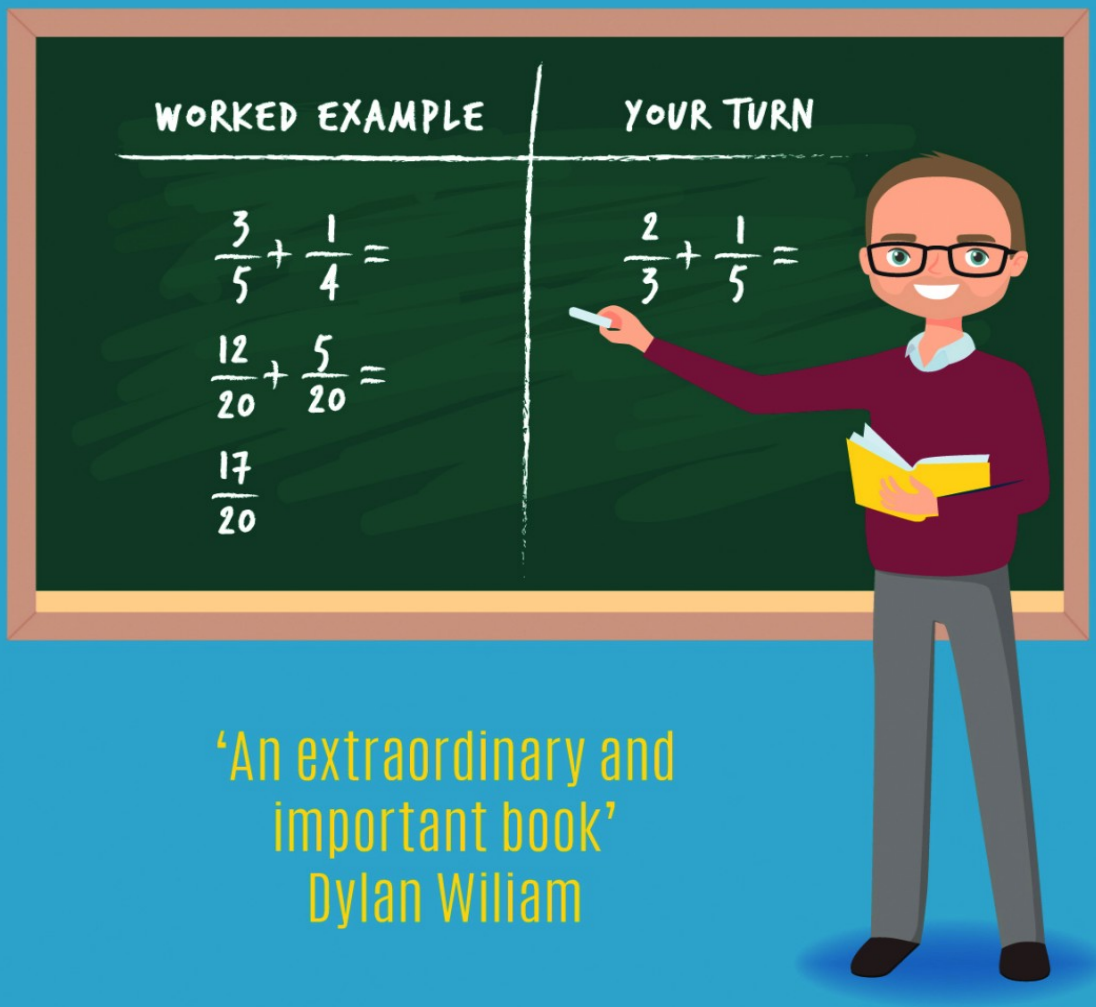
A JOHN CATT PUBLICATION

–CRAIG BARTON–

HOW I WISH I'D TAUGHT MATHS

LESSONS LEARNED FROM RESEARCH, CONVERSATIONS WITH EXPERTS, AND 12 YEARS OF MISTAKES



'An extraordinary and important book'
Dylan Wiliam

First Published 2018

by John Catt Educational Ltd,
15 Riduna Park, Station Road,
Melton, Woodbridge IP12 1QT

Tel: +44 (0) 1394 389850
Email: enquiries@johncatt.com
Website: www.johncatt.com

© 2018 Craig Barton

All rights reserved.

No part of this publication may be reproduced, stored in a retrieval system, transmitted in any form or by any means, electronic, mechanical, photocopying, recording, or otherwise, without the prior permission of the publishers.

Opinions expressed in this publication are those of the contributors and are not necessarily those of the publishers or the editors. We cannot accept responsibility for any errors or omissions.

ISBN: 978 1 913808 06 8

Set and designed by John Catt Educational Limited

OceanofPDF.com

ABOUT THE AUTHOR



Craig Barton has been teaching maths since 2004, predominantly in two comprehensive schools in the sunny North West of England – Range High School in Formby and Thornleigh Salesian College in Bolton. Four years into his career, Craig was appointed an Advanced Skills Teacher (AST), giving him the opportunity to work with and learn from many teachers and students in a wide variety of schools. Since 2009, he has been the Secondary Mathematics adviser for the Times Educational Supplement (TES), the largest professional network of teachers in the world, reviewing, creating and sharing resources with hundreds of thousands of teachers.

Craig is the creator of two of the country's most popular maths websites: www.mrbartonmaths.com, which provides free resources to teachers and students all around the world; and www.diagnosticquestions.com, a formative assessment website hosting the world's largest collection of high-quality maths diagnostic multiple-choice questions, which aims to help students and teachers identify, understand and resolve key misconceptions.

Craig is the host of the Mr Barton Maths Podcast, interviewing leading figures from the world of education, such as Dylan Wiliam, Doug Lemov, Daisy Christodoulou and Dan Meyer.

Craig has been fortunate enough to give talks, run workshops and work with teachers and students all over the world, from Bangkok to Basingstoke, and is the author of three (non-maths) novels. Fingers crossed he is also still married to Kate when you are reading this.

