

How to Study STEM

How to Study STEM

A Working System for
Problem-Solving Subjects

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 worked harder than the person
who got the better mark.*

Contents

Preface	x
I Why It Stopped Working	1
1 The Method Is the Unit	2
1.1 Three kinds of knowledge, and only one of them is yours	2
1.2 What a method actually contains	2
1.3 Why this makes reading feel productive and be useless	3
1.4 What follows from this	3
1.5 The two-minute version	4
2 Why School Study Skills Fail	5
2.1 Four things that changed	5
2.2 Rereading, specifically	6
2.3 Highlighting, and copying	6
2.4 Cramming, and what it actually costs	6
2.5 What school got right	6
2.6 The honest cost of switching	7
3 The Illusion of Fluency	8
3.1 What fluency is	8
3.2 Three tests that break the illusion	8
3.3 Why worked examples are the sharpest version of this	9
3.4 The confidence gap, and why it widens	9
3.5 Reading this book won't help either	9
II The Mechanics	11
4 Retrieval, When There's Nothing to Recall	12
4.1 What retrieval means for a method	12
4.2 Five retrieval activities that fit STEM	12
4.3 Retrieval needs feedback, and when	13
4.4 Making failure useful	13
4.5 The 20-minute retrieval session	13
4.6 What to retrieve, and how often	14
5 Spacing	15
5.1 Why spreading works	15
5.2 The schedule, and how much precision it needs	15
5.3 Fitting it into a real week	16
5.4 Interference, and why mixing helps here too	16

5.5	What spacing is not	16
5.6	The end-of-term arithmetic	16
6	Interleaving, and Why Blocked Practice Lies	18
6.1	Blocked versus mixed	18
6.2	When to block, and when to mix	18
6.3	How to build a mixed set in five minutes	19
6.4	The extra move: identification-only drills	19
6.5	Why it feels wrong, and what to do about that	19
6.6	Interleaving across courses	20
7	Desirable Difficulty	21
7.1	The idea, stated plainly	21
7.2	Desirable, not just difficult	21
7.3	Calibrating: the 70% rule	22
7.4	Where difficulty is genuinely undesirable	22
7.5	What this changes about how you plan a session	23
8	Worked Examples, Used Properly	24
8.1	The ordering error	24
8.2	The four-step protocol	24
8.3	Faded examples: the bridge	24
8.4	How many, and which	25
8.5	Where to find good ones	26
8.6	The transfer question	26
9	Elaboration: Asking Why It Works	27
9.1	What elaboration is	27
9.2	Four questions that work on anything	27
9.3	Self-explanation, out loud	27
9.4	Concrete examples, in both directions	28
9.5	Why this is not the same as "understanding"	28
9.6	Where it fits in the week	29
III	The Artefacts	30
10	The Problem Log	31
10.1	The format	31
10.2	Why it works	31
10.3	Running it	32
10.4	Using it	32
10.5	The variant for physics and engineering	32
10.6	What it looks like after twelve weeks	32
11	Method Cards	34
11.1	The card, inverted	34
11.2	Making them, in ten minutes after a lecture	34
11.3	Reviewing them: produce, don't read	35
11.4	Paper or software?	35
11.5	What not to put on a card	35

11.6 A worked stack	36
12 The Error Log	37
12.1 The four columns	37
12.2 Rules that work and rules that don't	37
12.3 Running it	38
12.4 Turning the log into practice	38
12.5 The distribution you'll find	38
12.6 Taking it to somebody	39
13 Summary Sheets and Method Maps	40
13.1 The method map	40
13.2 The summary sheet	40
13.3 What goes on a formula sheet, and what doesn't	41
13.4 How the four artefacts fit together	41
13.5 When to build the map	42
 IV The Week	 43
14 Before the Lecture	44
14.1 What it is	44
14.2 Why it pays so much	44
14.3 The list of unknown symbols	44
14.4 When you have no time	45
14.5 What not to do beforehand	45
15 In the Lecture	46
15.1 Split the page	46
15.2 If the slides are published	46
15.3 Marking your own confusion	47
15.4 Should you record it?	47
15.5 Handwriting or typing?	47
15.6 Asking a question in the lecture	47
16 After the Lecture: The Forty-Eight Hour Rule	48
16.1 The ten minutes	48
16.2 Why the window matters	48
16.3 The blank-page part is the important part	48
16.4 Fitting it into an actual day	49
16.5 What this replaces	49
16.6 When you missed the lecture	49
17 Problem Sets Without Suffering	50
17.1 The protocol	50
17.2 The twenty-minute rule	50
17.3 Getting unstuck, in order	51
17.4 How much time a set should take	51
17.5 Collaboration, done right	51
17.6 After it's marked	52

18 Study Groups That Work	53
18.1 Format 1: the teaching round	53
18.2 Format 2: the gap swap	53
18.3 Format 3: the mixed set race	53
18.4 What makes groups fail	54
18.5 The right size and cadence	54
18.6 Working alone, if that's your preference	54
19 Office Hours	55
19.1 What to bring	55
19.2 When to go	55
19.3 The questions worth asking that aren't about a specific problem	56
19.4 Tutorials and demonstrators	56
19.5 If you're anxious about going	56
V The Term, and the Years	57
20 Planning a Term	58
20.1 The weekly template, per course	58
20.2 Four courses, one week	58
20.3 The term, week by week	59
20.4 Weeks 1 to 3: build one habit at a time	59
20.5 Mid-term tests, and what they're for	59
20.6 Coursework-heavy and lab-heavy terms	59
21 Slumps, and Getting Back	61
21.1 The spiral, and its shape	61
21.2 The recovery protocol	61
21.3 Telling somebody, and why it isn't optional	62
21.4 When it's not a study problem	62
21.5 The everyday version: motivation, not crisis	62
21.6 Comparison, and why it misleads	63
22 Labs, Projects, and Coursework	64
22.1 Labs	64
22.2 Extended projects	64
22.3 Supervisor meetings	65
22.4 Coursework-heavy modules	65
22.5 Group projects	65
22.6 Placements and industrial work	65
23 Handing Off to the Exam	66
23.1 What you should be holding in week ten	66
23.2 The two weeks	66
23.3 Where this book stops	67
23.4 The last three days, in detail	67
23.5 Afterwards	67
23.6 What you're left with	67

VI Toolkit	68
A Three Worked Weeks	69
A.1 Week 1: four modules, first year, a normal week	69
A.2 Week 2: second year, heavy coursework, one bad day	70
A.3 Week 3: final year, project-dominated	70
A.4 What the three weeks have in common	71
B The Problem Log Template	72
C Method Card Templates	74
D The Error Log Template	76
E What the Evidence Actually Says	78
F Further Reading, Honestly Annotated	80

Preface

The method that got you here stops working, usually in the second term of first year, and nobody warns you.

At school you could read the chapter, do the exercises, revise the notes, and it worked. Now you read the chapter and understand it, revise the notes and feel fine, walk into the test, and find you can't start question two. The material makes sense when you look at it. It just isn't there when you need it.

That isn't a failure of intelligence and it isn't a failure of effort. It's a mismatch between a study method and a kind of knowledge.

School subjects are mostly made of facts and worked procedures you'll be asked to repeat. University STEM is made of *methods* you'll be asked to select and apply to something you haven't seen. Those are different things, and the study techniques that work for one are close to useless for the other. Rereading builds recognition; you need production. Highlighting marks importance; you need retrieval. Doing twenty questions of the same type builds execution; the exam tests selection.

So the fix isn't to work harder. Most students who are struggling are already working plenty. The fix is to change what the hours contain.

That's what this book is: not the science of learning, which has been written up well elsewhere, but the operational version. What to do after a lecture. What a study session actually consists of. What to write down and what to throw away. What a week looks like, and what a term looks like.

A word on the evidence. Retrieval practice, spacing and interleaving are about as well established as anything in the study of learning, and I'll say where the evidence is strong and where I'm extrapolating from teaching experience. Where I'm giving you my own practice rather than a research finding, I'll say so. Appendix E is an honest summary of what the research does and doesn't support, including the parts people overstate.

Two warnings.

Every technique in here feels worse than what you're doing now. Retrieval is harder than rereading. Spacing means you'll have forgotten some of it. Mixed practice makes you slower and less accurate in the session. All three of those are the mechanism working, not a sign it isn't, and that's the single hardest thing about switching.

And this isn't a productivity book. There's nothing here about apps, morning routines or motivation systems. It's about the content of the hours you already spend, which is where the difference actually lives.

Let's start with why your old method broke.

Gaurav Tiwari
gauravtiwari.org

Part I

Why It Stopped Working

1

The Method Is the Unit

Ask a student what they're revising and they'll name a topic: integration, thermodynamics, circuits.

Ask what they'll be tested on and the honest answer is different. You'll be given a problem you haven't seen, and asked to work out which method it wants and then run it. The topic is the label on the box. The method is what's inside, and the method is the thing that has to be learnable, retrievable, and selectable under pressure.

Almost everything in this book follows from taking that seriously.

1.1 Three kinds of knowledge, and only one of them is yours




Figure 1.1: Three kinds of knowledge. School examines mostly the top two; STEM degrees examine mostly the third, and it is the only one that can't be built by reading.

At school, an exam question told you which procedure to use, usually by appearing in the chapter about it. At university nobody tells you. The question just sits there, and the first thirty seconds are entirely about identification.

That's why a student can honestly say "I know how to do integration by parts" and still fail a question that needs it. Both statements are true. They know the procedure and they can't select it.

Knowing a method and being able to choose it are separate skills, learned in separate ways. Nearly all student revision builds the first and almost none of it builds the second.

1.2 What a method actually contains

If the method is the unit, it's worth being precise about what you have to store for each one. Four things, and most students store two.

The name	“Integration by parts.” Cheap, and nearly everyone has it.
The steps	How to run it. Also cheap, built by repetition, and where most revision stops.
The tell	What a problem looks like when this is the right method. A product where one factor simplifies on differentiation. This is the one that’s missing.
The boundary	When it doesn’t work, and what to use instead. Rarely taught, and it’s what stops you grinding for ten minutes on the wrong approach.

The tell is the difference between a student who recognises a question in five seconds and one who stares at it for four minutes. And it’s learnable: it’s just a list of surface features you’ve deliberately attached to the method.

Chapter 11 is how to build and store those four fields. For now, notice that your lecture notes almost certainly contain the first two and neither of the others.

1.3 Why this makes reading feel productive and be useless

Reading a worked example gives you the steps. It doesn’t give you the tell, because the example arrived pre-identified: it was in the section on integration by parts, in a list of integration by parts problems, in a lecture titled integration by parts.

Every context cue that will be absent in the exam is present while you study. That’s the whole problem in one sentence, and it explains why the gap between “I understood the lecture” and “I could do the exam” is so reliably large.

WHAT THE EVIDENCE SAYS

Retrieval practice and the testing effect are among the best-supported findings in the psychology of learning, replicated across decades and across subject areas. The specific claim here, that STEM study needs retrieval of *method selection* rather than facts, is my own emphasis from teaching, but it follows directly from the transfer literature.

1.4 What follows from this

The rest of the book is consequences of one idea, so here’s the map.

Part II rebuilds the three well-evidenced techniques, retrieval, spacing and interleaving, around methods rather than facts. Interleaving turns out to matter far more in STEM than the popular write-ups suggest, because it’s the only one that trains selection.

Part III is the artefacts: four things you build and keep. The problem log, method cards, the error log, and one-page method maps. None takes more than a few minutes at a time, and together they’re what you actually revise from.

Part IV is the week: what to do before a lecture, in it, and in the forty-eight hours after, plus problem sets, groups and office hours.

Part V is the term and the years, including what to do when it goes badly, which it will at some point for almost everyone.

You are not trying to cover the material. You are trying to build a small number of methods you can select and run under pressure. Those are different goals, and they lead

to completely different weeks.

1.5 The two-minute version

If you read nothing else in this book:

- After every lecture, write down the method and its tell. Ten minutes.
- Do problems from a blank page, not with the notes open.
- Every week, do a few problems from three weeks ago, mixed.
- Keep a one-line log of every problem you solve and the idea that unlocked it.

Four habits, under an hour a week per course, and they cover most of the value of everything that follows.

TRY IT THIS WEEK

Take the last method you were taught. Write down all four fields: the name, the steps, the tell, and the boundary. If you can't write the tell, that's what you'll be missing in the exam, and it takes two minutes with your notes to fix.

2

Why School Study Skills Fail

Every technique that worked at school is still available to you, and most of them have quietly stopped paying.

It's worth knowing exactly why, because the failure isn't obvious from the inside. All of them still *feel* like studying.

