

The STEM Exam Handbook

The STEM Exam Handbook

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

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 walked out of an exam knowing
they could have done better than that.*

Contents

Preface	xiii
----------------	-------------

I What You Are Actually Being Marked On	1
--	----------

1 The Paper Is a Machine	2
1.1 How a paper gets built	2
1.2 The arithmetic nobody does	3
1.3 Which questions are actually hard	4
1.4 What the machine can and can't see	4
1.5 The five-question paper audit	5
1.6 What's coming	6
1.7 The front page is not decoration	6
1.8 Three paper shapes, three strategies	7

2 How Marking Schemes Actually Work	9
2.1 The three kinds of mark	9
2.2 Follow-through, and why one slip needn't cost the question	10
2.3 Four conventions worth knowing by name	11
2.4 Reading a real scheme	11
2.5 When there's no scheme at all	12
2.6 What this changes about how you write	12
2.7 A scheme, read line by line	13
2.8 Where to find schemes, and what to do with three of them	14

3 Where Marks Leak	15
3.1 Leak 1: you never reached it	15
3.2 Leak 2: you ran out of time inside a question	16
3.3 Leak 3: no method shown	16
3.4 Leak 4: the uncaught arithmetic slip	16
3.5 Leak 5: you misread it	17
3.6 Leak 6: the cheap stuff at the end	17

3.7	Audit your own leaks	17
3.8	One drill per leak	18
3.9	What the leaks have in common	19
3.10	A worked audit	19
3.11	Why the audit has to use the script	20
 II The Exam Hour		21
4	The First Four Minutes	22
4.1	The protocol	22
4.2	Step 1: the arithmetic, written down	23
4.3	Step 2: read everything	23
4.4	Step 3: write the order, with clock times	24
4.5	Step 4: underline the verbs	24
4.6	Worked: four minutes on a real paper	25
4.7	What the four minutes feel like	25
4.8	When the paper has choice	26
4.9	When the triage comes back all C	26
4.10	Triage on the three common paper shapes	27
5	Order, and Clock Discipline	28
5.1	The running total	28
5.2	Checkpoints, not vigilance	28
5.3	Ordering inside a question	29
5.4	When to break your own order	30
5.5	The physical setup	30
5.6	Your checkpoints, computed once	31
5.7	The two-thirds rule for the last checkpoint	31
6	The Abandonment Decision	33
6.1	Why this is hard	33
6.2	The rule	33
6.3	How to leave a question	35
6.4	Coming back	35
6.5	The three things that feel like progress and aren't	36
6.6	The question you never start	36
6.7	Worked: an abandonment that saved eleven marks	36
6.8	The note to your future self	37

7	The Last Ten Minutes	39
7.1	The sweep, in order	39
7.2	Step 1: the blanks	39
7.3	Step 2: state the answers	40
7.4	Step 3: units and figures	41
7.5	Step 4: salvage the abandoned	41
7.6	Step 5: back to the hard one, maybe	41
7.7	The five minutes you should have reserved	42
7.8	Scaling the sweep to the paper	42
7.9	The two-alarm system	42
7.10	And when they call time	43
7.11	Worked: ten minutes, seven marks	43
7.12	The sweep on a paper that went well	44
III	Writing Under Pressure	45
8	Writing for Method Marks	46
8.1	The rule that decides most of it	46
8.2	Before and after	46
8.3	What “show your working” does not mean	47
8.4	Writing so a tired person can follow	47
8.5	The specific case of “show that”	48
8.6	The physical script	49
8.7	The forty-second insurance policy	49
8.8	Three more, from three subjects	50
8.9	What “visible” costs you, measured	51
9	State the Plan	52
9.1	What it looks like	52
9.2	Why it earns marks four different ways	52
9.3	The thirty seconds before the plan	53
9.4	When you can’t write a plan	54
9.5	Plans for proof questions	54
9.6	When the plan is wrong	55
9.7	The habit, compressed	55
9.8	Plan lines for the twenty questions you’ll actually meet	56
9.9	The tell, and why the plan is really about recognition	58

10 Salvage	59
10.1 Why blank pages are the most expensive thing on a script	59
10.2 The salvage stack	59
10.3 Worked salvage	60
10.4 Salvage on proof questions	61
10.5 Salvage as a first move, not a last one	61
10.6 What not to write	62
10.7 Budget it	62
10.8 Salvage in physics and engineering	62
10.9 The two-minute rule for salvage across a paper	63
11 Checking Work You Have No Time to Check	65
11.1 The five-second checks	65
11.2 Where checks pay most	67
11.3 The ugly-answer heuristic	67
11.4 Checks by subject	68
11.5 The check budget	69
11.6 Six checks, six real errors caught	69
11.7 Checking a multi-part question, in the right order	69
IV The Term Before	71
12 The Past-Paper System	72
12.1 Three modes, and they're not interchangeable	72
12.2 The schedule	73
12.3 The review is the point	73
12.4 Redo, don't just review	74
12.5 Question banks versus whole papers	74
12.6 Simulating the room	75
12.7 Where to get papers	76
12.8 A term, on a calendar	76
12.9 Marking yourself honestly	77
12.10 When there are no past papers	78
13 The Error Log	80
13.1 What goes in it	80
13.2 Why it works	81
13.3 How to run it	81

13.4	Turning the log into practice	82
13.5	Taking the log to somebody else	82
13.6	The last three days	83
13.7	A worked entry	83
13.8	Four weeks of a real log	84
13.9	Building the drill from the log	84
13.10	Retiring entries, and watching it shrink	85
14	Retrieval Practice for Problem Subjects	86
14.1	Retrieval, when there's nothing to recall	86
14.2	The three artifacts that make it work	87
14.3	Spacing, in practice	88
14.4	Interleaving, and why blocked practice lies	88
14.5	The problem log, in practice	89
14.6	What to stop doing	89
14.7	A workable week	90
14.8	A method map, built	90
14.9	A week that actually fits	91
15	The Final Two Weeks	93
15.1	What the two weeks are for	93
15.2	Days 14 to 10: close the real gaps	93
15.3	Days 9 to 4: rehearse	94
15.4	Days 3 to 1: patch, don't learn	94
15.5	A season, not an exam	95
15.6	The day before	96
15.7	The morning	96
15.8	If you're badly behind	96
15.9	What the last fortnight cannot do	97
V	Particular Kinds of Question and Paper	98
16	Units, Figures, and Sanity Checks	99
16.1	Units are marks, and they're also a free check	99
16.2	Significant figures, without the folklore	100
16.3	Significant figures, five worked cases	100
16.4	The sanity checks that belong to physics	101
16.5	Prefixes, the quiet killer	101

16.6	Estimation questions, and why they're gifts	102
16.7	The unit table you should know cold	103
16.8	Answering in the right form	103
16.9	Worked: five formulas checked in forty seconds	104
16.10	Worked: the prefix disaster	104
17	Data, Graphs, and Uncertainty	106
17.1	The graph checklist	106
17.2	Linearising	106
17.3	Uncertainty: the rules you actually need	108
17.4	The discussion marks	108
17.5	Uncertainty, worked	109
17.6	Anomalies	109
17.7	The tables you're given	110
17.8	Worked: a complete data question	110
17.9	Where this question usually loses marks	111
18	Multi-Part and "Show That" Questions	113
18.1	The ladder	113
18.2	How to read "hence"	114
18.3	Show that: what's actually being marked	114
18.4	When a later part is easier	115
18.5	The instruction verbs, complete	115
18.6	Marks tell you the size of the answer	117
18.7	Worked: a ladder, climbed and misclimbed	117
18.8	Reading the mark split as a map	118
19	Multiple-Choice and Short-Answer Papers	120
19.1	What changes, and what doesn't	120
19.2	The economics of guessing	121
19.3	The three-pass sweep	122
19.4	Techniques that only work here	122
19.5	Short-answer papers	123
19.6	Where MCQ marks actually go	123
19.7	The words that flip an item	124
19.8	The answer sheet itself	125
19.9	Four items, worked the fast way	125
19.10	Reading the distractors backwards	126

20 Open-Book, Take-Home, and Formula Sheets	128
20.1 Why open-book papers are harder	128
20.2 Preparing for open book	129
20.3 The lookup budget	129
20.4 Take-home papers	130
20.5 Formula sheets	131
 VI The Room	 132
21 Nerves, Blanking, and the Restart	133
21.1 What nerves actually do to you	133
21.2 The blank	133
21.3 The panic spiral, and how to break it	134
21.4 Physical things that help, unglamorously	135
21.5 The night before and the morning of	135
21.6 The kit	136
21.7 The first sixty seconds in the seat	136
21.8 If it goes badly anyway	137
21.9 Worked: a restart that worked	137
21.10 Things that make the room worse, and are avoidable . . .	138
21.11 Two exams in one day	138
 22 After the Paper	 140
22.1 The twenty minutes, within a day	140
22.2 What not to do	141
22.3 When the marks come back	141
22.4 What actually improves between papers	142
22.5 If the result is bad	142
22.6 The last thing	143
22.7 The post-mortem between papers	143
22.8 What improvement actually looks like	144
 VII Toolkit	 145
A Ten Annotated Scripts	146
A.1 Script 1: the mental substitution	146
A.2 Script 2: the unstated answer	147

A.3	Script 3: the ladder ignored	148
A.4	Script 4: never reached	149
A.5	Script 5: the calculator in the wrong mode	150
A.6	Script 6: the blank “show that”	150
A.7	Script 7: good salvage	151
A.8	Script 8: the answer to a different question	152
A.9	Script 9: no units, six times	152
A.10	Script 10: the graph rushed	153
A.11	What the ten have in common	154
B	Two Papers, Start to Finish	155
B.1	Paper 1: mathematics, 100 marks, 180 minutes	155
B.2	Paper 2: engineering, 90 minutes, mixed format	158
B.3	What both papers have in common	160
C	The Marker’s Shorthand	161
D	The Protocol Card	164
E	The Error Log Template	167
F	Marks-per-Minute Tables	169
G	The Two-Week Countdown	173
H	Where to Find Papers and Schemes	175

Preface

You knew that one.

You've had the walk home after an exam where the answers arrive one after another, twenty minutes too late. The integral you set up wrong, the question you never reached, the last part you'd have got in ninety seconds if you'd had ninety seconds.

That gap between what you knew and what you scored is the subject of this book, and it's usually much bigger than students think. In a paper marked mostly on method, a student can know eighty per cent of the material and walk out with fifty-five per cent of the marks, purely on execution.

The good news is that the gap is mechanical. It comes from a small number of repeatable mistakes: starting at question 1 because it's question 1, spending twenty-five minutes on a six-mark part, doing arithmetic in your head, leaving a page blank because you couldn't finish, checking nothing. Each one has a fix, and none of the fixes requires you to know any more mathematics or physics than you already do.

This isn't a revision guide and it isn't a study-skills book. There's nothing here about mind maps. It's about the three hours themselves, plus the term of practice that makes those three hours go well.

A word on who it's for. It's written for undergraduate maths, physics, and engineering students, because those papers share a structure: a small number of long, multi-part problems, marked with method marks, where partial credit is most of the game. It works just as well if you're drilling for JEE, GATE, or a university entrance paper, and there's a chapter on data and uncertainty questions that engineers and physicists will use more than mathematicians. Everything transfers to any exam where you have to *work something out* rather than recall it.

Two honest warnings.

First, this book cannot rescue you from not knowing the material. Nothing can. If you've done no work, no triage protocol will save the afternoon. What it will do is make sure you collect everything you have earned, and for most students that's between five and fifteen percentage points sitting on the table.

Second, some of the advice will feel wrong the first time. Leaving a question unfinished on purpose. Writing down a method you know won't reach the answer. Spending four minutes reading before you write anything at all. Every one of those feels like giving up, and every one of them is worth marks. Try them on a past paper before you decide.

Let's go and get the marks you've already earned.

Gaurav Tiwari
gauravtiwari.org

Part I

What You Are Actually Being Marked On

1

The Paper Is a Machine

Most students walk into an exam thinking of it as a test of what they know.

It isn't. It's a machine that converts three hours of your writing into a number, and it does that by a published, mechanical procedure. Somebody designed the machine. Somebody else operates it with a red pen and a rubric.

You can learn how it works. And once you have, a large chunk of your mark stops depending on how much you know and starts depending on decisions you make in the first four minutes.

1.1 How a paper gets built

Papers aren't assembled by picking hard questions. They're assembled to a specification, and the specification is boringly consistent across institutions.

A typical three-hour undergraduate paper has six to eight questions worth 15 to 25 marks each. Within each question there are three to five parts, labelled (a), (b), (c), and those parts are almost always ordered by difficulty. Part (a) is a definition, a statement, or a routine computation. Part (b) applies it. Part (c) is the one the examiner considers interesting, and the last part is often there to separate the top ten per cent from everyone else.

Here's what that means in practice, and it's the single most valuable fact in this chapter.

Marks are not spread evenly across difficulty. The first 60% of the marks in any question sit in its easiest 40% of work. The last part is worth the fewest marks per minute of anything on the paper.




Figure 1.1: One 20-mark question, four parts. Bar length is the time each part takes. Part (a) pays five times as many marks per minute as part (d), and they are printed one under the other.

Part (a) pays 1.3 marks a minute. Part (d) pays 0.25. They're on the same question, printed one under the other, and students work them in the order printed.

That ordering is the default, and the default is expensive.

1.2 The arithmetic nobody does

Take a three-hour paper, 100 marks, answer all questions. That's 1.8 minutes per mark, and you should write that number at the top of your script before you write anything else.

Now the numbers that follow from it:

- A 4-mark part deserves about 7 minutes. If you're 15 minutes in, you're losing.
- A 20-mark question deserves 36 minutes. Not 50, however interesting it is.
- Fifteen minutes spent on a 5-mark part you don't finish costs you roughly 8 marks elsewhere.

That last one is the whole book in a sentence. The marks you lose to bad time allocation don't come off the question you were stuck on. They come off the questions you never reached, and those were probably easier.

FROM THE MARKER'S SIDE

From my side of the desk, the most painful script isn't the one that gets things wrong. It's the one where questions 1 to 4 are near-

perfect, question 5 is a half-page of increasingly small handwriting, and questions 6 and 7 are blank. That student knew enough for a first. They walked away with a 2:2, and nothing in their subject knowledge caused it.

1.3 Which questions are actually hard

Students judge difficulty by how a question looks. Long questions look hard. Unfamiliar notation looks hard. A question about a topic you disliked looks hard.

None of those correlate well with how many marks you'll get.

What does correlate:

- **Have you seen this shape before?** Not this exact question, this shape. If you've done four questions that started "show that the sequence defined by...", you'll get marks on the fifth.
- **Can you start?** A question you can start is worth attempting even if you can't see the end. A question you can't start is worth four minutes and no more.
- **Does part (a) give you part (b)?** Very often it does, and spotting that is worth more than any amount of cleverness.

That third point deserves its own note, because it's the most commonly missed gift in exam papers. Multi-part questions are usually built as a ladder, where each part hands you the tool for the next. When part (a) says "show that $f'(x) = \dots$ " and part (b) asks for the stationary points, part (a) isn't a warm-up. It's the answer to part (b), handed to you in advance.

Chapter 18 is all about reading those ladders. For now, just notice that a question with four parts is usually four small questions with the hard thinking already done for you, and a question with one part is where the real difficulty lives.

1.4 What the machine can and can't see

The machine sees your script. That's all.

It can't see that you knew the method and ran out of time. It can't see that you meant to write $\pi/4$ and wrote $\pi/2$ because your hand was faster

than your head. It can't see the working you did on the page you crossed out.

But it also can't see that you had no idea what you were doing and copied a standard method that happened to fit. Method marks don't ask whether you understood. They ask whether the method is on the page.

That cuts both ways, and it's worth being clear-eyed about it. The exam is not a perfect measure of what you know, in either direction. It's a measure of what you wrote down in three hours, and the skill of writing down more of what you know is a separate skill from knowing it.

That skill is learnable in about two weeks. Most students never learn it, because nobody teaches it and it looks like cheating to think about it.

Nobody gets marks for what they knew. Everybody gets marks for what's on the paper. Closing the gap between the two is the entire subject of this book.

1.5 The five-question paper audit

Before you sit a paper, you can know most of its shape. Take three past papers from your course and answer five questions. Twenty minutes, once, and it informs everything else in this book.

1. **How many questions, and are they equally weighted?** If they're not, the heavy ones need their own budgets.
2. **How many parts per question, and how are the marks split across them?** Front-loaded, back-loaded, or flat, as in Chapter 18.
3. **Which topics appear on all three papers?** Those are your revision priority, and it's rarely the ones you find most interesting.
4. **What's the ratio of "show that" and "prove" to "find" and "calculate"?** It tells you how much of the paper is marked on argument rather than answers.
5. **Is there a formula sheet, and what's on it?** Everything on it is memory work you don't have to do.

Students who've done this walk into the exam knowing the shape of the thing in front of them. Students who haven't meet it for the first time with the clock running, which is the most expensive moment to be learning it.

FROM THE MARKER'S SIDE

Papers are far more stable year to year than students expect. The same topics, in a similar order, with a similar mark split. Not the same questions, but the same skeleton, and the skeleton is what you can prepare for.

1.6 What's coming

Part I is the machine: how papers are built, how mark schemes work, and where marks leak out of a script.

Part II is the exam hour itself. Triage, ordering, the clock, and the decision to abandon a question.

Part III is how to write when the clock is running: method marks, stating the plan, salvage, and checks that cost seconds.

Part IV is the term before, which is where the exam is actually won: past papers as a system, the error log, and what to do in the final fortnight.

Part V is for physicists and engineers specifically, and Part VI is about the room, the nerves, and what to do after.

Start anywhere. But if the next paper is close, read Chapter 4 and Appendix D tonight, and do one past paper with a timer tomorrow.

1.7 The front page is not decoration

Every paper opens with a rubric nobody reads. It tells you things that change what you should do, and it takes forty seconds.

What it says	What it changes
“Answer all questions”	no choice; triage decides order only
“Answer 4 from 6”	you choose; pick 4 plus a reserve
“Section A: answer all. Section B: answer 2 of 4”	two different strategies on one paper
“Marks for each question are shown”	you can compute marks per minute
“Formula sheet provided”	don’t waste memory on what’s printed
“Calculators must not be programmable”	check yours the week before
“Answers to 3 significant figures unless otherwise stated”	a standing instruction for every part
“Cross out work you do not want marked”	means uncrossed work is marked

That last row matters more than it looks. On most papers, work you leave uncrossed is marked, and if you produce two attempts at a question, the marker may take the worse one or may take the first. So cross out the attempt you don’t want, with a single line, and never scribble out anything you’d be happy to have marked.

1.8 Three paper shapes, three strategies

Papers come in a small number of shapes, and each one changes the plan.

The uniform paper. Six questions of equal marks, answer all. The simplest case. Triage for order, keep to the budget, sweep at the end.

The choice paper. Answer four from six. Now selection matters as much as execution, and the classic error is choosing a question because its part (a) looks easy, then discovering that part (d) is worth eight marks and impossible. Read *all* parts of a question before committing to it.

The two-section paper. A short section and a long section, usually with different mark densities. Compute marks per minute for each section separately, because they’re rarely the same, and the short section is usually the better rate.

Work out the marks-per-minute rate of each section, not just the paper. A short-answer section paying 1.2 marks a minute and a long section

paying 0.5 is a common design, and it tells you where the paper is won.

DO THIS BEFORE THE NEXT PAPER

Find a past paper for a subject you're taking. Don't do it. Just work out, for every part of every question, the marks and your honest estimate of the minutes. Then compute marks per minute for each part and circle the three worst deals on the paper. That's twenty minutes, and it'll change how you attack the real thing.

2

How Marking Schemes Actually Work

Almost nobody has read a mark scheme.

They exist, they're often published alongside past papers, and they're the single most useful document in your subject that you're not reading. A mark scheme tells you, line by line, exactly what a marker is instructed to reward. Reading three of them will teach you more about exam technique than any amount of advice, including mine.

Here's what you'll find inside.

2.1 The three kinds of mark

Most STEM mark schemes divide marks into three types. The labels vary by institution, but the structure is near-universal.

Read the M line again, because two consequences fall straight out of it and they're worth a grade between them.

You get method marks for a method that fails. Set up the integral correctly, integrate it wrongly, and the setup still scores. Which means an unfinished question with visible method is worth far more than a blank one, and often worth more than a finished question with no working.

You get no method marks for a method in your head. If you did the substitution mentally and wrote only the result, and the result is wrong, you score zero on a part where a slower student who wrote three lines scores three quarters.

M method *usually 50–70% of the paper*
Awarded for using a correct method, whether or not the arithmetic comes out. Visible on the page or not awarded.

A accuracy *20–40%*
Awarded for the correct value or expression. Usually *dependent*: no M mark, no A mark, even if the answer is right.

B independent *the rest*
Stand-alone: a stated definition, a sketch, a correct unit. Doesn't depend on anything else being right.

Figure 2.1: The three mark types. Method marks are the majority, and they don't require you to be right.

FROM THE MARKER'S SIDE

I can only mark what I can see. When a script jumps from the question to a wrong number with nothing in between, there's nothing for me to reward, and I promise you I've looked. The rubric says "M1 for a correct substitution", and if the substitution isn't written down, my hands are tied. Show me the substitution and I'll give you the mark even though your final answer is wrong.

2.2 Follow-through, and why one slip needn't cost the question

Here's the mechanism that surprises students most.

If you make an arithmetic error early and then do everything after it correctly, most schemes award the later marks anyway. It's called follow-through, or error carried forward, and it's usually written in the scheme as "ft" or "ecf".

So a slip in line two costs you one accuracy mark, not the whole question.

But follow-through only works if the marker can follow. Which gives you a concrete, mechanical instruction.

After an error, keep going with your wrong number and finish the method. Follow-through pays the rest of the marks, but only if the working is visible and the wrong number is carried consistently.

The student who spots a mistake, panics, scribbles everything out and restarts with four minutes left ends up with a blank space where a follow-through would have paid three marks. Cross out lightly, or better, don't cross out at all. A single line through work you've abandoned is enough, and crossed-out work is still marked at most institutions if nothing replaces it.

2.3 Four conventions worth knowing by name

Seen once. If a mark is available for stating a formula, you get it the first time you state it, not once per use. So state it explicitly, early, even if it feels obvious.

ISW: ignore subsequent working. Once you've written the correct answer, some schemes ignore anything you do to it afterwards. Some don't. Since you can't tell which, the safe habit is to end each part with a clear, final statement of the answer and stop writing.

Dependent marks. Written "dM1" or similar. The mark is only available if the earlier one was earned. This is why skipping the setup line to save time can cost you three marks, not one.

Any correct method. Most schemes accept alternative approaches and say so explicitly. You don't have to use the method from lectures. You do have to use one that works, and write it down.

2.4 Reading a real scheme

Take a past paper in your subject and its published scheme, and put them side by side for one question. You're looking for four things:

1. How many marks are method versus accuracy. Count them. The ratio will surprise you.
2. Which single line the scheme calls the key step, because that's the line the examiner cares about.
3. Where it says "ft" or "oe" (or equivalent) or "any valid method", because those are the places you have freedom.

4. What earns the B marks: units, a sketch, a stated assumption. Those are free and students skip them constantly.

Twenty minutes with one scheme, and you'll never again write a bare answer with no working.

2.5 When there's no scheme at all

Plenty of departments publish papers and no schemes. You can still infer most of what you need, from the mark allocation alone.

Count the marks and count the steps. A five-mark part with four natural steps is telling you those steps are the marks. Write all four, each on its own line.

Look for the answer-bearing step. Usually one mark, usually the last. Everything before it is method, which means most of the part is insured by writing the steps down.

Assume method dominates. Across institutions, method is 50 to 70 per cent of a STEM paper's marks. In the absence of information, write as though every part is mostly method, because it almost certainly is.

Ask your lecturer one question. Not "what's on the exam". Ask: "roughly what proportion of the marks are method marks?" It's a reasonable question, it isn't asking for anything confidential, and the answer changes how you write for the rest of the year.

FROM THE MARKER'S SIDE

It's a question almost nobody asks, and most of us would answer it happily and in detail, because a student writing for method is easier to mark and scores better. Nobody benefits from you writing bare answers.

2.6 What this changes about how you write

Three habits fall out of the whole chapter, and they're the ones Part III develops properly:

Write the method line even when you don't need it. The substitution, the formula, the setup. You lose fifteen seconds and buy insurance on every part.

Never leave a part blank if you can state its method. Two lines of correct approach on a question you can't finish is typically worth a third of that part.

Finish with an explicit answer, stated once, and stop. No trailing algebra after the box.

None of the three requires you to know more mathematics. All three are worth marks on every paper you'll ever sit.

2.7 A scheme, read line by line

Here is a five-mark part with its scheme beside it, which is what you'll find when you go and look at a real one.

Q. Solve $2 \sin^2 \theta = 3 \cos \theta$ for $0 \leq \theta \leq 2\pi$, giving answers in radians.
[5]

M1	use of $\sin^2 \theta = 1 - \cos^2 \theta$
A1	$2 \cos^2 \theta + 3 \cos \theta - 2 = 0$
M1	valid attempt to solve their quadratic in $\cos \theta$
A1	$\cos \theta = \frac{1}{2}$, and $\cos \theta = -2$ rejected
A1	$\theta = \pi/3, 5\pi/3$ (ft on their $\cos \theta$)

Four things to take from that, and they're all typical.