OceanofPDF.com

REVIEWS

How I Wish I'd Taught Maths *is an extraordinary and important book. Part guide to research, part memoir, part survival handbook, it's a wonderfully accessible guide to the latest research on teaching mathematics, presented in a disarmingly honest and readable way. I know of no other book that presents as much usable research evidence on the dos and don'ts of mathematics teaching in such a clear and practical way. No matter how long you have been doing it, if you teach mathematics – from primary school to university – this book is for you.*

Dylan Wiliam, Emeritus Professor of Educational Assessment, UCL
@dylanwiliam

How I Wish I'd Taught Maths *is a rare and wonderful book, one that could only have been written by someone with Craig's devotion to teaching and willingness to become immersed in the research literature on how people learn. In clear, concrete, and compelling terms, Craig illustrates evidence-based ways to upgrade mathematics instruction, ways that are often unintuitive and/or at odds with prevailing educational practices. It makes us wish that young people the world over might have the good fortune to find themselves in classes that incorporate Craig's insights. In fact, whereas Craig writes, 'I'll be honest – this book has been created for maths teachers', we think that Craig's 'lessons learned' can, with some creativity, enhance any teaching.*

Robert A. Bjork and Elizabeth L. Bjork, Department of Psychology,
University of California, Los Angeles

It's rare that we change our habits and beliefs once they are established – cognitive bias is strong in us. And that is what makes this book so exceptional. Craig describes not only what he's learned from a methodical study of cognitive science but how he's changed over time despite his initial success. There's a joyful relentlessness to Craig's study of teaching methods. He starts out telling us he wants to 'know every detail', and what makes the book so exceptional is just that – the way the story of how something he learned about teaching played out in a specific problem or lesson, was refined and improved. It's an incredibly useful book for maths teachers especially, but really for anyone who teaches and cares about getting it right.

Doug Lemov, former teacher and author of *Teach like a Champion*
@Doug_Lemov

This is a really thoughtful and thought-provoking discussion of a series of important and practical questions about the best way to teach.

Daisy Christodoulou, author of *Seven Myths about Education* and Director of Education at No More Marking @daisychristo

This book has the potential to have a huge impact on the way maths is taught. It is so refreshing to see a maths teacher honestly critiquing their own practice and suggesting alternative approaches based on sound research and analysis. Craig's warm and relatable style of writing is a pleasure to read. His book is brought to life by hilarious anecdotes and humble reflections. Craig summarises the key points of the relevant research succinctly and his advice to teachers is perfectly pitched and instantly transferable to any maths classroom. For the sake of our current and future students, I certainly hope that this book becomes essential reading for maths teachers.

Jo Morgan, maths teacher and creator of www.resourceaholic.com
@mathsjem

Teaching is on the cusp of a new era, bringing it out from the dark ages where 'don't smile before Christmas' and 'because we've always done it that way' is replaced with evidence-backed approaches based on insights from research into cognitive science and psychology. In a world where the highest-impacting teaching approaches are often counter-intuitive, experience alone cannot be relied upon; we need these research-informed insights in order to improve as practitioners.

History will look back on **How I Wish I'd Taught Maths** *as a seminal book leading mainstream teachers into this new world. And who else would you want to narrate us into the unknown than Craig Barton, the leading maths educator in the UK? A role model for excellence in professional development, Craig's enthusiasm, intellect and critical skills navigate readers through vitally important and complex concepts with ease.*

Pulling on contemporary research from hundreds of papers and books, **How I Wish I'd Taught Maths** *concisely summarises all that we know about high-impact teaching and gives specific examples of what it looks like in the classroom. This book helps you realise just how much you didn't know about teaching, and furthermore, how much we still have to learn as a profession.*

A must-read for experienced or newly qualified teachers alike. This book will quickly become a classic.

William Emeny, Teacher, Researcher, Author of *Great Maths Teaching Ideas*, Creator of Numeracy Ninjas @Maths_Master

How I Wish I'd Taught Maths *is an honest and insightful reflection of Craig's years as a maths teacher. Through experience, podcasting others and reading broadly, Craig carefully considers every assumption he used to make when*

teaching maths – assumptions we all make or made. In each chapter he forensically analyses a theme from Explicit Instruction to, my favourite, Choice of Examples, expounding his old approach and backing up his new approach with rich examples and scholarly references.

Written in the way he speaks – upbeat, humble and littered with 'flippin' 'ecks' – Craig brings together so many aspects of maths teaching and so many shared assumptions that there's something in here for anyone involved in maths education, including teacher trainers, early career teachers and those with many years at the chalkface. Having read **How I Wish I'd Taught Maths** *I'll now be a considerably stronger practitioner. No doubt this will become the defining book on maths pedagogy for generations of maths teachers.*

Bruno Reddy, Former Head of Maths at King Solomon Academy and creator of Times Tables Rockstars @MrReddyMaths

This is one of the most useful books on maths instruction that I have read. Craig's humility and honesty about his previous reasoning, and responses to received wisdom (and his own biases), invites the reader to examine their own preconceptions and blind spots. The volume of research behind his conclusions – and the down-to-earth summaries of its implications – make it accessible for anyone interested in maths teaching.

His stories ring horribly true, encapsulating predictable errors made early (and late!) in a maths teacher's career. I wish this book had existed when I trained; I might have avoided making so many mistakes, and for so long! The conversational style gives it the feel of a discussion with a trusted and inspiring maths mentor or university tutor, making you feel Craig's expectations of us, as teachers, are as high as those he has for his students!

I've made it a 'must-read' for anyone who wants to join our department. In addition to being accessible, it acts as a handy reference for discussions, giving a common language for maths teachers to discuss their practice. It is wonderful to have a go-to resource allowing maths teachers, and those who support them, to examine where their practice can be refined or, in some cases, altered.

As the debate about 'teacher standards' gathers pace, this book is a timely contribution in answer to the question, 'What should we expect all teachers, as professionals, to know about the craft of teaching maths to children?'. That said, there is a wealth of information and reflection in here that would be valuable for teachers across subjects and phases. I would not be surprised if this became compulsory reading for a range of PGCE courses and teacher inductions. Craig's conclusions about how to help pupils learn interrogates the most fundamental aspects of how we think about teaching, regardless of curriculum area.