2.1 Four things that changed

Volume. A school subject covered maybe forty topics in two years. A single university module covers thirty in twelve weeks, and you're taking four modules. Techniques that scale with time spent per topic stop fitting.

The unit of knowledge. From facts and drilled procedures to method selection, as in Chapter 1.

Who decides what matters. At school, the teacher told you what would be examined and drilled it. Now the lecturer presents the material once, and deciding what matters is your job.

The gap between contact and assessment. School tested you every few weeks. A university module might test you once, fourteen weeks after the first lecture, which means anything you don't deliberately maintain will be gone.




Figure 2.1: The same four hours, spent two ways. The left column is what most students do; the right is what the same time buys if it is spent producing rather than reviewing.

Every row on the left produces recognition. Every row on the right produces retrieval. That distinction is the whole of Part II.

2.2 Rereading, specifically

The most popular study technique in the world, and among the least effective per hour.

It fails for a precise reason: rereading builds *familiarity*, and familiarity feels exactly like knowledge from the inside. You finish the page thinking “yes, I know this”, and what you actually know is that you’ve seen it before.

The research here is unusually clean. Across many studies, students who reread score little better than students who read once, while students who tested themselves score substantially better, even when the testers spent less total time.

WHAT THE EVIDENCE SAYS

The comparison of rereading against self-testing is one of the most replicated results in this area. The size of the effect varies with material and delay; the direction doesn’t. And the finding that students *prefer* rereading and predict it will work better is equally robust, which is the awkward part.

2.3 Highlighting, and copying

Highlighting marks what’s important. It doesn’t move anything into your head, and it produces a document that looks studied.

Copying out notes is worse per hour, because it takes longer and requires even less thought. If you’re going to write, write something you had to produce: a method card, a summary from memory, a solution attempted before you looked.

There’s one legitimate use for both. If highlighting or copying is how you *select* what matters, and the selection then feeds a method card or a map, fine. As an end in itself, it’s an hour spent on transcription.

2.4 Cramming, and what it actually costs

Cramming works, in the narrow sense that it gets material into your head for tomorrow.

Its problem is that it doesn’t survive, and in a degree that matters more than at school. The second-year module assumes the first-year one. The final-year project assumes both. A student who crams each module in turn arrives at year three with three years of material they’ve each known for about a week.

Cramming is a loan against your future modules, at a high rate of interest. It is occasionally the right decision, and it is never a system.

2.5 What school got right

Two things worth keeping, because this chapter isn’t an argument that everything you learned is wrong.

Doing lots of problems works. It’s the single most valuable habit you brought with you. What changes is *which* problems: mixed rather than blocked, and attempted before the solution is read (Chapters 6 and 8).

Writing things down works. Also true, and also needs redirecting: write the things you had to produce, not the things you copied.

2.6 The honest cost of switching

Everything in Part II will feel worse for two or three weeks. Your recall in the session will be poorer, your problem sets will feel harder, and you'll be slower.

That's not a side effect. Retrieval is effortful, spacing means partial forgetting, and mixing means you can't coast. The discomfort is the mechanism, and if a study session feels easy and pleasant, it's usually because you're recognising rather than producing.

Give it three weeks before judging, and judge it on a cold test rather than on how the sessions felt.

TRY IT THIS WEEK

Look at your last three study sessions honestly and put each activity in the left or right column of Figure 2.1. Most students find four-fifths of their time is in the left column. That ratio, not the number of hours, is what this book is trying to change.

3

The Illusion of Fluency

The reason bad study methods survive is that they feel good, and the feeling is systematically misleading.

This chapter is about that specific error, because until you stop trusting the feeling you'll keep drifting back to rereading.

3.1 What fluency is

When material is in front of you and it's easy to follow, your brain reads that ease as understanding.

It isn't. It's fluency: a judgement about how smoothly the information is going in, which is affected by how recently you saw it, how clearly it's laid out, and how confident the lecturer sounded. None of those has much to do with whether you could reproduce it tomorrow.




Figure 3.1: Fluency against actual retention, for four common study activities. The two most comfortable activities are the two that leave least behind.

The dashed line isn't a measured curve, and I'm not going to pretend it is. It's the shape the evidence keeps producing: the activities students rate as most effective are reliably the ones that leave least behind.

3.2 Three tests that break the illusion

You can't fix a miscalibrated feeling by deciding to feel differently. You fix it by measuring, and there are three cheap measurements.

The blank page test. Close everything and write what you know about the topic: the methods, their tells, the key results. Two minutes. Then compare with your notes. The gap is the truth, and it's usually shocking the first time.

The cold problem test. Take a problem from three weeks ago that you solved successfully. Do it again, cold, with nothing open. If you can't start it, you never owned it; you followed it.

The explanation test. Explain the method out loud to somebody who doesn't do your course, or to a wall. The moment you stall is the moment your understanding actually ends, and it's usually two sentences in.

Never judge your understanding from inside a study session. Judge it from a cold start, on a different day, with nothing in front of you. Everything else is a measurement of familiarity.

3.3 Why worked examples are the sharpest version of this

Reading a worked solution produces the strongest illusion of fluency available, because every step is justified and follows the last. You nod along, it all makes sense, and you learn very little.

The fix isn't to avoid worked examples, which are genuinely valuable. It's the order: attempt first, then read. Chapter 8 is the full treatment, and it's the single most common repair I make to a student's week.

3.4 The confidence gap, and why it widens

There's a well-documented pattern where the students most confident about their preparation are often not the ones who do best, and it has a mechanism rather than being a joke about arrogance.

Confidence is built from fluency. A student who reread five times has a strong fluency signal and little retrieval. A student who tested themselves has a weak fluency signal, because they kept experiencing failure, and much better retention.

So the second student walks in feeling less prepared and performs better, and both of them are surprised.

WHAT THE EVIDENCE SAYS

The general finding that subjective confidence is a poor predictor of performance, and that difficulty during learning improves retention, is well supported. The specific size of the effect in university STEM courses is less studied than in laboratory settings, so treat the direction as solid and the magnitude as indicative.

3.5 Reading this book won't help either

Worth saying plainly, because it applies to the thing in your hands.

Reading about retrieval practice produces exactly the fluency illusion this chapter is about. You'll finish this chapter agreeing with it, and that agreement is worth nothing on its own.

Every chapter ends with a box called *Try It This Week* for that reason. The chapters are the argument; the boxes are the book. If you do none of them, you'll have had a pleasant few hours and changed nothing, and you'll be an excellent example of the effect described above.

TRY IT THIS WEEK

Do the blank page test right now, on whatever topic you studied most recently. Two minutes, nothing open, write what you know. Then open your notes and mark what was missing. That gap is your real starting point, and everything in Part II is aimed at closing it.

Part II

The Mechanics

4

Retrieval, When There's Nothing to Recall

Retrieval practice is the best-evidenced study technique there is, and every popular write-up of it uses flashcards and vocabulary.

Which leaves the STEM student nodding and then wondering what to do, because you don't need to recall that the chain rule exists. You need to be able to *use* it on something you haven't seen.

4.1 What retrieval means for a method

It means producing, from a blank page, with nothing open.

Not retrieval	Retrieval
Reading a worked example	Attempting it, then checking
Reviewing your solution	Redoing the problem cold, a week later
Looking at a method card	Writing the method card from memory
Following a derivation	Reproducing it with the book shut
Recognising which method it is because it's in that chapter	Identifying the method from a mixed set

The common feature on the right: at the moment of the activity, the answer isn't available to you. That's the whole mechanism. Effortful reconstruction is what strengthens the memory, and if it isn't effortful it isn't doing much.

If the answer is visible while you work, you are reviewing, not retrieving. Cover it, produce it, then check.

4.2 Five retrieval activities that fit STEM

Number 1 does most of the work. If you take one habit from this chapter, make it three cold problems from an older topic, twice a week, per course.




Figure 4.1: Five ways to retrieve a method rather than a fact. All five are things you produce; none involves looking at anything.

4.3 Retrieval needs feedback, and when

Retrieval without checking is just guessing into the void, so every attempt gets checked. The question is when.

For a short factual item, check immediately. For a problem, finish the whole attempt first, including the parts you're unsure about, and check afterwards. Checking mid-problem converts a retrieval attempt into a worked example, and you lose most of the benefit.

If you're properly stuck, use the twenty-minute rule: twenty minutes on one obstacle, then look. Grinding past that produces frustration rather than learning.

4.4 Making failure useful

You will fail a lot of retrieval attempts, especially at first. That's expected, and the failures are where the information is.

For each one, write a single line: what you couldn't produce, and what was missing. That line goes into the error log (Chapter 12). Across a term those lines cluster into four or five recurring gaps, and those gaps are your actual revision plan.

WHAT THE EVIDENCE SAYS

Retrieval practice works even when the attempt fails, provided feedback follows. Failed retrieval followed by the correct answer beats simply studying the answer, which is counter-intuitive and well replicated.

4.5 The 20-minute retrieval session

Concrete, so it's easy to start:

0–2 min	Blank page: write every method from the chosen topic, and its tell.
2–5 min	Check against your notes. Mark what was missing.
5–17 min	Three problems from that topic, cold, timed at four minutes each. Attempt fully before checking anything.
17–20 min	Check. One line in the error log per failure.

Twenty minutes, and it will feel much worse than twenty minutes of rereading. That's the point, and Chapter 7 is about why.

4.6 What to retrieve, and how often

Not everything, and not on a schedule you won't keep.

Per course, per week: one 20-minute session on a topic from about three weeks ago, plus the ten minutes after each lecture from Chapter 16. That's roughly forty minutes a week per course, and it's on top of the problem sets you're already doing.

Which topic? The one from three weeks ago that you'd least like to be tested on. That instinct is a good detector and students almost always override it.

TRY IT THIS WEEK

Pick the topic from three weeks ago that you'd least like a surprise test on. Do the 20-minute session above on it, today. Note how the first two minutes feel compared with rereading, and note how much came back by the end.

5

Spacing

Four twenty-minute sessions spread across a term beat one eighty-minute session the week-end before, by a distance that surprises people.


Same total time. Very different outcome.

5.1 Why spreading works

Because forgetting is part of the mechanism.

When you return to a topic after a gap, some of it has faded, and pulling it back is effortful. That effort is what strengthens it, exactly as in Chapter 4. Return too soon and there's nothing to retrieve, because it's all still sitting there; return after a gap and the retrieval is real.

Which produces the uncomfortable rule: the session that feels worst is usually the one doing most.



Closed book every time. Axis is ordered, not to scale.

Figure 5.1: Spacing one topic across a term. Five contacts, eighty minutes of revision after the first learning session, and each contact is a test rather than a read.

5.2 The schedule, and how much precision it needs

Very little. The gains come from spreading at all, not from hitting particular intervals.

A workable default: revisit at one day, one week, three weeks, six weeks. If you miss one, do it late; a late revisit is worth nearly as much as an on-time one and far more than a skipped one.

What doesn't work is the schedule you won't follow. Elaborate systems with daily queues collapse in week four, and a collapsed system does nothing at all.

Spread at all, and don't optimise. Four contacts at roughly increasing gaps captures most of the available benefit, and any system you'll actually run beats a better one you won't.

5.3 Fitting it into a real week

The practical problem is remembering to do it, so attach it to something that already happens.

Attach to the problem set. Before you start this week's sheet, do three problems from three weeks ago. Fifteen minutes, and the sheet is already in your calendar.

Attach to the lecture. The ten minutes after each lecture covers the one-day interval automatically.

One calendar block per course per week. Twenty minutes, named after the course, no topic specified until you sit down. Then pick the oldest topic you feel worst about.

That's the whole system. If you want something more structured, spaced-repetition software works well for method cards, and Chapter 11 covers what to put on them.

5.4 Interference, and why mixing helps here too

One complication worth knowing. Studying two similar topics in adjacent sessions can cause interference, where the newer one overwrites the older.

The fix is the same as the fix for selection: mix them deliberately rather than avoiding the clash. If integration by parts and substitution are getting confused, do a mixed set containing both and force yourself to choose. Practising the discrimination is what removes the confusion, and avoiding it preserves it.

That's Chapter 6, and it's the other half of the same idea.

WHAT THE EVIDENCE SAYS

The spacing effect is one of the oldest and most robust findings in the psychology of memory, going back well over a century and replicated across material types, ages and delays. The specific optimal intervals are much less settled, which is why the advice here is "spread, don't optimise".

5.5 What spacing is not

Not a substitute for the first learning session. You still have to understand it once, properly. Spacing maintains and strengthens; it doesn't install.

Not an excuse to do less per session. Twenty minutes of real retrieval, not twenty minutes of skimming.

