioms yourself.

It isn't cheating. It's the only way anyone works, because the alternative doesn't fit inside a life. What separates a professional isn't that they've proved everything. It's that they know precisely what they haven't, and could go and do it if it mattered.

21.6 Where to go from here

Three things to do with this book now that you've finished it.

Pick one technique and make it automatic. Not all of them. One. I'd pick the unpacking drill from Chapter 4, because everything else in mathematics rests on definitions. Run it on every definition for a month, until it stops feeling like a procedure.

Build the artifacts for one course. A notation table, a dependency graph per chapter, summary sheets, a monster page. One course, one term. Compare the revision experience with last term's and decide for yourself whether it was worth the three hours.

Write. The companion to this book is *How to Write Mathematics*, and the two skills are the same skill seen from opposite ends. Everything you now know about how mathematical prose compresses will make you better at writing it, because you know exactly what a reader has to do to unfold what you wrote.

Writing a proof for a reader you understand is a different act from writing one for a grader you're afraid of.

21.7 The last thing

Go back to the page from the preface. The one you'd read four times with nothing happening. It's probably still on your desk.

Read it again, but with a pen this time. Say every symbol out loud. List the types. Unpack the definition. Build the smallest example you can. Try to break it. Write the skeleton of the argument in five phrases.

Forty minutes instead of ten, and at the end you'll own it.

That trade, forty minutes for permanence instead of ten minutes for nothing, is the entire content of this book.

READ SOMETHING TODAY

Finish the page you were stuck on when you started this book. Then write one paragraph, in your own words, explaining what it says, and check it against the text. That paragraph is the first thing you've read properly. There will be a great many more.

Part VII

Toolkit

A

Five Reading Makeovers

Every technique in this book, applied to five real passages.

The BEFORE in each case isn't a lazy student. It's what careful, linear, well-intentioned reading actually produces. The AFTER is the same passage, the same reader, the same amount of intelligence, and twenty minutes of the right moves.

A.1 Makeover 1: a definition

The passage, as it turns up in a first analysis text:

> $f : A \to \mathbb{R}$ is *uniformly continuous* on A if for every $\varepsilon > 0$ there exists $\delta > 0$ such that $|f(x) - f(y)| < \varepsilon$ whenever $x, y \in A$ and $|x - y| < \delta$.

BEFORE

OK so uniformly continuous means continuous but uniformly. Stronger than continuity somehow. There's an epsilon and a delta, same as continuity. I'll come back to this when I see it used.

Honest, and worth nothing. There's no object in it, no test, and no way to tell uniform continuity from continuity, which is the only thing the definition exists to do.

AFTER

Say it. "For every tolerance, there's one gap that works everywhere on A at once."
Types. f is a function $A \to \mathbb{R}$; ε, δ are positive reals; x, y are points of A. Note δ is *not* indexed by a point.
Quantifiers. $\forall\varepsilon\ \exists\delta\ \forall x, y$. Compare ordinary continuity at each

point: $\forall x\ \forall \varepsilon\ \exists \delta$. In continuity, x comes before δ, so δ can depend on x. Here it can't. **That single swap is the entire definition.**
Smallest example. $f(x) = x$ on $\mathbb{R}$. Given ε, take $\delta = \varepsilon$. Works everywhere. Done.
Non-example. $f(x) = 1/x$ on $(0, 1)$. Near zero the function climbs arbitrarily fast, so a δ that works at $x = 0.5$ fails at $x = 0.001$. No single δ covers every point.
Boundary. $f(x) = x^2$ on $\mathbb{R}$: continuous, bounded on every bounded set, and still not uniformly continuous, because the slope grows without limit. Restrict it to $[0, 1]$ and it becomes uniformly continuous. So the failure is about the domain as much as the function, which is a real insight and is nowhere in the text.

Twelve minutes. The reader now has two objects, knows the exact quantifier swap that separates the concept from its neighbour, and has met the fact that compactness of the domain is what rescues it, which is the theorem three pages later.

A.2 Makeover 2: a theorem statement

> **Theorem.** Every bounded monotone sequence of real numbers converges.