Dani Quinn, Head of Maths at Michaela Community School @danicquinn

Evidence-fuelled, down-to-earth, and hugely practical. The book we maths teachers have all been waiting for.

Peps Mccrea, Associate Dean at the Institute for Teaching @pepsmccrea

OceanofPDF.com

FOREWORD

It's hard to imagine a maths teacher in the United Kingdom who hasn't heard of Craig Barton.

As the creator of both the *Mr Barton Maths* website and the *Mr Barton Maths Podcast*, the long-standing Maths Adviser to *TES*, the Diagnostic Questions phenomenon, and regular speaker at conferences throughout the country, and the world beyond, his contribution to the community of mathematics teachers has been extraordinary, as has the generosity with which he shares so much of his insight.

Given this, it is equally hard to imagine a reader who won't be shocked to hear Craig admit he thinks he's been getting it wrong the whole time, across more than a dozen years in the classroom.

I first came across Mr Barton when tutoring GCSE maths, in preparation to kick off a serious career in education by joining Teach First. His website provided a systematic and organised collection of resources that proved invaluable. I could hardly have imagined at the time that seven years later *the* Craig Barton would be interviewing me, to ask my thoughts on teaching.

But I'd been fortunate. In my first months in the classroom, friends in the teaching profession directed me to original research from the field of cognitive science, and what I read there completely redirected my thinking.

I entered the classroom obsessed with words like 'inspire', 'motivate' and 'engage', alongside their siblings 'challenge', 'understanding' and 'independence'. As I read the research, I realised that for the first time I was learning something about the human mind, about how we learn, and so how

I could teach in a way that pupils were more likely to follow successfully. I realised that I had been obsessed with 'motivation' and 'inspiration' because I knew how to give a pep talk, but didn't know how to teach, and I used 'challenge' and 'independence' as an excuse for my poor teaching: pupils were struggling because I was encouraging them to 'think for themselves', because I courageously refused to 'spoon-feed' them.

So, I got lucky, and over Christmas, after one term of teaching, I wrote a confession and sent it to 30 people, including family, friends, fellow student teachers and my mentors. It began: 'Everything I thought was wrong'.

I was lucky, and was able to dedicate my next five years in the classroom to figuring out how to apply the theory, constantly improving, with a false-start that limited me and my pupils for only four months. By contrast, Craig has come out after 12 years, ready to share his own confession, publicly, with thousands of fellow teachers, and he's done it because most of us aren't as lucky.

Seven years ago, almost no one in education that I spoke to had heard the words 'cognitive science'. Now, it seems everyone has heard of it; but it's easy to become trapped in the filter bubble. When I was interviewed by Craig, he frequently expressed surprise at some of the points I made, and asked follow-up questions whose answers I thought were obvious. There were moments when this was clearly a deliberate effort to add clarity for the audience, but in other moments I was left wondering if Craig was asking because he himself genuinely wasn't already familiar with the ideas I referred to. It must be a clever interview technique, I concluded, designed to create an informal atmosphere of conversation and discussion.

So, I was as shocked as anyone when I learnt, on reading *How I Wish I'd Taught Maths*, that it wasn't just a clever interview technique, that somehow the burgeoning revolution, and explosion in awareness of cognitive science, had completely passed Craig by. It wasn't until he started the podcast series, and interviewed the likes of Dylan Wiliam and Daisy Christodoulou, that he began to realise what he had been missing.

At this point it's worth saying that Ray Dalio is the of the world's most successful billionaires, and he makes the point that 'smart people are the ones who ask the most thoughtful questions, as opposed to thinking they have all the answers'. I believe Dalio has the right of it, and therefore Craig Barton is one of the smartest people I know. When first learning there was a body of research that might contradict his teaching recommendations, he could have done what most people feel compelled to do: dug his heels in and refused to question his existing beliefs. Instead, true to form, he ferociously set out to find and question people who could challenge and add to his thinking, and he freely shared his experiences with anyone who wanted to get involved and join in the journey.

People who are prepared to dedicate their public lives to a line of reasoning, and then in the face of new evidence say, equally publicly, that they had it wrong all that time, are rare, rare, rare. They must be treasured, and this book is a treasure for every one of us who hasn't been lucky enough to stumble upon the 'cheat codes to intelligence' for ourselves.

This is the book I wish I'd written. It is part confessional, part literature review, part pragmatic advice. It speaks directly to the mistakes we've all made in the classroom, and precisely to what we can change to do better. It is an uncommon book that is as practical as it is theoretical, and somehow retains its humility throughout.

It's been a while since I've read a book on education that I've enjoyed and taken as much from as I did here, so buckle up, you're in for a treat!

Kristopher Boulton

Former second in department at King Solomon Academy, Teach First maths ambassador and Director of Education at Up Learn

OceanofPDF.com

DEDICATION

For Kate, my (a)cute girl

OceanofPDF.com

INTRODUCTION

I thought I knew it all

As a maths teacher, I have been quite successful.

I know that is not exactly the most humble way to kick things off, but I feel it is important to establish the context in which this book is written. So please forgive me the indulgence of giving my own trumpet a hearty blow for just one paragraph.

At the time of writing, I am in my 13th year of teaching mathematics. I was one of the youngest Advanced Skills Teachers (ASTs) ever appointed in the UK, giving me the opportunity to work closely with teachers and students across dozens of schools. For the last 8 years I have been the *TES* Maths Adviser, with a responsibility for sharing the best practice and resources with the largest professional network of maths teachers in the world. I am a member of the AQA Expert Panel, responsible for giving advice on assessment and support materials at a national level. I am the creator of two popular maths websites: www.mrbartonmaths.com, which has received over 20 million visits from teachers, students and parents from across the globe; and www.diagnosticquestions.com, which contains the world's largest collection of multiple-choice diagnostic questions, and is used in over 80% of UK secondary schools and in more than 50 countries. My teaching has been judged as Outstanding in four successive Ofsted inspections in two different state schools, and each one of my classes has met or exceeded their challenging targets in national exams in every year I have taught. I have written two maths textbooks, had the honour of delivering workshops on teaching across the UK and all over the world, including Bangkok, Nanjing

and Kuala Lumpur, and I have worked directly on a national project with the Ministry of Education and the British Embassy to improve teaching standards in Cambodia. I have even been asked to pose for the odd selfie.