Not a way to avoid problem sets. The sets are your main practice. Spacing is the maintenance layer on top.

5.6 The end-of-term arithmetic

Twelve weeks, one course, one twenty-minute spaced session a week: four hours across the term.

Compare with the standard alternative, which is twelve to twenty hours of revision in the fortnight before the exam, on material that's mostly gone and has to be relearned rather than recalled.

The spaced version is less total time, spread thinly, and arrives at the exam with the material still present. That's the entire argument, and it's why Chapter 20 builds the term around it.

TRY IT THIS WEEK

Open your calendar and put one twenty-minute block per course per week for the rest of term, at a fixed time. Don't specify the topic. When the block arrives, pick the oldest topic you feel least confident about and do the 20-minute session from Chapter 4. That's the whole system, and it takes five minutes to set up.

6

Interleaving, and Why Blocked Practice Lies

This is the chapter that matters most in a STEM degree, and it's the one the popular write-ups treat as a footnote.

Do twenty integration-by-parts questions in a row and you'll feel fluent by question fifteen. You aren't. You've learned to execute a method somebody already selected for you, and selection is what the exam tests.

6.1 Blocked versus mixed




Figure 6.1: The same twenty problems, two ways. Blocked practice feels better and produces worse performance on a mixed test, which is what an exam is.

The awkward middle rows are real. In the session, mixed practice is slower, less accurate and less satisfying. On a delayed mixed test, it wins, often substantially.

Blocked practice builds execution. Mixed practice builds selection. Exams test selection first, and it's the half most revision skips entirely.

6.2 When to block, and when to mix

Blocking isn't wrong, it's a phase.

Stage	What to do
Meeting a method for the first time	Block. Five or six problems of the same type, to get the mechanics working.
Once you can execute it reliably	Stop blocking. Further blocked practice adds very little.
From then on	Mix, including older topics. Every set should require you to identify before you execute.
Revision	Always mixed, always with older material.

Most students never leave the first row, because problem sheets are organised by topic and so are textbooks. The material is arranged for teaching and you have to deliberately disarrange it for practice.

6.3 How to build a mixed set in five minutes

Nobody hands you one, so make it:

1. Take four topics you've covered, including one from at least a month ago.
2. Pick three problems from each, from your problem sheets or the textbook.
3. Write the twelve problem references on a sheet, in shuffled order.
4. Work them in that order, cold, writing the *tell* before the method each time.

That fourth step is the whole point. For each problem, write one line before you start: “product with a factor that simplifies on differentiation, so parts”. You're practising the identification explicitly rather than hoping it develops.

6.4 The extra move: identification-only drills

A drill I use with students who can execute everything and freeze in exams.

Take twenty problems from mixed topics. Don't solve any of them. For each, write only the method you'd use and the tell that told you. Twenty problems in twelve minutes.

It feels like cheating, because you're skipping the work. It's the most efficient selection practice available, and it isolates exactly the skill that's failing.

WHAT THE EVIDENCE SAYS

Interleaving effects are well demonstrated in mathematics practice specifically, which is unusual: much of the learning-science literature uses artificial material, but the interleaving work includes real maths problems. Session performance drops, delayed test performance rises.

6.5 Why it feels wrong, and what to do about that

Three things happen when you switch to mixed practice, and all three read as failure:

You get slower. Because you're now doing the identification step you were skipping.

You make more mistakes. Same reason.

You feel less confident afterwards. Because confidence tracks fluency, per Chapter 3.

The measurement that matters is a cold mixed test a fortnight later, not how the session went. Run mixed practice for three weeks before you judge it, and judge it on that.

6.6 Interleaving across courses

An extension worth trying in second year and beyond.

Mix problems from two modules that share machinery: the differential equations from your maths module and the ones in your circuits module, or the linear algebra in maths and the transformations in mechanics.

Students compartmentalise by module and are then unable to use a method in the other course's exam, which is one of the more frustrating things to watch. The fix is to practise them together at least occasionally.

TRY IT THIS WEEK

Build one mixed set this week: twelve problems, four topics, one of them at least a month old, shuffled. Before each problem, write the tell and the method. Then compare how it feels against your usual topic-by-topic sheet, and note that the difficulty is the point.

7

Desirable Difficulty

Three chapters in a row have now told you that the version which feels worse is the version that works.

That isn't a coincidence, and it has a name.

7.1 The idea, stated plainly

Conditions that make learning harder in the moment often make it stick better afterwards. Conditions that make it feel smooth often produce nothing that lasts.

Easy now	Harder now, better later
Reading the solution	Attempting it first
Same topic all session	Mixed topics
Revising the day after	Revising after a gap
Notes open	Notes closed
Recognising an answer	Producing an answer
One long session	Several short ones

The right-hand column is the whole of Part II. And the pattern is why students so reliably choose wrong: the left column is more pleasant, produces higher accuracy during the session, and generates more confidence.

In the moment, easier usually means less is being learned. That doesn't mean maximum difficulty is best, and the next section is about the difference.

7.2 Desirable, not just difficult

The qualifier matters, and it's the part that gets dropped when this idea is repeated.

A difficulty is desirable when it makes you *do the work of retrieval*. It isn't desirable when it just prevents progress.

Desirable	Just difficult
Attempting before reading the solution	Attempting a problem four topics beyond you
Notes closed	Notes lost
Mixed practice	Random material with no structure
A gap of a week	A gap of a year
A problem at the edge of your ability	A problem requiring a method nobody taught you
Working without worked examples available	Working without ever seeing a worked example

The test: after the effort, could you resolve it with a reasonable amount of checking? If yes, the difficulty was productive. If you simply had no route in, it was wasted time, and the answer is to go back a level rather than to grind.

7.3 Calibrating: the 70% rule

A rough target that works well in practice.

Aim for practice where you get roughly seven out of ten right. That's hard enough to require real retrieval and easy enough to keep you moving.

Getting	Do this
Nearly all right	Too easy. Add older topics, remove the notes, or move to harder problems.
About 7 in 10	About right.
Fewer than half	Too hard. Go back to worked examples for that topic, then return.
None	Wrong level entirely, or a missing prerequisite. Find the prerequisite.

That last row deserves emphasis. If nothing is going in, the problem is usually a missing earlier idea rather than insufficient effort, and no amount of desirable difficulty fixes a missing prerequisite.

WHAT THE EVIDENCE SAYS

The desirable difficulties framework comes from Robert and Elizabeth Bjork's work, and the individual effects under it, spacing, testing, interleaving, varied practice, are each well supported. The framework's practical limit is that "desirable" is easier to identify after the fact than before, which is why the 70% heuristic is a teaching rule of thumb rather than a research finding.

7.4 Where difficulty is genuinely undesirable

Three cases where making it harder helps nobody.

Bad materials. An unreadable textbook or an incoherent set of notes is a difficulty that teaches nothing. Get a better book.

Exhaustion. Effortful retrieval requires a functioning brain. After ninety minutes, or on four hours of sleep, difficulty stops being productive and starts being noise.

Anxiety. There's a level of stress above which retrieval simply fails. A student who is panicking is not experiencing a desirable difficulty, and the fix is to lower the stakes for that session, not to push harder.

7.5 What this changes about how you plan a session

Practically, four rules:

- Close the notes before you start, not after you get stuck.
- Attempt before reading, always.
- Include older material every time.
- Stop when the accuracy collapses for tiredness rather than for difficulty, and come back tomorrow.

TRY IT THIS WEEK

Take your next study session and change exactly one thing: close everything before you start, and only open it after you've produced something. Note how much worse it feels and how much more you can produce the next day. That contrast is the whole of Part II in one session.

8

Worked Examples, Used Properly

Worked examples are genuinely valuable and almost universally misused.

The misuse is a single ordering error, and fixing it is the cheapest improvement in this book.

8.1 The ordering error

BEFORE

Read the worked example. Nod. Understand each step. Move to the next worked example. Then attempt the problem set and discover you can't start.

AFTER

Cover the solution. Attempt it yourself, on paper, for a bounded time. Then read the solution, and specifically find the step you didn't have. Then close it and redo the problem from scratch.

Same example, three times the value, and about six extra minutes.

The reason is Chapter 3: reading a solution generates the strongest fluency signal available, because every step follows the last and nothing is missing. You experience understanding without producing anything.

Never read a solution to a problem you haven't attempted. The attempt is what creates the gap that the solution then fills, and without the gap the solution has nowhere to land.

8.2 The four-step protocol

Step 4 is non-negotiable and takes five minutes. Without it you've understood a solution; with it you've produced one, and only the second transfers.

8.3 Faded examples: the bridge

There's an intermediate technique between reading solutions and solving cold, and it's underused.

Take a worked solution and cover progressively more of it:




Figure 8.1: How to use a worked example. Step 4 is the one everyone skips, and it's where the learning is consolidated.

First pass	Cover the last step. Predict it, then check.
Second	Cover the last two steps.
Third	Cover everything after the setup line.
Fourth	Cover everything.

Four passes over one example, each a real retrieval attempt at increasing difficulty. It's especially good for long derivations, where jumping straight to a cold attempt fails so completely that nothing is learned.

WHAT THE EVIDENCE SAYS

Worked-example research shows they're most valuable early in learning a method, and that their advantage disappears, and can reverse, once a learner is competent. That's the expertise-reversal effect, and it's a good argument for the faded sequence above: reduce the support as competence grows.

8.4 How many, and which

Two or three per method, not twenty.

Choose them for variety rather than similarity. Two examples that differ only in numbers teach you one thing; two that differ in structure teach you what varies and what stays fixed, which is what the tell is made of.

And after the second one, stop reading and start attempting. Additional worked examples have sharply diminishing returns once you can start a problem yourself.

8.5 Where to find good ones

Lecture examples	Closest to your examiner's taste. Use first.
Textbook examples	Usually the best explained, and often more general.
Problem sheet solutions	Excellent, if you attempt before reading.
Video walkthroughs	Fine for a first exposure; the passivity is the risk. Pause and predict.
Solution manuals	Dangerous. Easy to read many and produce nothing.

The video row deserves a note. Watching somebody solve a problem is the most fluent, least productive study activity available, and it feels excellent. If you use videos, pause before every step and predict what comes next. That turns it into retrieval.

8.6 The transfer question

After every worked example, ask one question: *what would change if the problem were slightly different?*

- What if the limits were different?
- What if the coefficient were negative?
- What if the boundary condition were at the other end?
- What if n were even rather than odd?

Thirty seconds, and it's the difference between memorising an example and learning a method. It also generates your boundary field from Chapter 1, which is the field nobody has.

TRY IT THIS WEEK

Take the next worked example you'd normally read, and run the four steps instead: attempt, read hunting for the missing step, name the move, redo cold. Time it. It takes about twenty minutes rather than four, and you'll be able to do the problem tomorrow, which the four-minute version never achieves.

9

Elaboration: Asking Why It Works

The last of the mechanics, and the one that separates a student who can pass from one who can adapt.

Retrieval gets the method into your head. Spacing keeps it there. Interleaving lets you select it. Elaboration is what lets you use it on something genuinely new.

9.1 What elaboration is

Asking and answering *why* questions about material you already technically know.

Knowing	Elaborating
Integration by parts is $\int u dv = uv - \int v du$	Why does choosing u as the factor that simplifies make the second integral easier?
$F = ma$	Why is it mass times acceleration and not mass times velocity? What changes if mass varies?
The determinant tells you if a matrix is invertible	Why does a zero determinant mean the columns are dependent?
Resistors in parallel: $1/R = 1/R_1 + 1/R_2$	Why does adding a resistor in parallel always <i>lower</i> the total?

Each right-hand question takes a minute or two to answer, and each one attaches the method to something that makes it retrievable and adaptable.

9.2 Four questions that work on anything

The fourth is the one with the largest payoff in a degree, because STEM modules are taught in silos and the same mathematics turns up in four of them under four names. A student who notices that the RC circuit, the cooling cup of coffee and the radioactive decay are the same differential equation has learned one thing instead of three.

Every time you meet a method that feels new, ask what it resembles. Most of a STEM degree is a small number of structures wearing different vocabulary, and spotting that is most of what people mean by “getting it”.

9.3 Self-explanation, out loud

The cheapest version of elaboration, and the best evidenced.




Figure 9.1: Four elaboration questions. Ask them of any method and you will know it better than a student who has done twice as many problems.

While working through a solution, explain each step to yourself out loud: why this step, why now, what it achieves. Not what it says, why it's there.

It's slow, it feels strange, and it roughly doubles what you retain from a worked example. It also finds the gaps immediately, because the moment you can't explain a step is visible in a way that reading past it isn't.