The first mark is for an identity, not an answer. Writing $\sin^2 \theta = 1 - \cos^2 \theta$ before doing anything else claims it.

"Their quadratic" means follow-through. If you rearranged wrongly at A1 but then solved your wrong quadratic properly, the third mark is still yours.

Rejecting the impossible root is an explicit mark. Writing " $\cos \theta = -2$ has no solutions" is worth as much as finding the good root, and students omit it constantly because it feels obvious.

"ft on their $\cos \theta$ " rescues the last mark. Wrong value of $\cos \theta$, correctly converted to angles in the right range, still scores.

So a student who makes one sign error early and does everything else carefully scores 4 of 5. A student who does the whole thing in their head and writes " $\theta = \pi/3, 5\pi/3$ " scores 5 if they're right and 0 if they aren't.

2.8 Where to find schemes, and what to do with three of them

Departments vary. Some publish schemes with past papers, some publish neither, and some will hand them over if you ask the lecturer in week two rather than week eleven.

If yours publishes nothing, use school-level schemes. A-level, IB, and their equivalents publish full schemes free, and the marking logic, M1, A1, B1, ft, oe, ISW, is identical to what universities use. The mathematics is below your level, which doesn't matter, because you're studying the marking, not the material.

Do this once, with three schemes, and it takes an hour:

1. Count the M marks and the A marks on one long question. Note the ratio.
2. Find every "ft" and "oe" and work out what freedom each gives you.
3. Find every B mark and see what it's for. Usually a definition, a sketch, a unit, or a stated assumption, and usually something students skip.
4. Take your own attempt at that question and mark it with the scheme, strictly. Count the marks you'd have lost to invisible method.

That last number is the one that changes behaviour. It's usually between four and eight, on a single question.

DO THIS BEFORE THE NEXT PAPER

Find one published mark scheme for a subject you're taking. Take a single 20-mark question and count the M marks and the A marks. Then look at your own last attempt at a similar question and ask, honestly, how many of the M marks a marker could have awarded from what you actually wrote down. That number is your leak.

3

Where Marks Leak

Every script leaks marks in the same handful of places.

Not “you didn’t know the material” places. Places where the student knew the material and the marks went missing anyway, which is the only kind of loss this book can do anything about.

Here are the six, roughly in order of how much they cost.




Figure 3.1: The six leaks, with the share of recoverable marks each one typically accounts for. The two biggest are both time-allocation problems wearing different clothes.

3.1 Leak 1: you never reached it

The biggest single loss on most scripts, and the one students find hardest to believe, because it doesn’t feel like a mistake. It feels like running out of time.

But look at what it costs. An unattempted 18-mark question is 18 marks gone. Not partial credit. All of it. And the question you never reached was, on average, exactly as hard as the ones you did.

The cause is almost never slowness. It's spending 40 minutes on a question budgeted for 25 because you were close, you could feel it, and stopping felt like quitting.

Chapters 5 and 6 are entirely about this leak.

3.2 Leak 2: you ran out of time inside a question

Different failure, same root. You reached the question, started well, and the clock ran out with the last two parts half-done.

This one at least earns partial credit, so it's cheaper than leak 1. It's still expensive, and it's still a budgeting problem rather than a knowledge problem.

The fix is a rule you'll meet properly in Chapter 5: a checkpoint every thirty minutes, comparing where you are against where the marks say you should be.

3.3 Leak 3: no method shown

The one that hurts to look at as a marker. A correct approach, done mentally, with only the answer on the page. Answer wrong by a factor of two. Zero marks on a part where three quarters were available.

Students do this because writing feels slow and they're trying to save time. The arithmetic doesn't support them: writing the method line costs fifteen seconds and insures three or four marks.

If you take one habit from Part I, take that one.

3.4 Leak 4: the uncaught arithmetic slip

Everybody makes them. The question is whether you catch them, and the honest answer for most students is no, because they never look.

The good news from Chapter 2 is that follow-through means a slip usually costs one mark, not the question. The bad news is that an uncaught slip early in a multi-part question can make part (c) impossible, because the number you're carrying no longer behaves.

Chapter 11 is a set of checks that cost between two and fifteen seconds each. Magnitude, sign, units, limiting case. Not a full re-do, which you'll never have time for.

3.5 Leak 5: you misread it

Solving a different question beautifully is worth almost nothing, and it happens more than you'd think under time pressure. The classic versions:

- Missing a word: “show that” instead of “find”, “hence” instead of “otherwise”, “not” anywhere at all.
- Missing a constraint: “for $x > 0$ ”, “for real λ ”, “correct to three significant figures”.
- Answering part (b) with part (c)'s question in your head.
- Missing that part (c) says *hence*, which means it's telling you to use part (b), which you've just ignored.

The fix is mechanical and takes four seconds a part: underline the instruction verb and any constraint before you start writing. Do it every time and this leak closes almost entirely.

3.6 Leak 6: the cheap stuff at the end

Units. Significant figures. The final statement of the answer. A sketch that was asked for and never drawn. The stated assumption a B mark was waiting for.

Each one is worth a mark, sometimes two. Across a whole paper it's routinely five marks, and it's the purest waste on this list, because recovering it requires no thought at all. It requires a checklist and ten seconds a part.

For physics and engineering students this leak is bigger than for mathematicians, which is why it gets its own chapter (Chapter 16).

3.7 Audit your own leaks

The six leaks aren't equally distributed. Yours won't match the diagram, and the point of knowing the list is to find out which two are actually costing you.

So do this, once, properly. Take your most recent marked script, or a past paper you've done under timed conditions. Go through it question

by question, and for every mark you didn't get, assign it to one of seven buckets: the six leaks, or "I didn't know it".

Then total the buckets.

Two things happen. First, the "I didn't know it" column is usually much smaller than the feeling in the exam suggested. Second, one or two leaks dominate everything else, and those two are your entire improvement plan for the next six weeks.

You don't have six problems. You have one or two, repeated on every paper you've ever sat. Find them once and you'll be fixing something specific instead of resolving to "be more careful".

3.8 One drill per leak

Once the audit names your top two, here is what to practise. Each drill takes twenty to thirty minutes and targets the failure directly rather than the subject around it.

Leak	The drill
1 Never reached	Do a full paper where the <i>only</i> rule is that every question must be started. Score yourself on coverage, not on correctness.
2 Ran out inside	Six questions, hard timer per question. Success is stopping on time, finished or not.
3 No method shown	Ten parts where you write only the structural lines and no arithmetic at all. Then mark them for method.
4 Arithmetic slips	Twenty short calculations, each followed by the five checks. Score the checks, not the answers.
5 Misread	Ten questions where you only underline the verb and the constraints and state what's being asked. Don't solve any of them.
6 Units and form	Take a marked script and rewrite every final answer line properly. Fifteen minutes.

Notice that four of the six drills involve doing no mathematics at all. That’s deliberate, and it’s the part students resist: if the failure is in reading, timing, or presentation, then practising more mathematics trains the thing that isn’t broken.

Drill the failing skill in isolation, scored on the skill. A timing drill you “pass” by finishing every question has not tested timing.

3.9 What the leaks have in common

Look back at the six and notice that five are decisions, not knowledge. Which question to start. When to stop. Whether to write the method line. Whether to check. Whether to underline the verb.

That’s why the rest of this book is a set of protocols rather than a set of topics. You’re not going to out-think these leaks in the room, because the room is where your judgement is worst. You’re going to run a procedure you decided on in advance, when you were calm.

That procedure starts in the next chapter, four minutes before you write anything.

3.10 A worked audit

Here is the seven-bucket audit done properly, on one real script. The student scored 58 and was convinced the problem was that they hadn’t revised enough.

Bucket	Marks	Where
1 Never reached the question	18	Q6 entirely
2 Ran out of time in a question	9	Q5(c), Q5(d)
3 No method shown	6	Q2(b), Q4(a)
4 Arithmetic slip, uncaught	4	Q1(c), Q3(b)
5 Misread the question	3	Q3(d)
6 Units, figures, final statement	2	Q4(c), Q4(d)
7 Didn’t know it	0	—
Total lost	42	

Forty-two marks lost, and not one of them to not knowing the material.

Twenty-seven of the forty-two, nearly two thirds, are buckets 1 and 2: pure time allocation. This student's revision plan for the next paper was going to be "do more past papers". What they actually needed was a triage protocol and a rule for when to abandon Q5.

Do the audit before you decide what to revise. Most students plan their revision against bucket 7 and lose their marks in buckets 1 and 2.

3.11 Why the audit has to use the script

You cannot do this from memory, and it's worth understanding why.

After an exam, the marks you remember losing are the ones that felt bad: the question you couldn't do, the thing you'd forgotten. Those are bucket 7 events, and they're emotionally vivid.

The other six buckets are invisible from the inside. Nobody walks out of an exam thinking "I lost four marks by not writing down substitutions". It doesn't feel like anything at the time, because at the time you knew what you were doing.

So get the script back if your institution allows it, and if it doesn't, do the audit on timed past papers instead, where you have the evidence in front of you.

FROM THE MARKER'S SIDE

The single most common reaction to seeing a marked script is surprise at how much went to execution. Students expect to find gaps in their knowledge and instead find unstated answers, method done mentally, and one part missed on the back of a page.

DO THIS BEFORE THE NEXT PAPER

Take your last marked script and do the seven-bucket audit. Every lost mark gets a bucket. Total them, and write the two biggest at the top of the page. Those two are what you're fixing this term, and everything else on this list can wait.

Part II

The Exam Hour

4

The First Four Minutes

The bell goes. Around you, thirty people turn to page one and start writing. Don't.

The four minutes you spend not writing at the start of a paper are the highest-value four minutes of the whole three hours, and they're the hardest four minutes to spend, because everyone else is scribbling and your page is blank.

Do it anyway. Here's what goes in them.

4.1 The protocol




Figure 4.1: The four-minute triage. Nothing here requires you to solve anything, and it decides most of your mark.

Four minutes. Roughly two marks' worth of time, on a paper where this protocol routinely saves ten.

4.2 Step 1: the arithmetic, written down

Total marks. Total minutes. Divide.

A 100-mark, three-hour paper is 1.8 minutes per mark. A 75-mark, two-hour paper is 1.6. Whatever it is, write it at the top of your script and underline it, because in ninety minutes when you're absorbed in a differential equation you'll need to see it without thinking.

Then subtract. Take five minutes off the end for the final sweep (Chapter 7) and four for this. What's left is your working time, and that's the number the ordering uses.

4.3 Step 2: read everything

Two minutes, the whole paper, no solving.

You're sorting into three buckets, and you write the letter next to each question in the margin:

- **A: I can start this right now.** I know the first line.
- **B: I can probably start it.** I recognise the shape, I need a minute to find the entry point.
- **C: I can't start it.** Nothing comes.

Note what you're sorting on. Not whether you can finish. Whether you can *start*. Those are different questions and only the second one is answerable in twenty seconds.

Most students never do this, and it costs them in a specific way: they meet question 6 for the first time at minute 130, discover it's the easiest thing on the paper, and have eleven minutes for it.

FROM THE MARKER'S SIDE

You can tell instantly which scripts were triaged. They start with question 4, then 1, then 6, and the marks are spread evenly across everything attempted. The untriated ones go 1, 2, 3, 4 and stop, with the last two questions blank, and it's nearly always the case that question 6 was one the student could have done.

4.4 Step 3: write the order, with clock times

Now commit to a sequence and write it in the margin. Not in your head. On the paper.

The default order that works:

1. All the A questions, easiest first.
2. The B questions.
3. The C questions, with whatever's left.

Against each, write the clock time it has to end. Not the duration, the actual time on the wall. “Q4 by 10:22” is a decision you can check in one glance; “Q4 gets 25 minutes” requires arithmetic you won’t do at minute 80.

Convert every time budget into a wall-clock deadline before you start writing. Under pressure, you can compare two numbers. You cannot reliably subtract them.

Two refinements worth knowing.

Start with an A, not the hardest A. You want the first fifteen minutes to go well. The confidence is worth real marks, and a good first question settles your hands.

Don’t leave all the C questions to the end. Give one C question a slot in the middle, when you’re warm and there’s still time to reach a part (a). C questions still have part (a)s, and part (a)s are cheap.

4.5 Step 4: underline the verbs

Thirty seconds, on your first question only. You’ll repeat it for each question as you reach it.

Underline the instruction verb: *show that, find, prove, deduce, hence, state, sketch, estimate*. Each one is an instruction about what your answer has to look like, and they aren’t interchangeable.

Then circle every constraint: the domain, the range, the accuracy, the units, the “where n is a positive integer”.

This is leak 5 from Chapter 3, and thirty seconds a question closes almost all of it.

4.6 Worked: four minutes on a real paper

A 100-mark, three-hour paper, six questions, answer all. Here is what actually gets written in the margin, and when.

0:00	Top of script: <i>100 marks / 180 min. Less 9 = 171 working. 1.7 min/mark. RAD.</i>
0:30	Q1, integration. Knows the substitution on sight. A
0:50	Q2, series. Recognises the shape, not the route. B
1:15	Q3, group theory. Nothing. C
1:35	Q4, eigenvalues. Done twenty of these. A
2:00	Q5, ODE, 20 marks. Big, but the type is familiar. B
2:25	Q6, vector calculus. Divergence theorem, clearly. A
2:40	Order written: <i>Q1, Q4, Q6, Q3(a) only, Q2, Q5.</i>
3:10	Deadlines written beside each: <i>10:00, 10:26, 10:57, 11:04, 11:32, 12:07.</i> Sweep at 12:20.
3:40	Q1: underlines <i>evaluate</i> , circles <i>to 3 s.f.</i>
4:00	Starts writing.

Four minutes, about sixty words written, and no mathematics done at all.

Two decisions in there are worth noticing. The 20-mark question went *last*, despite being the biggest, because it was a B and the two As after it were more certain. And the C question got a seven-minute slot in the middle for its part (a) only, rather than being left to die at the end.

4.7 What the four minutes feel like

Bad. They feel bad.

Everyone around you is writing. You're reading. There's a voice saying you're wasting time, and at minute three that voice gets loud.

Two things help. First, know that it's four minutes out of a hundred and eighty, and that you've allowed for it in the arithmetic you did in step 1. Second, practise it. If the first time you try this protocol is in a real exam, the nerves will win and you'll start writing at minute one. Run it on three past papers and it becomes the thing you do without deciding.

Any protocol you first attempt in an exam will fail. Not because it's wrong, but because the room removes your capacity to follow new instructions. Rehearse it until it's boring.

4.8 When the paper has choice

Some papers say “answer four questions from six”. The four minutes matter even more, and the sorting changes slightly.

Read all six. Mark them A, B, C. Then pick your four, and here's the part students get wrong: **pick five**. Choose four to attempt and one reserve, because roughly a third of the time a question that looked like an A turns out to have a nasty part (c), and you want to know in advance where you're going instead of scanning the paper in a panic at minute 95.

Also, don't attempt five hoping the best four count. On most papers they do count the best four, but the fifth attempt costs you the time that would have improved the other four, and a rushed extra question is worth less than fifteen more minutes on your best one.

4.9 When the triage comes back all C

It happens, and it's the moment students panic hardest, so decide what to do about it now rather than then.

If you read all six questions and can't start any of them, three things are true, and only one of them is bad.

You're four minutes in and adrenaline is at its peak. Reading comprehension is genuinely impaired for the first few minutes of an exam. Questions that look impossible at minute three routinely open at minute twelve.

You sorted on “can I start”, which is a high bar under stress. Go round again and ask a lower-bar question: which part (a) can I do? Not which question, which part. Almost every part (a) on almost every paper is a definition, a statement, or a routine computation.

If it's genuinely all unfamiliar, the paper is hard for everybody. Scaling exists. Your job is unchanged: collect the most marks available, which now means sweeping every part (a) on the paper first.

The protocol for an all-C paper:

1. Take one more minute. Mark the *parts*, not the questions.

- 2. Do every part (a) on the paper, in order, quickly. That’s typically 20 to 25 per cent of the marks.
- 3. Then return to whichever question’s part (b) looked closest.
- 4. Salvage everything you can’t finish.

Students who do this on a hard paper routinely finish in the top quarter, because most of the room spent forty minutes on question 1.

4.10 Triage on the three common paper shapes

The protocol is the same; the numbers change.

Paper	Triage	What to sort
6 long questions, answer all	4 min	questions, then parts
Answer 4 from 6	5 min	pick 4 plus a named reserve
Section A short + Section B long	4 min	Section B only; Section A is worked in order

That third row is worth a note. On a paper with a short-answer section, don’t triage the short section: the items are worth the same and cost the same, so the sorting effort doesn’t pay. Work it in order, fast, and mark anything that resists for a second pass.

Where triage pays is wherever questions differ in mark value or in how long they’ll take, which is exactly the long-answer section.

DO THIS BEFORE THE NEXT PAPER

Take a past paper and run only the four-minute protocol. Set a timer for four minutes, do the arithmetic, sort every question A/B/C, write the order with wall-clock times, and underline the verbs on the first one. Then stop and put the paper away. You’ve just practised the highest-value four minutes of the exam without spending three hours.

5

Order, and Clock Discipline

You have your order. Now you have to hold it for three hours while your judgement quietly degrades.

That's the real problem. It isn't that students don't know they should manage time. It's that time management requires the exact faculty that absorption destroys, and being absorbed in a problem is otherwise a good thing.

So you don't manage time in the exam. You install checkpoints beforehand and obey them.

5.1 The running total

Here's the single habit that fixes most clock problems.

Every time you finish a question, write two numbers in the margin: the clock time, and the marks you've attempted so far. Not marks you think you've earned. Marks attempted.

Then compare against where you should be.

The diagram is the whole chapter. One question ran long, and the gap it opened never closed, because there's no mechanism in an exam that gives time back.

Catch it at the first checkpoint and you lose two marks. Catch it at the end and you lose eight.

5.2 Checkpoints, not vigilance

You can't watch the clock continuously. Nobody can, and trying makes you worse at the mathematics.




Figure 5.1: One script on a 100-mark, 180-minute paper. The solid line is the pace that finishes the paper; the dashed line is what this student actually attempted. The gap opens at the first checkpoint and never closes.

So set three checkpoints. On a three-hour paper: at 45 minutes, at 90, and at 135. Write those times at the top with the marks you should have attempted by each.

At a checkpoint you do exactly one thing, and it takes ten seconds: compare your running total against the target, and if you're more than eight marks behind, you move on from whatever you're doing right now. Not at the end of the part. Now.

Decide the checkpoints before the exam, and make the rule automatic: eight marks behind means move, immediately, mid-question if necessary. Judgement in the room will always vote to keep going.

Why eight? Because it's roughly fifteen minutes, which is recoverable, and because a number smaller than that will fire constantly and you'll start ignoring it. Tune it to your own paper structure, but pick the number beforehand and write it down.

5.3 Ordering inside a question

The A/B/C sort ordered the questions. Inside a question, there's a second ordering decision, and it's the one that separates good scripts from excellent ones.

Do all the part (a)s first, across every question you're attempting.

That sounds strange until you look at the marks. Part (a)s are typically 20 to 25 per cent of the paper, they take about ten per cent of the time, and they're mostly independent of the rest. A student who sweeps every part (a) on the paper in the first forty minutes has banked a quarter of the marks and can then spend the remaining time wherever it pays.

The pure version of this is aggressive and not for everybody. A softer and still very effective version: within each question you attempt, never skip a part because it looks trivial, and never start part (c) until (a) and (b) are written down.

FROM THE MARKER'S SIDE

The scripts that surprise me are the ones where a strong student has skipped part (a) because it was "obvious" and gone straight to the interesting part. Part (a) was three marks for stating a definition. They were never awarded, and the student would have written it perfectly in forty seconds.

5.4 When to break your own order

The order isn't sacred. Three legitimate reasons to change it mid-exam:

A question you marked C suddenly opens. It happens. Something in question 2 reminds you how question 5 works. Move it up, but write the new order down, because your memory of the plan is now unreliable.

You're two-thirds through a question and clearly winning. Finishing a question you understand is usually better value than starting a fresh one, even slightly over budget. Slightly. Not twenty minutes over.

You've stalled twice on the same question. That's a signal, not a challenge. Chapter 6 deals with it.

What isn't a legitimate reason: it's interesting, you're nearly there, or you've already invested twenty minutes. That last one is the sunk cost talking, and it's the most expensive sentence in any exam hall.

5.5 The physical setup

Small things, but they cost nothing and they remove decisions:

- Put your watch on the desk, face up, not on your wrist. Turning your wrist is a movement you'll skip.

- Write the finish time of the current question in the margin as you start it, not the duration.
- One question per fresh page where the rules allow. Cramped answers lose marks to illegibility and make it impossible to come back and add.
- Leave a third of a page after each question you're unhappy with, so returning is easy.

That last one matters more than it sounds. A student who abandons a question and leaves no room can't come back, so the abandonment becomes permanent. Leave the space and abandonment becomes temporary, which is the whole point of the next chapter.

5.6 Your checkpoints, computed once

Do this arithmetic for your own paper before the exam and write it on the protocol card. It takes two minutes and it removes all arithmetic from the room.

Paper	CP 1	CP 2	CP 3
100 marks / 180 min	45 min 22 marks	90 min 48 marks	135 min 74 marks
75 marks / 120 min	30 min 17 marks	60 min 36 marks	90 min 55 marks
60 marks / 90 min	25 min 15 marks	45 min 29 marks	65 min 43 marks
65 marks / 90 min (mixed format)	25 min 20 marks (MCQ done)	45 min 34 marks	70 min 55 marks

The targets aren't quite linear, because the first checkpoint has to absorb the triage minutes and the last one has to leave room for the sweep. The rule of thumb, if you'd rather not compute: your marks attempted should be a little *ahead* of the fraction of time elapsed, because the end of a paper is always slower than the middle.

5.7 The two-thirds rule for the last checkpoint

One refinement that matters at the third checkpoint, when there's still time to act but not much.

If at checkpoint 3 you're more than eight marks behind, don't try to catch up on the question you're holding. Go to the paper's remaining *parts* and ask which ones you can complete in under four minutes each. Those are almost always part (a)s and part (b)s on questions you haven't started.

Two-thirds of the remaining time on cheap parts, one-third on the expensive question you were fighting. That split recovers more marks than either extreme, and it stops the common failure where a student spends the last forty minutes on one question and collects six marks.

FROM THE MARKER'S SIDE

A script where the final forty minutes produced one heroic, complete answer to the hardest question on the paper usually scores worse than one where the same forty minutes produced four part (a)s and two part (b)s. The heroic version feels better and marks worse.

DO THIS BEFORE THE NEXT PAPER

On your next timed past paper, write the three checkpoint times and target marks at the top before you start, and keep the running total in the margin. At the end, look at where the gap opened. It'll be one question, and it'll be the one you were sure you were about to crack.

6

The Abandonment Decision

You're eighteen minutes into a question budgeted for twelve. You can feel the answer. One more substitution and it'll come out.

It won't come out. And the decision you make in the next thirty seconds is worth more marks than anything else you'll do today.

6.1 Why this is hard

Three forces push you to stay, and all three are lying.

Sunk cost. You've spent eighteen minutes. Leaving feels like wasting them. But the eighteen minutes are gone whatever you do, and the only live question is what the *next* ten minutes are worth here versus somewhere else.

The nearly-there feeling. It's real and it's uncalibrated. That feeling is equally strong when you're one line from the answer and when you're one line from a dead end, because it's produced by having a plan, not by the plan being right.

Identity. Leaving a question undone feels like admitting you can't do it, and in an exam that feels like a verdict. It isn't. It's a scheduling decision.

Abandoning a question isn't a judgement about you. It's an arithmetic result: the marks per minute here are now lower than the marks per minute somewhere else on this paper.

6.2 The rule

Make it mechanical, because in-the-room judgement is compromised.

Read the second branch again, because it's the one that catches people.




Figure 6.1: The abandonment decision. Every branch has a fixed action and none of them is “keep going and see”.

If you haven’t written anything for three minutes, you’re not working. You’re staring. Three minutes of no ink is the clearest signal available that the next ten minutes will produce nothing either, and it’s a signal you can check without any judgement at all.

The third branch is the useful nuance. “One specific thing to try” means a named next step: try the substitution $u = x^2$, try integrating by parts the other way round, try the case $n = 1$ to see the pattern. It does not mean “keep looking at it”. If you can name the thing, it’s worth three minutes. If you can’t, it isn’t.

6.3 How to leave a question

Leaving badly is nearly as expensive as not leaving. Sixty seconds of house-keeping protects everything you've done and makes the return cheap.

1. **Write the salvage lines.** What you know, what you were trying, where it broke. Chapter 10 is entirely about what to write and what it's worth, and this is where you spend it.
2. **Don't cross anything out.** Work you've abandoned still earns method marks if nothing replaces it. A page of scribbled-over algebra earns nothing.
3. **Leave space.** A third of a page, so returning is possible.
4. **Write one line to your future self.** "Stuck at the substitution, try parts." In forty minutes you'll have forgotten everything and that line saves you three minutes of re-reading your own work.

That fourth step is the one nobody does and everybody needs. Returning to an abandoned question cold costs three to five minutes of reconstruction, which is often more time than you have left. The one-line note cuts it to thirty seconds.

6.4 Coming back

Two rules for the return.

Come back only after you've swept everything else. The whole point of leaving was that the marks were cheaper elsewhere. If you come back before the cheap marks are collected, you've abandoned for nothing.