I do not list my achievements to make myself feel good (well, not entirely, anyway). I do so to illustrate just how far it is possible to come without really having a clue what you are doing.

If someone had asked me to write a book about maths teaching two years ago, it would have been simply unrecognisable from the one you are (hopefully) about to read. It would probably have been entitled *The Beauty of Rich Tasks*, *The Power of Mathematical Discovery*, or *The Ultimate Guide to Teaching Problem-Solving*, and no doubt featured a picture of me looking very happy with myself on the inside cover. It would have consisted of the open-ended tasks and ideas that I have developed over many years and used with thousands of students. It would have been full of me exclaiming how much my students enjoyed these activities, the insights they made, the problem-solving skills they developed, the independent learners they became, and the results they achieved. I would have extolled the benefits of discovery learning, inquiries, projects, puzzles, student-centred learning, and of me as a teacher taking a back seat (I probably would have used the phrase 'the teacher should be the facilitator of learning' more than once). The one noticeable absence from this hypothetical book, of course, would have been any research to backup my claims. And if someone had pointed out this tiny omission, I would have replied with a patronising smile and explained 'I don't need research, I know it works'.

If really pushed for some external justifications for the practices I had always employed, I could very quickly point to Ofsted. When I was observed as part of our school's inspection during the last week of term in December 2015 (Happy Christmas!), the inspector walked into my room to find a group of my Year 9 students stood at the front of the class, interactive whiteboard pen in hand, explaining the intricacies of rearranging algebraic equations to the remaining twenty-five students, whilst I sat at the back of the room playing on my phone. My 'teaching' was judged to be Outstanding.

Teacher-led lessons were dull, uninspiring and bad for learning. Everyone knew that – well, apart from those teachers who refused to turn on an interactive whiteboard, plan adequately for their kinesthetic learners, or engage with Kagan group structures. But they were a dying breed. I always found it a bit weird that those teachers seemed to consistently get good results and students seemed to really enjoy their lessons, but surely that was *despite* their archaic, ineffective teaching methods, not because of them? To use the old Ofsted language, such teacher-led lessons were 'satisfactory' at best. My lessons were better than that – flipping heck, *I* was better than that – and I saw it as part of my wider responsibilities to ensure I helped other teachers develop along similar lines.

My mid-career crisis

Towards the end of 2015, I began recording the *Mr Barton Maths Podcast*. This show gives me an opportunity to interview people from the world of education who interest and inspire me, and I have been fortunate enough to speak to the likes of Dylan Wiliam, Robert and Elizabeth Bjork, Daisy Christodoulou and Doug Lemov.

If my guest is a practising teacher, one of the first things I ask of them is to describe a lesson they have taught recently. I want to know every detail, from how they decide what to teach, how they plan, what resources they use, what happens as the students first enter the room, to how they assess, model and differentiate, how much independent and group work is involved, right through to the very end of the lesson. Above all, I want to know *why* they make these choices.

The types of lessons my guests describe vary widely. In 2017, Andrew Blair described a lesson built around an inquiry with the students determining the direction the lesson took, whereas that same year Dani Quinn talked in detail about a very teacher-led lesson encompassing drills and rigorous practice. The thought my guests put into the planning and delivery of their lessons is mind-blowing. Indeed, when I asked Kris Boulton to describe how he plans a lesson, it took him two hours to answer. And whilst there are key differences between the lessons themselves, they all have one thing in common – my guests can justify each and every decision they make.

These interviews made me realise that I could not. My only justification for the things I did was that it was obvious. It was obvious that teachers should talk less and place the responsibility of learning on the shoulders of the students. It was obvious that we should teach students to problem-solve, and to allow them to discover things for themselves. It was obvious that we should let students struggle. It was obvious because I had always done it that way. If I was a guest on my own show, I would have been reduced to a bumbling wreck.

I knew I needed to do something about this.

My guests often mentioned books and various bits of research they had read that informed their practices and methodologies. I had never really read any educational research, but I made a note of each reference and promised myself I would get around to them one day – maybe after I had planned another amazing lesson on problem-solving.

The day my life changed

What I am about to say next may sound ridiculous, but I promise if anything it is an understatement: the day I started reading educational research was the day my life changed.

It all began with Daniel Willingham's *Why Don't Students Like School*, which led me onto his ridiculously good series of *Ask the Cognitive Scientist* articles for the American Federation of Teachers and the associated research papers he quoted. I devoured them all with eyes wide open and jaw firmly on the floor. Next came Cognitive Load Theory, which I was introduced to via the incredible blog of Greg Ashman and our subsequent conversation on my podcast. I was familiar with the work of Dylan Wiliam through the development of my Diagnostic Questions website, but I read and re-read everything I could get my hands on from the great man. I didn't think life could get any better – and then I came across the work of Robert and Elizabeth Bjork. Good God.

One thing led to another, and before I knew it I had read well over 200 books and research papers. I would wake up in the middle of the night, my head buzzing with ideas – my wife is incredibly understanding, but even she

refused to indulge my enthusiasm for the Pretest Effect at 4*am* on a Thursday morning. I summarised all my reading and my practical takeaways on my website, and began slowly dropping them into the training and presentations I gave to teachers.

I started trialling ideas in my classroom, and the effects were immediate. My reading of the Redundancy Effect (Section 4.7) and my subsequent conscious effort to shut up more at key points in the lesson caused one concerned Year 11 to ask if I was feeling okay that day. Even after a few weeks I sensed a change. We were getting through more work in class. My students seemed more confident with the concepts we tackled. I felt like I was actually teaching.

But the more I read and the more I experimented, the more questions I had. So, I read more, spoke to more people, tried out more stuff with my students and colleagues, and wrote more takeaways. By the end I had over 100 hours of interviews, 1000 PowerPoint slides and more than 100,000 words of notes written.

This book is the result.

The intended audience

I'll be honest: this book has been created for maths teachers.