WHAT THE EVIDENCE SAYS

Self-explanation is among the better-supported techniques in this area, with effects shown specifically on physics problem-solving. It's also one of the few that helps most for students who are struggling, which is unusual: most techniques help the already-strong more.

9.4 Concrete examples, in both directions

Two moves, and students usually do only the first.

Abstract to concrete. You've met a general result; construct the simplest specific case. Set $n = 1$. Take the identity matrix. Make the damping zero. If the general result doesn't behave sensibly in the simple case, you've misunderstood it, and it takes ten seconds to check.

Concrete to abstract. You've done a specific problem; ask what class of problems it belongs to. This is the direction that builds the tell, and almost nobody does it deliberately.

9.5 Why this is not the same as "understanding"

A caution, because elaboration can turn into an excuse.

Students sometimes decide they need to understand a topic deeply before practising it, and spend three hours reading around a method they still can't execute. That's not elaboration, it's avoidance, and it's more comfortable than problems.

The order that works is: get the method running first, then elaborate. Execution, then understanding, then selection. Reversing it produces students who can discuss the physics beautifully and can't answer the question.

Elaborate on methods you can already run. Understanding before execution is usually procrastination wearing a respectable coat.

9.6 Where it fits in the week

It doesn't need its own session. Attach it:

- One *why* question per method card, when you make the card.
- The transfer question after every worked example (Chapter 8).
- One connection per week: find a method that appears in two of your modules, and write the one-line link.

Ten minutes a week, and it's the part that makes the difference between knowing your modules and knowing your subject.

TRY IT THIS WEEK

Take a method you can already execute confidently. Answer all four questions from Figure 9.1 in writing, in five minutes. The fourth one, where else this appears, is usually the one that produces the surprise.

Part III

The Artefacts

10

The Problem Log

One line per problem you solve, all term.

It's the artefact students are most sceptical about and most convinced by afterwards, and it costs about eight minutes a week.

10.1 The format

Three fields: the problem, the method, and the idea that unlocked it.

Sheet 4, Q3. Volume of revolution with a hole. Shells, not discs. **Key idea:** the radius is the distance from the axis, not from the origin.

Sheet 5, Q1(b). Second-order ODE, repeated root. Complementary function has a $te^{\lambda t}$ term. **Key idea:** a repeated root means the second solution is t times the first.

Sheet 6, Q4. Beam reactions. Moments about the left support. **Key idea:** choosing the support as the pivot kills one unknown immediately.

Tutorial, Q2. Series convergence. Comparison with $\sum 1/n^2$. **Key idea:** bound the numerator first, then compare.

Thirty seconds each. Note what the third field records: the *idea*, not the answer. In three months you won't need the answer. You'll need to know that taking moments about a support removes an unknown.

Log the idea that unlocked the problem, not the problem and not the answer. The idea is the transferable part, and it's the part you'll have forgotten by the exam.

10.2 Why it works

Three reasons, and the third is the one people don't anticipate.

It's revisable. By week twelve you have a two-page compressed index of every technique in the course. Your folder of full solutions never will be, because rereading a full solution takes as long as writing it did.

Writing the key idea is elaboration. Deciding what the unlocking move was forces you to abstract from the specific problem, which is the concrete-to-abstract direction from Chapter 9.

It shows you the shape of the course. After thirty entries you can see that eight of them are the same idea in different clothes. That consolidation happens automatically and it's how a module stops feeling like forty unrelated problems.

10.3 Running it

One notebook or one file, not scattered.

Log immediately after the problem set, not at the end of term. It takes thirty seconds per problem while the solution is fresh and about five minutes per problem three weeks later.

Log the ones you failed too. Same format, with the key idea being whatever you were missing. Those entries are usually the most valuable.

Don't log everything. Routine problems that applied a method you already had don't need an entry. If nothing unlocked, there's nothing to record.

10.4 Using it

Three uses, roughly in order of value.

Before a mixed practice session, read the log for the topics you're about to mix. Two minutes, and it loads the tells into working memory.

Before an exam, read the whole thing. For a twelve-week module that's two pages and about eight minutes, and it's the highest-density revision material you own.

When you're stuck on a new problem, scan the log for a similar shape. Students underuse this: your own log is a searchable index of methods you've personally made work.

WHAT THE EVIDENCE SAYS

There's no direct research literature on "problem logs" as such. The components are well supported, self-explanation, abstraction of the underlying structure, and spaced review, and this is a format I've found works for students. Treat the specific format as my practice rather than a finding.

10.5 The variant for physics and engineering

For applied subjects, add a fourth field: the physical insight.

Sheet 3, Q5. Block on an incline with friction. Resolve along and perpendicular to the slope. **Key idea:** choosing the axes along the slope makes one equation trivial. **Physically:** the normal force isn't mg once the surface is tilted, which is the thing that keeps catching me.

That last clause, "which is the thing that keeps catching me", is the bridge to the error log in Chapter 12. When the same physical misunderstanding appears twice in the problem log, it belongs in the error log as a rule.

10.6 What it looks like after twelve weeks

Thirty to fifty entries per module. Two or three pages. Perhaps twenty minutes of total writing spread across the term.

And when you sit down to revise, you have a document that says: here is every method this course used, here is what a problem looks like when it wants that method, and here is the idea that makes it work.

That's the thing students spend the fortnight before the exam trying to reconstruct from scratch, badly, from a folder of solutions.

TRY IT THIS WEEK

Start the log today with your last problem set. Go back through it and write one line per problem: problem, method, key idea. It'll take fifteen minutes for a set you've already done, and about thirty seconds per problem from now on.

11

Method Cards

Flashcards, rebuilt for methods.

The standard flashcard has a term on the front and a definition on the back, which is exactly the wrong shape for a subject where the difficulty is selection rather than recall.

11.1 The card, inverted

A method card has the *tell* on the front, not the name.




Figure 11.1: A method card. The front carries the situation, because in an exam nobody tells you the method's name.

Five fields on the back: method, the choice it requires, the steps, the boundary, and the reason. That's the four-field structure from Chapter 1 plus the elaboration from Chapter 9.

Put the situation on the front and the method on the back. A card that asks "what is integration by parts?" tests a skill the exam never assesses.

11.2 Making them, in ten minutes after a lecture

One to three cards per lecture, no more.

0–3 min	Which methods did this lecture introduce? Usually one or two.
3–7 min	For each: write the tell, from memory. What does a problem look like when this is the answer?
7–9 min	The steps, briefly. Not a full derivation.
9–10 min	The boundary: when does it fail? Check the notes if you must, but try first.

Writing the tell from memory is the retrieval; copying it from the notes is transcription. If you can't produce it, that's your gap, and it's better found now than in the exam.

11.3 Reviewing them: produce, don't read

Reading through a stack of cards is not retrieval. The protocol:

1. Read the front.
2. **Say or write the method and the first step**, out loud or on paper.
3. Turn the card. Check.
4. Sort into "got it" and "didn't".
5. Redo the "didn't" pile at the end, and again tomorrow.

Five minutes for a stack of twenty. And the sorting is what makes it spaced automatically: cards you fail come back sooner.

11.4 Paper or software?

Both work. The choice matters less than the discipline of producing rather than reading.

Paper	Faster to make, especially with symbols and diagrams. Easy to sort physically. Better for anything you'd have to typeset.
Software	Handles the scheduling for you, syncs across devices, and nags. Worth it if you'll use the scheduling; a liability if the setup becomes the activity.

If you use spaced-repetition software, resist making hundreds of tiny cards. Fifty good method cards per module beat four hundred fragments, and the fragments are usually facts, which is the wrong unit.

11.5 What not to put on a card

Definitions alone. "What is a derivative?" isn't going to be asked. What a problem looks like when you need one might be.

Whole derivations. Too long for a card. Put them in the reconstruction rotation instead.

Anything you've never used. Cards for methods you haven't practised are cards you can't answer, and they poison the stack.

Numbers you can look up. Constants and standard values belong on a formula sheet, not in your retrieval rotation.

WHAT THE EVIDENCE SAYS

Flashcards are among the best-studied study tools, and the evidence strongly favours *producing* the answer over recognising it, and favours spaced review over massed. The specific inversion here, situation on the front, is my adaptation for method-based subjects rather than a research finding.

11.6 A worked stack

What one lecture's worth looks like, from a first-year mechanics course:

Front: A body on a slope, in equilibrium or accelerating along it.

Back: Resolve along and perpendicular to the slope, not horizontally and vertically. Because it makes one equation trivial and the normal force falls out directly.

Front: Two masses connected over a pulley.

Back: Newton's second law for each mass separately, same magnitude of acceleration, opposite senses. Fails if the string stretches or the pulley has significant mass.

Front: A collision, and you're asked about velocities afterwards.

Back: Conserve momentum first, always. Then check whether the question says elastic before assuming energy is conserved.

Three cards, one lecture, ten minutes. And each one is a tell, which is the thing your notes don't contain.

TRY IT THIS WEEK

Make three method cards from your most recent lecture, with the tell on the front. Write the tells from memory before checking your notes. Then try them tomorrow: read the front, produce the method out loud, and see how many you can actually do.

12

The Error Log

Your mistakes repeat, and you don't know it.

Every student believes their errors are scattered and various. They aren't. Log for four weeks and a pattern appears that's usually three or four items long, with one of them accounting for a third of everything.

12.1 The four columns

Date	What went wrong	Why	The rule
12 Oct	sign wrong in parts	swapped u and dv	write both, boxed, before starting
19 Oct	limits not changed	substituted in my head	new line for new limits, always
26 Oct	used sin in degrees	calculator mode	check mode at the start of every session
2 Nov	sign wrong in parts	(repeat)	

The fourth column is the only one that changes anything. "What went wrong" is a diary. The rule is a change in behaviour, phrased so that future you can follow it while tired.

An error log without a rule column is a list of regrets. And the rule has to name the moment as well as the action: "check units" doesn't work, "write the unit when I box the answer" does.

12.2 Rules that work and rules that don't

Not a rule	A rule
Be more careful with signs	Check the sign of every answer before moving on
Read the question properly	Underline the verb before starting each part
Show more working	Write the substitution on its own line, always
Practise more	Three problems from this topic every week until it retires
Don't rush	Stop at the timer, finished or not

12.3 Running it

One page, physical, visible. A file on your laptop gets opened twice.

Log immediately, while you remember why you did it. Ninety seconds.

Log technique errors as well as content errors. “Spent forty minutes on one problem” belongs here as much as “forgot the chain rule”.

Read it before every problem set and every practice test. Thirty seconds, top to bottom. That’s the mechanism: you’re loading your known failure modes into working memory right before you’re exposed to them.

Retire entries. When a rule has held for three sets, cross it off and date it. Watching the list shrink is the most motivating thing in a term, and it keeps the log short enough to read.

12.4 Turning the log into practice

Every fortnight, take the top two entries and build a drill around them.

Top entry	The drill
Sign errors in parts	Fifteen parts problems; write u and dv boxed; check <i>only</i> the signs. Don’t finish them.
Timing overruns	Six problems with a hard timer; success is stopping, not finishing.
Wrong method chosen	Twenty identification-only problems (Chapter 6).
Algebra slips	Ten problems where you rewrite every line rather than working across.

Notice that three of those four involve doing no complete problems at all. If the failure is in a specific step, practising whole problems trains mostly the parts that already work.

Drill the failing skill in isolation, and score the skill rather than the answer. A timing drill you “pass” by finishing everything hasn’t tested timing.

12.5 The distribution you’ll find

After four weeks, typically:

Look at where “genuinely didn’t know it” sits. Fifth, with two occurrences out of thirty.

That’s the finding that changes how people revise. Most students plan their work against that row and lose their marks in the rows above it.

WHAT THE EVIDENCE SAYS

Error analysis of this kind is standard practice in coaching and in competitive exam preparation, and it’s under-studied in the academic literature compared with retrieval and spacing. I use it because it works with students, and because the alternative, “be more careful”, has never worked for anybody.




Figure 12.1: One student's error log after four weeks, sorted. Thirty entries, and the top two account for nineteen of them.

12.6 Taking it to somebody

The log is the best thing you can put in front of a tutor.

"I'm struggling with this module" gets you reassurance. "Here's four weeks of log; three of these sixteen entries are sign errors in parts and I can't see what I keep doing, can you watch me do one?" gets you somebody watching your hand move and spotting in ninety seconds that you're writing dv before deciding what u is.

Same appointment, completely different value.

TRY IT THIS WEEK

Start the log today with your last three problem sets. Go through them, find every lost mark, and write the four columns. You'll have six to ten entries in twenty minutes, and two of them will already be repeats. Those two are what you're fixing this month.