BEFORE

Bounded and monotone gives convergent. Makes sense: if it's increasing and can't go past a ceiling, it has to settle. Got it.

The intuition is correct, and that's exactly what makes this reading dangerous. It feels complete. There's no audit in it, no test, and no awareness of what the proof is going to need.

AFTER

Split. H1: bounded. H2: monotone. C: converges. Note the conclusion doesn't say what the limit is.
Restate. "An increasing sequence with a ceiling converges, and so does a decreasing one with a floor."

Drop-one audit. Drop boundedness: $a_n = n$ is monotone and diverges. Drop monotone: $a_n = (-1)^n$ is bounded and diverges. Both hypotheses are load-bearing, and each has a one-line counterexample.
Strengthen? Could the conclusion name the limit? For an increasing sequence, yes: the limit is the supremum of the set of terms. So the theorem as stated says less than is true, and noticing that predicts the proof exactly.
Predict the proof. The conclusion is an existence claim, so the proof has to produce a limit. The only limit available is the supremum, which needs the completeness of $\mathbb{R}$. So this theorem can't be true of the rationals. Test: truncate the decimal expansion of $\sqrt{2}$ at n places. Bounded, monotone, rational, doesn't converge in $\mathbb{Q}$. Confirmed.

Six minutes, and the reader has predicted the proof's mechanism, found the reason the theorem belongs to the real numbers, and picked up two counterexamples. The proof, when they get to it, takes a third as long.

A.3 Makeover 3: a short proof

Proposition. Every convergent sequence is bounded.
Proof. Let $a_n \to L$. Take $\varepsilon = 1$. Then there is N with $|a_n - L| < 1$ for all $n > N$. Hence $|a_n| < |L| + 1$ for those n. Let $M = \max\{|a_1|, \dots, |a_N|, |L| + 1\}$. Then $|a_n| \leq M$ for all n. $\square$

BEFORE

Follows fine. Take epsilon equals one, get an N, take the max. Yes, that works.

A clean follow-level read. It'll be gone by Thursday, because nothing got interrogated and nothing got rebuilt.

AFTER

Shape. Direct. The conclusion is an existence claim (there exists M), so the proof has to construct M. It does, in the second-to-last line. That line is the destination.

The load-bearing step. Not the max. It's *take* $\varepsilon = 1$. The definition of convergence is a promise for every ε, and the proof spends that promise once, at a convenient value, to split the sequence into a finite head and a controlled tail. That move, choose a convenient ε and buy a split, recurs everywhere in analysis.
Interrogate. Why does $|a_n - L| < 1$ give $|a_n| < |L| + 1$? Triangle inequality: $|a_n| \leq |a_n - L| + |L|$. The proof doesn't say so. Margin. Why is the max legitimate? Because it's a maximum of finitely many numbers, and the finiteness comes from N being a fixed integer. If the head were infinite there'd be no max. That's the hinge of the whole argument and the text doesn't mention it.
Rebuild, book closed. One sentence: *spend the convergence promise at* $\varepsilon = 1$ *to split into a finite head and a bounded tail, then take a max.* From that, the proof regenerates in two minutes.
Change the theorem. Is the converse true? No: $(-1)^n$. Does the argument show a convergent sequence is eventually inside any fixed band? Yes, and that's what convergence says. Does it still work with $\varepsilon = 7$? Yes, nothing needed the 1, which tells you the 1 was a convenience and not a constraint.

A.4 Makeover 4: the sentence that changes the goal

> …It suffices to show that E has no isolated points. Suppose $x \in E$ were isolated. Then there is $r > 0$ with $B(x, r) \cap E = \{x\}$…

BEFORE

They're proving something about isolated points now. I thought we were proving E is perfect. I'll keep going and hope it connects.

This is the commonest way a reader silently loses a proof, and the damage happens at the phrase everyone skims.

AFTER

Stop at "it suffices". The goal just changed. Write both goals down: Old goal: E is perfect. New goal: E has no isolated points.