Come back with a time limit, and write it down. "Q3 until 11:40." Otherwise the question that ate twenty minutes the first time will eat fifteen more.

Something genuinely useful happens in the gap, by the way. A surprising number of abandoned questions solve themselves while you're working on something else, and you'll get the answer arriving unbidden at minute 95 in the middle of a different problem.

That isn't luck. It's what happens when you stop applying pressure to a stuck problem. It's also an argument for abandoning *early*, while there's still time for the gap to do its work.

Abandon early enough that you can come back. A question dropped at minute 20 is a question you might finish. A question dropped at minute 45 is gone.

6.5 The three things that feel like progress and aren't

Worth naming, because each one can absorb ten minutes while feeling productive.

Rewriting the question in your own notation. Sometimes useful. Usually it's stalling with a pen, and you can tell the difference by asking whether the rewrite changed anything you know.

Trying the same method again more carefully. If the method was wrong, care doesn't fix it. If the method was right and the algebra failed, you should be checking the algebra, not restarting.

Simplifying an expression that's already stuck. Algebraic tidying feels like motion and generates no new information. If you don't know what you're heading towards, a tidier version of not knowing is still not knowing.

The test for all three: has anything appeared on the page in the last three minutes that you didn't know three minutes ago? If not, you're not working, and the no-ink rule has already fired.

6.6 The question you never start

One more version of the same decision, and it's cheaper.

If a question is a C after triage, it still has a part (a). Give it four minutes near the end, do the part (a), state whatever definition it wants, and leave. Four minutes for three marks is a better rate than most of the paper.

Students skip C questions entirely because they can't do the whole thing. That's leak 1 from Chapter 3 in its purest form: whole questions worth zero because the hard parts were unreachable, when the easy parts were sitting right there.

6.7 Worked: an abandonment that saved eleven marks

A 100-mark paper. The student is on Q4, a 16-mark question budgeted for 27 minutes.

Clock	What happens
+0 min	Starts Q4. Parts (a) and (b) go cleanly, 9 marks.
+18 min	Part (c) needs a substitution that isn't coming.
+27 min	Budget gone. Feels close. Keeps going.
+31 min	Over budget by 50%. First test fires.
+34 min	Realises nothing has been written for three minutes. Second test fires: stop.
+34 to +35.5	Salvage: names the substitution family being tried, states the integral in its current form, writes " $\text{try } t = \tan(x/2)$ next" and leaves a third of a page.
+35.5 min	Moves to Q5.

Q5 and Q6 then get their full budgets and score 14 and 16.

Later, in the final sweep, the student returns to Q4(c) for four minutes, tries the substitution they noted, and it works. Part (c) scores 4 of 7.

The counterfactual. Had they stayed on Q4(c) for the twenty more minutes it would have taken to grind out, Q6 would have had eleven minutes instead of thirty, and would have scored perhaps 5 instead of 16.

The abandonment was worth eleven marks, and at the moment it happened it felt exactly like giving up.

FROM THE MARKER'S SIDE

Almost every strong script I mark has an abandoned question with salvage lines and a gap left for the return. Almost every disappointing one has a question that clearly consumed forty minutes, and blank pages afterwards.

6.8 The note to your future self

The one-line note in that walkthrough deserves its own emphasis, because it's the step everyone skips and it's what makes the return cheap.

Coming back to your own abandoned working cold costs three to five minutes of reconstruction: reading your own algebra, remembering what you were attempting, working out which line was the last good one. With fifteen minutes left, that's most of your remaining time.

The note cuts it to thirty seconds. Write it in the margin, in the imperative, and make it specific:

Useless	Useful
“stuck”	“stuck at the integral; try $t = \tan(x/2)$ ”
“come back”	“last good line is the one boxed”
“check this”	“check the sign in line 4 first”
“didn’t finish”	“need I_0 ; it’s elementary”

Thirty seconds of writing at the moment you leave, buying three minutes at the moment you return, when minutes are worth the most they’ll ever be worth.

DO THIS BEFORE THE NEXT PAPER

Next timed paper: the moment you notice you’re over budget on a question, start a stopwatch on the desk and see how long you actually stay. Most students are shocked. Then run the decision tree and leave properly, with the salvage lines, the space, and the note to yourself. Compare the score with your usual.

7

The Last Ten Minutes

Ten minutes left. Most students use them to keep grinding on whatever they're holding.

That's almost always the worst available use of the time, because the question you're holding at minute 170 is, by definition, the hardest thing left. Meanwhile there are two or three marks lying loose on pages you've already written.

The last ten minutes have their own protocol, and it's worth between three and eight marks on a typical script.

7.1 The sweep, in order

The order is by marks per minute, and it's close to the reverse of what students actually do.

7.2 Step 1: the blanks

Flip through the whole paper looking only for parts with nothing under them.

For each one, ask a single question: is there a definition, a formula, or a one-line setup I can write in forty seconds? If yes, write it. If no, move on.

You are not trying to answer these. You're collecting the B marks and the first M mark, and on a paper with six questions there are usually three or four of these sitting unclaimed.

FROM THE MARKER'S SIDE

The last ten minutes are visible in the handwriting. I can see the sweep happen: the writing gets faster, and suddenly there are three




Figure 7.1: The final sweep. Cheapest marks first, and the question you're stuck on comes last, if at all.

or four short additions scattered across earlier questions. Those additions are almost always worth marks, and the scripts that have them tend to be four or five marks above the ones that don't, from students of the same ability.

7.3 Step 2: state the answers

This one is embarrassing and extremely common.

Half a page of correct working, and no sentence at the end saying what the answer is. On many schemes that's an accuracy mark lost per part, and it's the cheapest mark on the paper to recover.

So: every part you attempted, find the answer, and write one clear line.

Therefore the stationary points are at $x = 1$ and $x = -3$.

Box it if you like, though the box is decoration. What matters is that the answer is stated, unambiguous, and not buried in the third line of a paragraph of algebra.

7.4 Step 3: units and figures

Sweep every numerical answer on the paper. Two checks, three seconds each:

Is there a unit? Physics and engineering scripts bleed marks here. A number without a unit is often worth zero even when it's right.

Is the precision sane? If the data was given to three significant figures, an answer to nine is wrong-looking and sometimes explicitly penalised. And if the question said "to 2 d.p.", it said it for a reason.

Chapter 16 covers this properly. In the last ten minutes, you're just sweeping.

7.5 Step 4: salvage the abandoned

For each question you left: two or three lines saying what you know, what method you'd use, and where it broke.

Chapter 10 has the full treatment. In the sweep, the compressed version is enough:

By the divergence theorem this reduces to a volume integral of $\nabla \cdot \mathbf{F}$. Computed $\nabla \cdot \mathbf{F} = 3$, so the integral is $3 \times$ volume of the region. Ran out of time evaluating the volume.

That's forty seconds and it's frequently worth two or three method marks, because the method is exactly what the scheme rewards.

7.6 Step 5: back to the hard one, maybe

If minutes remain, now you can return to the question you were fighting.

But notice what's happened. You've collected everything cheap first, so whatever you spend here is genuinely spare, and if it produces nothing you've lost nothing.

That's the entire logic of the sweep. It converts your last ten minutes from a gamble into a guaranteed return with a gamble on top.

7.7 The five minutes you should have reserved

Back in Chapter 4, step 1 said to subtract five minutes from the working time for this sweep. That subtraction is the thing that makes the sweep possible, and it's the first thing students cut when they're behind.

Don't cut it. When you're behind is exactly when the sweep pays most, because being behind means more blanks, more unstated answers, and more abandoned questions.

The final sweep is worth more when the exam has gone badly. Protect it hardest on the papers where you most want to keep writing.

7.8 Scaling the sweep to the paper

Five minutes is right for a three-hour paper. Shorter papers need a shorter sweep, but never a missing one.

Paper	Sweep	What fits
180 min	10 min	all five steps, plus a return
120 min	7 min	all five steps
90 min	5 min	blanks, answers, units
60 min	3 min	blanks and units only

On the shortest papers, cut steps 4 and 5, not steps 1 to 3. Blank parts and missing units are the cheapest marks on any paper, and they stay cheapest regardless of how long the paper is.

7.9 The two-alarm system

Most students set no alarm and discover the end of the paper by hearing somebody stand up. Set two.

Alarm one, at $T - 20$. Not to start the sweep, but to make the decision: is there a question I should abandon now so the last twenty minutes go somewhere useful?

Alarm two, at $T - 10$. Sweep, immediately, mid-line if necessary. The whole point is that this alarm overrides your judgement, because your judgement at $T - 10$ is exhausted and will always vote to keep going.

If your watch has no alarm, write both times at the top of the script next to your checkpoints, and check them at every natural break.

The sweep alarm exists to overrule you. If you find yourself negotiating with it, that's the alarm working exactly as designed.

7.10 And when they call time

Stop writing. Immediately.

Not because of the rule, though the rule matters, but because the last sentence written in a panic after “pens down” is worth nothing and occasionally causes real trouble. Put the pen down, and let the sweep you already did be the thing that saved the paper.

7.11 Worked: ten minutes, seven marks

A real sweep on a paper that had gone moderately badly. The student had attempted five of six questions and was five marks behind where they wanted to be.

Time	What was found	Marks
0:00–2:30	Q3(a) blank: stated the definition of a conservative field. Q6(a) blank: wrote down the standard form of the ODE.	+3
2:30–4:30	Q1(d) had a full page of working and no final line. Q4(b) the same. Wrote both answers out.	+2
4:30–6:30	Six numerical answers checked: two had no units, one was quoted to eight figures.	+2
6:30–8:30	Q5, abandoned at minute 70, got three salvage lines naming the method and the setup.	+0 *
8:30–10:00	Back to Q5(c). Didn't finish.	0
Total recovered		+7

*The salvage lines scored 2 when the paper was marked; the student couldn't know that at the time.

Seven marks in ten minutes is 0.7 marks per minute. For comparison, the question they wanted to keep working on was paying about 0.15.

Notice also that nothing in the sweep required any thinking. It's clerical work, which is exactly what you want from your last ten minutes, because your capacity for anything else is gone by then.

7.12 The sweep on a paper that went well

Worth saying, because the instinct differs.

If you've finished everything with fifteen minutes left, the sweep is still the first thing you do, and it's still worth two or three marks: unstated answers, missing units, a part (c) you skipped and forgot to come back to.

Only after the five steps should you go back and check a long calculation, and even then, check the ones that feed later parts rather than the ones that stand alone.

Finish early and you sweep first, check second. Almost everyone does it the other way round and spends twelve minutes re-verifying an answer that was already right while three unstated answers sit on earlier pages.

DO THIS BEFORE THE NEXT PAPER

On your next timed paper, set an alarm for ten minutes before the end and run the five sweep steps in order, even if you're mid-line on something. Count the marks you pick up in those ten minutes. For most students it's between three and eight, and it takes no additional knowledge whatsoever.

Part III

Writing Under Pressure

8

Writing for Method Marks

Chapter 2 said that method marks are 50 to 70 per cent of most STEM papers, and that you only get them if the method is visible.

This chapter is what “visible” actually means, line by line.

8.1 The rule that decides most of it

Write the line where the method changes. Skip the lines where it only computes.

That’s the whole thing, and it resolves the false trade-off students worry about. You don’t have to choose between showing your working and having time. You show the *structural* lines and skip the arithmetic in between.

Structural lines are where a decision happened:

- the formula you’re invoking, written out before you use it;
- the substitution, stated as $u = \dots$, with $du = \dots$;
- the setup of the integral or sum, with its limits;
- the equation you’re about to solve, before solving it;
- the theorem you’re citing, by name or statement;
- the case split, announced.

Non-structural lines are the algebra that follows mechanically. Expanding brackets, collecting terms, dividing through. Those are where you can move fast, and no scheme in the world has a mark for line three of an expansion.

8.2 Before and after

Here’s a five-mark part, written two ways by the same student.

BEFORE

$$\int_0^2 xe^{x^2} dx = \frac{1}{2}(e^4 - 1) = 26.8$$

If the number is right, this scores full marks. If the number is wrong, and it is wrong about a third of the time under exam conditions, it scores zero. There’s nothing on the page for a marker to reward.

AFTER

Let $u = x^2$, so $du = 2x dx$ and $x dx = \frac{1}{2} du$. [M1]

When $x = 0$, $u = 0$; when $x = 2$, $u = 4$. [M1]

$$\int_0^2 xe^{x^2} dx = \frac{1}{2} \int_0^4 e^u du = \frac{1}{2} [e^u]_0^4$$

$$[M1] = \frac{1}{2}(e^4 - 1) = 26.8 \text{ (3 s.f.)} \quad [A1, A1]$$

Twenty extra seconds of writing. Three of the five marks are now guaranteed regardless of what happens to the arithmetic, and the two accuracy marks are still available.

Note the changed limits on their own line. That’s a classic M mark, and it’s the one students most often do in their heads.

8.3 What “show your working” does not mean

It doesn’t mean write more words. Long prose explanations of what you’re doing cost time and earn nothing.

The right-hand column is where students who “show all their working” actually lose time. Copying the question out is the worst offender: it feels like starting, it takes forty seconds, and it earns nothing at all.

8.4 Writing so a tired person can follow

Your marker has ninety scripts and twelve minutes for yours. Three habits make their job possible, and all three convert into marks:

One idea per line. Not three equalities chained across a line with the reasoning implied.

Earns marks	Earns nothing
the substitution, stated	“Now I will integrate”
the formula, before use	every line of expansion
changed limits	re-copying the question
the theorem cited	apologies and notes
the case split	decorative boxes

Figure 8.1: What earns method marks and what only costs time.

Left-hand margin discipline. Start each new line at the margin. Working that wanders diagonally across the page is genuinely hard to follow, and a marker who can’t follow it can’t award it.

Label your parts. Write (a), (b), (c) clearly. If you do parts out of order, say so: “(c) below (b) on next page”. Markers work question by question and a part they can’t find is a part that scores zero.

FROM THE MARKER’S SIDE

I’ve hunted for a part (c) that turned out to be on the back of the sheet for question 6, upside down, unlabelled. I found it because I was looking. I don’t always look.

8.5 The specific case of “show that”

When a question says *show that* and gives you the answer, the marks are entirely in the method. You cannot get them by writing the answer, because the answer is printed in the question.

So the standard is higher: every step needs to be visible, because the step *is* the mark. And there’s a trap worth naming. If you get stuck halfway, don’t work backwards from the printed answer to fake a connection. Markers spot it instantly and it usually scores less than an honest partial attempt.

Chapter 18 covers the structure of these questions, including the fact that a “show that” you can’t complete still lets you use the given result in the next part.

8.6 The physical script

Twelve minutes, ninety scripts, and yours is somewhere in the middle of the pile. Legibility isn't politeness, it's a marking condition.

Five things that cost nothing and change what can be awarded:

Write bigger than feels necessary. Cramped work is misread, and a misread digit is a lost accuracy mark. Paper is provided, and using more of it is free.

One question per fresh page, where the rules allow. It makes adding to a question later possible, and markers work question by question.

Label every part, clearly, at the left margin. If you answer out of order, say where things are: "(c) continued on page 11". A part the marker can't find scores zero, and nobody is obliged to hunt.

Keep the working vertical. Each new line starts at the margin. Work that drifts diagonally across a page, with three parallel columns going at once, is genuinely hard to follow and the marks go with the difficulty.

Box or underline the final answer of each part. Not decoration: it tells the marker where to stop looking, which matters when your working runs to half a page.

FROM THE MARKER'S SIDE

I've never deducted a mark for bad handwriting, and I've certainly failed to award marks I couldn't read. Those are different things, and only one of them is in the rubric.

8.7 The forty-second insurance policy

Here's how to think about all of this when the clock is loud.

For any part worth four or more marks, spend forty seconds writing the structural lines before you compute anything. Formula, substitution, setup, limits.

That forty seconds insures roughly half the marks on that part against every arithmetic slip you're about to make. Nothing else on this list of techniques has that return.

8.8 Three more, from three subjects

The rule is the same everywhere, but what counts as a structural line changes with the subject. Here it is in three.

Mechanics. The structural lines are the free-body diagram, the axis convention, and the equation you're about to apply.

BEFORE

$$a = 2.4, \text{ so } T = 5 \times 2.4 + 49 = 61 \text{ N}$$

AFTER

Taking upwards as positive for the 5 kg mass:

$$T - mg = ma \quad \Rightarrow \quad T = m(g + a)$$

$$= 5(9.8 + 2.4) = 61 \text{ N.}$$

The sign convention line is a mark on most schemes, and it's the line that prevents half the errors in the topic.

Circuits. The structural lines are the law being applied and the loop or node it's applied to.

AFTER

Applying KVL around the left-hand loop, clockwise:

$$12 - 4I_1 - 6(I_1 - I_2) = 0$$

Naming the loop and the direction costs four words and makes an otherwise unreadable set of simultaneous equations markable.

Pure mathematics. The structural line is the logical move.

AFTER

Suppose, for contradiction, that $\sqrt{2} = p/q$ in lowest terms.

That one sentence claims the mark for identifying the proof structure, and it does so before any mathematics has happened.

8.9 What “visible” costs you, measured

Students resist this because writing feels slow. Here is the actual arithmetic on a 100-mark, three-hour paper.

	Time	Marks
Structural lines, 20 parts at 40 s	13 min	—
Method marks thereby insured	—	15–25
Time available after triage and sweep	171 min	
Cost as a share of working time	8%	

Eight per cent of your time, insuring somewhere between fifteen and twenty-five marks against arithmetic you’re going to get wrong somewhere.

There is nothing else on the paper with that return, and the students who skip it are almost always the ones who then run out of time anyway, because the time went on redoing calculations that failed.

DO THIS BEFORE THE NEXT PAPER

Take your last past-paper attempt and pick any five-mark part you got wrong. Rewrite it, showing only the structural lines and skipping the arithmetic. Then estimate, against a real mark scheme, how many marks the rewrite would have scored. The difference is what this habit is worth to you per part, and you can multiply it by the number of parts on a paper.

9

State the Plan

One line at the top of an answer, before any working, saying what you're about to do.

It costs fifteen seconds. It's the highest-return sentence you can write in an exam, and hardly anyone writes it.

9.1 What it looks like

I'll use integration by parts with $u = \ln x$ and $dv = x^2 dx$.

Substituting $y = vx$ makes this separable.

By induction on n . Base case $n = 1$, then assume for n and prove for $n + 1$.

Take moments about the left support to eliminate the unknown reaction.

Each one is a single sentence, written before the mathematics, naming the method.

9.2 Why it earns marks four different ways

It's often the M1 itself. Many schemes award the first method mark for "correct approach identified". Your plan line is that identification, written in the clearest possible place.

It survives your arithmetic. If everything after it goes wrong, the plan line is still on the page and still correct.

It stops you starting the wrong method. You cannot write "I'll use integration by parts" without noticing, half a second later, whether parts is actually the right move. The sentence forces a decision that otherwise happens implicitly three lines in, by which point you've committed.

It makes returning cheap. Come back to the question at minute 140 and the plan line tells you what you were doing. Without it you re-read your own algebra to reconstruct your own intention.

Write the method before you use it, in a sentence, in English. It is frequently worth a mark on its own, and it's always worth the thirty seconds it saves you later.

9.3 The thirty seconds before the plan

To write the plan you have to have one, which means the first thirty seconds of a question aren't writing at all.




Figure 9.1: The thirty-second opening. Three questions, then the plan line, then the mathematics.

Question 2 deserves emphasis. *Including anything part (a) just handed you.* Multi-part questions are ladders, and the single most common way to waste ten minutes is to prove something in part (b) that part (a) already gave you.

Before starting any part after the first, look back at what the previous parts produced. If the question says *hence*, that's not a suggestion, it's an instruction, and a method that ignores it may score nothing even if it reaches the right answer.

9.4 When you can't write a plan

Sometimes thirty seconds pass and no method comes. That's information, and it's cheap information.

Three moves, in order:

Name the type of object. What kind of thing is being asked for? An integral over a region, an eigenvalue, a maximum, a proof of uniqueness. The type usually implies a short list of standard methods, and one of them is probably it.

Look at what's given that you haven't used. If the question mentions that a function is continuous, or that a matrix is symmetric, or that n is prime, that fact is in the question because the answer needs it. Unused hypotheses are the biggest hint in any exam question.

Write the plan you'd use if the numbers were easier. Often it's the same plan.

If none of those produce anything in thirty seconds, this is a C question. Do its part (a), write a salvage line, and move on. You've spent ninety seconds finding that out, which is exactly what triage is for.

9.5 Plans for proof questions

In pure mathematics the plan line does even more work, because the structure of a proof is dictated by the statement's shape.

The question says	Your plan line
prove for all n	"By induction on n ."
prove there exists	"Take $x = \dots$ and verify."
prove exactly one	"Existence, then uniqueness."
prove $A = B$ for sets	"Show $A \subseteq B$, then $B \subseteq A$."
prove no such thing exists	"Suppose one existed."
prove $P \iff Q$	"Both directions, $\Rightarrow$ first."

Every one of those is worth writing down even if you can't complete the proof, because it demonstrates that you know the required structure, and structure is what proof questions are marked on.

FROM THE MARKER'S SIDE

On a proof question, a student who writes "by induction: base case, then assume for n " and then fails at the inductive step gets meaning-

fully more than a student who writes half a page of correct-looking algebra with no visible structure. The first one has shown they know what a proof of this shape requires.

9.6 When the plan is wrong

You write a plan line, work for four minutes, and it isn't going anywhere. This is normal, and there's a right way to handle it.

Don't erase the plan line. It's still a method mark on most schemes, even if the method turned out to be the wrong one for this question. A wrong-but-reasonable approach, clearly stated, is worth more than a blank.

Write a second plan line underneath. "Parts isn't working; the integrand is a rational function, so partial fractions instead." Now the marker can see two attempts, and you've claimed whatever the second approach is worth as well.

Cap it at two. If two named methods have both failed in six minutes, this is a question to leave, and Chapter 6 takes over.

That cap matters. Cycling through five approaches, each abandoned after ninety seconds, is one of the recognisable shapes of a lost question, and it produces a page of fragments that scores about as well as a blank.

Two named methods, then leave. A third attempt almost never works, and by then you're eight minutes into a question you should have abandoned at four.

9.7 The habit, compressed

Every part, every question, for the rest of your degree:

1. Underline the verb and the constraints. (5 seconds)
2. Ask what you've been given, including from earlier parts. (10 seconds)
3. Write one sentence naming the method. (15 seconds)
4. Now work.

Thirty seconds a part. On a paper with twenty parts that's ten minutes, and it buys a method mark on most of them plus every minute you'd have lost to starting the wrong way.

9.8 Plan lines for the twenty questions you'll actually meet

Most exam questions are one of a small number of shapes. Here are the plan lines, ready to use.

The question	The plan line
Integral with a product	"By parts, $u = \dots$, $dv = \dots$ "
Integral with $\sqrt{a^2 - x^2}$	"Trig substitution $x = a \sin \theta$ "
Rational function integral	"Partial fractions first"
First-order ODE, messy	"Check for separable; if not, integrating factor"
Second-order linear ODE	"Complementary function, then particular integral"
Series convergence	"Comparison with $\sum 1/n^p$ " or "ratio test"
Find a maximum	"Differentiate, set to zero, classify with the second derivative"
Matrix, repeated use	"Eigenvalues from $\det(A - \lambda I) = 0$ "
Volume or flux	"Divergence theorem to convert to a volume integral"
Prove for all n	"Induction on n "
Particle on a slope	"Resolve along and perpendicular to the slope"
Two connected masses	"Newton's second law for each mass separately"
Collision	"Conserve momentum; check whether energy is conserved"
Circuit with two loops	"KVL around each loop, then solve simultaneously"
Beam with a load	"Take moments about the support to remove one reaction"
Thermodynamic cycle	"First law per stage, then sum"
Estimate something	"State assumptions with numbers, then multiply"
Data and a straight line	"Linearise to $Y = mX + c$, name Y , X , m , c "
Uncertainty asked for	"Percentage uncertainties, added"
Sketch a function	"Intercepts, asymptotes, stationary points, then draw"

Copy the ten that apply to your course onto the back of your protocol card. In the exam, the sentence you need will be one of them about eighty per cent of the time, and having it ready turns thirty seconds of searching into five seconds of recognition.

9.9 The tell, and why the plan is really about recognition

Notice what the right-hand column is doing. Each plan line is triggered by a *tell* in the question: a product inside an integral, a difference of squares under a root, two masses joined by a string.

That's the skill the exam is testing, and it's the one blocked practice never builds (Chapter 14). When you do twenty integration by parts questions in a row, the tell is redundant because somebody already told you the method.

So when you build method cards, put the tell on the card. Front: the tell. Back: the plan line and the steps. That's the direction the exam asks the question in.

FROM THE MARKER'S SIDE

The scripts that stall aren't usually the ones that can't execute a method. They're the ones that spent six minutes deciding which method, and then had nine minutes left for a twelve-minute calculation.

DO THIS BEFORE THE NEXT PAPER

On your next past paper, write a plan line on every single part before you start it, including the ones that look trivial. It'll feel slow for the first three questions and then stop feeling like anything. Afterwards, check how many of your plan lines correspond to an M mark in the scheme.

10

Salvage

You can't finish it. The clock says move on. There's half a page of dead algebra and a question you're about to leave.

What you write in the next ninety seconds is worth between zero and five marks, and the difference is entirely mechanical.