13

Summary Sheets and Method Maps

The last artefact, and the one you'll revise from.

13.1 The method map

One page per topic. Two columns: the tell, and the method.

AFTER

First-year integration

When you see	Use
A product, one factor simplifying	by parts, u = the simplifier
$\sqrt{a^2 - x^2}$	$x = a \sin \theta$
$\sqrt{a^2 + x^2}$	$x = a \tan \theta$
A rational function, denominator factorises	partial fractions
Inner function and its derivative both present	substitution
Odd power of sine or cosine	peel one off, use the identity
Even powers only	double-angle formulas
$\int_{-a}^a$, odd integrand	zero, by symmetry
A recurrence between I_n and I_{n-1}	reduction formula; find the base case

Nine rows, one page, and it's what you read for ninety seconds before a mixed practice session or a paper.

Notice the left column is written in the exam's language, not the textbook's. The textbook says "integration by parts" and gives you twenty by-parts problems. The exam gives you a product and says nothing.

Build the map from the tell inwards. A map organised by method name is your contents page, which you already have. A map organised by situation is the thing you don't have and need.

13.2 The summary sheet

One page per topic, written *from memory* at the end of a topic, containing:

- every definition, in your own words, with one example and one non-example;
- every result you'd quote, stated precisely;

- the method map rows for that topic;
- the two or three derivations you're expected to reproduce, as skeletons rather than in full;
- your monsters: the counterexamples and gotchas from the problem log.

And the rule that makes it work: **write it from memory in one pen, then open the notes and fill the gaps in a second colour.**

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

WHAT THE EVIDENCE SAYS

Generating a summary from memory is retrieval; copying one is transcription, and the difference in outcome is large. The two-pen version also gives you an accurate self-assessment, which Chapter 3 says you otherwise won't have.

13.3 What goes on a formula sheet, and what doesn't


If your exam provides a formula sheet, get it in week one and work with it all term. Two consequences.

Stop memorising what's printed. That's real capacity released.

Learn its geography. A sheet you've used for twelve weeks is a tool; one you meet in the exam is three pages of dense unfamiliar notation at the worst moment.

If no sheet is provided, build your own and revise from it. What goes on it: the results you'd quote rather than derive. What doesn't: anything you're expected to derive, and anything whose *tell* matters more than its statement, which belongs on the method map instead.

13.4 How the four artefacts fit together



At revision, you read the map, test with the cards, drill from the error log, and use the problem log to remind you what unlocked each technique. The textbook stays shut.

Figure 13.1: The four artefacts. Two are built continuously, two at the end of a topic, and together they replace the textbook at revision time.

Total cost across a twelve-week module: roughly three hours, spread in five and ten minute pieces.

Total value at revision: it replaces the textbook, and it's written in your own words, which means reading it is retrieval rather than recognition.

13.5 When to build the map

At the end of a topic, not at the end of term.

The end-of-term version is a reconstruction exercise on material you've half forgotten, which is slow and produces a worse map. The end-of-topic version takes twenty-five minutes while everything is still available and doubles as the consolidation session for that topic.

Put it in the calendar as the last item of each topic, the way you'd schedule a problem set.

TRY IT THIS WEEK

Take the topic you finished most recently and build its method map: one page, two columns, tells on the left. Write as much as you can from memory, then fill the gaps from your notes in a different colour. The proportion in the second colour tells you where you actually stand on that topic.

Part IV

The Week

14

Before the Lecture

Twenty minutes, and it's the highest-return habit in this book.

14.1 What it is

Not reading the chapter. A survey.

0–5 min	Read the lecture title and the section headings in the notes or textbook. What is this about?
5–12 min	Read the definitions and the statements of the main results. Skip every proof and every derivation.
12–17 min	Write down the names of the new objects and any symbols you don't recognise.
17–20 min	One sentence: what do I think this lecture is for?

You're not trying to understand it. You're trying to arrive with the vocabulary already loaded.

14.2 Why it pays so much

Because of what a lecture costs you when you go in cold.

In an unprepared lecture your attention is spent decoding: what does that symbol mean, what did she just call that, is this the same thing as last week. All of that happens in real time, at the speaker's pace, while the argument moves on without you.

In a prepared lecture the vocabulary is familiar and your attention is free for the argument, which is the only thing worth being in the room for.

Twenty minutes of survey before a lecture is worth more than an hour of reading after it. Going in cold and planning to catch up later is the default strategy and the expensive one.

14.3 The list of unknown symbols

The one artefact from this session, and it's more useful than it sounds.

Write down every symbol and term you couldn't say out loud. Ten items, maybe. Then in the lecture, when the lecturer introduces each one, you're listening for something specific rather than being surprised by it.

And anything still on the list at the end of the lecture is your question for the lecturer, ready-made.

14.4 When you have no time

The five-minute version, which is worth most of the twenty.

Read the headings. Read the definitions only, ignoring everything else. Write down the two or three new words.

That's it. Five minutes, on a bus, and it still converts a large part of the decoding cost.

WHAT THE EVIDENCE SAYS

Pre-exposure effects, sometimes called pretraining or advance organisers, are reasonably well supported: learners who have seen the key terms before an explanation get more from it. The size depends heavily on how dense the material is, and STEM lectures are dense.

14.5 What not to do beforehand

Don't try to learn it. You'll spend two hours on something the lecturer is about to explain, and you'll arrive tired.

Don't read the proofs. They're the slowest part and the part the lecture is most useful for.

Don't do the problem set first. It comes after, when you have the methods.

The one exception: if the lecture explicitly builds on a topic you know you never understood, spend the twenty minutes on *that* instead. Arriving without the prerequisite is a worse position than arriving without the preview.

TRY IT THIS WEEK

Do the twenty-minute survey before your next lecture, and write the list of unfamiliar symbols. In the lecture, notice how different it feels to meet a word you've already seen once. If twenty minutes is too much, do the five-minute version, which still works.

15

In the Lecture

The commonest lecture mistake is trying to write 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*.

15.1 Split the page




Figure 15.1: Split the page. The left column is already written down somewhere; the right column exists nowhere else and is what you’ll reread.

Sentences worth catching word for word:

- “The reason we need this condition is...”
- “This is the step people always get wrong.”
- “You should think of this as...”
- “We’ll use this again in week nine.”
- Anything beginning “intuitively”.
- Any aside about how to spot when this method applies. That’s a tell, free.

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

Copy the board lightly and the speech fully. Most students do the exact opposite and end up with a worse copy of a document they already own.

15.2 If the slides are published

Then the left column mostly doesn’t need writing at all.

Print the slides beforehand, or have them open, and use the whole page for the right column: the reasoning, the asides, the tells, your own questions.

Students with published slides often stop taking notes entirely, which wastes the lecture. The slides are the board; the value is still in the speech.

15.3 Marking your own confusion

Two symbols, used consistently.

A question mark beside anything you didn't follow. Don't stop to work it out during the lecture; you'll lose the next three minutes and gain nothing.

A star beside anything the lecturer emphasised. Emphasis is the only signal you get about what matters, and it's easy to lose by the evening.

Then in the ten minutes after (Chapter 16), the question marks become your list of things to resolve, and the stars become your method cards.

15.4 Should you record it?

Sometimes, with a caveat.

Recordings are genuinely useful for a fast lecturer, a language barrier, a missed session, or a topic you know you struggled with. They're a trap when they become a reason not to attend or not to engage, because rewatching is the most fluent, least productive study activity there is.

If you rewatch, do it actively: pause before each step and predict what comes next. Otherwise you'll spend fifty minutes recognising material and call it revision.

15.5 Handwriting or typing?

For STEM, mostly handwriting, and for a practical reason rather than an ideological one: equations, diagrams and free-body sketches are far faster by hand.

If you type, use something that handles mathematics quickly, and expect to sketch on paper anyway.

There's some evidence that handwriting encourages summarising rather than transcribing, which is the behaviour you want. The evidence is more contested than the popular write-ups suggest, so treat this as a practical recommendation for symbol-heavy subjects rather than a law.

15.6 Asking a question in the lecture

Worth doing, and worth doing well.

A good in-lecture question is short and specific: "why does that require f to be continuous?" It gets a thirty-second answer that helps everyone.

A question that needs five minutes belongs at the end or in office hours. And if you're lost enough that you can't formulate a question, that's a question mark in the margin and a visit to Chapter 19, not an interruption.

TRY IT THIS WEEK

In your next lecture, split the page and write the right-hand column in full sentences. Afterwards, read only that column. Most students find it's the part they'd normally have lost entirely, and it's the part that makes the notes make sense in November.

16

After the Lecture: The Forty-Eight Hour Rule

Ten minutes, within two days, and it's the habit that does the most per minute in this book.

16.1 The ten minutes

0–3 min	Close the notes. Write, from memory, what the lecture was about and which method it introduced.
3–5 min	Open the notes. Fill what you missed, in a different colour.
5–8 min	Make one to three method cards, tell on the front (Chapter 11).
8–10 min	Resolve or record the question marks from the margin. Anything unresolved goes on the list for office hours.

Ten minutes. Not thirty, not an hour. The size is deliberate: a habit that takes ten minutes survives a bad week, and one that takes an hour doesn't.

The forty-eight hour window is where lecture notes are still decodable. After that you stop remembering what your own abbreviations meant, and the notes lose most of their value.

16.2 Why the window matters

Two mechanisms, and both are practical rather than mysterious.

Your notes decay faster than your memory. Abbreviations, half-finished diagrams, a line you meant to complete. Within two days you can repair them; after two weeks you're reading somebody else's handwriting.

The first spaced interval lands here. From Chapter 5, the one-day revisit is the cheapest and one of the most valuable, and this session is it.

16.3 The blank-page part is the important part

Three minutes, notes closed, writing what you remember.

It will feel bad. You'll produce less than you expect, and you'll be tempted to open the notes at minute one.

Don't. The gap between what you produce and what's in the notes is the most accurate measurement of your understanding you'll get all week, and it's available for free ninety minutes after the lecture.

WHAT THE EVIDENCE SAYS

This is the testing effect applied at short delay, and it's about as well supported as anything in this book. The specific ten-minute format is my own; the underlying claim, that a brief retrieval attempt shortly after learning substantially improves retention, is solid.

16.4 Fitting it into an actual day

The habit fails on scheduling, not on effort. Three attachments that work:

The gap between lectures. If you have back-to-back sessions, the ten minutes before the next one starts is free time that already exists.

The walk or the bus. The three-minute recall part can be done in your head, out loud if you're alone. The writing can wait until evening.

Same time every day. Six-thirty, all four courses' lectures from that day, forty minutes total. Less good than doing each one immediately, much better than not doing it.

16.5 What this replaces

Not extra work. It replaces something.

Most students, at some point in the evening, reread that day's notes. That takes fifteen or twenty minutes and produces recognition. Replace it with this: three minutes of recall, two of repair, five of card-making.

Same slot, less time, much more retained.

16.6 When you missed the lecture

Different protocol, and worth having.

1. Get the notes and do the twenty-minute survey from Chapter 14 *on them*, as though preparing.
2. Work through one worked example properly, using the four steps from Chapter 8.
3. Make the method cards.
4. Take the question marks to office hours or a coursemate.

That's about forty minutes and it recovers most of a missed lecture. What doesn't work is reading somebody else's notes end to end, which produces recognition of their understanding rather than any of your own.

TRY IT THIS WEEK

After your next lecture, do the ten minutes before you do anything else, starting with three minutes of blank-page recall. Note how much less you can produce than you expected. That gap is exactly what would have been missing in the exam, and you've just found it in three minutes.

17

Problem Sets Without Suffering

Problem sets are the main event. They're where most of your learning actually happens, and most students run them in a way that wastes about half the value.

17.1 The protocol




Figure 17.1: How to run a problem set. Steps 1 and 5 are the ones students skip, and they're where most of the value is.

Step 2 is where most of the value lives, and it's the step students compress by having the notes open from the start.

First pass with everything closed. The point of a problem set isn't to produce correct answers, it's to find out which methods you can select and run without help. Notes open, you learn nothing about that.

17.2 The twenty-minute rule

Twenty minutes on one problem, then move on and come back.

Past that, returns go negative: you're not thinking any more, you're anxious, and anxiety rereads the same line indefinitely.

When you leave a problem, write one line about where you stopped and what you'd try next. Coming back cold costs three minutes of reconstruction otherwise, which is often more time than you have.

And a surprising number of problems solve themselves in the gap while you're working on something else. That's not luck; it's what happens when you stop applying pressure.