I don't know if maths teachers are a bit of a funny breed (in fact, I definitely do know that we are), but every time I sit through some generic cross-curricular training session, I always end up thinking to myself, 'Yeah sure, but how would that work in maths?'. I know this reflects incredibly badly on me, but as soon as I hear my colleague from geography tell me about a problem-solving strategy that I could use when teaching quadratic equations, or a colleague from PE offer up wisdom on using group work that I could employ when investigating the finer points of conditional probability, I tend to switch off.

So, it would be incredibly hypocritical of me to suggest that this book has any worth whatsoever outside of the boundaries of maths teaching. English teachers may find strategies in here that they can use when studying Shakespeare, and there may well be nuggets that history teachers can use

when contrasting the validity of sources of the First World War; but I genuinely do not know if that is true, and I am not going to claim it is so. I believe maths is different (*special*, you might say), and this book has been written by a maths teacher, for maths teachers.

Whilst I am successfully losing readers, I will also confess that the majority of my 13 years' teaching experience has been spent in secondary schools working with students between the ages of 11 and 18. Whilst I have spent a significant amount of time in primary schools, working with both students and teachers, I am by no means an expert in that field. I hope the strategies I write about can be transferred into a primary setting – and the feedback I have had when teachers have trialled these is that they can be – but I want to be open and honest and say that they have mainly been written from a secondary perspective.

That said, I hope that no matter your age, experience, mathematical background, type of school, role, responsibility, or favourite number, there will be plenty in this book that you find useful.

The structure of the book

I have tried to make this book as user-friendly as possible, appealing to the time-poor teacher who has five minutes spare on a weeknight, through to the rare occasion where we might have a chunk of time free to do some reading on a Sunday morning or a holiday.

The book covers 12 key themes, with each theme further broken down into ideas. Each idea consists of four sections:

What I used to do

Here I explain my previous beliefs, where they came from, and how they impacted what I did in the classroom. At times this makes for painful reading, and I can assure you it was equally painful to write. Please don't think too badly of me!

Sources of inspiration

Here I will give my sources of inspiration, which will be research papers, books, blog posts or interviews. The original versions of almost all quoted

research papers are linked to directly from the following page on my website: mrbartonmaths.com/teachers/research.

My takeaway
This is my summary of the key points from the sources of inspiration that are pertinent to the issue under discussion. You may well disagree with my interpretation, or simply want to dive in for greater detail, so I encourage you to study the original source and reach your own conclusions.

What I do now
This explains the ways I have changed my approach based on my takeaway. It is how I wished I had taught maths throughout my career, and how I plan on teaching it going forward. I want this book to be full of practical strategies that you can use right away, and this is the section where you will find these.

It is possible to dip in and dip out of this book as you please, jumping to the sections and issues that interest you most. However, I must warn you that I have thought carefully about the order I present these ideas, and it is often the case that later sections of the book will refer to ideas introduced and developed earlier on.

Ten things I used to believe

In order to whet your appetite for what follows, here are ten things I used to believe were true for at least the first ten years of my career, and which I no longer do. Each of these will be covered at some point in the book.

1. The best lessons have little teacher-talk and lots of student-talk.
2. Where possible, students should 'discover' things for themselves.
3. We can teach problem-solving.
4. Effective differentiation means giving students different work to do.
5. The maths we teach should be relevant to our students' lives.
6. Students should always know why they are doing something before they learn how to do it.

7. The more feedback we give students, the better.
8. Tests are predominantly tools of assessment.
9. Doing lots of past papers is the best way to prepare for an exam.
10. If students are struggling, then they are learning.

Hopefully I am not alone in having believed these things.

A word of warning

Despite what many of my students may tell you, I am a human being, and as a human being I am susceptible to a number of cognitive biases that are directly relevant to what you are about to read. Despite my best efforts to avoid them, the selection of sources of inspiration and my interpretation of them have undoubtedly been influenced by the following:

The Confirmation Bias

The tendency to search for, interpret, favor, and recall information in a way that confirms one's preexisting beliefs or hypotheses.

It turns out that not everybody agrees with what is the best way to teach in order to enable students to learn. This is quite annoying when you are trying to write a book with a coherent narrative. When I came across a book or paper that made a really interesting claim, an unavoidable urge pulsed through my veins to search for other sources that reached a similar conclusion. Likewise, any time I came across something that called that claim into question, my instincts where to look for flaws in that particular source's methodology (small sample size, no control group, written by an English teacher, *etc*). I have tried my best to overcome this failing by attempting to objectively judge every paper on its merits, getting second and third opinions where need be. This has led to me completely rewriting or even binning various sections as the findings were just too contradictory.

The Halo Effect

When an observer's overall impression of a person, company, brand, or product influences the observer's feelings and thoughts about that entity's character or properties.

I simply love the writings of Daniel Willingham, John Sweller, Dylan Wiliam and the Bjorks. I think those people are geniuses. They are my heroes. I hang on their every word. Much like the parent who has been told their beloved child has been caught doing something wrong at school, my first instinct when someone criticises their work is to go into denial, because how could people as good as they are ever do or say something wrong? In an attempt to combat my infatuation, I have gone out of my way to read and understand their critics, and where I feel they have a valid point, I have addressed this directly in the book.

The Dunning-Kruger Effect
Where persons of low ability suffer from illusory superiority, mistakenly assessing their ability as greater than it is.

For about ten years I was happily teaching in ignorant bliss. I firmly believed I was an Outstanding teacher, I knew what was best for my students, and I had no real interest in anyone telling me otherwise. The more I have read, the more people I have talked to, and the more things I have tried out, the more I realise just how clueless I really was. Perhaps more importantly, in the words of Paul Weller, the more I know, the less I understand (and hopefully I am now a 'Changingman'). I am a firm believer that the only way to improve is to confront your weaknesses head-on, acknowledge them for what they are, and strive to improve. This book is my attempt to do just that.

Anyway, I have built this up far too much. It is time to get cracking. I really hope you enjoy this book. At times it may prove uncomfortable reading, and not just because of my poor grammar, but because it may challenge techniques and ideas you have held for years. You will certainly not agree with all of it – flipping heck, I am not sure I agree with all of it myself – but I hope it will make you think.