10.1 Why blank pages are the most expensive thing on a script

A blank part scores zero. Always. There is no scheme in which nothing earns something.

But almost nothing earns something surprisingly often, because method marks don't require completion. They require the method to be identifiable on the page.

There is no such thing as a question you can't write anything on. If you can name the topic, you can write something worth a mark. The only question is how much.

10.2 The salvage stack

Write as many of these as you can, in this order. Stop when you run out of things you know.

Lines 1 to 3 are the ones that pay. They're method, they're marked as method, and none of them requires you to be able to finish.

Line 5 pays nothing directly. Write it anyway: it costs eight seconds and it's what lets you resume in thirty seconds instead of three minutes if you come back.




Figure 10.1: The salvage stack. Each line is worth something on its own, and you write as far down as your knowledge goes.

10.3 Worked salvage

A question you can't do, salvaged three different ways depending on how much you know.

You know almost nothing:

This is a surface integral over a closed surface, so the divergence theorem applies.

One line. Frequently one mark.

You know the method but not the computation:

By the divergence theorem, $\oint_S \mathbf{F} \cdot d\mathbf{S} = \int_V (\nabla \cdot \mathbf{F}) dV$.
Here $\nabla \cdot \mathbf{F} = 2x + 2y + 2z$, so I need $\int_V (2x + 2y + 2z) dV$ over the given region.

Two or three marks, on a question you did not answer.

You got most of the way:

...so the integral becomes $\int_0^1 \int_0^{1-x} (2x + 2y + 2z) dy dx$ after substituting the plane. Evaluated the inner integral to $2xy + y^2 + 2zy$ at $y = 1 - x$. Ran out of time on the outer integral.

Four marks or so, and the only thing missing is arithmetic.

10.4 Salvage on proof questions

Pure questions salvage even better, because their marks are structural.

If you can write the shape, you're claiming a real share of the marks:

By induction on n .

Base case $n = 1$: LHS = 1, RHS = 1, so it holds.

Assume true for $n = k$, so $\sum_{i=1}^k i = \frac{1}{2}k(k+1)$.

Then $\sum_{i=1}^{k+1} i = \frac{1}{2}k(k+1) + (k+1)$, and I need to show this equals $\frac{1}{2}(k+1)(k+2)$.

Even if the final algebra never appears, that's the base case, the hypothesis, and the correct target, and most schemes pay for all three.

10.5 Salvage as a first move, not a last one

There's a version of this that most students never consider, and it's worth more than the rest of the chapter.

Write the salvage lines *before* you attempt the question.

On any part worth four or more marks, spend forty seconds writing the method, the relevant formula, and the setup, exactly as you would if you were abandoning it. Then start solving.

Two things follow. If the solve works, you've written the structural lines you were supposed to write anyway (Chapter 8), and they cost you nothing. If the solve fails, the salvage is already on the page and you can leave immediately, with no ninety-second exit cost at the moment when you're most stressed.

Write the salvage first and it stops being salvage. It becomes the setup, and it earns the same marks either way.

This also removes the psychological problem with salvage, which is real: writing it at the end feels like formally conceding, and students avoid it for that reason. Written at the start, it's just how you begin a question.

10.6 What not to write

Three things that cost you time or goodwill and earn nothing:

Apologies. “Sorry, ran out of time, I do know how to do this.” The marker cannot award anything for it, and it reads badly.

Vague gestures. “Use the standard method and simplify.” If you can't name the method, this is worth nothing. If you can name it, name it.

Fake completions. Working backwards from a printed answer to manufacture a connection, or writing a final answer with no support because you know what it should be. Markers see it, and on “show that” questions it typically scores less than an honest partial attempt.

FROM THE MARKER'S SIDE

Honest partial work is easy to reward. Every marker would rather find three correct lines and a note saying where it stopped than a page engineered to look complete. And the engineered version usually breaks somewhere visible anyway.

10.7 Budget it

Salvage is a scheduled activity, not a mood.

When you abandon a question (Chapter 6), spend sixty to ninety seconds on the stack before you leave. And in the final sweep (Chapter 7), step 4 exists precisely so that anything you abandoned in a hurry gets its salvage lines before the paper ends.

Two minutes of salvage across a paper is routinely worth four or five marks. There is no other two minutes on the paper with that return.

10.8 Salvage in physics and engineering

The stack is the same, but what fills it is different, and in applied subjects there's more of it available than students expect.

Draw the diagram. A free-body diagram with the forces named and the axes marked is frequently a mark on its own, and it's available even when the algebra defeats you entirely.

State the governing equation. Newton's second law, KVL, the first law of thermodynamics, the continuity equation. Write it in full before applying it.

List what you know, with units. Converting the given data to base SI units and listing it is a real step, and on some schemes it carries a mark.

Name the assumptions. Frictionless, incompressible, steady state, small angles. On modelling questions the assumptions are explicitly marked, and students skip them because they feel like preamble.

Estimate the answer. If you can't compute it, say what order of magnitude it should be and why. It shows physical understanding, and on estimation questions it's most of the marks anyway.

Here's the whole stack on a question the student couldn't finish:

AFTER

Assumptions: pulley frictionless and massless, string inextensible.
Taking downwards positive for the 3 kg mass:

$$3g - T = 3a, \quad T - 2g = 2a.$$

Adding: $g(3 - 2) = 5a$, so $a = g/5 = 1.96 \text{ m s}^{-2}$.

Then $T = 2(g + a) = 2(9.8 + 1.96) = 23.5 \text{ N}$.

Could not complete part (c), the time to fall 2 m, but it follows from $s = ut + \frac{1}{2}at^2$ with $u = 0$.

Even the unfinished part carries a method mark, because the equation of motion is named and the initial condition is stated.

10.9 The two-minute rule for salvage across a paper

Salvage doesn't have a budget line in most students' heads, so it doesn't happen. Give it one.

When	Time	Typical value
On abandoning a question	60–90 s	2–3 marks
On a part you can’t start	30 s	1 mark
In the final sweep, per abandoned question	40 s	1–2 marks
Across a whole paper	≈ 4 min	4–6 marks

Four minutes, four to six marks, and it’s the only technique in this book that works better the worse the paper is going.

Salvage is worth most on your worst paper. When you’re behind and everything is unfinished, the four minutes of naming methods are the highest-return minutes available, and they’re the first thing panic deletes.

DO THIS BEFORE THE NEXT PAPER

Find a question from a past paper you genuinely can’t do. Set a timer for ninety seconds and write as far down the salvage stack as your knowledge reaches. Then mark it against the scheme. Most students score two or three on a question they’d otherwise have left blank, and the shock of that is what makes the habit stick.

11

Checking Work You Have No Time to Check

“Check your answers” is the advice everyone gives and nobody follows, for a good reason: there’s no time, and re-doing a calculation you’ve already done tends to reproduce the same error anyway.

So forget re-doing. What follows is a set of checks that cost between two and fifteen seconds each and catch a large share of real exam errors.

11.1 The five-second checks

Run these on every numerical answer, without exception.




Figure 11.1: Five checks, none longer than fifteen seconds. Run them per answer, not per paper.

Sign

A probability came out negative. An area came out negative. A distance is negative. A definite integral of a positive function over an increasing interval is negative.

Each one is instantly wrong, and each is a two-second look. Sign errors are the most common arithmetic slip in every subject, and they're the easiest to catch because they don't need arithmetic to detect.

Magnitude

Estimate the answer roughly, in your head, before you look at what you wrote.

The car is going 4,000 metres per second. The resistor dissipates 30 megawatts. The angle is 400 degrees. The probability is 4.

Order-of-magnitude sanity is nearly free, and the errors it catches, a factor of ten from a misplaced decimal, a factor of two from a missing half, a factor of π from a radians-degrees mix-up, are exactly the errors that exam conditions produce.

Units

Covered properly in Chapter 16, but the short version belongs here because it's a check, not a decoration.

Carry units through your working and see whether they come out right. An answer in metres per second when you were asked for an energy tells you the formula is wrong, before you've checked a single number. This is the cheapest error-detection method in physics and engineering, and it's usually the fastest way to find that you've used the wrong equation entirely.

Special case

For an algebraic answer, substitute something trivial and see whether it behaves.

Set $n = 1$. Set $x = 0$. Let the mass go to zero, or the resistance go to infinity, or the angle go to zero. Does your general formula collapse to the obvious answer?

Every general result has a special case whose answer you already know. Ten seconds spent putting $n = 1$ into your formula catches more algebra

errors than five minutes of re-deriving it.

If the question had an earlier part that computed a specific case, you've been handed a free check: does your general formula reproduce part (a)'s number?

Did I answer the question?

Read the underlined verb again. You found x ; it asked for x^2 . You found the speed; it asked for the velocity. You solved for t ; it asked when the particle returns, which is $2t$.

This one costs three seconds and catches the single most annoying way to lose four marks.

11.2 Where checks pay most

Not everywhere equally. Prioritise:

- **The first part of a multi-part question.** An error here poisons every later part. Sixty seconds checking part (a) can be worth more than the rest of your checking combined.
- **Anything that feeds a later part.** Same logic.
- **Answers that came out ugly.** Exam answers are usually designed to be clean. If you get $x = 17/43$, you're probably wrong, and if you're right it cost you ten seconds to wonder.
- **Anything you rushed.** You know which ones.

And where checks don't pay: re-reading a long correct-looking derivation from the top. You'll read what you meant to write, not what you wrote. That's not a discipline problem, it's how reading works.

11.3 The ugly-answer heuristic

Worth its own note, because students distrust it and it's remarkably reliable.

Exam questions are constructed backwards from nice answers. Someone chose the numbers so the algebra resolves. So when your answer is $\frac{3}{7}\sqrt{2} - \frac{19}{5}$, the prior probability that you slipped somewhere is high.

It isn't proof. Some answers genuinely are ugly, particularly in applied questions with real data. But ugliness is a flag, and flags are what you want when you have ninety seconds.

FROM THE MARKER'S SIDE

When I see an ugly answer with a note beside it saying “this looks wrong, expected something simpler”, the student has told me they understand the subject even though the arithmetic failed. It doesn't earn a mark by itself. It does mean I read the working more carefully to find where the follow-through starts.

11.4 Checks by subject

The five general checks apply everywhere. Each subject adds two or three of its own, and they're worth having as reflexes.

Pure maths	Does the formula work for $n = 1$? Is the answer defined at the endpoints? Did I lose a solution when I divided by something that could be zero?
Statistics	Is the probability in $[0, 1]$? Do the probabilities sum to 1? Is the variance positive?
Mechanics	Does the direction make sense? Is the tension positive? Would the answer change sign if I flipped my axis convention, and did I state that convention?
Circuits	Do the currents at a node sum to zero? Is the power dissipated positive? Is a resistance in the right decade?
Thermo	Is efficiency below the Carnot limit? Is entropy increasing for the isolated system?
Materials	Is the strain dimensionless? Is the stress below yield, and if it isn't, did the question intend that?

Pick the three from your own subject that match your error log, and put them on the protocol card. Not all of them. Three, chosen because they catch the mistakes you actually make.

11.5 The check budget

Ten seconds per part. That’s it.

On a paper with twenty parts, that’s a little over three minutes total, which you can afford, and it’ll catch most of what’s catchable. Anything more than that competes with the final sweep, which pays better.

Do the checks as you finish each part, not at the end. At the end you’ll have neither time nor the memory of what you were doing.

11.6 Six checks, six real errors caught

Each of these took under fifteen seconds and saved between two and five marks.


The answer	The check	The error
Probability = 1.4	range	added instead of multiplying
$\int_0^1 x^2 dx = -\frac{1}{3}$	sign	limits swapped
Efficiency = 118%	possible?	used input where output belonged
$g = 0.025 \text{ m s}^{-2}$	magnitude	forgot to divide t_{20} by 20
Resistance in volts	units	wrong formula entirely
$\sum_{k=1}^n k = n^2$	put $n = 1$	gives 1, but $n = 2$ gives 4, not 3

The last one is the special-case check doing its job. The formula looks plausible, the student half-remembered it, and one substitution killed it in four seconds.

11.7 Checking a multi-part question, in the right order

Not all checks are worth the same, because errors propagate downhill.

So check part (a) hardest, and check anything whose answer gets used later. On a question where part (a) produces a constant that the rest of the question multiplies by, forty seconds on part (a) is worth more than the other three checks combined.



An error in (a) is inherited by (b), (c) and (d). An error in (d) is inherited by nobody.

Figure 11.2: Where to spend a ninety-second check budget on a four-part question. An error in (a) costs four times what the same error costs in (d).

Check upstream, not downstream. The value of a check is proportional to how many later marks depend on the number you're checking.

DO THIS BEFORE THE NEXT PAPER

Go through your last three past-paper attempts and classify every arithmetic error you made against the five checks. You'll find one check would have caught most of them, and it's usually magnitude or sign. That's the one to run automatically for the rest of your degree.

Part IV

The Term Before

12

The Past-Paper System

Everyone says do past papers. Almost nobody says how, so students do them the way that feels productive and isn't: open one, work through it slowly with notes to hand, check the answers, feel reasonably good, do another.

That's not practice. It's reading with extra steps.

Here's the system that works, and the difference is mostly about conditions and about what happens *after* each paper.

12.1 Three modes, and they're not interchangeable




Figure 12.1: Three ways to use a past paper. Most students only ever use the first, which is the one that teaches least.

The third mode is the one students avoid, because it's uncomfortable and takes a whole afternoon. It's also the only one that practises triage,

clock discipline, abandonment, and the final sweep, which is to say everything in Parts II and III of this book.

Untimed past papers test your knowledge. Timed past papers test your exam, and the exam is what you're sitting. Do at least four full timed rehearsals before any paper that matters.

12.2 The schedule

Work backwards from the exam. For a twelve-week course with the exam at the end:

When	Mode	What for
Weeks 1–8	Open, by topic	meet the question types as you learn them
Weeks 9–10	Closed, untimed	find out what you actually know
Week 11	Two full timed	rehearse the protocols
Week 12	Two full timed	rehearse again, on the newest papers
Final 3 days	Error log only	no new papers

Two things in that table are deliberate and get ignored.

Save the two most recent papers for week 12. They're the closest match to what you'll sit, and they're wasted in week 3 when you don't know the material yet.

No new papers in the last three days. A paper attempted three days before the exam has no time to feed back into anything. It only generates anxiety. Those days belong to the error log.

12.3 The review is the point

Doing the paper is maybe a third of the value. The review is the rest, and it takes about as long as the paper did.

For every mark you didn't get:

1. **Bucket it**, using the seven buckets from Chapter 3: the six leaks, or "didn't know it".
2. **If it was a leak**, note which one and what you'd do differently. That's a protocol fix.
3. **If you didn't know it**, write it in the error log (Chapter 13) as a topic, not as a question.

4. **Mark against the real scheme**, not against your own opinion of your answer. Students are systematically generous to their own working and systematically surprised by method marks they didn't earn.

That fourth point is worth insisting on. Mark yourself with the published scheme, line by line, and be brutal. If the scheme wanted the substitution written down and you did it in your head, you didn't get that mark, no matter how well you knew the method.

FROM THE MARKER'S SIDE

Students routinely tell me they got 70 on a past paper and then score 55 on the real one. Almost always the difference is self-marking. They gave themselves method marks for methods that weren't on the page, which is exactly the mark I can't give them.

12.4 Redo, don't just review

The step almost everyone skips.

Three days after you review a paper, redo the questions you got wrong. From scratch, closed book, timed to the original budget.

Not the whole paper. Just the failures.

This is what converts "I understand where I went wrong" into "I can now do this", and those are very different states. Understanding an error takes five minutes and lasts a week. Re-solving the problem takes fifteen and lasts.

12.5 Question banks versus whole papers

Both, at different times.

Topic sets early. When you're learning integration by parts, do twenty parts questions in a row from wherever you can find them. Massed practice on one method is how you get fluent with it.

Whole papers late. Once you're fluent, massed practice stops teaching you anything, because in a real exam nobody tells you which method to use. Mixed papers force you to identify the method, which is the skill that actually gets tested.

Practising twenty questions of the same type teaches you the method. Practising twenty questions of mixed types teaches you to choose the method, and choosing is what the exam is testing.

That transition, from blocked to mixed, is the single biggest upgrade available to most students’ revision, and Chapter 14 goes into why.

12.6 Simulating the room

A timed paper at your own desk is not the same test, and the differences are worth engineering away before it matters.

At your desk	In the hall
your chair, your light	a hall chair and strip lighting
you can stand up	you can’t
water, snacks, phone	water only
your own handwriting	the same, but with adrenaline
speed	
you know the paper is	you don’t
practice	
nine in the evening	nine in the morning

You can’t remove all of it, and you don’t need to. Fix the three that change performance most:

Time of day. Do at least one full rehearsal at the exam’s actual start time. A student who has only ever worked at 10 pm will find a 9 am paper genuinely harder.

Location. One rehearsal somewhere that isn’t your desk. A library room, an empty lecture theatre. Feeling slightly wrong is part of what you’re practising.

No interruptions. Phone off, in another room. One three-hour block. If you can’t do the paper in a single sitting, you haven’t rehearsed the thing that’s hard about it, which is the third hour.

FROM THE MARKER’S SIDE

The most common thing students discover in their first genuine three-hour rehearsal is that their handwriting and their arithmetic both deteriorate badly after about a hundred minutes. That’s worth finding out in week 11 rather than in the exam.

12.7 Where to get papers

In order of usefulness:

- Your own department's past papers, with schemes if published. Nothing else matches the style.
- The same course from a different institution. The material is the same and the question style is a useful stress test.
- Textbook chapter reviews, which are often built from old exam questions.
- Any published exam in the subject: professional bodies, entrance exams, olympiad-adjacent sets for pure material.

Appendix H lists what's reliably available and free by subject.

If your department publishes no schemes, get one from any institution. Mark schemes are more similar across institutions than papers are, because they all follow the same M/A/B logic.

12.8 A term, on a calendar

The schedule table earlier is the shape. Here it is as actual appointments, for a twelve-week course with the exam in week 14.

Week	Session	Mode
3	Topic set, 20 questions	open book, by topic
5	Topic set, 20 questions	open book, by topic
7	Topic set, mixed	closed book, untimed
9	Half paper	closed book, untimed
10	Review + redo	—
11	Paper 1	closed, timed, one sitting
11	Review against scheme	—
12	Paper 2	closed, timed
12	Review + redo paper 1 failures	—
13	Paper 3 (exam time of day)	closed, timed
13	Review + redo paper 2 failures	—
14	Paper 4 (most recent past paper)	closed, timed
14	Review; then error log only	—

Thirteen appointments per course. That’s the difference between having a system and having a pile of papers, and the appointments are the part that makes it real: a plan to “do past papers” produces two in week 13.

12.9 Marking yourself honestly

The single biggest distortion in self-study, and it’s worth naming precisely. Students mark themselves generously in a specific way: they award method marks for methods they used but didn’t write. It feels fair, because they did in fact use the method. It isn’t fair, because the marker can’t see it, and the gap between your self-mark and your real mark is exactly the size of your leak 3.

Generous marking	Honest marking
“I knew to substitute, so M1”	Is the substitution written? No mark.
“Only an arithmetic slip”	Is the working visible for follow-through? If not, no ft.
“I’d have finished with time”	Marks what is on the page at the whistle.
“That’s basically the answer”	Is the answer stated? No? No A mark.

Mark yourself the way a stranger with ninety scripts would. The number will be four to eight lower than your instinct on your first honest attempt, and that difference is the most useful piece of information you’ll get all term.

FROM THE MARKER’S SIDE

When a student tells me they scored 70 on the past paper and 55 on the real one, the first thing I ask is who marked the past paper. It’s almost always self-marked, and almost always generous in exactly this way.

12.10 When there are no past papers

New module, new lecturer, or a department that publishes nothing. It happens, and the answer isn’t to skip rehearsal.

Build a paper from the problem sheets. Take the problem sets from the whole course, pick twelve questions across the topics in proportion to lecture time, assign marks in proportion to length, and sit it under timed conditions. It won’t match the real paper’s style, and it will rehearse everything in Parts II and III perfectly well.

Use another institution’s paper. The same module exists elsewhere. The style will differ; the material won’t.

Ask what the paper looks like. Not what’s on it. How many questions, whether there’s choice, how the marks split, whether a formula sheet is provided. Lecturers answer this readily and it’s most of what triage needs.

Use the lecturer’s own examples. The questions worked in lectures are the closest thing you have to the examiner’s taste, because they came from the same person.

A rehearsal with the wrong questions still rehearses the right protocol. The four minutes, the checkpoints, the abandonment rule and the sweep are all subject-independent, and they're most of what you're practising.

DO THIS BEFORE THE NEXT PAPER

Open a calendar and put four full timed papers on it, with dates, between now and your next exam. Then put a review session two days after each one and a redo session three days after that. Twelve appointments. That's the system, and the appointments are the whole difference between having a system and having a pile of papers.

13

The Error Log

One page. A table. It’s the highest-leverage revision artifact there is, and it takes ninety seconds a day.

13.1 What goes in it

Every mistake you make, all term, in 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 head	new line for new limits, always
26 Oct	ran 18 min over on Q3	thought I was close	3 min no ink means stop
2 Nov	used degrees	calculator mode	check mode at the start of every paper

The fourth column is the one that matters, and it’s the one people leave out. “What went wrong” is a diary. “The rule” is a change in behaviour, phrased so that future you can actually follow it.

An error log without a rule column is a list of regrets. The rule is the only part that changes what happens next time.

13.2 Why it works

Because your errors repeat, and you don't know it.

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




Figure 13.1: One student's log after four weeks, sorted by frequency. Thirty entries, and the top two account for nineteen of them. "Didn't know it" is fifth.

Look at where "genuinely didn't know" sits. Fifth. For most students, most of the time, not knowing the material is a minority cause of lost marks by the time exams come round, and it's the only one they revise for.

13.3 How to run it

One page, physical, visible. Front of your folder or stuck to the wall. A file on your laptop will be opened twice.

Log immediately. After every problem set, tutorial, and past paper. Ninety seconds while you still remember why you did it.

Log exam-technique errors as well as content errors. The overrun, the unwritten method, the uncaught sign. Those belong here just as much as "forgot the chain rule".

Read it before every practice paper. Thirty seconds, top to bottom. That's the whole mechanism: you're loading your own known failure modes into working memory right before you're going to be exposed to them.

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

13.4 Turning the log into practice

Every fortnight, look at the top two entries and build a targeted session around them.

If the top entry is sign errors in integration by parts, do fifteen parts questions and check only the signs. Not the whole solutions. The signs.

If the top entry is time overruns, do six questions with a hard timer, where the only success criterion is stopping when the timer says stop, regardless of whether you finished.

That second one feels absurd and it's the most effective drill in this book. You're not practising mathematics. You're practising stopping, which is the skill that's failing.

Practise the thing that's failing, not the subject it's failing in. If you lose marks to timing, do drills where timing is the only thing being scored.

13.5 Taking the log to somebody else

The log turns out to be the best thing you can put in front of a tutor, and it changes what you get back.

Compare two students at office hours in week 8.

"I'm not doing very well in this module."

"Here's my log for four weeks. Three of the sixteen entries are sign errors in integration by parts, and I can't work out what I keep doing. Can you watch me do one?"

The first student gets reassurance. The second gets somebody watching their hand move, spotting in ninety seconds that they're writing dv before deciding what u is, and fixing a recurring error permanently.

The same applies to a study group. "Bring your logs" is the most productive instruction you can give a group meeting, because everyone arrives with specific material and the overlap is usually large: three people with the same top entry can drill it together.

FROM THE MARKER'S SIDE

A student who brings evidence gets a different conversation. I can't do much with "I'm struggling". I can do a great deal with four weeks of data and a specific question about one recurring mistake.

13.6 The last three days

Back in Chapter 12: no new papers in the final three days. Here's what the log does with them.

Read the whole log. For each entry, do one question that exposes you to exactly that error and see whether the rule holds. An hour a day, three days, and you go into the exam with your specific failure modes freshly patched.

Compare that with the standard alternative, which is re-reading notes you already know or attempting a fresh paper you have no time to learn from.

FROM THE MARKER'S SIDE

The students who improve most between mock and final almost never learn a lot of new content in between. They fix two or three recurring execution errors, which is a smaller job than it sounds and worth more marks than another topic.

13.7 A worked entry

Here's what a good entry looks like when you write it properly.

9 Nov. Lost 4 marks on Q2(b): computed the eigenvector for $\lambda = 3$ but the question asked for $\lambda = -1$.

Why: I found both eigenvalues, then wrote out the working for the first one out of habit without re-reading the question.

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

Twenty seconds to write. And notice that the rule is specific enough to follow under pressure, which "be more careful" never is.

13.8 Four weeks of a real log

Here is what one looks like once it's running. This is a second-year engineering student, four weeks, sixteen entries compressed to nine.

Date	What went wrong	The rule
5 Oct	sign in parts	write u and dv boxed before starting
7 Oct	mm and m mixed	convert everything to SI before the formula
9 Oct	sign in parts	(repeat)
12 Oct	22 min on a 9-min part	3 min no ink means stop
14 Oct	no units, 3 answers	unit check on every answer, at the part
16 Oct	sign in parts	(repeat, third time)
19 Oct	mm and m mixed	(repeat)
23 Oct	30 min on a 12-min part	(repeat)
26 Oct	efficiency > 1 , submitted	possible? check on every ratio

Sixteen entries, and look at what they actually are: three sign errors in integration by parts, two unit conversions, two time overruns. Seven of the sixteen are three causes.