Ask why the swap is legal. Perfect means closed and with no isolated points. So the author must already have established closedness, probably in the previous paragraph. Go back and check.
If they have, the substitution is honest and the proof is now single-purpose. If they haven't, you've found a gap, which happens more often than textbooks admit.
Then catch the second shift. "Suppose $x \in E$ were isolated" means the rest is indirect. So the destination is an absurdity, and every line from here is walking toward one.
Two lines of margin. *goal changed: perfect → no isolated points, legal because closed was shown on p. 61. Now indirect.*

That's the whole technique. Every *it suffices*, every *without loss of generality*, every *suppose not* gets a written margin note recording what changed. Three seconds each, and long proofs stop being places you get lost.

A.5 Makeover 5: a research paper, first contact

Abstract. We prove that every quasi-regular Λ-module over a locally finite ring admits a canonical decomposition, answering a question of Ellis and Baptiste. The proof uses a refinement of the transfer method of [12].

BEFORE

I don't know what a quasi-regular Λ-module is, or a locally finite ring, or the transfer method. Clearly I'm not ready for this. I should go and learn module theory properly first.

Every clause of that is understandable, and the conclusion is wrong. This reader is about to spend three weeks on prerequisites and never come back.

AFTER

Triage, twenty minutes. Abstract, then the introduction twice, then the references, then decide.
Shape, extracted without knowing the objects. The main the-

orem is an *existence and canonicity* result: a class of objects decomposes, and the decomposition is unique in some sense. It answers a named question, so it settles something the field cared about. The method refines an existing technique, so this is an incremental-technique paper rather than a new-framework one.
Borrowed results list. quasi-regular Λ-module (undefined, borrowed); locally finite ring (undefined, borrowed); transfer method (borrowed, from reference [12]).
Density test. Three unknown objects in two sentences. That's breadth-first territory, so don't read section 3.
The route. Check the references for a survey. There's one: reference [12] is a 40-page expository article on the transfer method. Read that instead and come back in a month.
The three questions, written down. (1) Main claim: this class of modules decomposes canonically. (2) The one object to learn: what a Λ-module is. (3) What I'd need: a graduate course in module theory, or the survey in [12].

Twenty minutes, no despair, a written route, and a specific next action.

That's what first contact with a research paper is supposed to produce. Understanding it was never on the table today, and nobody expected it to be.

A.6 Makeover 6: an induction proof

Theorem. For every positive integer n,

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

Proof. The result is true for $n = 1$. Assume it is true for n. Adding $2n + 1$ gives $n^2 + 2n + 1 = (n + 1)^2$. The result follows by induction. $\square$

BEFORE

Base case, assume it, add the next term, done. This is the standard induction proof and I understand it.

The reader followed the lines. But ask them which expression is the inductive hypothesis, why $2n + 1$ is the next term, or which statement has been proved at the end, and the proof evaporates.

AFTER

Index the statements. Let

$$P(n): \quad 1 + 3 + \cdots + (2n - 1) = n^2.$$

The proof must establish $P(1)$ and $P(n) \Rightarrow P(n + 1)$.
Base. $P(1)$ reads $1 = 1^2$.
Hypothesis. For one fixed but arbitrary $n \geq 1$, assume

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

The word "one" matters. We have not assumed the theorem for every n.
Target. $P(n + 1)$ is

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

because substituting $n + 1$ into $2n - 1$ gives $2n + 1$.
Bridge. Start with the left side of the target, isolate the old sum, use the hypothesis once, then simplify:

$$\big[1 + 3 + \cdots + (2n - 1)\big] + (2n + 1) = n^2 + 2n + 1 = (n + 1)^2.$$

Rebuild. Base, name $P(n)$, write $P(n + 1)$ before proving it, use the old sum inside the new one. That's the whole skeleton.

The proof did not change. The reading did. The target was written out, the substitution creating the next term was checked, and the hypothesis was located at the exact line where it enters.

A.7 Makeover 7: a definition that depends on an earlier one

> Let (X, d) be a metric space. A sequence (x_n) in X is *Cauchy* if for every $\varepsilon > 0$ there exists N such that $d(x_m, x_n) < \varepsilon$ whenever $m, n > N$. The space X is *complete* if every Cauchy sequence in X converges to a point of X.