This is how I wish I'd taught maths.

OceanofPDF.com

1. HOW STUDENTS THINK AND LEARN

We begin our journey by taking a look inside the minds of our students, considering how they think and learn, and the implications for matters such as planning, lesson content and the introduction of mathematical concepts.

OceanofPDF.com

1.1. A simple model of thinking and learning

What I used to think

The next 15 words I am about to write are more than a little worrying: *in 12 years of teaching, I never really considered how my students think and learn*. A teacher not considering how their students think and learn is kind of like a doctor not being overly concerned about the workings of the body, or a baker taking only a casual interest in the best conditions for bread to rise. If we don't know how students think and learn, how on earth can we know how to teach them effectively?

I had vague notions of concepts like working memory and schema, but *vague* is certainly the operative word. So, I went about my business, blissfully unaware of my own ignorance, going off nothing more than intuition to try to teach students to the best of my abilities.

Sources of inspiration

- Anderson, J. R. (1996) 'ACT: A simple theory of complex cognition', *American Psychologist* 51 (4) p. 355.
- Anderson, R. C. (1977) 'The notion of schemata and the educational enterprise: General discussion of conference' in Anderson, R. C., Spiro, R. J. and Montague, W. E. (eds) *Schooling and the acquisition of knowledge*. Hoboken, NJ: John Wiley & Sons, pp. 415-431.
- Bartlett, F. C. (1932) *Remembering: a study in experimental and social psychology*. Cambridge: Cambridge University Press.
- Bjork, R. A. (1975) 'Retrieval as a memory modifier' in Solso, R. L. (ed.) *Information processing and cognition: the Loyola symposium*. Mahwah, NJ: L. Erlbaum Associates, pp 123-144. Coe, R. (2013)

Improving education: a triumph of hope over experience, CEM Inaugural Lecture. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf

- Cowan, N. (2010) 'The magical mystery four: how is working memory capacity limited, and why?', *Current Directions in Psychological Science* 19 (1) pp. 51-57.
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.
- Mayer, R. E. (2008) 'Applying the science of learning: evidence-based principles for designing multimedia instruction', *American Psychologist* 63 (8) pp. 760-769.
- Mccrea, P. (2017) *Memorable teaching: leveraging memory to build deep and durable learning in the classroom*. CreateSpace Independent Publishing Platform.
- Piaget, J. (1928) *Judgement and reasoning in the child*. Harcourt, NY.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.
- Wiliam, D. (2017) 'Memories are made of this', *TES Magazine*.
- Willingham, D. T. (2009) *Why don't students like school?: A cognitive scientist answers questions about how the mind works and what it means for the classroom*. Hoboken, NJ: John Wiley & Sons.
- Willingham, D. T. (2017) 'On the Definition of Learning…', *Science and Education Blog*. Available at: http://www.danielwillingham.com/daniel-willingham-science-and-education-blog/on-the-definition-of-learning

My takeaway

As thinking and learning are clearly such important parts of teaching, they are also going to be crucial parts of this book. With that in mind, it is important early on to establish a model of how students think and learn. Any attempt to reduce something as complex as human cognition to a model that can be explained in a couple of pages is fraught with danger. But as Dylan Wiliam explains in a 2017 article for *TES*: 'what makes a model valuable is not how accurate it is – any model can be made more accurate by making it more complex – but rather the trade-off between simplicity and power. This is particularly important when we look at the human brain, which is probably the most complex thing in the universe'.

Fortunately, there are sufficient commonalities across leading research to enable us to form a basic model for how thinking and learning occurs that will serve us well for what follows.

How students think

For Willigham (2009), thought occurs when you combine information in new ways, and successful thinking relies on four factors: information from the environment, facts in long-term memory, procedures in long-term memory, and space in working memory. If any one of these factors is deficient, thinking will likely fail.

Let's take a closer look at those two components of memory.

Long-term memory

For Mccrea (2017), long-term memory may be thought of as our mental model of the world: a map we construct ourselves, holding facts, procedures, beliefs, mindsets and dispositions. **Long-term memory represents what we know and who we are and informs how we act**. It has no known limits. All of the information stored in long-term memory resides outside of awareness, lying there patiently until it is needed, before entering working memory and becoming conscious.

Throughout this book, the facts and procedures stored in long-term memory will be grouped together under the term *knowledge*. When I say I want my students to have a good knowledge of fractions, I mean I want them to know relevant facts, such as what a numerator is and that three-

quarters is more than one-quarter, as well as to be able to carry out relevant procedures, such as how to add, simplify and divide fractions. I also want as much of this knowledge as possible to be automated, so students know instantly that a half of 24 is 12 and that to add fractions you need a common denominator, without imposing any strain on their working memories. Later on in this book there will be the need to be more specific with the nature of this knowledge, in particular when we come to consider procedural fluency and conceptual understanding in Section 3.9 and the eternal question about teaching the How before the Why. But for now, think of knowledge as the interconnected facts and procedures stored and organised in long-term memory that allow us humans to operate.

This knowledge is stored and organised in long-term memory in *schemas* (see Piaget, 1928; Bartlett, 1932; and Anderson, 1977). Schemas do this by incorporating or chunking multiple elements of related information into a single element with a specific function. So, you may have a schema for adding fractions, which contains all the relevant knowledge you have acquired over many years. There is no limit to how complex schemas can become, or how many can be stored in long-term memory. Indeed, for the proponents of Cognitive Load Theory (*eg* Sweller *et al*, 1998) to prevent cognitive overload, the ability to solve problems demands the acquisition of tens of thousands of these domain-specific schemas, together with the automation of key knowledge following extensive practice. Therefore, long-term memory is not just a vast databank of knowledge, but an integral component of all cognitive activity.