That's the whole point of logging. Before the log, this student would have told you their problem was "being careless". After four weeks, their problem is a specific, fixable list of three items.

13.9 Building the drill from the log

Take the top two and design a session that practises the failure, not the subject.

Top entry: sign errors in integration by parts. The drill is fifteen parts questions where you write u and dv in a box first, and check only the signs. Don't evaluate them fully. Twenty minutes.

Second entry: time overruns. The drill is six questions with a hard timer, where the success criterion is *stopping when the timer sounds*, whether or not you finished. If you finish three and abandon three cleanly, that's a perfect score.

That second drill feels absurd the first time. You're practising putting your pen down, and it's the most effective session in this book, because stopping is the skill that's failing and no amount of extra mathematics practice will fix it.

Drill the failure, not the topic. If your log says timing, do timing drills. Doing more mathematics is the answer to exactly one of the seven buckets.

13.10 Retiring entries, and watching it shrink

The log should get shorter, and watching it shrink is what keeps you running it.

Retire an entry when its rule has held for three consecutive papers. Draw a line through it, write the date, and leave it visible. A page with six crossed-out entries and two live ones is a picture of a term going well, and it's a much better motivator than any amount of resolve.

If an entry never retires, it isn't a rule yet. Rewrite it more specifically. "Check units" doesn't retire; "write the unit at the moment I box the answer, before moving to the next part" does, because it names the moment as well as the action.

DO THIS BEFORE THE NEXT PAPER

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 within twenty minutes, and two of them will already be repeats. Those two are what you're fixing.

14

Retrieval Practice for Problem Subjects

The research on how to study is unusually clear, unusually consistent, and almost never applied properly in STEM.

Testing yourself beats rereading. Spacing beats massing. Mixing beats blocking. All three are well-established, and all three are written up using examples like vocabulary and capital cities, which is why students in maths and physics read them, nod, and carry on rereading their notes.

Here's what the three look like when the unit of knowledge is a method rather than a fact.

14.1 Retrieval, when there's nothing to recall

Flashcards work by prompting recall of a stored fact. But 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.

So retrieval in a problem subject means solving, from a cold start, with nothing in front of you.

Which gives a concrete rule.

A worked example you read is worth almost nothing. The same example attempted from a blank page, then checked, is worth a great deal. The difference is whether you produced it.

The practical version: for every worked example in your notes, cover the solution, attempt it, then compare. That's it. Same material, same time, several times the retention.

14.2 The three artifacts that make it work

The problem log

Different from the error log, and just as valuable.

One line per problem you've solved this term: the problem, the method, and the single idea that unlocked it.

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

Thirty seconds per problem. Over a term it becomes a compressed index of every technique in the course, and it's revisable in an hour, which your folder of full solutions never will be.

Method cards

Flashcards, but for procedures. Front: a problem type. Back: the method, in four or five steps, plus the tell that identifies it.

Front: Integral with $\sqrt{a^2 - x^2}$ in it.

Back: Trig substitution $x = a \sin \theta$. Tell: a difference of squares under a root. Steps: substitute, replace dx , change limits, simplify with $1 - \sin^2 = \cos^2$, integrate, convert back if indefinite.

The tell is the part that matters and the part students leave off. In an exam nobody says "use trig substitution". Recognition is the skill, so the card has to train recognition, not just execution.


The one-page method map

Per topic, one page listing every method with its tell. Build it from the cards at the end of the topic.

That map is what you read before a timed paper, and it's the thing that converts "I know all these methods" into "I can pick the right one in twenty seconds".

14.3 Spacing, in practice

The rule is simple: revisit a topic at increasing intervals, and revisit it by testing, not reading.



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

Figure 14.1: Spacing one topic across a term. Five contacts, eighty minutes of revision in total after the first learning session, and every contact after the first is a closed-book test.

Twenty minutes, four times, spread out. Total eighty minutes per topic across a term, and it beats the four-hour session the weekend before by a distance that surprises people.

The catch is that spacing feels worse. Recall is harder at day 7 than at day 1, and that difficulty is exactly what makes it work. If it feels easy, it isn't doing anything.

14.4 Interleaving, and why blocked practice lies

Do twenty integration-by-parts questions in a row and you'll feel fluent by question 15.

You aren't. You've learned to execute a method you were told to use. In the exam nobody tells you, and the students who practised in blocks routinely can't identify which method a question wants, despite being able to perform all of them.

So after the initial blocked practice that gets you fluent, switch to mixed sets. Ten questions drawn from four different topics, in random order, where the first decision on each is *which method is this?*

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

Mixed practice feels worse too. Your accuracy drops, you make more mistakes, and it's less satisfying. Both of the effective techniques feel worse

than the ineffective one, which is the whole reason ineffective revision is so popular.

14.5 The problem log, in practice

The problem log gets one line per problem you solve, all term, and it's the artifact students are most sceptical about and most convinced by afterwards.

Three entries, so the format is clear:

Sheet 4, Q3. Volume of revolution with a hole. Method: 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. Method: CF has the $te^{\lambda t}$ term. Key idea: repeated root means the second solution is t times the first.

Sheet 6, Q4. Beam reactions. Method: moments about the left support. Key idea: choosing the support kills one unknown immediately.

Thirty seconds each. Notice what the third field is doing: it 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.

Over a term this becomes a compressed index of every technique in the course, revisable in an hour. Your folder of full solutions never will be, because rereading a full solution takes as long as it took to write.

FROM THE MARKER'S SIDE

Ask a student in week 12 what they did on problem sheet 4 and they usually can't tell you. Ask them to find it and they'll produce four pages of algebra with no indication of what the point was. The log is the fix, and it costs about eight minutes a week.

14.6 What to stop doing

Three habits that consume enormous amounts of time and produce very little:

Rewriting notes. It feels like work, it's nearly passive, and the output is a tidier version of something you already have. If you're going to write, write a method map or a summary from memory.

Reading worked solutions. Only after attempting. Reading a solution to a problem you haven't tried produces the strong feeling of understanding and no ability.

Highlighting. No.

FROM THE MARKER'S SIDE

The gap between the students who do well and the students who work equally hard and don't is almost never effort. It's that one group spends their hours producing answers and the other spends them looking at answers.

14.7 A workable week

For one course, in a normal week:

- **After each lecture, 10 minutes.** Close the notes and write what the lecture's method was, and its tell. Add a method card.
- **Problem set, as normal.** Then log each problem in one line.
- **One 20-minute spaced session** on a topic from three weeks ago. Closed book, three problems, mixed.
- **Ninety seconds** updating the error log.

Under an hour a week, on top of the work you're already doing, and it's the difference between arriving at the exam with a term of material and arriving with a term of material you can retrieve.

14.8 A method map, built

Here's what one page looks like at the end of a topic. This is first-year integration, which is the topic students most often "know" and most often can't select from.

The tell	The method
A product, one factor easy to differentiate	by parts, u = the one that simplifies
$\sqrt{a^2 - x^2}$	$x = a \sin \theta$
$\sqrt{a^2 + x^2}$	$x = a \tan \theta$
$\sqrt{x^2 - a^2}$	$x = a \sec \theta$
A rational function, denominator factorises	partial fractions
Inner function and its derivative both present	substitution
Odd power of sin or cos	peel one off, use the identity
Even powers of sin or cos	double-angle formulas
$\int_{-a}^a$ with an odd integrand	zero, by symmetry
Recurrence between I_n and I_{n-1}	reduction formula, find the base case

Ten rows, one page, and it’s the sheet you read for ninety seconds before a timed paper.

Notice the left column. That’s the exam’s question, not the textbook’s. The textbook says “integration by parts” and then gives you twenty by-parts problems; the exam gives you a product and says nothing.

14.9 A week that actually fits

Everything in this chapter, for one course, in a normal week:

When	Time	What
After each lecture	10 min	close the notes, write the method and its tell; add a card
Problem set	as usual	then one line per problem in the problem log
Midweek	20 min	spaced session on a topic from three weeks ago, closed book, mixed
End of week	90 s	update the error log
Total added	≈ 50 min	on top of work you already do

Fifty minutes a week per course. Over a twelve-week term that's ten hours, and it replaces the thirty hours of rereading that most students do in week 13 with something that works.

FROM THE MARKER'S SIDE

The students who arrive at the exam able to select methods quickly are almost never the ones who worked the most hours. They're the ones who spent ten minutes after each lecture writing down what the method was for.

DO THIS BEFORE THE NEXT PAPER

Take the topic you covered three weeks ago. Close everything, set twenty minutes, and do three problems from it cold. Whatever happens, that's a more accurate picture of what you know than any amount of rereading, and the gaps it exposes go straight into the error log.

15

The Final Two Weeks

Two weeks out, most students do the same thing: panic gently, and try to cover everything again.

Covering everything again is impossible and wouldn't help. The last fortnight has a different job, and it's mostly about execution and selection rather than content.

15.1 What the two weeks are for




Figure 15.1: The last fourteen days, counting down to the exam. Content learning stops on day 10; the whole back half is rehearsal and patching.

15.2 Days 14 to 10: close the real gaps

Not all the gaps. The ones that pay.

Take your error log and your course outline, and rank topics by two things: how often they've appeared on past papers, and how badly you did last time you were tested.

A topic that appears on every paper and that you can't do is worth four days of anyone's fortnight. A topic that appeared once in nine years and that you can't do is worth nothing, and you should write it off explicitly so it stops occupying you.

Write down, in the first week, the two topics you are choosing not to learn. Naming them stops them costing you attention for a fortnight, and frees the time for topics that actually appear.

That decision feels reckless and is simply arithmetic. You have finite hours. Spending them uniformly across topics of wildly different exam frequency is the least efficient allocation available.

15.3 Days 9 to 4: rehearse

Four full papers, timed, in one sitting each, in the mode from Chapter 12.

Each one gets: the paper (3 hours), the review against the real scheme (1–2 hours, two days later), and a redo of the failures (45 minutes).

Run the full protocol every time. The four minutes. The wall-clock deadlines. The checkpoints. The abandonment rule. The final sweep. If you're not running the protocols in rehearsal, you won't run them in the exam, because the room doesn't grant you new capabilities.

Do at least one of them at the actual exam time of day, in a room that isn't where you study, with no phone. Feeling slightly wrong is part of what you're rehearsing.

15.4 Days 3 to 1: patch, don't learn

No new papers. No new topics.

- **Read the error log end to end**, once a day.
- **One question per log entry**, to check the rule holds.
- **Read your method maps**, one topic at a time.
- **Rewrite the formula sheet from memory**, then check it. Twice.
- **Sleep**. Properly. This is not a wellness note; sleep deprivation costs more marks than any topic you could learn in the time.

The instinct in these three days is to keep learning, because stopping feels like giving up. But new content learned 48 hours before an exam is almost never retrievable under pressure, and the hours spent on it come out of the sleep and the patching, both of which are worth more.

FROM THE MARKER'S SIDE

I've never seen a script where the difference was a topic learned two days before. I've seen many where the difference was obvious exhaustion: arithmetic slips everywhere, questions half-read, timing collapsed after ninety minutes.

15.5 A season, not an exam

Most students don't sit one exam. They sit four or five inside a fortnight, and the fortnight has to be planned as a unit or the later papers get whatever's left of you.

Situation	What to do
Two papers, three days apart	Full post-mortem after the first is fine; one day of recovery, one of patching
Two papers, one day apart	Ten-minute post-mortem only. Everything else goes to the second subject
Two papers, same day	No post-mortem at all. Eat, walk, read the second subject's method map
Four papers, ten days	Front-load the revision <i>before</i> the season starts. Inside it there is no learning time, only patching

The mistake that costs most: allocating revision time inside the season proportionally to how the papers are going. A bad first paper pulls attention backwards to a subject that's finished, and the third and fourth papers pay for it.

Once the season starts, every hour belongs to the next paper, not the last one. The finished paper's mark is already fixed; the next one isn't.

Plan the whole fortnight before the first paper, on one sheet, with the patching hours allocated per subject. Then follow it regardless of how any individual paper feels, because your judgement inside a season is worse than your judgement before it.

15.6 The day before

Short and specific:

- Two hours of review, maximum, and only material you already know.
- Assemble the physical things: pens that work, calculator with fresh batteries and the right mode, ID, water, watch.
- Write your protocol on a card: minutes per mark, the four minutes, the three checkpoints, the sweep. Read it in the morning.
- Stop by early evening.

The two-hour cap is deliberate. The marginal value of hour seven the day before is negative, and everybody who's done it knows the feeling of arriving at an exam already tired.

15.7 The morning

Eat. Arrive early enough that you're not rushing, not so early that you stand around absorbing other people's panic.

Read your protocol card. Nothing else. Do not do a last-minute question, because a question you get wrong forty minutes before the exam does real damage to how you start.

Then go and run the protocol you rehearsed four times.

15.8 If you're badly behind

The plan above assumes you've been working all term. Sometimes you haven't, and a fortnight out you're looking at a course you barely know.

The instinct is to start at week 1 and work forwards. Don't. You will reach week 6 of 12 and run out of days, having learned the easy half.

Do this instead.

1. **Get four past papers and count topic frequency.** Which three topics appear on every single paper? Those are your fortnight.
2. **Learn those three from worked examples backwards.** Take a past exam question with a solution, read the solution, then close it and reproduce it. Repeat with four more questions on the same topic. This is faster than learning the theory first and it's aimed exactly at what's examined.

3. **Accept the rest.** Write down the topics you are not going to learn, and stop thinking about them.
4. **Protect the protocol.** A student who knows three topics and executes perfectly beats a student who knows five and manages their time badly. Do at least two timed papers.

That's a triage plan, and it feels like giving up on the course. It isn't. It's the same marks-per-minute logic as the exam hour, applied to a fortnight.

When you're short of time, depth on the three topics that always appear beats breadth across all twelve. Exam papers are not uniform samples of the syllabus, and pretending they are is the expensive assumption.

15.9 What the last fortnight cannot do

Two honest limits, so the plan doesn't promise what it can't deliver.

It can't build fluency you never had. If you've never solved a second-order ODE, ten days won't make you quick at it under pressure. It can make you able to start one, recognise the shape, and collect method marks, which on a mark scheme weighted to method is a large fraction of the available credit.

It can't replace sleep. Every hour after about the sixth in a day is worth much less than the hour of sleep it costs, and the effect compounds across a fortnight. Students routinely arrive at the first paper of the season already depleted and then wonder why the third one went badly.

What the fortnight can do, reliably, is convert knowledge you already have into marks you actually collect. That's usually worth more than it sounds: the audits in Chapter 3 keep landing on the same number, somewhere between five and fifteen percentage points sitting on the table for execution reasons.

DO THIS BEFORE THE NEXT PAPER

Two weeks out, sit down with your error log and the past-paper topic frequencies, and write three lists: topics to fix, topics to rehearse, and topics you are choosing not to learn. Then put the four timed papers in your calendar with dates and times. Fifteen minutes of planning, and the fortnight stops being a fog.

Part V

Particular Kinds of Question and Paper

16

Units, Figures, and Sanity Checks

For physicists and engineers this chapter is worth more than any other in the book, because these marks are pure loss. They don't require any physics. They're given away by students who understood the question perfectly.

16.1 Units are marks, and they're also a free check

Two separate facts, and students underrate both.

Units are marks. A numerical answer without a unit is incomplete, and most schemes say so explicitly. Across a paper that's routinely three or four marks, and it takes two seconds per answer.

Units are a check. This is the bigger one. Carry units through the working and the algebra tells you when you've used the wrong formula, before you've checked a single number.

Dimensional analysis is the cheapest error detection in physics and engineering. If the units of your answer are wrong, the formula is wrong, and you knew it without computing anything.

Worked

You want a period and your working gives $T = 2\pi\sqrt{g/L}$.

Check: g is m s^{-2} , L is m , so g/L is s^{-2} , and the square root is s^{-1} . That's a frequency, not a period.

So the formula is upside down and it's $T = 2\pi\sqrt{L/g}$. Four seconds, no memory required, and it catches one of the most common inversions in first-year mechanics.

Get in the habit and you'll stop needing to remember which way up half of your formulas go, because you can rebuild them.

16.2 Significant figures, without the folklore

The rules that actually matter in an exam:

- **Match the data.** If the question gives values to three significant figures, answer to three. Giving nine is wrong-looking and sometimes penalised.
- **Obey explicit instructions exactly.** “To 2 decimal places” is not “to 3 significant figures”. These are different and the scheme checks.
- **Never round in the middle.** Carry full precision through the working and round once, at the end. Mid-calculation rounding produces answers that miss the scheme's accepted range.
- **Show the unrounded value once.** Write “ $= 4.68219 \dots = 4.68$ (3 s.f.)”. It proves you didn't round early and it protects you if your final rounding is off.

That last habit costs three seconds and settles most accuracy disputes in your favour.

16.3 Significant figures, five worked cases

The rules are short; applying them is where the marks go.

Given	Answer	Why
$m = 2.5 \text{ kg}$, $a = 3.75 \text{ m s}^{-2}$	$F = 9.4 \text{ N}$	data to 2 and 3 s.f.; quote to the <i>lowest</i> , 2 s.f., so 9.4
“to 2 decimal places”	4.07	not 2 s.f., which would be 4.1
πr^2 , $r = 3.0 \text{ cm}$	28 cm^2	π is exact, so precision comes from r : 2 s.f.
Counted: 12 components	exact	counts are exact and don’t limit precision
4.68219 ..., no instruction	4.68	default to 3 s.f. and show the unrounded value once

The first row is the rule students most often invert: a chain of multiplications carries the precision of the *least* precise input, not the most.

And the fourth is worth knowing because it looks like a trap and isn’t. If you counted something, or the question specifies an exact quantity, it imposes no limit at all.

16.4 The sanity checks that belong to physics

Beyond the five general checks in Chapter 11, physics and engineering answers have their own quick tests.

The last row deserves a sentence of its own. Check your calculator’s angle mode in the first minute of the exam and write “RAD” or “DEG” at the top of your script. A mode error doesn’t cost you one mark, it silently corrupts every trigonometric answer on the paper, and students discover it during the walk home.

16.5 Prefixes, the quiet killer

Kilo, milli, micro, nano, mega, giga. Almost every engineering script loses marks here, and the fix is a habit rather than knowledge.

Convert everything to base SI units at the start of the question. Write the converted values down. Do the whole calculation in metres, kilograms, seconds, and only convert back at the very end if the question asks for a particular unit.




Figure 16.1: Physics-specific sanity checks. All are faster than recomputing, and all catch real exam errors.

Mixed prefixes mid-calculation, millimetres in one term and metres in another, is one of the highest-frequency errors in the subject and it's entirely eliminated by that one rule.

FROM THE MARKER'S SIDE

An enormous number of scripts have the right method and an answer out by exactly 10^3 or 10^6 . Follow-through usually saves most of the marks, but the accuracy mark is gone, and on a multi-part question the wrong magnitude often makes the next part look absurd, which costs more.

16.6 Estimation questions, and why they're gifts

"Estimate the number of..." questions look scary and are usually the cheapest marks on a physics paper.

The marks are for the method: stating assumptions, choosing sensible values, and getting the arithmetic to the right order of magnitude. The answer isn't expected to be exact, and the scheme usually accepts a wide range.

So the structure that scores:

1. State every assumption explicitly, with a number.
2. Show the formula or the chain of reasoning.

3. Compute to one or two significant figures.
4. State the answer to an order of magnitude and say it's an estimate.

Students lose marks here by trying to be precise, or by making assumptions silently. The assumptions *are* the marks.

16.7 The unit table you should know cold

Not for memorising formulas, for checking them. These are the ones that catch the most errors in first and second year.

Quantity	SI unit	In base units
force	N	kg m s^{-2}
energy, work	J	$\text{kg m}^2 \text{s}^{-2}$
power	W	$\text{kg m}^2 \text{s}^{-3}$
pressure, stress	Pa	$\text{kg m}^{-1} \text{s}^{-2}$
charge	C	A s
voltage	V	$\text{kg m}^2 \text{s}^{-3} \text{A}^{-1}$
resistance	Ω	$\text{kg m}^2 \text{s}^{-3} \text{A}^{-2}$
spring constant	N m^{-1}	kg s^{-2}
Young's modulus	Pa	$\text{kg m}^{-1} \text{s}^{-2}$
frequency	Hz	s^{-1}

Two facts from that table do a lot of work. Energy and work share units, so any formula producing one can be checked against the other. And pressure, stress and Young's modulus are all pascals, which is why the stress-strain relation has to have a dimensionless strain, and why a strain quoted with units is always wrong.

16.8 Answering in the right form

Last one, and it's cheap. Read what form the answer should take:

- “in terms of m , g and θ ” means no numbers in the answer;
- “in surd form” means don't decimalise;
- “correct to the nearest whole number of components” means round up if it's a physical count, because 3.2 cables means 4 cables;
- “state the direction” means a vector answer isn't complete without it.

Each of those is a mark, each is available for free, and each is lost constantly by students who solved the physics perfectly.

16.9 Worked: five formulas checked in forty seconds

Dimensional analysis is only useful if it's fast, so here it is at speed. Each of these is a formula a student wrote in an exam. Three are wrong, and you can tell without knowing any physics beyond the units.

Written	Check
$T = 2\pi\sqrt{g/L}$	$\sqrt{\text{m s}^{-2}/\text{m}} = \text{s}^{-1}$. That's a frequency. Upside down.
$v = \sqrt{2gh}$	$\sqrt{\text{m s}^{-2} \cdot \text{m}} = \text{m s}^{-1}$. Fine.
$E = \frac{1}{2}mv$	kg m s^{-1} is momentum, not energy. Missing a v.
$R = \rho A/L$	$\Omega \text{ m} \cdot \text{m}^2/\text{m} = \Omega \text{ m}^2$. Should be $\rho L/A$.
$P = I^2 R$	$\text{A}^2 \Omega = \text{A}^2 \cdot \text{V/A} = \text{VA} = \text{W}$. Fine.

Eight seconds each. And notice what the check gives you beyond error detection: once you've done this a few dozen times you stop needing to remember which way up $\rho L/A$ goes, because you can rebuild it from the units in the time it takes to write it down.

Dimensional analysis isn't only a check. It's a way of reconstructing a formula you half remember, which is a much better position to be in than guessing between two versions.

16.10 Worked: the prefix disaster

This one is worth walking through slowly, because it's the highest-volume error in engineering exams and it's entirely mechanical.

Q. A steel rod of diameter 12 mm and length 1.5 m carries a tensile load of 25 kN. Young's modulus is 210 GPa. Find the extension. [5]

The formula is $\delta = FL/AE$. Here's the script that loses marks:

BEFORE

$$A = \pi r^2 = \pi(6)^2 = 113$$
$$\delta = \frac{25 \times 1.5}{113 \times 210} = 0.00158$$

Radius in millimetres, area therefore in mm^2 , load in kilonewtons, modulus in gigapascals, and the answer in nothing at all. The number is wrong by a factor of 10^6 , and there's no unit to reveal it.

Convert everything first, before touching the formula:

AFTER

$r = 6 \text{ mm} = 6 \times 10^{-3} \text{ m}$, so $A = \pi r^2 = \pi(6 \times 10^{-3})^2 = 1.131 \times 10^{-4} \text{ m}^2$.
 $F = 25 \text{ kN} = 25 \times 10^3 \text{ N}$. $E = 210 \text{ GPa} = 210 \times 10^9 \text{ Pa}$. $L = 1.5 \text{ m}$.

$$\delta = \frac{FL}{AE} = \frac{(25 \times 10^3)(1.5)}{(1.131 \times 10^{-4})(210 \times 10^9)} = 1.58 \times 10^{-3} \text{ m}$$

$= 1.58 \text{ mm}$ (3 s.f.)

Same arithmetic, and now the units come out as metres on their own, which is the check. An extension of 1.58 mm on a 1.5 m rod is also physically sensible, which is the second check, and both took five seconds.

FROM THE MARKER'S SIDE

When an answer arrives out by exactly 10^3 or 10^6 , follow-through usually rescues most of the method marks. What it can't rescue is the next part, where a stress of 200 MPa reported as 200 Pa makes the student's own safety-factor calculation come out at 0.0000012 and they write it down anyway.

DO THIS BEFORE THE NEXT PAPER

Take your last physics or engineering script and count only the marks lost to units, prefixes, significant figures, calculator mode, and answer form. Total them. For most students the number is between three and seven, and none of it required any additional knowledge to recover.

17

Data, Graphs, and Uncertainty

Data questions have a reputation for being fiddly, and they're actually the most predictable marks on a physics or engineering paper, because they're marked against a near-identical checklist every time.

Learn the checklist and these become the questions you want.

17.1 The graph checklist

Every “plot a graph and find the gradient” question is marked on roughly the same six things.

Five of those six are awarded before any physics happens. Students who treat graph questions as an annoyance and rush them are throwing away the most reliable marks available.

Two specific errors worth naming, because they're near-universal:

Reading the gradient off two data points. The whole reason you drew a best-fit line is to average the noise. Take two points *on the line*, far apart, and preferably not data points at all.

A triangle that's too small. A gradient triangle spanning a quarter of your line multiplies your reading error by four. Use the biggest triangle the line allows, and draw it on the graph so the marker can see you did.

17.2 Linearising

The single most common design in data questions: you're given a relationship that isn't linear, and you have to make it linear so a straight-line graph gives you the constant.




Figure 17.1: The six graph marks. They are the same on nearly every paper, and five of them require no physics at all.