17.3 Getting unstuck, in order

-
- | | |
|---|---|
| 1 | Re-read the question. Underline the verb and every constraint. |
| 2 | What kind of object is being asked for? That usually shortlists the method. |
| 3 | What in the question haven't I used? Unused information is the biggest hint there is. |
| 4 | What's the simplest case? Set $n = 1$, or make a constant zero. |
| 5 | Which method card has this tell? Check your own map. |
| 6 | Have I done anything like this? Check the problem log. |
| 7 | Now open the notes. |
-

Six moves before the notes, and each takes under a minute. Students go straight to step 7, which resolves the immediate problem and teaches nothing.

Step 3 deserves emphasis. If a question tells you a function is continuous, or a matrix is symmetric, or n is prime, that fact is there because the answer needs it. Unused hypotheses are the most reliable hint in any problem.

17.4 How much time a set should take

Longer than you think, and the length isn't the problem.

A weekly set for a demanding module is three to five hours, done in two or three sittings rather than one. That's normal, and students who assume something is wrong with them for taking four hours are usually taking a normal amount of time.

What's not normal is four hours in one sitting the night before it's due, because the last two of those hours are worth very little.

WHAT THE EVIDENCE SAYS

The claim that distributed practice beats massed practice applies to problem sets as much as to revision. Two two-hour sittings beat one four-hour sitting, and the difference shows up on later tests rather than on the set itself.

17.5 Collaboration, done right

Working with others is legitimate and valuable, and it has a shape that keeps it that way.

Attempt alone first. Everything, properly, before the group meets. Otherwise the fastest person in the room does the thinking and everyone else watches, which feels productive and teaches nobody.

Compare approaches, not answers. "I used parts, what did you use?" is worth far more than checking whether you got the same number.

Explain, don't dictate. If you've solved it, ask them what they've tried and give them the next move, not the solution. And if you're stuck, say precisely where.

Write your own submission. Whatever the module's rules are, read them, and then follow them. Where they're unwritten, shared discussion and individual writing is the safe default.

17.6 After it's marked

The step almost nobody takes, and it's fifteen minutes.

Take the marked set, go through every lost mark, and put each in the error log with a rule. Then redo the questions you got wrong, from scratch, three days later.

Understanding an error takes five minutes and lasts a week. Re-solving the problem takes fifteen and lasts.

TRY IT THIS WEEK

On your next problem set, do step 1: three problems from three weeks ago, cold, before you start. Then do the whole first pass with the notes shut. Note how different the experience is, and how much more precisely you can say what you don't know afterwards.

18

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, and each does something you can't do alone.

18.1 Format 1: the teaching round

Each person is assigned one method or one problem before the meeting and teaches it at the board for ten minutes.

That's it, and it's by a distance the most effective format.

Teaching is retrieval under pressure with an audience who will interrupt. Every gap in your understanding surfaces within about ninety seconds, and the questions you get are exactly the ones you couldn't ask yourself.

If nobody in the group is standing up and explaining, 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.

18.2 Format 2: the gap swap

Everyone brings their error log and their list of unresolved question marks.

Go round the table, five minutes per question. Most get answered by somebody in the room. 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.

Twenty minutes, and it clears more confusion per minute than any amount of working through problems together.

18.3 Format 3: the mixed set race

Somebody builds a mixed set (Chapter 6), twelve problems from four topics, shuffled. Everyone works it independently, timed, then compares.

Not the answers. The *tells*: what made you think it was that method? Where people diverge is where the discrimination is weak, and that conversation is worth an hour of solo practice.

18.4 What makes groups fail

Failure	Fix
Nobody prepared	Assign the teaching round in advance
One person does the work	Rotate who's at the board, every time
It becomes social	Fixed start, fixed end, agenda in one line
Too big	Three or four people. Six is a seminar
Copying, not learning	Attempt alone first, always
Different levels	That's fine; the stronger student teaching is the one learning most

That last row surprises people. Explaining something to somebody who doesn't get it is more valuable to the explainer than to the listener, which means the "I'm carrying this group" complaint is usually backwards.

WHAT THE EVIDENCE SAYS

The learning-by-teaching effect is well supported, including the finding that expecting to teach improves learning even before the teaching happens. It's one of the more reliable ways to convert passive understanding into retrievable knowledge.

18.5 The right size and cadence

Three or four people, ninety minutes, once a week, same time.

Longer than ninety minutes and attention goes. More than four people and somebody is always passive. Less than weekly and it never builds the habit of preparing for it.

And it works best after the individual attempt, not instead of it. A group that meets before anybody has tried the set spends its time watching one person think.

18.6 Working alone, if that's your preference

Groups aren't compulsory, and plenty of good students prefer to work alone. Two substitutes for what the group provides:

For the teaching effect, explain out loud to nobody. Record a two-minute explanation of a method on your phone, then listen to it. The stalls are audible and the embarrassment is private.

For the gap swap, take the question marks to office hours (Chapter 19) or to a subject forum, where a precise question usually gets a precise answer.

What you can't easily substitute is the comparison of approaches, so if you work alone, it's worth attending tutorials properly for exactly that.

TRY IT THIS WEEK

Propose a teaching round to three coursemates for next week: each person takes one method from the current topic and gets ten minutes at the board. It's the one format that reliably produces the "I thought I knew this" moment, and that moment is worth the whole session.

19

Office Hours

The most underused resource in any degree, and the one with the best ratio of value to effort.

The barrier is almost never that lecturers are unwilling. It's that students don't know what to bring, so they don't go.

19.1 What to bring

A question with three parts.




Figure 19.1: The anatomy of a question that gets a useful answer. The third part is the one students omit, and it's the one that makes the answer specific.

Compare that with “I don’t understand Tuesday’s lecture”. 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.

Bring evidence and a specific gap. A student with a marked-up page and a precise question gets a two-minute answer that fixes the problem permanently; a student with a general complaint gets the lecture repeated.

There’s a bonus effect, and it isn’t small: constructing the three-part question forces you to locate your own confusion precisely, and perhaps a third of the time the question answers itself before you ask it. That isn’t wasted effort, that’s the technique working.

19.2 When to go

Early in the term, not the week before the exam.

Two reasons. The queue is shorter, so you get longer. And a misunderstanding fixed in week three stops corrupting everything built on it, whereas the same fix in week eleven arrives after the damage.

Go with one or two questions, not a list of ten. And go having tried: “I attempted it three ways and here’s where each stopped” gets you a different conversation.

19.3 The questions worth asking that aren't about a specific problem

Students think office hours are only for stuck problems. Four questions that are worth more:

- “Which results in this course does everything else depend on?” Tells you where to spend your depth.
- “What do students usually find hardest here?” You get the list of traps before you fall into them.
- “Is there a better book for this topic?” Lecturers usually have a strong opinion and rarely get asked.
- “What does a really good answer to this kind of question look like?” Directly about assessment, entirely reasonable, and almost nobody asks it.

None of those is asking for the exam. All four change how you spend the next month.

WHAT THE EVIDENCE SAYS

No research here, just teaching experience: the students who come to office hours early with specific questions are, over and over, the ones whose marks move most across a year. Whether that’s the office hours or the habits that lead to them, I can’t separate.

19.4 Tutorials and demonstrators

For many courses the tutorial is where the real help is, and it’s often treated as a place to receive answers.

Two changes that transform it:

Arrive having attempted everything. A tutorial where you’ve tried the set is a conversation; one where you haven’t is a lecture with fewer people.

Ask about the choice, not the execution. “Why parts rather than substitution here?” is a better use of a demonstrator than “what’s the answer to 4(b)?”, and it’s the question that builds tells.

19.5 If you’re anxious about going

Common, and worth addressing directly.

Nobody thinks less of a student for asking a specific question. The opposite: a precise question is evidence you’ve engaged, and it’s the most interesting part of most academics’ week.

If it still feels difficult: go with a coursemate, go with a written question so you don’t have to compose it on the spot, or send it by email first. The point is the answer, not the performance.

TRY IT THIS WEEK

Write one three-part question this week, from a question mark in your lecture notes, and take it to office hours. If you can’t construct part three, the precise gap, that’s the work: go back and find exactly where your understanding stops. Either way you win.

Part V

The Term, and the Years

Planning a Term

Everything so far has been a technique. This chapter is where they go in the week, and what a term looks like when they're all running.

20.1 The weekly template, per course

When	Time	What
Before each lecture	20 min	Survey: headings, definitions, unknown symbols
In each lecture	—	Split page; stars and question marks
After each lecture	10 min	Blank-page recall, repair, 1–3 method cards
Problem set	3–5 h	Cold warm-up, closed-book first pass, log
One spaced session	20 min	Oldest topic you feel worst about, mixed
End of topic	25 min	Method map and summary sheet, from memory
Weekly	2 min	Error log update and read

For a course with two lectures a week, the additions beyond the problem set come to about ninety minutes. That's the whole system.

The additions are small and they replace something. The ten minutes after a lecture replaces rereading that evening; the spaced session replaces part of the pre-exam scramble. This isn't extra hours, it's different hours.

20.2 Four courses, one week

The obvious objection: ninety minutes per course times four is six hours, on top of everything else.

Two answers.

Most of it is attached to things already in your diary. The pre-lecture survey and the post-lecture ten minutes happen around lectures you're already attending. The cold warm-up is the first fifteen minutes of a problem set you were doing anyway.

And you can run a reduced version. If four courses is too much, run the full system on the two hardest and the minimum version on the others: post-lecture ten minutes and the problem log, nothing else.

A reduced system that runs beats a complete one that collapses in week five.

20.3 The term, week by week




Figure 20.1: A twelve-week term. Notice that revision doesn't start in week ten, because by then it's relearning rather than revision.

The shape that matters: by week eleven you're consolidating something you already have, rather than learning eleven weeks of material in a fortnight.

20.4 Weeks 1 to 3: build one habit at a time

Do not try to install the whole system in week one. It fails, reliably, and then the conclusion is that the system doesn't work.

Week 1	The ten minutes after each lecture. Nothing else.
Week 2	Add the problem log.
Week 3	Add the cold warm-up before each problem set.
Week 4	Add the weekly spaced session.
Week 5	Add method maps at the end of each topic.

One new habit a week, and by week five the whole thing is running without requiring decisions.

20.5 Mid-term tests, and what they're for

Treat every mid-term as data rather than as a verdict.

Afterwards, do the audit: every lost mark into a bucket, either a technique failure or genuinely not knowing it. The proportion tells you what to change, and it's usually not what you assumed.

Then put the technique failures into the error log as rules, and the knowledge gaps into the next fortnight's spaced sessions.

20.6 Coursework-heavy and lab-heavy terms

If your term is dominated by labs, projects and reports rather than exams, the system barely changes. The artefacts stay; what shifts is the proportion.

Chapter 22 covers this properly, and its short version is that a project has the same failure modes as an exam module, they just arrive on a longer timescale.

TRY IT THIS WEEK

Write the weekly template for one course, with actual times, into your calendar for the next three weeks. Only the first row of the habit table above: the ten minutes after

each lecture. Add the next habit in week two, and only if the first one held.

21


Slumps, and Getting Back

At some point in a STEM degree you will fall behind, badly enough that catching up looks impossible.

It happens to nearly everyone, and how you handle the first two weeks determines whether it costs you a module or a year.

21.1 The spiral, and its shape

It has a mechanism, and knowing the mechanism helps.



The exit is at the second box, and it is not “catch up on everything first”.

Figure 21.1: The falling-behind spiral. Each step is a reasonable response to the last, which is why it’s so easy to enter.

The critical error is at box three, and it’s a reasonable-sounding decision: *I’ll stop attending until I’ve caught up*. That guarantees the gap grows, because lectures keep happening.

Keep attending, always, even when you understand nothing. The lecture is the cheapest exposure available and stopping is what turns a two-week gap into a lost module.

21.2 The recovery protocol

Not “go back to week one and work forwards”. You’ll reach week four of nine and run out of time, having relearned the easy part.

-
- 1 **Keep attending.** Non-negotiable, from today.
 - 2 **Work backwards from the current lecture.** What does today depend on? Learn that. Then what that depends on. Stop as soon as you hit something you know.
 - 3 **Do the most recent problem set,** not the oldest. It's the one that's marked and the one the next lecture builds on.
 - 4 **Tell somebody.** Lecturer, tutor, or personal tutor. Early, factually, without dramatising.
 - 5 **Accept that some of it stays missed.** Pick the topics that appear on every past paper and let the rest go.
-

Step 2 is the important one. Backwards from where you are, depth-first, stopping at the first thing you know, is much faster than forwards from where you stopped.

21.3 Telling somebody, and why it isn't optional

Two reasons, one practical and one administrative.

Practically, a lecturer can tell you which three topics matter and which you can skip, which turns an impossible recovery into a plan.