According to Anderson's (2012) Adaptive Character of Thought (ACT-R) theory, complex cognition arises from an interaction of declarative and procedural knowledge. Declarative knowledge is factual knowledge that can be reported or described, and its most basic unit is a chunk. Procedural knowledge is dynamic and involves rules, or productions, that guide how thinking occurs. Declarative knowledge can be acquired quickly from direct encoding of the environment, while procedural knowledge takes longer and must be compiled from declarative knowledge through practice. After a certain amount of practice, the path of production becomes stable and procedural learning has occurred.

It is worth noting that psychologists and maths teachers are likely to have different interpretations of the term *procedure*. Psychologists mainly think in terms of procedures derived from implicit memories, such as tying a shoelace or driving a car. Whilst maths contains examples of such procedures – experts fluently adding fractions or rearranging equations, for example – there are also procedures that rely on more explicit memories, such as working through a multi-step trigonometry question slowly and methodically. But the key point remains the same – the conditions under which we learn procedures are determined by existing declarative knowledge.

There are two key implications from Anderson's model that we will revisit many times throughout this book:

1. Existing knowledge makes thinking and learning easier.
2. Once something has been learned – whether it is correct or not – it is very difficult to unlearn. Hence, practice does not make perfect, **practice makes permanent**.

Working memory

Working memory is best viewed as the place where thought occurs. It is all about the here and now. Unlike long-term memory, working memory has a finite capacity, with Cowan (2010) estimating the number of items that can be held and processed at any one time to be around four. Cognitive Load Theory (*eg* Sweller *et al*, 1998; and Chapter 4 of this book) is primarily concerned with the limits of working memory. The theory is centred around the way in which a learner's cognitive resources are focused and used during problem-solving, suggesting that for instruction to be effective, care must be taken to not overload the mind's capacity for processing information. If working memory experiences cognitive overload, no learning may take place. We can get around working memory's restrictive limit by chunking related information, carefully designing instruction and automating key knowledge.

Putting these related approaches together, we are able to form a very simplified model of how students think. This will be constantly revisited and

expanded upon throughout this book, but the fundamentals will remain the same. The model looks like this:

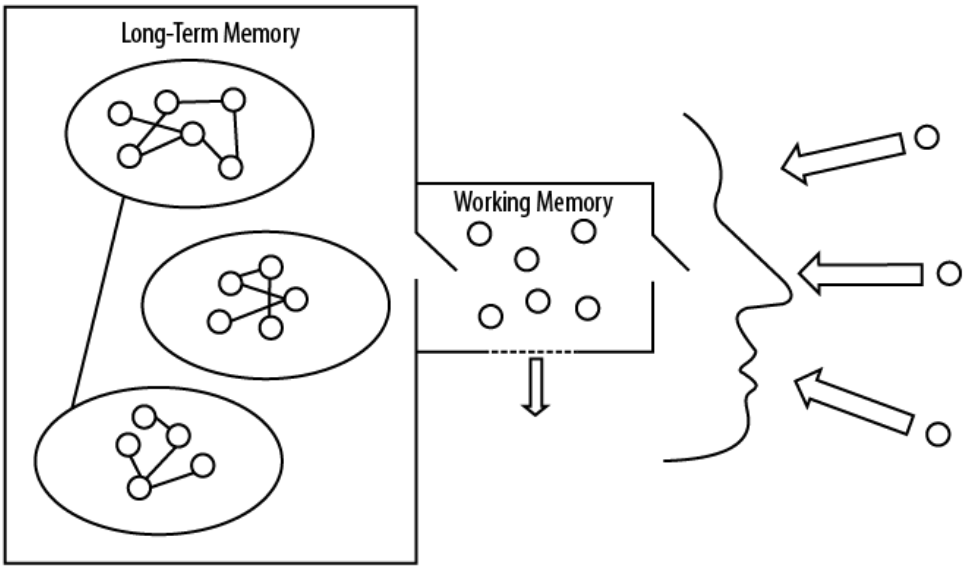


Figure 1.1 – Source: Craig Barton

The circles represent units of knowledge, the ovals are schemas, and the lines represent connections. Thinking takes place in working memory, focused on the interplay between the environment and what we retrieve from long-term memory. The more knowledge we have stored and organised in long-term memory, and the more of this knowledge that is automated, the easier thinking is and the more we can think about. Knowledge helps students take in more information, think about new information, and remember new information.

This simple model has huge implications for teaching and learning that we will delve into in the remainder of this chapter and throughout this book.

How students learn

There are many different definitions about what learning is, but the one I am going to opt for throughout this book is provided by Kirschner, Sweller and Clark (1998), who define learning as '**a change in long-term memory**', going on to say that if nothing has changed in long-term memory, then nothing has been learned. Working memory is the vehicle that instigates this change.

It is worth noting that not everyone agrees with this definition. In a 2017 blog post, Daniel Willingham points out that this definition does not specify that the change in long-term memory must be long-lasting (so does that mean that a change lasting a few hours qualifies?), nor does it specify that the change must lead to positive consequences (does a change in long-term memory that results from Alzheimer's disease qualify as learning?). In Chapter 12 we will address the first point – the durability of learning, and how we can improve it. The second point is beyond the scope of this book, but in Section 3.8 we will consider the acquisition of incorrect knowledge, and the difficulties of resolving it.

Our definition of learning implies that knowledge in long-term memory is not static. We acquire brand new knowledge, we adapt, change or accommodate existing knowledge based upon new experiences and information, and knowledge may become more or less accessible. This is all summed up beautifully by Mccrea (2017) who explains that knowledge 'is constantly evolving and decaying as a result of our thinking and interaction with the environment. Our long-term memory is more like a forest than a library'.

For Coe (2013), **learning happens when people have to think hard**, and indeed making changes to long-term memory is likely to be effortful. There are two main pathways through which such a change may occur, and – in a sense – they travel in opposite directions to each other:

From working memory to long-term memory
Students learn new ideas by reference to ideas they already know. A learner holds information in working memory, and then makes connections between that information and knowledge already stored in long-term memory (assuming such knowledge is present). For Willingham (2009), '**understanding is remembering in disguise**' – it is taking correct old ideas from long-term memory, getting them into working memory, and rearranging them in a new order to make new connections. So, if a student encounters algebraic fractions for the first time, they will (hopefully) make connections between their existing organised knowledge of non-algebraic fractions, rules of algebra, factors and so on. If they are unable to do this – either due to cognitive overload or because such knowledge does not exist in

their long-term memory – then learning is unlikely to take place. However, if new ideas, information, facts and procedures are successfully processed in working memory, they may become assimilated into an existing schema or create a brand-new connected one, thus changing long-term memory.