Relationship	Plot	Gradient is
$y = kx^n$	ln y against ln x	n, intercept ln k
$y = Ae^{kx}$	ln y against x	k, intercept ln A
$T = 2\pi\sqrt{L/g}$	T^2 against L	$4\pi^2/g$
$1/v + 1/u = 1/f$	$1/v$ against $1/u$	-1, intercept $1/f$

The method that generates all of these: rearrange until the equation has the form $Y = mX + c$, then say explicitly what Y , X , m and c are. Write that line down. It’s nearly always a method mark, and it’s the line students do in their heads.

17.3 Uncertainty: the rules you actually need

Uncertainty is taught as arithmetic and marked as writing. Three rules cover almost everything an exam asks.

Adding or subtracting: add absolute uncertainties. If $c = a + b$ then $\Delta c = \Delta a + \Delta b$.

Multiplying or dividing: add percentage uncertainties. If $c = ab$ then the percentage uncertainty in c is the sum of the percentage uncertainties in a and b .

Powers: multiply the percentage uncertainty by the power. If $c = a^3$ then the percentage uncertainty in c is three times that in a .

Then two presentation rules that carry their own marks:

- Quote the uncertainty to one significant figure, and match the answer's precision to it. " 4.68 ± 0.03 ", not " 4.6821 ± 0.0312 ".
- State it in the same units as the quantity.

17.4 The discussion marks

The last part of a data question usually asks something like "suggest an improvement" or "comment on the reliability", and it's worth two or three marks that students answer with generalities.

Generalities score nothing. "Repeat the experiment" and "be more careful" and "use better equipment" are on every marker's list of non-answers.

What scores is specific and causal: name the source of error, say what it does to the result, and say what would fix it.

Weak: There may be human error in the timing.

Scores: Reaction time in starting and stopping the stopwatch adds an uncertainty of roughly 0.2 s to each measurement. Since each swing is about 1 s, timing 20 swings and dividing reduces this to about 0.01 s per swing.

Notice the shape: source, effect with a number, fix. Three clauses, three marks, and it's a template you can carry into every data question you ever sit.

Improvement marks are for causal statements with numbers, not for good intentions. Name the source, quantify the effect, state the fix.

17.5 Uncertainty, worked

The rules are short. Applying them under time pressure is where marks go, so here they are on one calculation.

A resistor carries $I = 2.40 \pm 0.05$ A with a voltage $V = 12.0 \pm 0.2$ V across it. Find the power dissipated and its uncertainty.

Value. $P = VI = 12.0 \times 2.40 = 28.8$ W.

Percentage uncertainties.

$$\frac{\Delta V}{V} = \frac{0.2}{12.0} = 1.67\%, \quad \frac{\Delta I}{I} = \frac{0.05}{2.40} = 2.08\%.$$

Combine. Multiplication, so add the percentages: $1.67 + 2.08 = 3.75\%$.

Convert back. 3.75% of 28.8 is 1.08, so $\Delta P = 1$ W to one significant figure.

Quote it properly. $P = 29 \pm 1$ W.

That last step is where the marks are lost. Writing “ 28.8 ± 1.08 W” is wrong on both halves: the uncertainty should be one significant figure, and the value’s precision should match it. If the uncertainty is in whole watts, the value cannot be quoted to a tenth of a watt.

Round the uncertainty first, to one significant figure. Then round the value to the same decimal place. Never the other way round.

And if the question had asked for $P = I^2 R$ instead, the power rule applies: the percentage uncertainty in I^2 is twice that in I . Powers multiply, they don’t add.

17.6 Anomalies

If a point is obviously off the line, say so in writing, exclude it from the best-fit line, and give a reason if you can.

Circle it on the graph. Write “anomalous, excluded from line of best fit”. That’s usually a mark, and silently ignoring the point is not.

17.7 The tables you're given

Two habits worth having before you plot anything:

Add the derived column. If you're plotting T^2 against L , compute a T^2 column in the table and show it, with its unit in the header. That table is marked.

Keep consistent precision. Every value in a column to the same number of decimal places. It looks trivial and it's an easy presentation mark on many schemes.

17.8 Worked: a complete data question

One question, run end to end, with every mark identified.

Q. A simple pendulum is timed for 20 oscillations at five different lengths. Plot a suitable straight-line graph and use it to determine g . Estimate the uncertainty in your value. [12]

L / m	0.200	0.400	0.600	0.800	1.000
t_{20} / s	17.9	25.4	31.1	35.9	40.1

Step 1: linearise, and say so. $T = 2\pi\sqrt{L/g}$, so $T^2 = \frac{4\pi^2}{g}L$. Plotting T^2 against L gives a straight line through the origin with gradient $4\pi^2/g$.

Write that sentence down. On most schemes it's a method mark, and it's the line students do in their heads.

Step 2: build the derived column, in the table.

L / m	0.200	0.400	0.600	0.800	1.000
T / s	0.895	1.270	1.555	1.795	2.005
T^2 / s^2	0.801	1.613	2.418	3.222	4.020

Three decimal places throughout, consistently, and the unit in every header. That table is marked.

Step 3: plot and take the gradient. Axes labelled " T^2 / s^2 " and " L / m ", scale using most of the grid, small crosses, best-fit line drawn with a straight edge.

Gradient triangle spanning the full line, using points *on the line* rather than data points:

$$\text{gradient} = \frac{4.02 - 0.80}{1.000 - 0.200} = \frac{3.22}{0.800} = 4.03 \text{ s}^2 \text{ m}^{-1}.$$

Step 4: extract the quantity, with units.

$$\frac{4\pi^2}{g} = 4.03 \quad \Rightarrow \quad g = \frac{4\pi^2}{4.03} = 9.80 \text{ m s}^{-2}.$$

Step 5: the uncertainty, quoted properly. Draw the steepest and shallowest lines that still fit the points. Suppose they give gradients of 4.11 and 3.95. Then

$$\Delta(\text{gradient}) = \frac{1}{2}(4.11 - 3.95) = 0.08,$$

which is 2%. Since $g \propto 1/\text{gradient}$, the percentage uncertainty carries straight across, giving $\Delta g = 0.2 \text{ m s}^{-2}$.

Quote it as $g = 9.8 \pm 0.2 \text{ m s}^{-2}$: uncertainty to one significant figure, and the answer's precision matched to it.

FROM THE MARKER'S SIDE

Look at how many of those twelve marks needed no physics: the derived column, the axes, the scale, the line, the triangle, the units, the uncertainty format. That's eight of them. The physics was two lines, at steps 1 and 4.

17.9 Where this question usually loses marks

Same question, the four ways students routinely drop four or five marks on it:

- **Plotting T against L** and drawing a curve, then trying to find g from it. The word *suitable* in “a suitable straight-line graph” was the instruction to linearise.
- **Forgetting to divide t_{20} by 20.** The gradient comes out 400 times too large and g arrives at 0.025 m s^{-2} , which the magnitude check from Chapter 11 would have caught in two seconds.
- **Taking the gradient from two data points.** It defeats the purpose of the line and is explicitly penalised on many schemes.

- **Quoting** $g = 9.7983$ with no uncertainty. Both halves of that are wrong: the precision is unjustified, and the uncertainty was asked for.

DO THIS BEFORE THE NEXT PAPER

Find a data question from a past paper and mark your own attempt against the six-item graph checklist, item by item. Most students find they lose two or three of the six every time, always the same two or three. Those are now log entries.

18

Multi-Part and “Show That” Questions

The structure of a multi-part question is a gift, and most students unwrap it upside down.

18.1 The ladder

A question with parts (a), (b), (c) is almost never three separate questions. It’s one problem, broken into steps by somebody who wanted you to succeed, where each part produces the tool for the next.




Figure 18.1: The ladder. Each part hands the next one its input, and the printed result in a “show that” is yours to use even if you couldn’t derive it.

Now the fact that makes this worth a chapter.

If part (a) says “show that X ”, you may use X in every later part even if you couldn’t prove it. The result is printed in the question. It is given to

you.

Students who can't do part (a) routinely abandon the entire question. That is 11 marks left on the table in the example above, on parts they could have done, because part (a)'s answer was sitting there in print.

Whenever you meet a “show that” you can't complete: write “unable to show this; using the given result” and carry straight on to part (b). Nobody deducts for it, and every later part is now fully available.

18.2 How to read “hence”

Hence is an instruction, not a connective.

It means: use the thing you just did. A method that ignores it may score nothing at all, even if it reaches the right answer, because the mark scheme is written around the intended route.

Related words worth knowing precisely:

- **Hence** use the previous result. Compulsory.
- **Hence or otherwise** the previous result is the easy route, but any correct method scores. Use the previous result anyway; it's quicker.
- **Deduce** a short step from what's just been established. If you're doing heavy work, you've missed something.
- **Verify** or **check** substitute and confirm. You do not need to derive it.
- **Write down** this should take one line. If it's taking you five, you're overcomplicating a question that's testing recall.

That last one saves real time. “Write down the value of...” is the examiner telling you no work is expected.

18.3 Show that: what's actually being marked

Since the answer is printed, none of the marks are for the answer. They're all for the route, which changes how you write.

Every step visible. No mental algebra. The step is the mark.

Work forwards, not backwards. Starting from the printed result and manipulating until you reach the start is usually accepted only if the steps are reversible and you say so. It's safer to work forwards.

End at exactly the printed form. If they ask you to show that $y = \frac{3x}{x+2}$, don't stop at $y = 3x(x+2)^{-1}$ and leave the last step implied. Land exactly on the printed expression.

Never fabricate the connection. If you're two lines short, say so and use the result anyway. A visible gap costs one mark. A fake bridge between your work and the printed answer costs the marker's trust and usually more marks than the gap.

FROM THE MARKER'S SIDE

The tell for a fabricated "show that" is a step that does nothing except change the expression into the printed one, with no justification. It's obvious, and it's a shame, because the honest version, three correct lines and a note saying where it stopped, usually scores more.

18.4 When a later part is easier

It happens constantly, and the ladder makes it feel forbidden. It isn't.

If you can't do (b) but you can do (c) using (b)'s printed or implied result, do (c). Label it clearly, state what you're assuming from (b), and carry on.

(c) Assuming from (b) that the stationary points are at $x = 0$ and $x = 4$:

One line, and part (c) is now fully markable. Students leave (c) blank because (b) is unfinished, which converts one lost part into two.

18.5 The instruction verbs, complete

This is the list the whole book keeps referring to. Learn what each one obliges you to produce, because they are not interchangeable and the scheme is written around them.

Verb	What the answer must contain
State	one line, no working. Recall only.
Write down	one line. If it's taking five, you've misread it.
Find, Calculate	a value, with units, plus enough method to be marked
Determine	a value, usually with a supporting argument
Show that	every step visible, ending exactly on the printed result
Prove	a complete argument with its structure explicit
Verify	substitute and confirm; no derivation needed
Deduce	a short step from what was just established
Hence	compulsory use of the previous part
Hence or otherwise	the previous part is the easy route; others score
Sketch	shape, intercepts, asymptotes, labels. Not to scale.
Draw	to scale, with instruments, on the grid
Estimate	assumptions stated with numbers, one or two significant figures
Explain, Justify	prose, causal, usually one mark per distinct reason
Comment on	a judgement plus the evidence for it
Compare	both things, explicitly, in the same sentence or table
Derive	from stated first principles, not quoted
Simplify	a specific target form is expected
Solve	all solutions in the stated range, not just one

Three of those cause most of the losses.

Solve means all solutions in the range. A trigonometric equation with one answer given where four exist loses three quarters of the accuracy marks.

Sketch versus *draw* is a real distinction. A sketch that took you six minutes with a ruler wasted five of them; a “draw” answered freehand loses the accuracy marks.

Explain wants prose. A calculation in place of an explanation scores nothing, however correct, because the mark is for the reason.

The verb tells you the form of the answer, and the marks are allocated to that form. Underlining it takes two seconds and is the cheapest insurance on the paper.

18.6 Marks tell you the size of the answer

The mark allocation is a message about how much writing is expected.

- **1 mark:** one line. State something.
- **2–3 marks:** a short method plus an answer.
- **4–6 marks:** a multi-step method. If you've written one line, you've misread it.
- **7+ marks:** a substantial argument, usually with a structure the question implies.

If your answer to a 6-mark part is two lines, stop and reread the question. Either you've missed a requirement, or you've solved a simpler problem than the one asked. Both are common and both are catchable in fifteen seconds by comparing the length of your answer with the marks available.

Compare the size of your answer with the marks. It's the fastest misreading detector there is, and it costs one glance per part.

18.7 Worked: a ladder, climbed and misclimbed

Q(a). Show that $\int_0^1 x^n e^{-x} dx = n \int_0^1 x^{n-1} e^{-x} dx - e^{-1}$ for $n \geq 1$.
[4]

Q(b). Hence evaluate $\int_0^1 x^2 e^{-x} dx$. [5]

Q(c). Deduce the value of $\int_0^1 x^3 e^{-x} dx$. [3]

The ladder. Part (a) builds a reduction formula. Part (b) applies it twice, from a base case you can compute directly. Part (c) applies it once more, using (b)’s answer.

Each rung needs the one below it, and the printed result in (a) means every rung is available whether or not you can climb the first.

How it goes right. In (b), write the plan line: “using the reduction formula from (a) with $n = 2$, then $n = 1$ ”. Then

$$I_2 = 2I_1 - e^{-1}, \quad I_1 = 1 \cdot I_0 - e^{-1}, \quad I_0 = \int_0^1 e^{-x} dx = 1 - e^{-1}.$$

So $I_1 = 1 - 2e^{-1}$ and $I_2 = 2 - 5e^{-1} \approx 0.160$.

Then (c) is a single application: $I_3 = 3I_2 - e^{-1} = 6 - 16e^{-1} \approx 0.113$. Three marks for two lines, because you climbed the ladder in order.

How it goes wrong, three ways.

Abandoning after (a). The student can’t do the integration by parts in (a), so leaves the whole question. Eight marks in (b) and (c) gone, on parts that only needed the printed formula.

Ignoring “hence”. The student does (b) by integrating by parts twice from scratch. It works, it takes nine minutes instead of three, and on schemes that specify the reduction route it may score two of the five.

Missing the base case. The student applies the formula down to I_0 and then doesn’t know what I_0 is, having never noticed that $\int_0^1 e^{-x} dx$ is elementary. The ladder needs a floor, and the floor is always something you can compute directly.

When a reduction formula appears, find the base case before you start descending. It is always a direct, elementary integral, and without it the recursion has nowhere to land.

18.8 Reading the mark split as a map

The allocation across a ladder tells you where the work is.

Split	Shape	What it tells you
4 / 5 / 3	rising then falling	(b) is the main event; (c) is a one-line consequence
2 / 3 / 10	back-loaded	(a) and (b) are setup; the real question is (c), budget accordingly
8 / 3 / 3	front-loaded	(a) is a substantial proof; the rest is application
5 / 5 / 5	flat	three near-independent parts, so a failure in one need not sink the others

Read the split during the four minutes, not at minute forty. A back-loaded question that you enter thinking it's flat will eat your schedule.

DO THIS BEFORE THE NEXT PAPER

Take a past paper and, for every multi-part question, draw the ladder: which part feeds which. Then find every “show that” and write next to it “result available from here”. Ten minutes, and you’ll never again abandon a question because part (a) defeated you.

19

Multiple-Choice and Short-Answer Papers

Everything so far assumed a paper of long, multi-part problems marked on method. Plenty of engineering and physics courses don't set one. They set forty short questions, or a machine-marked multiple-choice section worth a third of the module.

Those papers reward a different set of decisions, and some of the habits from earlier chapters actively hurt on them.

19.1 What changes, and what doesn't




Figure 19.1: Long-answer papers versus short-answer and multiple-choice. Two of the earlier habits reverse completely.

The two dark rows are the reversals, and they're the ones that catch students who've trained on long papers.

There is no partial credit. A ninety per cent correct answer to a multiple-choice item scores exactly what a blank scores. This changes the entire economics of the paper.

Working earns nothing by itself. Write working because it stops you making errors, not because a marker will reward it. Nobody reads it.

What doesn't change: triage, the clock, and reading the question. Those matter more here, not less.

19.2 The economics of guessing

Work this out before the exam, not during it, because it decides how you treat every item you can't do.

Marking rule	What to do
No penalty for wrong answers	Answer every item. Never leave a blank.
$-\frac{1}{4}$ for wrong, 5 options	Guess if you can eliminate one option; skip if you can't.
$-\frac{1}{3}$ for wrong, 4 options	Guess if you can eliminate one; skip if you can't.
Full negative marking	Answer only what you can justify, plus eliminations of two or more.

The arithmetic behind the middle rows: with five options and a $-\frac{1}{4}$ penalty, a blind guess has expected value $\frac{1}{5}(1) + \frac{4}{5}(-\frac{1}{4}) = 0$. Break-even. Eliminate a single option and the expected value turns positive, so guessing is now correct. That's the whole calculation, and it means one crossed-out option is worth real marks.

Eliminating one option is worth more than it feels. Under the standard negative-marking schemes it flips a guess from break-even to profitable, so cross options out on the paper as you rule them out.

Most university papers, and most professional exams, have no penalty at all. On those, a blank answer is strictly worse than a random one, and leaving any item unanswered is simply throwing away expected marks.

19.3 The three-pass sweep

Don't work an MCQ section front to back. Sweep it three times.




Figure 19.2: Three passes over a 40-item section in 60 minutes. Nothing hard is attempted until everything easy is banked.

The reason for the passes is that every item is worth the same, so an item you can do in twenty seconds and an item that takes four minutes pay identically. Doing them in printed order means spending four minutes on item 7 while items 31 to 40, which you could have done in twenty seconds each, sit unattempted.

That's the same argument as Chapter 4, sharpened, because on a long-answer paper questions at least differ in mark value. Here they don't.

19.4 Techniques that only work here

Five moves that are legitimate on machine-marked papers and useless on a long-answer script.

Work backwards from the options. If the item asks for the solution of an equation and gives you five candidates, substituting is often faster than solving. Especially with ugly algebra.

Eliminate on dimensions. In physics and engineering, wrong options frequently have the wrong units. Check the dimensions of each option before doing any calculation and two of them usually die.

Eliminate on limiting behaviour. What should happen as $x \rightarrow 0$, or as the mass goes to infinity? Options that behave wrongly at the extremes are wrong everywhere.

Estimate first. An order-of-magnitude estimate frequently picks the answer without the calculation, because exam distractors are usually spread across magnitudes.

Watch the distractors. Wrong options aren't random. They're built from the specific mistakes students make: the sign error, the missing factor of two, the degrees-radians confusion, the reciprocal. So if your answer matches an option exactly, that's weak evidence you're right, and if it matches the option that's exactly double another one, look again.

FROM THE MARKER'S SIDE

Every distractor on a well-built paper is somebody's mistake, written down. When your answer lands on one, the question worth asking is which mistake it corresponds to.

19.5 Short-answer papers

Between the two formats sits the short-answer paper: forty questions, two or three marks each, answers written in a box.

The rules are a hybrid, and the important ones are:

- **The box is the answer.** Anything outside it may not be marked. If working is wanted, there'll be space for it.
- **Some marks are still method marks.** On a three-mark short-answer item there's often one M mark available. Write one line of setup even here.
- **Units and form still count**, exactly as in Chapter 16, and they're the commonest loss on these papers.
- **No salvage.** There's nowhere to put it and nobody reading for it. If you can't do it, guess something dimensionally sensible and move.

19.6 Where MCQ marks actually go

An audit of a twenty-item section, the same exercise as Chapter 3 but for this format. One student, four practice sections, eighty items, twenty-three wrong.

Why it was wrong	Items	
Ran out of time, guessed blind	7	30%
Walked into a designed distractor	6	26%
Misread the item (“not”, “least”)	4	17%
Arithmetic slip	3	13%
Didn’t know it	3	13%

Same shape as every other audit in this book. Thirteen per cent was knowledge; everything else was execution, and the biggest single cause was spending too long on hard items early and having to guess the last seven.

That’s the three-pass sweep, and it’s why the sweep matters more here than on a long paper, not less.

19.7 The words that flip an item

Short items hide their instruction in one word, and it’s usually mid-sentence where the eye skips it.

Word	What it does
not	inverts the whole item; circle it
least, smallest	you’ll compute the largest by habit
except	same as “not”, harder to see
approximately	an estimate is expected; don’t compute exactly
must	the answer has to hold always, so one counterexample kills an option
could	the answer only has to be possible
initially, finally	selects which state is being asked about
per	a rate, so a division you might have skipped

Circle the word as you read the item. It costs a second, and it is the single highest-frequency error on short-item papers after running out of time.

19.8 The answer sheet itself

A mechanical risk that has nothing to do with the subject, and destroys scripts every year.

If you're filling a separate machine-marked sheet, transfer answers in blocks of five, checking the item number each time. A single row offset partway through a section turns every subsequent answer wrong, and the machine has no way to notice you knew the material.

Two habits that prevent it:

- Answer on the question paper first, then transfer in blocks. Never transfer one at a time from memory.
- Every ten items, check that the number on the sheet matches the number on the paper. Ten seconds, five times, and the failure mode is eliminated.

Leave five minutes at the end for the transfer if you haven't been doing it as you go. Running out of time with thirty correct answers on the question paper and a blank answer sheet is the single most painful way to fail an exam, and it happens.

19.9 Four items, worked the fast way

Each of these is solvable the long way. Each is faster the short way, and on a machine-marked paper the short way is worth exactly the same.

Item 1. Substitute, don't solve.

Which value of x satisfies $x^3 - 6x^2 + 11x - 6 = 0$? A: -1 B: 1 C: 4
D: 5 E: 6

Factorising takes two minutes. Substituting B takes ten seconds: $1 - 6 + 11 - 6 = 0$. Done.

Item 2. Kill options on dimensions.

The period of a mass m on a spring of stiffness k is: A: $2\pi\sqrt{k/m}$ B:
 $2\pi\sqrt{m/k}$ C: $2\pi mk$ D: $2\pi\sqrt{mk}$ E: $2\pi m/k$

k is $\text{N m}^{-1} = \text{kg s}^{-2}$. So m/k has units s^2 , and only B gives seconds. Fifteen seconds, no physics recalled.

Item 3. Test the limit.

Two resistors R_1 and R_2 in parallel have combined resistance: A: $R_1 + R_2$ B: $\frac{R_1 R_2}{R_1 + R_2}$ C: $\frac{R_1 + R_2}{R_1 R_2}$ D: $\sqrt{R_1 R_2}$ E: $\frac{1}{R_1} + \frac{1}{R_2}$

Let $R_2 \rightarrow \infty$: an open branch, so the answer should tend to R_1 . B does. Nothing else does. Ten seconds.

Item 4. Estimate.

A 2 kW kettle raises 1.5 litres of water from 20 °C to 100 °C. Roughly how long? A: 8 s B: 50 s C: 4 min D: 25 min E: 2 hours

Energy $\approx 1.5 \times 4200 \times 80 \approx 500$ kJ. At 2 kW that's 250 s, so C. And notice the distractors are spread across magnitudes, which is exactly what makes the estimate sufficient.

19.10 Reading the distractors backwards

On a well-built paper, the wrong options aren't filler. Each one is a specific mistake, and recognising which mistake is a way of checking yourself.

The distractor	The mistake it catches
Exactly double the right answer	forgot a factor of $\frac{1}{2}$, or counted one way not both
The reciprocal	inverted a fraction, or read a gradient upside down
Right number, wrong unit	didn't convert prefixes
Right magnitude, wrong sign	dropped a minus, or wrong direction convention
The intermediate value	stopped one step early

So when your answer lands exactly on the option that's double another option, spend ten more seconds. You may have just walked into the trap the question was built around.

DO THIS BEFORE THE NEXT PAPER

Find out, this week, whether your paper carries a penalty for wrong answers, and how many options each item has. Then work out your

own guessing rule from the table above and write it on your protocol card. That's a decision you never want to be making at minute 55.

20

Open-Book, Take-Home, and Formula Sheets

Three formats that look like a relief and reliably produce worse marks than closed-book papers.

They produce worse marks for the same reason in each case: students prepare less, because the material is available, and then discover that availability is not the same as access.

20.1 Why open-book papers are harder

An open-book exam is not a lookup test. If it were, everybody would score full marks.

Examiners know you have the notes, so they don't ask questions the notes answer. They ask questions that require you to *select* and *apply*, and they set them at a length that assumes you aren't reading anything.




Figure 20.1: Where the time goes in an open-book paper, for a prepared and an unprepared student on the same three-hour paper.

Same paper, same notes. The prepared student spends 25 minutes looking things up and 155 solving. The unprepared one spends 95 looking up, has 85 left, and doesn't finish.

Open book rewards knowing where things are, not knowing that they exist. The preparation isn't reading the notes. It's indexing them.

20.2 Preparing for open book

Three artifacts, built in the week before, and all three are things you'd want anyway.

The index. One page listing every method, formula, and worked example, with the page number it's on. Sorted by *tell*, not by topic name, exactly as in Chapter 14. This is the single highest return item, and building it is itself excellent revision.

Tabs. Physical tabs on the six or eight pages you'll actually need. Colour-code by topic if you like. The goal is that no lookup takes more than fifteen seconds.

The worked-example sheet. For each major method, one complete worked example, on its own page, in your handwriting. In the exam you will not read theory. You will pattern-match against a worked example, so make sure the examples exist and are findable.