BEFORE

Complete means all Cauchy sequences converge. Cauchy means the terms get close together. This sounds like convergence, so a complete space is one where sequences behave well.

The paraphrase is warm and unusable. "Behave well" cannot test an example, and the difference between Cauchy and convergent has been smoothed away.

AFTER

Unpack the dependency. "Complete" cannot be read until "Cauchy" and "converges in X" are available. Put those two on the desk first.

Cauchy, said aloud. Far enough along the sequence, every pair of terms is close to each other. The definition does not name a proposed limit.

Convergent, contrasted. A sequence converges when its terms get close to one particular point $x \in X$. That point must belong to the space.

Complete, rebuilt. X is complete when internal evidence of settling (the terms crowding together) always corresponds to an actual destination inside X.

Test the boundary. The rational numbers with the usual distance are not complete. Rational truncations

$$1, \quad 1.4, \quad 1.41, \quad 1.414, \ldots$$

form a Cauchy sequence in $\mathbb{Q}$, but their destination $\sqrt{2}$ is not rational. The terms settle; the space has a hole where the limit should be.

Dependency card.

$$\text{metric} \longrightarrow \text{Cauchy sequence} \longrightarrow \text{complete space.}$$

If the first object is still folded, the third will remain words.

This is what a dependent definition asks from you. Do not paraphrase the top line while leaving its ingredients opaque. Walk down the dependency chain until you reach objects you can construct, then rebuild upward.

A.8 What the seven have in common

Look back at the AFTER boxes and notice how little cleverness is in them.

Not one needed insight. Every improvement came from a fixed move: say it out loud, list the types, mark the quantifiers, build an object, drop a hypothesis, name the shape, write down the goal change, record what you borrowed.

That's the argument of this whole book, in seven worked cases. Reading mathematics isn't a talent you have or lack. It's a set of moves, and you have them now.

B

The Unpacking Drill

One page. Photocopy it, or copy it onto the inside cover of your notebook. Run it on every definition that matters.

THE UNPACKING DRILL *8–15 minutes*

↓

1 Say it out loud in English, no symbols.
If you stumble, the problem is notation, not the concept. Build the symbol list first.

↓

2 List every object and its type.
Two columns: symbol, and what kind of thing it is. Ninety seconds. Eliminates half of all confusion.

↓

3 Mark the quantifiers, in order.
Each variable may depend on everything before it and nothing after it. Draw the arrows. Read it as a game: they hand you the $\forall$s, you answer with the $\exists$s.

↓

4 Build the smallest example.
The dullest object that works. Verify it clause by clause, with the definition in view.

↓

5 Build a non-example.
Negate the definition first, flipping quantifiers, and read the negation as instructions. Name the clause that fails.

↓

6 Push to the boundary.
What is the closest thing that *just* fails? That object is the definition.

If step 4 defeats you, one of three things is true, and none of them is that you aren't clever enough:

1. An earlier definition is still folded. Go back one level and run the drill there. Recurse until you can build something.
2. The examples are genuinely hard, and the author knows it. Read ahead; the example is almost certainly in the next paragraph.
3. You're missing notation. Build the symbol list.

The two failure modes and their fixes:

Symptom	Name	Fix
"I thought N bounded a_n"	type confusion	step 2, 90 seconds
"I let N depend on n"	scope confusion	step 3, 2 minutes

C

The Stuck Protocol

For the moment you've read the same paragraph three times. One page, designed to be run rather than read.

FIRST: NAME THE STUCK *30 seconds*

I can't say this symbol → notation table

I don't know what this word means → go back, unpack it

I don't see why this step is legal → name the licence, or flag it

The algebra doesn't work out → do it yourself on paper

I follow it but don't believe it → test it on an example

Where did this idea come from? → legitimate. park it and read on

Then run the protocol

1. Can you say every symbol on the page? If not, build the list.
2. Is every term defined, and did you unpack it? If not, go back one level.

3. Have you tried the smallest example? If not, try it now.
4. Have twenty minutes passed? If yes, **mark it and read on**. The obstacle is often explained by text you haven't reached.
5. Still stuck after coming back? Ask a person, with a precise question.