Administratively, most institutions have mechanisms, extensions, mitigating circumstances, resit arrangements, and nearly all of them have deadlines. Students discover this after the deadline, routinely, because they waited until they were coping again before mentioning it.

WHAT THE EVIDENCE SAYS

No research here, and I'd rather say so: this is what I've seen work with students. The consistent pattern is that the ones who tell somebody in week four recover the module and the ones who tell somebody in week ten usually don't, and the difference is administrative as much as academic.

21.4 When it's not a study problem

Worth naming plainly.

If you can't concentrate for twenty minutes, if you've stopped sleeping, if the work has become frightening rather than difficult, or if this has been going on for weeks, that isn't a study-technique problem and no chapter in this book will fix it.

Every institution has support services, they're free, and they're used by far more students than you'd guess from the outside. Using them early is the same decision as telling a lecturer early: it converts a problem with a deadline into a problem with options.

21.5 The everyday version: motivation, not crisis

Most weeks the issue is smaller. You know what to do and you don't want to do it.

Three things that work better than resolve:

Reduce the starting cost. "Do three problems" is startable. "Revise thermodynamics" isn't. Every session in this book is defined in minutes and quantities for that reason.

Attach to something fixed. The ten minutes after a lecture happens because the lecture happened. Habits attached to existing events survive; free-floating intentions don't.

Make the record visible. The problem log filling up and the error log shrinking are both visible evidence of progress, which is otherwise completely absent in a twelve-week module with one exam.

That third one is why the artefacts matter beyond their content. In a subject where you can work for six weeks with no feedback, a page that demonstrably grew is worth something.

21.6 Comparison, and why it misleads

One more, because it drives a lot of unnecessary misery.

The student who seems to find it easy is usually either doing more work than they admit, or has seen the material before, or is struggling somewhere you can't see. Everyone in a STEM degree is confused most of the time; some are just quieter about it.

The only useful comparison is with your own error log four weeks ago.

TRY IT THIS WEEK

If you're behind right now: keep the next lecture, and spend one hour working backwards from it rather than forwards from where you stopped. Then send one email, today, to whoever runs the module, saying factually where you are. That's the whole recovery, started.

Labs, Projects, and Coursework

Not every module is a lecture course with an exam at the end. Labs, extended projects and coursework-heavy modules have their own failure modes, and the system adapts rather than changes.

22.1 Labs

The lab session itself is not the learning. It's data collection. The learning happens in the write-up, which is why Chapter 22's companion volume exists.

Three habits that make labs pay:

Prepare the lab like a lecture. Twenty minutes on the script beforehand: what's being measured, what theory is being tested, what the apparatus does. Students who arrive cold spend the first hour working out what they're supposed to do, which is time taken from the measuring.

Write the method during the session. Not afterwards. Model numbers, resolutions, settings and dimensions are available in front of you and gone by Thursday.

Log the lab like a problem. One line: what technique did this teach, what would you do differently. Lab skills are methods with tells, exactly like anything else.

22.2 Extended projects

A project fails in slow motion, which makes it harder to notice than a module you're behind on.

Failure mode	The fix
No visible progress for weeks	A weekly artefact: a plot, a page, a working script. Something that exists
Writing left to the end	Write the method chapter while you do the work
Supervisor meetings with nothing to show	Bring the artefact, even if it's negative
Scope grows quietly	Write the aims down and reread them monthly
Data and code chaos	One folder, dated or version-controlled, from day one

A project needs a weekly artefact, not a weekly effort. Something that exists at the end of each week and can be shown to somebody. Effort with nothing to show is how eight weeks disappear.

22.3 Supervisor meetings

The same principle as office hours, at a different cadence.

Bring three things: what you did since last time, what you're doing next, and what's blocked. The third is the one students soften, and softening it is expensive: a supervisor who hears about a problem in week four can redirect the work, and one who hears in week ten can't.

Send a two-line summary after each meeting confirming what was agreed. It takes ninety seconds and it prevents the common failure where you and your supervisor remember different decisions.

22.4 Coursework-heavy modules

Where the assessment is four pieces of coursework rather than one exam, two things change.

The deadlines do the spacing for you, which is genuinely helpful. Each piece forces retrieval of that block of material.

But nothing forces retention afterwards. A student can complete four pieces of coursework and retain almost nothing, because each was completed with the notes open and never revisited.

So the spaced sessions matter more here, not less, particularly if the module is a prerequisite for something next year.

22.5 Group projects

Everything in Chapter 18 about study groups applies, plus three things specific to assessed group work:

Agree the split by strand, not by section. Each person owning a coherent piece of work produces something mergeable; four people writing two pages each produces four voices.

Name a document owner. One person responsible for consistency of notation, structure and numbering. Without it the report will contradict itself somewhere.

Set an internal deadline three days early. That's the merge and read-through time, and a group planning to finish on the deadline will submit four reports stapled together.

22.6 Placements and industrial work

If your degree includes a placement, one habit transfers and is worth keeping: the log.

A weekly line about what you did, what technique it used, and what you'd do differently is the raw material for your placement report, your CV, and your final-year project ideas. Almost nobody keeps one, and everybody wishes they had when the report is due.

TRY IT THIS WEEK

If you have a project running: define this week's artefact now, in one sentence, and make it something that will exist by Friday. If you have a lab this week: spend twenty minutes on the script beforehand and write the method during the session rather than after it.

23

Handing Off to the Exam

If the system has been running, the last fortnight is a handover rather than a rescue.

This chapter is that handover, and it's short, because most of the work already happened.

23.1 What you should be holding in week ten

Artefact	State
Problem log	30–50 entries, two or three pages
Method cards	30–60 per module, mostly in the “got it” pile
Method maps	One page per topic, written from memory
Summary sheets	One per topic, with the gaps marked in a second colour
Error log	Short, because entries have been retiring

If you're holding those, the fortnight is consolidation. If you're not, Chapter 21's recovery protocol applies, and the honest version is that you're doing triage rather than revision.

23.2 The two weeks

Days 14–10	Close the real gaps. Rank topics by past-paper frequency and by your own error log, then work the top three. Write down the two topics you're choosing not to learn.
Days 9–4	Timed practice papers, reviewed properly. This is where the exam-technique work lives.
Days 3–1	Maps, cards and error log only. No new material. Sleep.

The instruction to name the topics you're skipping sounds reckless and is arithmetic. You have finite hours; spending them uniformly across topics of wildly different exam frequency is the least efficient allocation available.

Decide what you're not learning, in writing, ten days out. Naming it stops it consuming attention for a fortnight, and frees the time for topics that actually appear.

23.3 Where this book stops

Deliberately, at the door of the exam hall.

Everything from the moment the paper is turned over, triage, timing, partial credit, salvage writing, the final sweep, is a separate skill set and it's a whole book of its own. If your marks are lower than your understanding, and particularly if you routinely fail to finish papers, that's where the missing marks are rather than here.

What this book is responsible for is arriving at the door with the material actually available: methods you can select, tells you recognise, and a set of pages in your own handwriting that you can read in an hour.

23.4 The last three days, in detail

- **Read every method map**, one topic at a time. Ninety seconds each.
- **Run the cards**, producing rather than reading. The pile you fail is your list.
- **Read the error log end to end**, once a day, and do one problem per entry to check the rule holds.
- **Write the formula sheet from memory**, then check it. Twice.
- **Sleep properly, all three nights**. Not a wellness note: sleep deprivation costs more marks than any topic you could learn in the time.

No new topics. Content learned 48 hours before an exam is rarely retrievable under pressure, and the hours come out of the sleep and the consolidation, both of which are worth more.

23.5 Afterwards

Two things, and only one is useful.

The walk home, where every question you got wrong arrives in sequence with the answers, is compulsory apparently and worth nothing.

The twenty-minute post-mortem the next day is worth a great deal: what happened, why, and what you'll do differently. Put the answers in the error log before you open a book for the next paper.

23.6 What you're left with

If you run this system for a year, three things are different.

You'll know your subject rather than your modules, because elaboration and mixed practice keep connecting them.

You'll have a written record of how you actually learn, which is worth more in year three than any single module's content.

And the thing that changes most: the gap between what you understand and what you can produce closes, which is the gap this whole book has been about.

That gap is where most students lose marks, and almost none of it is about intelligence. It's about what the hours contained.

TRY IT THIS WEEK

Ten days before your next exam, write three lists: topics to fix, topics to rehearse, and topics you're choosing not to learn. Then book the timed papers into the calendar with dates. Fifteen minutes, and the fortnight stops being a fog.

Part VI

Toolkit

A

Three Worked Weeks

The system, running, at three different loads. Every entry is an actual block of time with an actual activity in it.

A.1 Week 1: four modules, first year, a normal week

Two lectures per module, one problem set per module per fortnight, one lab.

Day	Slot	Activity
Mon	08:40–09:00	Survey before Maths lecture
	11:00–11:10	After Maths: recall, repair, 2 cards
	13:40–14:00	Survey before Physics lecture
	16:00–16:10	After Physics: recall, repair, 1 card
Tue	09:50–10:00	After Mechanics: recall, 2 cards
	14:00–17:00	Maths problem set: 15 min cold warm-up (topic from wk 1), then closed-book first pass
	17:00–17:05	Log 6 problems
Wed	10:00–13:00	Lab. Method written during the session
	20:00–20:20	Spaced session: Physics, week-2 topic, 3 problems mixed
Thu	08:40–09:00	Survey before Maths
	11:00–11:10	After Maths
	15:00–16:30	Cards for all four modules, produced not read
Fri	09:50–10:00	After Mechanics
	14:00–16:30	Physics problem set, same protocol
	16:30–16:40	Log, and error log update
Sat	10:00–11:30	Study group: teaching round, Maths
Sun	—	Nothing. Deliberately.

Totals. Problem sets 5.5 h. Everything else in this book: about 2 h 40 m, nearly all of it in ten and twenty minute pieces attached to lectures that were happening anyway.

A.2 Week 2: second year, heavy coursework, one bad day

Same student, a week with two deadlines and a day lost to illness.

Day	Slot	Activity
Mon	11:00–11:10	After lecture: recall, 1 card
	14:00–18:00	Coursework 1
Tue	—	Ill. Nothing done.
Wed	09:00–09:10	After lecture (attended despite being behind)
	10:00–10:30	Catch-up on Tuesday's lecture: survey the notes, 2 cards, question marks listed
	13:00–18:00	Coursework 1, submitted
Thu	11:00–11:10	After lecture
	15:00–15:20	Spaced session, oldest topic
	16:00–17:00	Office hours with Tuesday's question marks
Fri	10:00–13:00	Coursework 2, first pass
	15:00–17:30	Problem set, reduced: attempt everything, log only the failures
Sat	14:00–16:00	Coursework 2, finished
Sun	11:00–11:25	Method map for the topic that just ended

What was cut. Pre-lecture surveys, the study group, and most of the card review.

What survived. The ten minutes after every lecture, one spaced session, and the logs. That's the minimum viable version, and it's what you protect when a week goes wrong.

When a week collapses, protect the ten minutes after each lecture and one spaced session. Everything else can be missed for a week without lasting damage; those two are what stop the gap opening.

A.3 Week 3: final year, project-dominated

One taught module, a large project, and a supervisor meeting.

Day	Slot	Activity
Mon	09:00–12:00	Project: run the parameter sweep
	14:00–14:10	After lecture: recall, 1 card
	14:10–14:40	Project log: what ran, what settings, what came out
Tue	09:00–13:00	Project: analysis, produce this week's artefact (one plot, with its script)
	15:00–16:00	Write the method section for what was done last week
Wed	10:00–10:30	Prepare supervisor meeting: done, next, blocked
	11:00–11:30	Supervisor meeting
	11:30–11:35	Two-line summary of what was agreed, sent
	14:00–17:00	Act on the meeting
Thu	09:00–12:00	Project
	14:00–14:10	After lecture
	15:00–15:20	Spaced session on the taught module
Fri	09:00–12:00	Project
	14:00–16:00	Taught module problem set
	16:00–16:10	Logs

What's different. The project needs a weekly artefact rather than weekly effort, and the method section gets written while the work is fresh rather than in April. The taught-module habits shrink to the minimum and keep running.

A.4 What the three weeks have in common

- The ten minutes after every lecture, in all three.
- At least one spaced session, in all three.
- Logs updated, in all three, even in the bad week.
- Problem sets attempted closed-book first, in all three.
- One deliberately empty day in week 1, and no guilt about it.

Everything else flexed. That's the system: five things that hold, and the rest scales with the week you're having.

B

The Problem Log Template

One line per problem. Thirty seconds each. Photocopy the grid or rule your own.