What not to do: bring everything. A student with four textbooks, a full set of notes and no index is slower than a student with one page and no books, because unbounded material means unbounded searching.

20.3 The lookup budget

Decide it in advance, like everything else.

Situation	Rule
You know the method, want to check a constant	look it up, 20 s
You know the method, can't recall the form	look it up, 40 s
You don't know which method	don't look think first, 60 s
Nothing after 60 s of thinking	one lookup in the index, then move on

The third row is the trap. Searching your notes for inspiration feels productive and is the main way open-book time disappears. The notes

cannot tell you which method a question wants, because the question doesn't say.

FROM THE MARKER'S SIDE

Open-book scripts split into two kinds. One has the same answers a closed-book student would give, only better presented. The other has three questions attempted, each one preceded by what is clearly twenty minutes of reading. The second student almost always says afterwards that they ran out of time.

20.4 Take-home papers

A take-home is an open-book paper with the clock removed, which changes two things and only two.

The standard rises. You have hours, so presentation, completeness and rigour are all expected to be higher. The marks that would have been partial credit in the room are now marks for a full argument.

Time management becomes self-imposed. Nobody will stop you spending nine hours on it, and plenty of students do, and the extra six hours earn almost nothing.

So impose the structure yourself:

- Budget the same hours the equivalent exam would have taken, plus half again for writing up properly.
- Do the triage exactly as you would in the room, and set wall-clock deadlines per question.
- Write the answers out once, properly, at the end. Don't submit your working.
- Stop at the budget. A take-home that consumes your whole week has cost you the other module you were supposed to be revising.

And read the collaboration rules carefully, twice, before you talk to anybody. They vary between courses, they're usually stricter than students assume, and the consequences are of a completely different kind from losing marks.

20.5 Formula sheets

If a formula sheet is provided, get a copy in week one and work with it all term. Two reasons.

You stop memorising what's printed. That's real capacity released, and students routinely spend hours learning things the exam was always going to hand them.

You learn where things are on it. A sheet you've used for twelve weeks is a tool. A sheet you meet in the exam is three pages of dense unfamiliar notation at the worst possible moment.

If no sheet is provided, build your own during the term and revise from it. One page per module, formulas grouped by the tell that triggers them, and the act of choosing what goes on it is most of the value.

Whether or not the exam provides a formula sheet, build one. If they provide it, learn its geography. If they don't, the page you built is the thing you revise from.

Either way, know what *isn't* on it. The gaps are where the memory work actually is, and they're a much shorter list than the whole syllabus.

DO THIS BEFORE THE NEXT PAPER

Find out this week whether each of your papers is open book, and whether a formula sheet is provided. If any is open book, start the index today: one page, sorted by tell, with page numbers. It doubles as revision and it's the difference between 25 minutes of lookup and 95.

Part VI

The Room

21

Nerves, Blanking, and the Restart

Every technique in this book assumes you can execute it in the room, and sometimes you can't, because your hands are shaking and question 1 has stopped making sense.

This chapter is about that.

21.1 What nerves actually do to you

Not “make you forget everything”. Something more specific and more manageable.

Exam stress narrows working memory. You can still do everything you could do yesterday, but you can hold fewer things in your head at once, and anything that requires juggling several items, a multi-step method with three intermediate values, is the first thing to go.

Which tells you the fix, and it's mechanical rather than emotional.

Under stress, write more down. Every intermediate value on the page rather than in your head. You're not being slow, you're moving the load off a faculty that's temporarily reduced.

This is also why the protocols in this book are written down rather than remembered. The four minutes, the checkpoints, the sweep. They're external because in the room your internal capacity is the thing that's degraded.

21.2 The blank

You read question 1 and nothing comes. Not “this is hard”. Nothing.

It's common, it's temporary, and it has a protocol.




Figure 21.1: The restart. Ninety seconds, and the point is to get your hand moving on something you can definitely do.

Step 3 is the one that does the work. Blanking is partly a stall in starting, and writing anything true breaks the stall. It doesn't have to be relevant. A student who writes down the definition of a derivative while blanked on a mechanics question has restarted their own hand, and the mechanics usually arrives within a minute.

Step 4 matters just as much. Don't sit with the question that broke you. Go and do something you can do, get twenty marks on the board, and return to it with the mood entirely changed.

21.3 The panic spiral, and how to break it

Different from a blank. This is the one where you're thirty minutes in, behind schedule, and the awareness of being behind is now consuming the capacity you needed to catch up.

The spiral has a shape: behind → anxious → slower → further behind. And it feeds on itself because the anxiety is about the clock, so every glance at the clock makes it worse.

Break it with a decision rather than with calm. Calm isn't available on demand; a decision is.

1. Stop the current question, mid-line if necessary.

2. Look at the paper and find the easiest unattempted part on it. Any question, any part.
3. Do that one.

Ten minutes of collecting marks resets both the schedule and the mood, in that order. What doesn't work is deciding to feel calmer, or trying to speed up, which produces the arithmetic slips that make everything worse.

21.4 Physical things that help, unglamorously

- **Long exhale.** Longer out than in, four or five times. It's the only quick physiological intervention that reliably works.
- **Unclench your hand.** People grip the pen hard enough to cramp, and the cramp then becomes another thing to worry about.
- **Feet flat, sit back, shoulders down for five seconds.** Costs nothing.
- **Water.** Take some in, drink it.
- **Don't look at other people.** They're writing fast because they're writing fast, not because they're right. Some of them are producing three pages of confident nonsense.

21.5 The night before and the morning of

Most exam anxiety is manufactured in the twelve hours before, and it's the most controllable part.

Stop revising early. Get the sleep, because tiredness and anxiety are nearly indistinguishable from the inside, and the one you can fix is the tiredness.

In the morning, don't do a practice question. If it goes badly you've started the day with evidence you're not ready, and forty minutes later you're in the room carrying it.

And arrive with your protocol card, because having a first move is most of what stops the blank in the first place. A student who sits down knowing that the first four minutes are triage has something to do with their hands immediately, and the nerves have nowhere to land.

FROM THE MARKER'S SIDE

The scripts that show real panic have a signature: the first question is laboured and over-long, then everything speeds up and the errors multiply. Almost always the student was fine and lost the first twenty minutes to the feeling of not being fine.

21.6 The kit

Unglamorous, and every item on this list has cost somebody marks.

Item	The failure it prevents
Two working pens	one dies at minute 40 and you spend six minutes borrowing
The calculator you rehearsed with	unfamiliar menus, wrong default mode
Spare batteries, or a spare calculator	it dies, and some halls won't lend
A watch, on the desk	wall clocks are often absent or behind you
ID and entry card	twenty minutes of your morning at the door
Water, in the permitted bottle	
Ruler, for graphs	a freehand best-fit line loses a mark
Protocol card, read then put away	the first move is decided before you sit down

Assemble it the night before, not the morning of. The morning is for eating and reading the card.

21.7 The first sixty seconds in the seat

Before the paper is turned over, there's a minute of dead time. Use it.

- Watch on the desk, face up.
- Calculator on, mode checked, and the mode written on your hand or the desk card if the rules allow.

- Two long exhales.
- Recall the first move: *four minutes, arithmetic, A/B/C, order, verbs.* Nothing else.

That last one is what stops the blank. A student who sits down knowing exactly what their hands will do for the first four minutes has nowhere for the panic to land, and the difference between that and staring at question 1 is the difference between a good start and twenty lost minutes.

21.8 If it goes badly anyway

Sometimes it does.

Finish the paper. Run the sweep. Collect the cheap marks, because a paper that's gone badly is exactly where the sweep is worth most and exactly where students give up with fifteen minutes left.

Then go and do the post-mortem in the next chapter, once, and stop. A bad paper is not a verdict on you, and it's certainly not a reason to skip preparing for the next one, which is what the walk home usually suggests.

21.9 Worked: a restart that worked

A student, first paper of the season, opens the booklet and cannot read question 1. The words are English and nothing means anything.

Clock	What happens
0:00	Reads Q1. Nothing. Reads it again. Still nothing. Heart rate climbing.
0:02	Stops reading it. Recognises the blank.
0:02–	Four long exhales, longer out than in. Puts the
0:03	pen down, sits back for five seconds.
0:03–	Writes something true at the top of the page:
0:04	the definition of a derivative, from the course. Nothing to do with Q1.
0:04–	Runs the four-minute triage as normal. Q1 is
0:08	still opaque. Q3 and Q5 are clearly A.
0:08–	Does Q5. It goes fine. 16 marks on the board.
0:31	
0:31	Returns to Q1 out of curiosity. It's a substitution question and the substitution is obvious.

The blank lasted about two minutes and cost nothing, because the protocol existed and the student didn't sit in it.

What would have happened without the protocol: fifteen minutes on Q1, increasing panic, an eventual move to Q2 at minute twenty with the day already framed as a disaster.

The blank is a stall in starting, not a loss of knowledge. Writing anything true restarts your hand, and doing a question you can do restarts everything else.

21.10 Things that make the room worse, and are avoidable

Small, unglamorous, and each one costs marks on somebody's paper every year.

- **Arriving at the last minute.** Adrenaline is already high; rushing adds to it and there's no way to bleed it off once you're seated.
- **Arriving forty minutes early.** You'll absorb the corridor, where somebody is always confidently explaining a topic you've never heard of.
- **A last-minute practice question.** If it goes badly you've started the day with evidence you're not ready. There is no upside.
- **An unfamiliar calculator.** Borrowing one, or buying a new one for the exam, is a genuine risk. Use the one you rehearsed with.
- **Comparing papers at the interval.** Between two exams on the same day, the corridor conversation is pure downside for the second one.
- **Watching the room.** People write fast for lots of reasons, including being wrong at speed.

21.11 Two exams in one day

Increasingly common, and it needs its own plan.

The first paper is over. Do not do the post-mortem. Do not discuss it. The post-mortem in Chapter 22 is explicitly a next-day activity, and the reason is that a bad post-mortem between papers will cost you more in the second exam than any revision could earn.

Between them: eat something, walk, read your protocol card and the method map for the second subject. Nothing else. Twenty minutes of quiet

is worth more than twenty minutes of revision, because the limiting factor in the afternoon is attention, not knowledge.

DO THIS BEFORE THE NEXT PAPER

Write your restart protocol on the same card as your timing protocol. Four lines. Then, in your next timed practice, deliberately trigger it: start with the hardest question on the paper, let yourself blank, and run the four steps. Practising the recovery is what makes it available when you need it.

After the Paper

Two things happen after an exam, and only one of them is useful.

The first is the walk home, where every question you got wrong arrives in sequence, along with the answers, and you conduct a detailed prosecution of yourself. That one is compulsory, apparently, and worth nothing.

The second is a twenty-minute post-mortem, done once, that makes the next paper better. Almost nobody does it.

22.1 The twenty minutes, within a day

Not immediately after. Give it a few hours, or until the next morning, when the material is still fresh and you're no longer in the room.

Then, on one page:




Figure 22.1: The post-mortem. Three columns, twenty minutes, and the right-hand one is the only one that changes anything.

Then put the right-hand column straight into your error log, and you're done. Twenty minutes.

22.2 What not to do

Don't reconstruct your whole script. You can't, accurately, and the attempt produces a version of the paper that's worse than what you wrote, because you only remember the parts that went badly.

Don't compare answers with everyone. A short, factual comparison with one person is fine if you find it useful. A group post-mortem in a corridor is a machine for manufacturing certainty about marks nobody has, and the confident student in the middle is often wrong.

Don't estimate your mark. Students are terrible at it in both directions, and the number you invent will follow you around for six weeks until the real one arrives.

Don't skip the next preparation because of this one. The single most expensive consequence of a bad paper is the two days of preparation lost to feeling bad about it, especially when the next exam is in seventy-two hours.

The purpose of the post-mortem is the next paper, not this one. If a question doesn't change what you'll do next time, it doesn't belong in it.

22.3 When the marks come back

Weeks later, and this is the more valuable review, because now you have the actual script and the actual marks.

Get your script back if your institution allows it. Then do the seven-bucket audit from Chapter 3 properly: every lost mark assigned to one of the six leaks or to "didn't know it".

That's the audit you can't do from memory, because from memory everything looks like "didn't know it". With the script in front of you, it usually isn't.

FROM THE MARKER'S SIDE

The most common reaction when a student sees their own marked script is surprise at how many marks went to execution rather than knowledge. Unstated answers. Method done mentally. A whole part missed on the back of a page. It's rarely the disaster they remembered, and it's rarely for the reasons they assumed.

22.4 What actually improves between papers

Look back over the book and notice what's been claimed and what hasn't.

Nothing here makes you better at your subject. That's what your courses, your problem sets, and your reading do, and there's no shortcut.

What this book does is close the gap between what you know and what you score. For most students that gap is five to fifteen percentage points, and it closes through a small number of decisions made in advance:

- Four minutes of triage before writing.
- Wall-clock deadlines and three checkpoints.
- Method lines written down, always.
- A plan sentence before each part.
- The three-minute no-ink rule and a clean abandonment.
- Salvage on everything you leave.
- Ten seconds of checking per part.
- Five minutes reserved for the final sweep.

Eight things. None of them requires you to know any more than you know right now, and all eight are on the card in Appendix D.

22.5 If the result is bad

Sometimes the mark comes back and it's genuinely poor. Three things are worth knowing, in order.

Get the script and do the audit. Not for comfort. Because the difference between "I need to learn this subject properly" and "I lost twenty marks to timing" is the difference between six months of work and two weeks, and you cannot tell which from the mark alone.

Find out what's actually available. Resits, capped or uncapped. Whether the module can be repeated. Whether there were circumstances the department should have known about. These have deadlines, often short ones, and students routinely miss them while feeling bad.

Then treat it as data. One paper is a noisy measurement of what you know, taken on one morning, under conditions designed to be stressful. It's information about your execution and your preparation. It isn't information about whether you belong in the subject, however much it feels like it on the day.

A bad mark tells you a lot about a process and almost nothing about a person. The audit turns it into a task list, which is the only form in which it's any use.

22.6 The last thing

The walk home will still happen. The answers will still arrive twenty minutes too late, and it'll still be annoying.

But there's a version of that walk where the answers arriving late are answers to questions you *reached*, on a paper where you spent your time in proportion to the marks and wrote down what you knew. That's the best any of us manages, and it's a long way from where most students start.

Go and get the marks you've already earned.

22.7 The post-mortem between papers

If the next exam is in three days, the post-mortem changes shape. It gets shorter and it gets ruthlessly forward-looking.

Ten minutes, not twenty, and only two questions:

1. **Which protocol step did I skip?** Not which question went badly. Which step. Triage, checkpoints, salvage, sweep.
2. **What will I do differently in seventy-two hours?** One thing. Written on the protocol card.

Then stop, and go and prepare for the next one.

The reason for the compression is that a full post-mortem three days before another paper produces rumination, and rumination costs more than the insight is worth. The full version, with the script, happens when the marks come back.

FROM THE MARKER'S SIDE

The most expensive consequence of a bad paper is almost never the paper itself. It's the two days of preparation lost to feeling bad about it, with the next exam already scheduled.

22.8 What improvement actually looks like

Across a year, tracked properly, the pattern is consistent and worth knowing in advance so you can recognise it.

Paper	What changes	Typical gain
1st	you run the four minutes for the first time	3–5 marks
2nd	you abandon a question cleanly	4–8 marks
3rd	salvage and sweep become automatic	3–6 marks
4th	the log’s top two errors stop happening	3–5 marks
cumulative		10–20 marks

Those aren’t independent, and you won’t get the full sum. But the shape is right: the gains come in steps, each one from a single protocol becoming automatic, and each one is available without learning any more of the subject.

Improvement between papers comes from one protocol at a time becoming automatic. Trying to install all eight at once is why the first attempt usually fails and the second usually works.

DO THIS BEFORE THE NEXT PAPER

Do the twenty-minute post-mortem after your next exam, on one page, three columns. Then put the right-hand column into your error log before you open a book for the next paper. That single habit compounds across every exam you’ll ever sit.

Part VII

Toolkit

A

Ten Annotated Scripts

Ten real answers, each with the marks it scored, the marks it didn't, and the mark scheme line that decided it.

None of these students lacked the knowledge. Every one of them lost marks to execution, and the total across the ten comes to about sixty marks.

The scripts are compressed. The structure of what was written, and every mistake in it, is exactly as it appeared.

A.1 Script 1: the mental substitution

Q. Evaluate $\int_0^2 xe^{x^2} dx$, giving your answer to three significant figures. [5]

BEFORE

$$\begin{aligned}\int_0^2 xe^{x^2} dx &= \frac{1}{2}e^{x^2} \text{ evaluated 0 to 2} \\ &= \frac{1}{2}(e^4 - 1) = 26.3\end{aligned}$$

Scored 2 of 5.

The scheme:

M1	correct substitution stated, $u = x^2$, $du = 2x dx$
M1	limits changed to 0 and 4, or reverted correctly at the end
M1	$\frac{1}{2} [e^u]$ or equivalent
A1	$\frac{1}{2}(e^4 - 1)$
A1	26.8 (3 s.f.)

The antiderivative is right, so the third M mark is safe. The first two were done in the student's head and never appeared, so they can't be awarded. And the arithmetic is wrong: $\frac{1}{2}(e^4 - 1) = 26.799$, which is 26.8, not 26.3. Both A marks go.

What makes this painful is that the student knew the method completely. Here is the same understanding, written for the scheme:

AFTER

Let $u = x^2$, so $du = 2x dx$, giving $x dx = \frac{1}{2} du$.

When $x = 0$, $u = 0$. When $x = 2$, $u = 4$.

$$\int_0^2 x e^{x^2} dx = \frac{1}{2} \int_0^4 e^u du = \frac{1}{2} [e^u]_0^4 = \frac{1}{2}(e^4 - 1)$$

$$= 26.799 \dots = 26.8 \text{ (3 s.f.)}$$

Twenty seconds more writing, and now three marks are guaranteed before the calculator is touched. Note the unrounded value shown once, which protects the final accuracy mark.

FROM THE MARKER'S SIDE

This is the commonest failure in the whole appendix, and the cheapest to fix. Two lines of substitution would have banked three marks regardless of what happened to the arithmetic afterwards.

A.2 Script 2: the unstated answer

Q. Find the stationary points of $f(x) = x^3 - 6x^2 + 5$. [4]

BEFORE

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

$$= 3x(x - 4)$$

$$3x(x - 4) = 0$$

Scored 3 of 4.

M1 for differentiating, A1 for the correct derivative, M1 for setting it to zero. Then the script stops.

The final A mark wanted the stationary points, and there aren't any on the page. The student solved the equation in their head, decided it was obvious, and moved to the next question.

Two separate things went wrong, and the second one is worth more than the first.

AFTER

$$f'(x) = 3x^2 - 12x = 3x(x - 4) = 0, \text{ so } x = 0 \text{ or } x = 4.$$

$$f(0) = 5 \text{ and } f(4) = 64 - 96 + 5 = -27.$$

Stationary points at (0, 5) and (4, -27).

The question said *points*. Points have two coordinates. A script that stops at " $x = 0$ or $x = 4$ " has answered a different question, and on many schemes that's the difference between the A mark and nothing.

Ten seconds. Read the noun in the question, not just the verb.

A.3 Script 3: the ladder ignored

Q(a). Show that $\frac{d}{dx}(x \ln x) = \ln x + 1$. [2]

Q(b). Hence find $\int \ln x \, dx$. [4]

Part (a) was done correctly and scored 2 of 2.

In part (b) the student started integration by parts from scratch, wrote $u = \ln x$, $dv = dx$, got three lines in, made a sign error, ran out of time, and left it.

Scored 2 of 6.

Part (b) was four marks for rearranging part (a). Integrate both sides of the result you were just given:

AFTER

From (a), $\frac{d}{dx}(x \ln x) = \ln x + 1$. Integrating both sides,

$$x \ln x = \int \ln x \, dx + \int 1 \, dx = \int \ln x \, dx + x.$$

Therefore $\int \ln x \, dx = x \ln x - x + C$.

Ninety seconds, four marks.

The word that carried all of this was *hence*. It isn't a connective, it's an instruction, and it means the previous part is the tool for this one. A student who reads it does part (b) in a minute and a half. A student who skims it does it from scratch, slowly, and often doesn't finish.

A.4 Script 4: never reached

A three-hour paper, six questions, worked strictly in the printed order.

	Scored	Available	
Q1	15	15	clean
Q2	14	16	one slip
Q3	16	16	clean
Q4	13	15	clean
Q5	3	20	a page and a half, unfinished
Q6	0	18	blank
	61	100	

Questions 1 to 4 scored 58 out of 62, which is first-class work.

Question 5 took fifty minutes and produced three marks. Question 6 was never seen, and it was a standard eigenvalue problem this student could have done in about twenty minutes.

FROM THE MARKER'S SIDE

This is the script that made me want to write this book. Nothing in the subject knowledge produced that mark. The student knew enough for a first and walked away with a 2:2, and the entire difference is that they met question 6 for the first time with no minutes left. Four minutes of triage would have been worth roughly fifteen marks.

The fix is Chapter 4 and Chapter 6 together: read everything first, and leave question 5 at the twenty-five minute mark with salvage lines instead of fighting it for fifty.

A.5 Script 5: the calculator in the wrong mode

A mechanics paper. Six answers required trigonometric evaluation, and the calculator was in degrees when every formula on the script assumed radians.

Scored 22 of 40 on the affected parts.

The method marks all survived, which is why the damage was contained rather than catastrophic. Every accuracy mark that depended on a numerical value was lost, six times.

Cost of prevention: three seconds at the start of the paper, and two letters written at the top of the script.

*Check the calculator mode in the first minute and write **RAD** or **DEG** at the top of your script. A mode error doesn't cost one mark. It silently corrupts every trigonometric answer on the paper.*

The student found out on the walk home, which is when everybody finds out.

A.6 Script 6: the blank “show that”

Q(a). Show that the series $\sum_{n=1}^{\infty} \frac{1}{n(n+1)}$ converges. [5]

Q(b). Given that it converges, find its sum. [6]

Q(c). Hence deduce the value of $\sum_{n=1}^{\infty} \frac{1}{n(n+1)(n+2)}$. [4]

The student couldn't construct the convergence proof in (a), so they left the whole question.

Scored 0 of 15.

Parts (b) and (c) were both doable. Part (b) is a telescoping sum: write $\frac{1}{n(n+1)} = \frac{1}{n} - \frac{1}{n+1}$, and the partial sums collapse to $1 - \frac{1}{N+1} \rightarrow 1$. That's a standard manipulation this student had done in class.

And crucially, part (b) hands you convergence anyway. It says *given that it converges*.

The one line that unlocks fifteen marks:

AFTER

(a) I could not complete the convergence proof.

(b) Using the given result that the series converges: $\frac{1}{n(n+1)} = \frac{1}{n} - \frac{1}{n+1}$,
so the partial sum is $S_N = 1 - \frac{1}{N+1}$, and $S_N \rightarrow 1$

Nobody deducts for admitting you couldn't do part (a). The result is printed in the question, and it's yours to use.

A.7 Script 7: good salvage

The counter-example, and the reason salvage gets a chapter.

Q. Use the divergence theorem to evaluate the flux of $\mathbf{F} = (x^2, y^2, z^2)$ out of the closed surface of the tetrahedron bounded by $x + y + z = 1$ and the coordinate planes. [8]

AFTER

By the divergence theorem, $\oint_S \mathbf{F} \cdot d\mathbf{S} = \int_V (\nabla \cdot \mathbf{F}) dV$.

$\nabla \cdot \mathbf{F} = 2x + 2y + 2z$.

V is the tetrahedron bounded by $x + y + z = 1$ and the coordinate planes, so

$$\int_0^1 \int_0^{1-x} \int_0^{1-x-y} (2x + 2y + 2z) dz dy dx.$$

Inner integral: $2xz + 2yz + z^2$ evaluated at $z = 1 - x - y$.

Ran out of time evaluating the outer integrals.

Scored 5 of 8, on a question that was never finished.

M1 for naming the theorem, M1 for stating it correctly, A1 for the divergence, M1 for the correct region and limits, M1 for beginning the evaluation. Only the two accuracy marks for the final value were lost.

That's the salvage stack from Chapter 10 working exactly as designed, and it took the student about ninety seconds to write once they knew they weren't going to finish.

A.8 Script 8: the answer to a different question

Q. A particle is projected from ground level with speed 20 m s^{-1} at 30° to the horizontal. Find the time at which it returns to its starting height. [5]

BEFORE

$$u_y = 20 \sin 30^\circ = 10$$

$$\text{At the top, } v_y = 0: 0 = 10 - 9.8t$$

$$t = 1.02 \text{ s}$$

Scored 3 of 5.

Every line is correct. The student computed the time to *maximum height*, which is half the answer, literally.

The method marks survived because resolving the velocity and applying the equation of motion were both correct and both visible. The accuracy marks went, because the number doesn't answer the question.

The fix is four seconds long: underline the verb and its object. *Find the time at which it returns to its starting height.* With that underlined, the doubling is obvious, and the final line becomes

AFTER

Time to maximum height = 1.02 s. By symmetry, the total time to return to the starting height is $2 \times 1.02 = 2.04 \text{ s}$.

FROM THE MARKER'S SIDE

It's a strange thing to mark. The physics is entirely right and the answer is entirely wrong, and there's nothing in the rubric that lets me be generous about it. The student solved a question the paper didn't ask.

A.9 Script 9: no units, six times

An engineering paper where every numerical answer was correct and not one of them carried a unit.