Escalate in this order

	Move	Cost
1	A different book on the same subject	free, 15 min
2	A different example	free, 10 min
3	Explain your stuck to a peer	cheap, 10 min
4	Ask the lecturer, with the three-part question	their time
5	Park it and move on, without guilt	free, and allowed

The three-part question

1. **Where you are.** "In the proof of Theorem 4.6, at the line where we set $\delta = \min(1, \varepsilon/5)\ldots$"
2. **What you tried.** "I checked that $\varepsilon/5$ makes the second term work…"
3. **The precise gap.** "…but I can't see what the 1 is protecting against."

Before you escalate anything

How long have you been working? Mathematical reading degrades sharply after about ninety minutes, and the degradation presents exactly as a hard passage. If you've been at it two hours, the page is probably fine. Stop. It will open in four minutes tomorrow.

D

Time Budgets and Honest Speeds

Numbers to plan with. They're drawn from what students actually manage when they work properly, not from what anyone wishes were true. Use them to build a term plan and to stop treating your own pace as a verdict.

Reading speeds, by gear

Gear	Rate	What you get	Use for
Survey	a chapter in 10 min	the map	everything, always
Orient	a section per hour	what is true	the week ahead
Working	2–6 pages per hour	it sticks	current courses
Verify	1 page per hour	mastery	the core results

Costs of the individual moves

Move	Time
Unpacking drill on one definition	8–15 min
Building the four-object example bank	10 min
Splitting and auditing a theorem statement	4 min
Extracting a proof skeleton	5 min
Verifying a one-page proof line by line	20–40 min
Reconstructing a proof, first attempt	20 min
Reconstructing it, third attempt	4 min
Pass 1 over a chapter	20–30 min
Pass 2 over a chapter	2–4 hours
Pass 3, exercises and summary sheet	1–2 hours
A summary sheet, two-pen method	25 min
Triage of a research paper	20 min
Notation table for a whole book	15 min

A term plan, worked

A twelve-week course, fifteen pages of dense text a week.

- **Total reading:** 180 pages.
- **If all of it were done at working pace** (4 min/page): 12 hours, which you don't have on top of lectures and problem sets.
- **So triage.** Everything gets a survey (10 min per chapter, 2 hours per term). The core four or five results get verify depth (about 4 hours). Everything the problem sets touch gets a working read (about 5 hours). The rest gets orientation, and you accept knowing what is true without knowing why.
- **Realistic total:** 11–12 hours per course per term of reading, plus exercises. That's under an hour a week per course, done deliberately, and it beats four hours a week done by rereading.

The two-hour block

Minutes	Activity
0–10	Survey, or resurvey last session's material
10–55	Working read, pen moving; mark the first proof that resists
55–65	Break, away from the desk, genuinely
65–105	The marked proof at verify depth, or exercises
105–120	Book closed: write what you met, from memory

The last fifteen minutes is the highest-value block on the page and the first thing everyone cuts. Protect it.

What is normal

- Fifty minutes on one page, when the page is a hard definition: normal.
- Not finishing a hard exercise after twenty minutes: normal, and it is why the twenty-minute limit exists.
- Rebuilding a proof and stalling three times on the first attempt: normal, and the stalls are the useful output.
- Reading a research paper and understanding the introduction only: normal, and often the correct stopping point.
- Needing three passes over a chapter before it makes sense: normal. Nobody understands a chapter on the first pass, including the people who write them.

E

Say It Out Loud

The rule from Chapter 6: if you can't pronounce a line as an English sentence, it's not yet available to you. This is the starter table. Add every symbol you meet, with the meaning *this book* gives it, because conventions differ between authors.