The format

Source	Method	Key idea
--------	--------	----------

Rules

- **Log the idea, not the answer.** The idea transfers; the answer doesn't.
- **Log failures too**, with the key idea being whatever you were missing.
- **Don't log routine problems** that applied a method you already had. If nothing unlocked, there's nothing to record.
- **Write it immediately** after the set, not at the end of term.

Worked entries

Sheet 4 Q3	shells not discs	radius is distance from the axis, not the origin
Sheet 5 Q1b	repeated root	second solution is t times the first
Sheet 6 Q4	moments about the support	choosing the pivot kills an unknown
Tutorial Q2	comparison test	bound the numerator first, then compare
Sheet 7 Q6	<i>failed</i>	needed the substitution $t = \tan(x/2)$; had never seen it

The applied-subject variant

Add a fourth field for the physical insight:

Sheet 3, Q5. Block on an incline with friction. Resolve along and perpendicular to the slope. **Key idea:** axes along the slope make one equation trivial. **Physically:** the normal force isn't mg once the surface is tilted, which is what keeps catching me.

When the same physical misunderstanding shows up twice, promote it to the error log as a rule.

Using it

Before mixed practice	Read the entries for the topics you're about to mix. Two minutes, loads the tells.
When stuck on something new	Scan for a similar shape. It's a searchable index of methods you personally made work.
Before an exam	Read the whole thing. Two pages, eight minutes, the densest revision material you own.

C

Method Card Templates

Situation on the front. Method on the back. Five fields.

The blank

FRONT	<i>When you see...</i>
BACK	
Method	
Choose	<i>(the decision the method requires)</i>
Steps	
Fails when	
Why it works	

Filled: mathematics

Front: An integral containing a product, where one factor gets simpler when differentiated.

Method: integration by parts, $\int u dv = uv - \int v du$.

Choose: u = the factor that simplifies.

Steps: pick u and dv ; differentiate and integrate; apply; check the new integral is easier than the old one.

Fails when: neither factor simplifies. Try substitution, or a reduction formula if there's a recurrence.

Why: it's the product rule, integrated.

Filled: mechanics

Front: A body on a slope, in equilibrium or accelerating along it.

Method: resolve along and perpendicular to the slope.

Choose: the positive direction, and state it.

Steps: draw the free-body diagram; resolve; write $\sum F = ma$ in each direction; solve.

Fails when: the surface is curved, or the body leaves the surface.

Why: choosing axes along the motion makes the perpendicular equation an equilibrium one, so the normal force falls straight out.

Filled: circuits

Front: A circuit with two or more loops and unknown currents.

Method: Kirchhoff's voltage law around each independent loop.

Choose: loop directions, and state them.

Steps: label currents; write KVL per loop; solve simultaneously; check with KCL at a node.

Fails when: there are dependent sources you haven't accounted for, or the circuit is better handled by node voltages.

Why: energy conservation round a closed path.

Review protocol

1. Read the front.
2. **Produce** the method and the first step, out loud or on paper.
3. Turn the card. Check.
4. Sort into "got it" and "didn't".
5. Redo the "didn't" pile at the end, and again tomorrow.

Five minutes for a stack of twenty. Reading through them without producing anything is not review, however comfortable it feels.

What not to make a card for

- Definitions with no associated situation.
- Whole derivations. Use reconstruction instead.
- Methods you've never practised.
- Constants you can look up.

Fifty good cards per module beat four hundred fragments, and the fragments are usually facts, which is the wrong unit.

D

The Error Log Template

One page, physical, visible. Ninety seconds a day. Read it before every problem set and every practice paper.

The grid

Date	What went wrong	Why	The rule
------	-----------------	-----	----------

Rules that work

The fourth column is the only one that changes anything. Name the moment as well as the action.

Not a rule	A rule
Be more careful with signs	Check the sign of every answer before moving on
Read the question properly	Underline the verb before starting each part
Show more working	Write the substitution on its own line, always
Manage time better	20 minutes on one problem, then move
Practise more	Three problems on this topic every week until it retires
Don't rush	Stop at the timer, finished or not

Running it

- **Log immediately**, while you still remember why you did it.
- **Log technique errors as well as content errors.**
- **Read it before every set.** Thirty seconds, top to bottom.
- **Retire an entry** when its rule has held for three sets. Cross it off and date it.
- **Build a drill** every fortnight around the top two entries.

Drills, by entry type

Top entry	The drill
Sign errors	Fifteen problems; check only the signs; don't finish them
Time overruns	Six problems, hard timer; success is stopping
Wrong method chosen	Twenty identification-only problems: name the method, don't solve
Algebra slips	Ten problems, rewriting every line rather than working across
Misreading	Ten problems where you only underline the verb and state what is being asked

Three of those five involve solving nothing. If the failure is in one step, practising whole problems trains mostly the parts that already work.

A worked entry

9 Nov. Lost the last part of Q2: computed the eigenvector for $\lambda = 3$ when the question asked for $\lambda = -1$.

Why: found both eigenvalues, then wrote out the first one from habit without re-reading the question.

Rule: after finding candidates, underline which one the question asked for before computing anything.

Twenty seconds to write, and the rule is specific enough to follow while tired, which “be more careful” never is.

E

What the Evidence Actually Says

An honest summary, including the parts that get overstated.

I've kept the chapters free of citation clutter, so this is where the bookkeeping lives: which claims in this book rest on strong research, which on weaker research, and which are my own practice.

Strong: use these with confidence

Retrieval practice (the testing effect). Producing information from memory beats restudying it, across material types, ages and delays. Among the most replicated findings in the field. It works even when the retrieval attempt fails, provided feedback follows.

Spacing (distributed practice). Spreading study over time beats massing the same total time, robustly, over more than a century of work. The *optimal* intervals are much less settled, which is why this book says spread rather than optimise.

Interleaving, for mathematics specifically. Mixed practice produces worse session performance and better delayed test performance. Unusually for this literature, some of the strongest studies use real mathematics problems rather than artificial material.

Self-explanation. Explaining each step to yourself while working improves learning, with good evidence in physics problem-solving. It also tends to help struggling students most, which is uncommon.

Rereading and highlighting are weak. Consistently among the least effective techniques per hour studied, and consistently among the most popular. That combination is itself well documented.

Moderate: directionally sound, details uncertain

Desirable difficulties. The framework is well motivated and its components are individually supported. Its practical weakness is that “desirable” is easier to identify afterwards than in advance, so the 70% heuristic in Chapter 7 is a teaching rule of thumb, not a measured optimum.

Worked examples and the expertise-reversal effect. Worked examples help most early in learning and their advantage shrinks or reverses with competence. Well supported; the exact crossover point isn't something you can look up for your situation.

Pre-exposure before instruction. Advance organisers and pretraining generally help, with effect sizes that vary a lot with how dense the material is.

Learning by teaching. Well supported, including the finding that merely expecting to teach improves learning. Most studies are short-term.

Handwriting versus typing. More contested than popular write-ups suggest. The recommendation in Chapter 15 rests mainly on the practical point that symbols and diagrams are faster by hand.

Weak or absent: my practice, not a finding

The problem log. No direct literature on this format. Its components, abstraction of structure, self-explanation, spaced review, are supported. The format is mine.

The error log. Standard practice in coaching and competitive-exam preparation, thinly studied academically. I use it because it works with students.

Method cards with the tell on the front. Flashcards are well studied; this specific inversion for method-based subjects is my adaptation.

The ten-minute post-lecture routine. The underlying claim, brief retrieval shortly after learning, is solid. The specific format and timing are mine.

Everything in Chapter 21 about recovery. Teaching experience, stated as such in the chapter.

Things people claim that this book doesn't

Learning styles. There's no good evidence that matching teaching to a preferred "style" improves outcomes. It's one of the most durable myths in education. Study the way the material demands, which for STEM means producing and problem-solving.

Fixed intervals as magic numbers. Spaced repetition software often implies precise scheduling matters. Spreading matters; the exact day rarely does at student timescales.

Multitasking, and background media. Divided attention reliably costs you on demanding material. No amount of preference changes it.

Any single technique as a complete answer. Retrieval without spacing fades. Spacing without interleaving leaves selection untrained. The combination is the point.

How to read a study-skills claim

Three questions, and they'll filter most of what you'll meet:

1. **What was the delay?** Techniques that look good on an immediate test often lose on a delayed one, which is the test that matches an exam.
2. **What was the material?** Word pairs are not differential equations. Prefer studies using material like yours.
3. **Compared with what?** "Better than nothing" is a weak claim. "Better than the same time spent restudying" is the useful one.

F

Further Reading, Honestly Annotated

Short list, with who each one is actually for. Books I haven't used aren't here.

On how learning works

Peter Brown, Henry Roediger and Mark McDaniel, *Make It Stick*. The best general introduction to the evidence, written by two researchers and a novelist. Strong on why the effective techniques feel wrong. Light on what to do on a Tuesday night in a maths degree, which is the gap this book is trying to fill.

Barbara Oakley, *A Mind for Numbers*. Aimed squarely at STEM students and much better known than anything else on this list. Its strength is the psychology and the reassurance; its limit is that it explains the mechanisms and leaves you to design the system. Read it if this book's Part I resonated and you want the science at greater length.

Ulrich Boser, *Learn Better*. Broader and more journalistic. Useful on the motivational and metacognitive side.

On practice and expertise

Anders Ericsson and Robert Pool, *Peak*. The deliberate practice literature, from the person who built it. The mapping onto academic study is looser than onto music or chess, and the book overstates the case in places, but the core idea, that practice must target specific weaknesses with feedback, is exactly what the error log implements.

George Pólya, *How to Solve It*. Seventy years old and still the best book on what to do when you're stuck on a problem. The heuristics chapter repays rereading every couple of years.

On the subjects themselves

Lara Alcock, *How to Study as a Mathematics Major* (also published as *How to Study for a Mathematics Degree*). Written by a mathematics education researcher, and the best account of what changes in the transition to proof-based courses.

Lara Alcock, *How to Think About Analysis*. If analysis is the module that's breaking you, this is the one.

The companion volumes

***The STEM Exam Handbook*.** Where this book stops. Timing, triage, method marks, salvage writing, and ten annotated scripts. If your marks are lower than your understanding,

and especially if you don't finish papers, that's where the missing marks are.

How to Read Mathematics. A procedure for definitions, theorems, proofs and papers, plus honest reading speeds. The natural companion to Part II of this book if your subject is mathematics-heavy.

How to Write Mathematics. For proofs and problem sets: writing solutions that collect the marks the work deserved.

How to Write Physics and Engineering. Lab reports, derivations, design documents and the final-year project. The other half of Chapter 22.

Where not to spend your time

General productivity systems. Elaborate task-management setups have a way of becoming the activity. The system in this book is four artefacts and a calendar block; that's deliberate.

Videos about studying. There's an enormous amount of it, most is the same handful of ideas, and watching it is the most fluent way to feel like you're improving without changing a single hour of your week.

Note-taking apps, past a point. Use whatever writes fastest. Nobody has ever failed a module for using the wrong app, and plenty have lost a weekend to configuring one.

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 what your week contains, 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 whoever in your year is working the hardest for the least return.

The appendices are working tools rather than reading. The problem log, the method cards and the error log are meant to be photocopied and used with a pen. You'll find printable companions and more of my writing for students at **gauravtiwari.org**, and corrections and questions are welcome there.

If you try the system for a term and something in it doesn't survive contact with your course, I'd genuinely like to know which part and why. Future editions will owe their sharpest fixes to those messages.

Also by Gaurav Tiwari

The STEM Exam Handbook

Timing, Triage, and Partial Credit in Maths, Physics, and Engineering Exams

Where this book stops. The four-minute triage, how method marks really work, salvage writing, and ten annotated scripts where every lost mark is traced to its cause. If your marks are lower than your understanding, start there.

How to Read Mathematics

Definitions, Theorems, Proofs, and Papers for Students

A six-step drill for opening any definition, three ways to read a proof, a stuck protocol that always ends in an action, and honest reading speeds.

How to Write Mathematics

Proofs, Papers, and Problem Sets for Students

Writing solutions that collect the marks the work deserved: proof makeovers, full-credit problem sets, and the grader's-eye view of partial credit.

How to Write Physics and Engineering

Lab Reports, Derivations, Design Documents, and Your Final-Year Project

The report side of a STEM degree, from the four-sentence abstract to the final-year thesis, including uncertainty written as prose rather than arithmetic.

Other titles

- *LaTeX for Students*
- *LaTeX for Theses and Dissertations*
- *The Student Research Workflow*
- *Indian Polity: Complete Notes for UPSC Prelims & Mains*

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