Part	Written	Wanted
2(a)	4.7	4.7 kN
2(c)	12.3	12.3 mm
3(b)	0.86	0.86 W
3(d)	240	240 V
4(a)	1.6	1.6 m s^{-2}
4(c)	55	55 %

Lost 6 marks, one per part.

There is no cheaper six marks anywhere in this book, and no part of it requires any additional knowledge. The final sweep in Chapter 7 exists mostly for this: two minutes at the end, running down every numerical answer on the paper, asking one question.

Worth noting that the last one, the efficiency, needed a per cent sign rather than an SI unit. A dimensionless answer still has a form the scheme expects.

A.10 Script 10: the graph rushed

A data question on a physics paper. The student measured well, plotted correctly, computed the gradient from the line, and interpreted it correctly as $1/g$.

Scored 5 of 9.

Everything lost was presentation, and every item is on the six-point checklist in Chapter 17:

What the scheme wanted	What the script had
axes labelled with quantity <i>and unit</i>	“time” and “extension”
scale using at least half the grid	points crammed into a third
gradient triangle drawn on the graph	gradient computed off-page
gradient quoted with its unit	a bare number

Four marks, none of which required any physics at all. The two marks for the actual result were both earned.

FROM THE MARKER'S SIDE

Data questions are the most predictable marks on any physics paper, because they're marked against nearly the same checklist every year. Students treat them as a chore and rush them, which is exactly backwards.

A.11 What the ten have in common

Sixty marks, across ten scripts. Here is where they went.

Cause	Scripts	Marks
Time allocation: never reached, or ran out	4, 6	32
Method not written down	1, 5	21
Answer not stated, or the wrong one	2, 8	3
Ladder or instruction word ignored	3	4
Units, figures, presentation	9, 10	10
Didn't know the material	none	0

That last row is the point of the appendix, and it's why this book exists.

Ten scripts. Sixty marks. Not one of them lost because the student didn't know enough.

B

Two Papers, Start to Finish

Everything in this book, run over two complete papers, minute by minute.

The first is a 100-mark, three-hour mathematics paper, answer all questions. The second is a 90-minute engineering paper with a multiple-choice section, a data question, and two long questions. Neither goes perfectly, because papers don't.

Read these once now, and again the week before your next exam.

B.1 Paper 1: mathematics, 100 marks, 180 minutes

Six questions. Answer all. The student is competent and has done four timed rehearsals.

0:00 to 0:04 The four minutes

Writes at the top of the script:

100 marks / 180 min. Less 4 triage, less 5 sweep = 171 working. **1.7 min per mark.** RAD. Checkpoints: 45 / 90 / 135 → 22 / 48 / 74 marks.

Then reads all six questions, marking each in the margin:

Q	Marks	First read
1	15	A standard integration, know the first line
2	16	B series, recognise the shape
3	16	C nothing comes
4	15	A eigenvalues, done twenty of these
5	20	B ODE, needs a minute to see the route
6	18	A vector calculus, divergence theorem

Order written in the margin, with wall-clock finishes, exam starting at 9:30:

Q1 by 10:00 • Q4 by 10:26 • Q6 by 10:57 • Q3(a) by 11:04 • Q2 by 11:32 • Q5 by 12:07 • sweep 12:20

Note the C question sitting in the middle rather than at the end, with seven minutes allocated for its part (a) only.

0:04 to 0:30 Q1, and it goes well

Underlines the verbs. Writes a plan line on each part: “substitution $u = 1 + x^2$ ”, “by parts, $u = \ln x$ ”. Structural lines written, limits changed on their own line, answers stated with units where relevant.

Finishes at 9:58, two minutes early. Writes in the margin: **9:58, 15 marks.**

0:30 to 0:56 Q4, one slip

Eigenvalues come out fine. In part (c) the student computes an eigenvector and gets a vector that doesn’t satisfy the original equation when checked.

Ten-second check caught it: substitution back into $A\mathbf{v} = \lambda\mathbf{v}$ fails. Sign error two lines up. Fixed in forty seconds, which is far cheaper than the follow-through damage would have been in part (d).

10:24, 30 marks. First checkpoint passed at 10:15 with 26 marks attempted against a target of 22. Slightly ahead.

0:56 to 1:31 Q6, and the first real problem

The divergence theorem part goes smoothly. Then part (d) asks for the flux through an open surface, and the student can't see how to close it.

At 11:57 on the wall clock, eighteen minutes into a question budgeted for thirty-one, they've written nothing for three minutes.

The three-minute no-ink rule fires. They stop.

Salvage lines, ninety seconds:

(d) The surface is not closed, so the divergence theorem does not apply directly. Intend to close it with the disc $z = 0$, apply the theorem to the closed region, then subtract the flux through the disc. Could not evaluate the disc integral.

Leaves a third of a page. Writes "*close with disc at $z = 0$, then subtract*" as a note to self. Moves on at 10:57, on schedule.

That salvage later scored 3 of the 6 marks on part (d).

1:31 to 1:38 Q3, part (a) only

The C question. Seven minutes, part (a) only, which asks for a definition and a small verification. Writes both. Three marks, banked, on a question they can't otherwise do.

11:04, 51 marks. Second checkpoint (90 min, target 48) passed at 11:00 with 48. On pace.

1:38 to 2:12 Q2, which turns out to be an A

The series question, marked B on the first read, opens immediately once they start. Twenty-eight minutes for sixteen marks, all parts complete.

11:32, 67 marks.

2:12 to 2:47 Q5, the twenty-marker

The ODE. Parts (a) to (c) go well. Part (d) requires a particular integral the student half-remembers.

At 12:04 they're three minutes from the deadline they set, with part (d) unfinished. They write the trial solution they believe is right, note that they couldn't determine the constants, and stop.

12:07, 84 marks attempted. Third checkpoint (135 min, target 74) passed at 11:45 with 79. Comfortable.

2:47 to 3:00 The sweep

Thirteen minutes, more than planned, because Q2 went fast.

3 min	Blank parts. Q3(b) and Q3(c) get one line each: the relevant theorem stated. Two marks.
2 min	Unstated answers. Q1(c) and Q5(b) had working and no final line. Two marks.
2 min	Units and figures. Q4(d) answer given to seven figures; rounded to three. One mark.
2 min	Salvage. Q5(d) gets two more lines on the method.
4 min	Back to Q6(d). Sets up the disc integral. Doesn't finish, but the setup was worth two marks.

Final: 78 of 100.

Worth noticing where that came from. The student's raw subject knowledge was worth perhaps 70 on this paper. The protocol added roughly eight marks: three from the salvage on Q6(d), three from Q3's part (a), and several from the sweep. And it prevented the much larger loss that would have come from spending forty minutes on Q6(d) and never reaching Q5.

B.2 Paper 2: engineering, 90 minutes, mixed format

Structure: twenty multiple-choice items (20 marks, no penalty), one data question (15 marks), two long questions (15 each). Total 65 marks in 90 minutes.

0:00 to 0:03 Triage

65 marks / 90 min, less 3 triage, less 5 sweep = 82 working. **1.26 min per mark.** DEG for this paper. No penalty for wrong MCQ answers, so **answer all twenty.**

Budgets: MCQ 25 min, data question 19, long questions 19 each.

Reads the two long questions only. Q22 is an A, Q23 is a B. Decides the order: MCQ first (because it's mechanical and warms you up), then Q22, then the data question, then Q23.

0:03 to 0:28 The MCQ section, three passes

Pass 1, 14 minutes. Answers thirteen items on sight. Marks five as needing work and two as no idea.

Pass 2, 9 minutes. Works the five. Two of them fall to a dimensional check on the options: three of the five candidates had the wrong units, leaving a coin flip that a quick estimate settled.

Pass 3, 2 minutes. The remaining two. Eliminates one option in each on limiting behaviour (as $R \rightarrow \infty$, the answer must go to zero, which kills two options), then guesses. One of the two guesses was right.

Answers transferred to the machine sheet in blocks of five, with the item number checked at each block.

17 of 20.

0:28 to 0:47 Q22, the A question

Straightforward. Structural lines written, units carried through the working, and a dimensional check at the end caught one wrong formula before it propagated: an answer in newtons where the question wanted joules.

14 of 15.

0:47 to 1:06 The data question

Runs the six-item graph checklist deliberately:

-
- | | |
|---|--|
| 1 | Axes labelled " T^2 / s^2 " and " L / m ". |
| 2 | Scale chosen to use most of the grid. |
| 3 | Points plotted as small crosses. |
| 4 | Best-fit line drawn with a straight edge, not through the origin. |
| 5 | Gradient triangle drawn on the graph, spanning most of the line, using points <i>on the line</i> . |
| 6 | Gradient interpreted as $4\pi^2/g$, with units, and g computed. |
-

The improvement question at the end gets the three-clause answer: source, effect with a number, fix. "Reaction time adds about 0.2 s to each

timing; over a single swing of roughly 1 s that is a 20% uncertainty, so time 20 swings and divide.”

14 of 15.

1:06 to 1:25 Q23, and it goes badly

The B question. Part (a) is fine. Part (b) requires a transformation the student can't recall, and at 1:14 the no-ink rule fires.

Salvage: names the method, states the governing equation, sets up the boundary conditions, and writes where it broke. Then moves to part (c), which turns out not to depend on (b) at all, and completes it.

8 of 15, on a question that felt like a disaster.

1:25 to 1:30 The sweep

Two blank MCQ items found and filled (no penalty, so never leave blanks). One unstated answer in Q22 written out. Every numerical answer checked for a unit, and two were missing.

Final: 55 of 65, about 85%.

B.3 What both papers have in common

Neither student was brilliant on the day. Both hit a question they couldn't do, and both had at least one arithmetic scare.

The difference between these scripts and the ones in Appendix A is entirely in the decisions:

- Four minutes spent reading before writing, on both papers.
- Every question given a wall-clock deadline before it was started.
- The three-minute no-ink rule obeyed, twice, both times against the instinct to keep going.
- Salvage written on everything abandoned, worth about five marks across the two papers.
- The final sweep protected and used, worth another seven.
- Checks that cost ten seconds catching two errors that would have cost far more.

None of it was subject knowledge. All of it was decided in advance and rehearsed, which is the only reason it survived contact with the room.

C

The Marker's Shorthand

Mark schemes are written in a compressed notation, and it's the reason most students bounce off them. Here it is, decoded.

Conventions vary slightly between institutions and boards. The structure doesn't.

Mark types

M1	Method mark. Awarded for a correct method, whether or not the arithmetic works out. Must be visible on the page. This is where most of the paper lives.
A1	Accuracy mark. For the correct value or expression. Almost always dependent on the corresponding M mark: no method, no accuracy, even if the answer is right.
B1	Independent mark. Stands alone. A stated definition, a correct unit, a sketch, an assumption named. Doesn't depend on anything else being right, and is the most commonly skipped mark on any paper.
dM1	Dependent method mark. Only available if the previous M mark was earned. Skipping the setup line to save time can therefore cost several marks, not one.
M1A1	Both together on one line: correct method <i>and</i> correct result.

Instructions to the marker

ft	Follow through. Award the mark using the candidate's own earlier (wrong) value. This is why you keep going after an error instead of scribbling everything out.
ecf	Error carried forward. The same thing, different name.
oe	Or equivalent. Any algebraically equivalent form scores. You don't have to match the printed form exactly.
cao	Correct answer only. No follow-through here. The exact value or nothing.
awrt	Anything which rounds to. The scheme accepts a range. Common on numerical answers.
ISW	Ignore subsequent working. Once the right answer appears, later mess doesn't lose it. Not universal, so still finish with a clean statement and stop.
SC	Special case. A partial award for a specific common near-miss, written into the scheme in advance.
NB	A note to the marker about something to watch for. Usually the trap the question was built around.

Phrases, and what they permit

The scheme says	Which means
"any valid method"	the lecture method isn't required
"their quadratic"	follow-through on your own earlier expression
"must be seen"	doing it mentally scores nothing
"allow one slip"	a single arithmetic error is tolerated here
"condone missing units"	unusual; usually they aren't condoned
"do not accept"	the specific wrong form the question targets
"seen once"	state the formula explicitly, early, and it counts

What this adds up to

Read those three tables together and the strategy of this whole book falls out of them.

- Method is most of the paper, and method must be *seen*.
- One arithmetic error usually costs one mark, not a question, as long as the working continues visibly.
- B marks are free and they're the ones students skip: the unit, the definition, the assumption, the sketch.
- "cao" is rare. Most of the paper forgives you, if you give it something to forgive.

The shorthand is the examiner telling you, in advance, exactly what they will pay for. It's published, it's free, and almost nobody reads it.

D

The Protocol Card

One page. Copy it onto a card, read it the morning of the paper, and run it.

BEFORE YOU WRITE ANYTHING 4 minutes

0:00 Write the arithmetic.

Total marks. Total minutes. Minutes per mark. Subtract 5 for the final sweep. Check the calculator mode and write RAD or DEG.

0:30 Read every question. Mark A / B / C.

A = can start now. B = probably. C = can't start. You are sorting on *starting*, not finishing.

2:30 Write the order, with wall-clock end times.

A questions first, easiest first. One C in the middle. "Q4 by 10:22", never "Q4 gets 25 min".

3:30 Underline the verb and every constraint.

Show / find / prove / hence / deduce / state / sketch.

Every part, before you start it (30 seconds)

1. Underline the verb and the constraints.
2. What have I been given, including from earlier parts?
3. Write one sentence naming the method.

Every part, while you work

- Write the **structural** lines: formula, substitution, setup, changed limits, theorem cited, case split.
- Skip the mechanical algebra.
- End with the answer stated explicitly, once, with its unit.

Every part, before you move on (10 seconds)

Sign. Magnitude. Units. Special case. Did I answer the question that was asked?

Checkpoints

At **45**, **90**, and **135** minutes: compare marks attempted against target. **More than 8 marks behind means move on now**, mid-question if necessary.

When to abandon

- Over budget by 50%? Go to the next test.
- Written nothing for 3 minutes? **Stop**. You're staring, not working.
- One specific thing left to try? 3 more minutes, then stop.
- Otherwise: salvage and move.

Leaving properly: salvage lines, no crossing out, leave a third of a page, one note to yourself saying where you got stuck.

Salvage stack

1. The method, named.
2. The formula or theorem, stated in full.
3. The setup, with limits.
4. Anything you computed.
5. One line saying where it broke.

Last ten minutes, in this order

1. Blank parts: state the definition or formula. (3 min)

2. Missing final answers: write the answer line. (2 min)
3. Units and significant figures, every numerical answer. (2 min)
4. Salvage lines on everything abandoned. (2 min)
5. Back to the hard question, with whatever is left.

If you blank

Stop reading it. Two long exhales. Write something true. Move to your easiest A question and come back.

E

The Error Log Template

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

The table

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

The fourth column is the only one that changes anything. Phrase it as an instruction you could follow while tired, not as an intention.

Not a rule	A rule
be more careful with signs	check the sign of every answer before moving on
manage time better	3 minutes with no ink means stop
show more working	write the substitution on its own line, always
read the question properly	underline the verb before starting each part
don't rush the end	alarm at $T - 10$; run the five sweep steps

The seven buckets

When reviewing a paper, every lost mark goes in exactly one bucket:

1. Never reached the question.
2. Ran out of time inside a question.
3. No method shown.
4. Arithmetic slip, uncaught.
5. Misread the question.
6. Units, figures, or final statement.
7. Didn't know it.

Total them after every paper. Buckets 1 to 6 are protocol fixes and go in the log. Bucket 7 is a study topic and goes in your revision plan.

Rules for running it

- **Log immediately**, while you still remember why you did it.
- **Log technique errors as well as content errors.** The overrun, the unwritten method, the uncaught sign.
- **Read the whole log** before every timed paper. Thirty seconds.
- **Retire an entry** when its rule has held for three papers. Cross it off. The list should shrink.
- **Build a drill** every fortnight around the top two entries. Practise the failing skill, not the subject it fails in.

Worked entry

9 Nov. Lost 4 marks on Q2(b): computed the eigenvector for $\lambda = 3$ when the question asked for $\lambda = -1$.

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

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

F

Marks-per-Minute Tables

Do this arithmetic before the paper, not during it.

Minutes per mark

	60 marks	75 marks	100 marks	120 marks
90 min	1.5	1.2	0.9	0.75
120 min	2.0	1.6	1.2	1.0
150 min	2.5	2.0	1.5	1.25
180 min	3.0	2.4	1.8	1.5

Those are raw figures. Subtract nine minutes first, four for triage and five for the sweep, then divide.

Worked: 100 marks in 180 minutes

Usable time: $180 - 9 = 171$ minutes. Per mark: 1.71, call it 1.7.

Part is worth	Budget	Walk away at
2 marks	3.5 min	5 min
4 marks	7 min	10 min
6 marks	10 min	15 min
10 marks	17 min	25 min
20 marks	34 min	50 min

The third column is the 50% overrun trigger from Chapter 6. Write it in the margin beside each question when you set the order.

Checkpoint targets

For a 100-mark, 180-minute paper, marks that should be attempted by each checkpoint:

Clock	Marks attempted	Action if 8+ behind
45 min	22	move on now
90 min	48	move on now, drop your weakest question
135 min	74	move on now, go straight to easy unattempted parts
170 min	100	begin the sweep regardless

Scale proportionally for other papers. The useful property is that the targets are roughly linear, so “marks attempted should be about 60% of the fraction of time elapsed, times the total” is close enough to check in your head.

Choice papers

“Answer 4 from 6” on a 100-mark, 180-minute paper means four questions of 25 marks each.

Activity	Minutes
Triage, all six questions	5
Four questions at 39 min	156
Reserve for the fifth if one collapses	14
Final sweep	5

Note the reserve. On choice papers, budget for one question going wrong, because on choice papers one question usually does.

Four more paper shapes, budgeted

75 marks, 120 minutes, answer all. Usable = 111. Per mark = 1.48.

Worth	Budget	Walk away
3 marks	4.5 min	7 min
6 marks	9 min	13 min
15 marks	22 min	33 min

60 marks, 90 minutes, answer all. Usable = 82. Per mark = 1.37.

Worth	Budget	Walk away
4 marks	5.5 min	8 min
10 marks	14 min	21 min
20 marks	27 min	41 min

Section A 30 marks short + Section B 45 marks long, 150 min.

Budget the sections separately, because the rates differ:

	Minutes	Rate
Triage	4	—
Section A (30 marks)	42	0.71 marks/min
Section B (45 marks)	99	0.45 marks/min
Sweep	5	—

Section A pays over 1.5 times what Section B pays per minute. Do it first, do it fast, and don't let it overrun into the long section.

20 MCQ + 45 long, 90 minutes. MCQ items are worth 1 mark each and take 60 to 90 seconds.

	Minutes	Rate
Triage	3	—
MCQ, three passes	25	0.80 marks/min
Long questions	57	0.79 marks/min
Sweep	5	—

Roughly equal rates here, which is unusual and means order is a matter of preference. Do the MCQ first anyway: it's mechanical, it warms you up, and it banks a fifth of the paper before the hard thinking starts.

The rule of thumb, if you remember nothing else

Minutes available divided by marks available. Write it at the top. Then no part gets more than 1.5 times its share, ever, and the last five minutes belong to the sweep.

G

The Two-Week Countdown

Print it, tick it off. Adjust the dates, keep the structure.

Days 14–10: close the gaps that pay

- ☐ Rank every topic by past-paper frequency and by your last test performance.
- ☐ Write the three topics you will fix.
- ☐ Write the two topics you are choosing **not** to learn. Name them so they stop costing attention.
- ☐ Work the three, using problems rather than notes.
- ☐ Update the error log daily.

Days 9–4: four full rehearsals

- ☐ Paper 1, timed, one sitting, full protocol.
- ☐ Review paper 1 against the published scheme, brutally.
- ☐ Paper 2, timed. Redo paper 1's failures.
- ☐ Review paper 2. Seven-bucket audit.
- ☐ Paper 3, timed, at the real exam's time of day, in an unfamiliar room.
- ☐ Review paper 3. Redo paper 2's failures.
- ☐ Paper 4, timed, using the most recent past paper you have.
- ☐ Review paper 4.

Every rehearsal runs the whole protocol: the four minutes, wall-clock deadlines, three checkpoints, the abandonment rule, the sweep. A rehearsal without the protocol is just a long test.

Days 3–1: patch, don't learn

- ☐ Read the error log end to end, once a day.
- ☐ One question per log entry, checking the rule holds.
- ☐ Read the method maps, one topic at a time.
- ☐ Write the formula sheet from memory, then check it. Twice.
- ☐ No new papers. No new topics.
- ☐ Sleep, properly, all three nights.

The day before

- ☐ Two hours of review, maximum, on material you already know.
- ☐ Pens, spare pens, calculator, spare batteries, ID, water, watch.
- ☐ Check the exam time and the room. Twice.
- ☐ Write the protocol card (Appendix D).
- ☐ Stop by early evening.

The morning

- ☐ Eat.
- ☐ Read the protocol card. Nothing else.
- ☐ No last-minute practice question.
- ☐ Arrive early, but not so early that you absorb the corridor.

In the first four minutes

- ☐ Minutes per mark, written at the top.
- ☐ Calculator mode checked, written at the top.
- ☐ Every question read and marked A / B / C.
- ☐ Order written, with wall-clock end times.
- ☐ Checkpoint times written: 45, 90, 135.
- ☐ Verb and constraints underlined on question one.

Afterwards, within a day

- ☐ Twenty-minute post-mortem: what happened, why, next time.
- ☐ Right-hand column straight into the error log.
- ☐ Then stop, and prepare for the next paper.

H

Where to Find Papers and Schemes

Ranked by usefulness, with what each is actually good for.

Papers

Your own department's past papers. Nothing else matches the style, the phrasing, or the examiner's habits. Usually on the department website, the library system, or the student portal. If they're not published, ask the lecturer directly; a surprising number will hand over old papers if asked politely and early in the term.

The same course at another institution. The material is the same and the phrasing is different, which makes this a useful stress test once you're fluent. It's also where to go when your own department has only three papers.

Textbook chapter reviews. Often assembled from old exam questions, and usually with answers. Good for blocked practice while you're learning a method.

Professional and entrance exams. For engineers, the professional body papers. For anyone drilling problem-solving under time pressure, the big entrance exams (JEE Advanced, GATE, and their equivalents) are an enormous free bank of multi-part problems with published solutions, and the timing pressure is the point even if the syllabus doesn't match exactly.

Olympiad and competition sets. For pure mathematics, useful for building the habit of starting a problem you don't recognise. Not representative of exam papers, so don't use them for rehearsal.

Mark schemes

Worth more per page than the papers, and much harder to find.

Your department's, if published. Read three of them properly.

Any institution's, if yours doesn't publish. Schemes are more similar across institutions than papers are, because they all use the same M/A/B structure. A scheme from elsewhere still teaches you what method marks look like and where they sit.

School-level schemes (A-level, IB, and equivalents) are published in full and free, and they're the clearest documentation of marking convention you'll find anywhere: M1, A1, B1, ft, oe, ISW, all annotated. The mathematics is below your level. The marking logic is identical to what your university uses.

That last one is the tip most students never hear, and an hour with a school-level scheme will teach you more about how marks are awarded than anything else on this list.

Solutions, and how to use them

Only after a genuine, written attempt. Reading a solution to a problem you haven't tried produces a strong sense of understanding and no ability, and it's the most common way revision time disappears.

When you do read one, read it the way Chapter 12 describes: bucket every lost mark, then redo the question from scratch three days later.

Building your own bank

Keep every problem sheet, every class test, and every past paper you attempt, in one folder, with your attempt attached.

By the final year that folder is the best-matched question bank in existence for you specifically, because it's built entirely from the material your own department chose to set.

The companion volumes

How to Write Mathematics (Gaurav Tiwari). The other side of the same problem: how to write a proof or a solution so a grader can follow it, with proof makeovers and the grader's-eye view of partial credit.

How to Read Mathematics (Gaurav Tiwari). For the term before the term before. Definitions, theorems, proofs, and papers, with a procedure for each, plus honest reading speeds.

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 happened in a paper, two minutes of your help matter more than you'd guess: leave a short, honest review wherever you bought it, or lend your copy to the friend who always runs out of time on question 6.

The appendices are working tools, not reading. The protocol card and the error log are meant to be photocopied, carried, and written on. You'll find printable companions and more of my writing for students at **gauravtiwari.org**, and corrections and questions are welcome there. Future editions will owe their sharpest fixes to readers who wrote in, and to markers who tell me what they're really seeing on scripts.

Also by Gaurav Tiwari

How to Write Mathematics

Proofs, Papers, and Problem Sets for Students

The companion to this book, one level deeper. Where this one is about the three hours, that one is about the writing itself: proof makeovers, full-credit solutions, and the grader's-eye view of partial marks, from Tuesday's problem set through to your first paper.

How to Read Mathematics

Definitions, Theorems, Proofs, and Papers for Students

For the term before the term before. 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 so slowness stops reading as failure.

LaTeX for Students

A Practical Guide to Beautiful Documents and Mathematical Typesetting

Everything you write, from Tuesday's problem set to the thesis, and how to make it look as good as it reads.

Other titles

- *The Student Research Workflow*
- *LaTeX for Theses and Dissertations*
- *Indian Polity: Complete Notes for UPSC Prelims & Mains*
- *Website and Blog Copywriting*
- *Marketing Essentials for Bloggers*
- *The Affiliate Bloggers System*

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