Logic and quantifiers

Symbol	Say	Watch out for
$\forall x$	for every x	order relative to $\exists$
$\exists x$	there exists an x	may depend on earlier variables
$\exists! x$	there's exactly one x	two obligations, not one
$\neg P$	not P	flips quantifiers when pushed inward
$P \wedge Q$	P and Q	
$P \vee Q$	P or Q	inclusive or, always
$P \Rightarrow Q$	P implies Q	not the same as its converse
$P \iff Q$	P if and only if Q	two proofs required
$\square$	end of proof	

Sets

Symbol	Say	Watch out for
$x \in A$	x is in A	element, not subset
$A \subseteq B$	A is contained in B	
$A \subset B$	A is contained in B	proper, for some authors
$A \cup B$	A union B	
$A \cap B$	A intersect B	
$A \setminus B$	A minus B	
A^c	the complement of A	complement in what?
$\emptyset$	the empty set	often a hidden hypothesis
$\{x \in A : P(x)\}$	the set of x in A such that P	
$\bigcup_{n=1}^{\infty} A_n$	the union over n from 1 to infinity	
$\mathcal{P}(A)$	the power set of A	
$\lvert A\rvert$	the cardinality of A	same bar as absolute value

Numbers and functions

Symbol	Say	Watch out for
$\mathbb{N}$	the naturals	contains 0? Author-dependent
$\mathbb{Z}, \mathbb{Q}, \mathbb{R}, \mathbb{C}$	the integers, rationals, reals, complexes	
$f : A \to B$	f maps A into B	into, not onto
$x \longmapsto x^2$	x goes to x squared	the rule, not the function
$f\vert_S$	f restricted to S	
$f^{-1}(B)$	the preimage of B	exists without an inverse
$f \circ g$	f composed with g	g acts first
(a, b)	open interval, or ordered pair, or gcd	type decides
$[a, b]$	closed interval	
$\lvert x\rvert$	the absolute value of x	
$\lVert x\rVert$	the norm of x	which norm?
$a \mid b$	a divides b	not a fraction bar
$n!$	n factorial	$0! = 1$
$\binom{n}{k}$	n choose k	

Analysis

Symbol	Say	Watch out for
$a_n \to L$	a_n tends to L	
$\lim_{n\to\infty} a_n$	the limit as n goes to infinity	
$\limsup a_n$	the lim sup	exists when the limit doesn't
$\liminf a_n$	the lim inf	
$\sup A$	the supremum of A	need not be in A
$\inf A$	the infimum of A	
$\sum_{n=1}^{\infty} a_n$	the sum from n equals 1 to infinity	
ε, δ	epsilon, delta	always positive
$\int_a^b f(t)\,dt$	the integral from a to b of f of t dee t	
$f', \dot{f}$	f prime, f dot	prime also means "derived set"
$\overline{A}$	the closure of A	or conjugate, or mean
$B(x,r)$	the ball about x of radius r	open or closed?

Algebra and elsewhere

Symbol	Say	Watch out for
$\lvert G \rvert$	the order of G	same bar again
$H \leq G$	H is a subgroup of G	not an inequality
G/H	G mod H	needs H normal
$a \equiv b \pmod{n}$	a is congruent to b mod n	
$\cong$	is isomorphic to	
$\sim$	is equivalent to, or is asymptotic to	context decides
$\langle x, y \rangle$	the inner product, or the generated subgroup	
$\oplus$	direct sum	
$\times$	product, or cross product	
$\det M$, $\lvert M \rvert$	the determinant of M	
$\lfloor x \rfloor$	the floor of x	
$\lceil x \rceil$	the ceiling of x	

The Greek letters students mispronounce

ε	epsilon	ξ	ksi
ζ	zeta	φ	phi
η	eta	χ	kai
θ	theta	ψ	psi
ι	iota	ω	omega
κ	kappa	Λ	capital lambda
μ	mu	Σ	capital sigma
ν	nu	Π	capital pi
ρ	rho	Ω	capital omega
σ	sigma	$\aleph$	aleph
τ	tau	∂	partial, or "del"

Nobody will laugh if you say a Greek letter wrongly, but you'll hesitate before speaking, and hesitation in a tutorial costs you the question you were about to ask. Learn the sounds once.

F

Further Reading, Honestly Annotated

Short list, and each entry says who it's actually for. I have left out books I haven't used.

On reading and studying mathematics

Lara Alcock, *How to Study as a Mathematics Major* (also published as *How to Study for a Mathematics Degree*). The closest thing to a companion for this book, written by a mathematics education researcher. Stronger than anything else available on the psychology of the transition to proof-based courses, lighter on procedure. Read it alongside this one, not instead.

Lara Alcock, *How to Think About Analysis*. A guided tour of the objects of a first analysis course, aimed at the reader who has the definitions and can't see what they mean. Excellent as the "gentle book" in the two-book strategy of Chapter 16.

Paul Halmos, "How to Write Mathematics" (1970 essay, in *How to Write Mathematics*, AMS). About writing rather than reading, and fifty-five years old, and still worth an evening. Halmos's spiral method of exposition explains a good deal about why textbooks are ordered the way they are.

On proof, for the reader who wants more

Daniel Velleman, *How to Prove It*. The standard bridge book. Slow, thorough, and full of exercises with answers, which makes it one of the few genuinely self-study-friendly options. If the logical structure material in Part III of this book was new to you, go here next.

Kevin Houston, *How to Think Like a Mathematician.* Broader than Velleman and more conversational, with a good chapter on reading proofs specifically. The exercises are the weak point.

George Pólya, *How to Solve It.* About solving rather than reading, and the single most quoted book in mathematics education. The heuristics chapter repays rereading every couple of years. Dated in its examples and undimmed in its content.

Gentle first books, by subject

These are the "one level below" texts referred to in Chapter 16. Each is worth reading fast before the standard text.

- **Analysis:** Stephen Abbott, *Understanding Analysis*. Written to explain why the definitions are what they are. Then go to Rudin if you need the terseness, or Tao if you want the construction from the ground up.
- **Linear algebra:** Sheldon Axler, *Linear Algebra Done Right* for the structural view; Gilbert Strang for the computational and applied one. They're different books about different subjects and both are worth having.
- **Abstract algebra:** Fraleigh for gentleness, Dummit and Foote as the reference you'll keep. Dummit and Foote isn't a first read; it's where you look things up for the next ten years.
- **Topology:** Munkres, and it's unusually well written for a standard text. Read the first two chapters slowly and the pace picks up on its own.
- **Number theory:** Silverman, *A Friendly Introduction*, which is friendly in fact and not only in title.

On the mathematical life

Timothy Gowers (ed.), *The Princeton Companion to Mathematics.* The best possible object for the shape-reading of Chapter 18. Every article gives you the map of a field in twenty pages. Read it at random for pleasure; you'll absorb an enormous amount of context.

Terence Tao's blog, and his "career advice" pages. Free, and the writing on how to read, how to learn, and what rigour is for is better than most of what is in print.

The companion volume

Gaurav Tiwari, *How to Write Mathematics: Proofs, Papers, and Problem Sets for Students.* The other half of this book. Reading is decompressing what somebody else compressed; writing is compressing your own thinking so that a reader can unfold it. Each makes the other easier, and the two share a vocabulary deliberately: the skeletons, shapes, and structure words in Part IV here are the same objects you will build when the proof is yours to write.

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 changed how an evening with a textbook goes, two minutes of your help matter more than you'd guess: leave a short, honest review wherever you bought it, or lend your copy to the classmate who has been on page 41 for a week.

The appendices are working tools, not reading. The unpacking drill and the stuck protocol are meant to be photocopied, taped inside a notebook cover, and used with a pen in your hand. You'll find printable companions and more of my writing for students at **gauravtiwari.org**. Questions and corrections are welcome there too, and future editions will owe their sharpest fixes to readers who wrote in.

Also by Gaurav Tiwari

How to Write Mathematics

Proofs, Papers, and Problem Sets for Students

The other half of this book. Reading mathematics means unfolding what somebody else compressed; writing it means compressing your own thinking so a reader can unfold it. Proof makeovers, full-credit solutions, the grader's-eye view of partial marks, and the architecture of a first paper or thesis.

LaTeX for Students

A Practical Guide to Beautiful Documents and Mathematical Typesetting

Everything you 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.

The Student Research Workflow

How to Find, Read, and Organize Research, and Write Your First Literature Review

What this book does for a textbook page, that one does for a pile of forty papers: searching, screening, note systems, reference managers, and the literature review that has to come out of them.

Other titles

- *LaTeX for Theses and Dissertations*
- *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**.