From long-term memory to working memory
When information is successfully retrieved from long-term memory into working memory, its representation in long-term memory is changed such that it becomes more accessible in the future. Using our memories changes our memories – or as Bjork (1975) put it, '**retrieval is a powerful memory modifier**'. Thus, the act of retrieval can result in learning. This rather surprising pathway to learning will be the focus of Chapter 12, and I promise it is worth waiting for.

What I do now

I think a lot about…well, thinking. Specifically, how can I design my teaching to ensure:

1. …my students are thinking about the right things?
2. …what they think about is transferred to long-term memory?
3. …once it is in long-term memory, it can be accessed?
4. …it can be applied successfully to different situations?

I will try my very best to answer these questions in this book.

OceanofPDF.com

1.2. Experts and Novices

What I used to think

I believed I could help novice learners become experts simply by treating them as experts. I would model how experts think, give them the same work as experts, and encourage them to struggle. I would do this in much the same way as I always say to my new crop of Year 12 students each September: 'I am going to treat you like adults, so you had better behave like adults'.

Astonishingly, neither my Year 12s nor my novice learners responded in quite the way I had hoped.

Sources of inspiration

- Barton, C. (2017) 'Andrew Blair', *Mr Barton Maths Podcast*.
- Didau, D. (2017) 'A Novice→Expert Model of Learning', *The Learning Spy Blog*.
- Hill, N. M. and Schneider, W. (2006) 'Brain changes in the development of expertise: neuroanatomical and neurophysiological evidence about skill-based adaptations' in Ericsson, K. A., Charness, N., Feltovich, P. J. and Hoffman, R. R. (eds) *The Cambridge handbook of expertise and expert performance.* Cambridge: Cambridge University Press, pp. 653-682.
- National Research Council (2000) *How people learn: brain, mind, experience, and school: expanded edition*. Washington, DC: National Academies Press.

- Wiemann, C. (2007) 'The "curse of knowledge", or why intuition about teaching often fails', *APS News* 16 (10) p. 9.

My takeaway

If I am ever to get an education-related tattoo (it is only a matter of time), it would say '**experts and novices think differently**'. The distinction between expert learners and novice learners, and the implications for teaching, is at the heart of this book. It may well be the single most important thing I have learned from my reading and speaking to guests on my podcast. Experts do not simply know more than novices, they think in a fundamentally different way.

What is an expert?

Before we go any further, we need to define what we mean by *expert* and *novice*. They are terms used frequently throughout the literature, but as Andrew Blair pointed out on my podcast, how do you know when someone has reached a level of expertise?

Oxforddictionaries.com define an expert as '*a person who is very knowledgeable about or skilful in a particular area*'. You could make a convincing case that students never achieve such status in the subject of mathematics whilst at school, perhaps only becoming experts at Degree or Masters level. However, I think we can observe expertise in *narrow domains* within both primary and secondary school mathematics. Students can become experts in specific areas such as times tables, adding fractions and solving linear equations. Of course, this raises questions such as which times tables, what kind of fractions, and how complex are the equations, all of which are perfectly valid. So, as a rule, the more narrow the domain, the more confidence we can have in attributing expertise.

Expertise is domain-specific and not person-specific. I am an expert at circle theorems (in fact, I cannot get enough of them), but absolutely useless at cooking, fixing things, and a whole host of other skills that would probably be more useful in my life than the ability to spot a subtle use of the alternate segment theorem. This is also true for students within our subject. It is too easy to label all top-set students as experts and all bottom-set students as novices when it comes to maths. I have taught plenty of students who excel

at all things algebraic, and yet really struggle with certain geometrical aspects of mathematics, such as identifying 2D right-angled triangles in 3D trigonometry problems. Even within a relatively narrow domain such as mathematics, everyone is a novice in something, especially the first time they encounter a concept. As we shall see throughout this book, labelling and treating someone as an expert when they are not could have serious consequences for their learning.

Likewise it is not true that people are either an expert or a novice in any particular domain, and nothing in between. The journey from novice to expert is continuous and fuzzy, with no clear indication of when you have reached your destination. It is certainly not the case that during one particular period of instruction, or midway through a certain train of thought, a light bulb illuminates and a student magically achieves the status of expert in the way they might be awarded a badge in *Call of Duty*. People will be at different stages of the journey at different points in time. Just as importantly, factors such as a lack of practice, poor instruction and the development of misconceptions – all of which will be addressed in this book – may mean that students are not always travelling the right way.

Given these difficulties, it may be tempting to conclude that the terms *expert* and *novice* are nothing more than meaningless labels, and hence abandon them. But I feel that would be a mistake. The differences between how those with strong domain-specific knowledge think and learn compared to how those with less domain-specific knowledge do so are significant – and have huge implications for how we teach that will be recurring time and time again throughout this book. Hence, whilst no one is truly an expert or a novice, as teachers we should attempt to determine the relative expertise of our students in specific areas wherever possible. This is difficult, but the techniques of formative assessment (Chapter 11) and low-stakes tests (Chapter 12) can provide valuable information to assist us in making such judgements.

So, when you encounter the terms *expert* and *novice* throughout this book, treat them with caution, but try not to dismiss them. Think of expert learners not as those belonging to some discrete, absolute category, but as

those further along the path to expertise. Likewise, think in terms of domain-specific expertise, making that domain as narrow as possible.

How experts and novices think differently

Hill and Schneider (2006) explain that expert learners differ from novices in terms of their knowledge, effort, recognition, analysis, strategy, memory use, and monitoring, and that these differences are due to the structure of their long-term memories. As we learn, our brain architecture changes and thoughts are processed differently. This means that as we move to mastery of a given skill or concept, our brains form different links in long-term memories, and it is actually possible to observe different activation patterns during problem-solving. They conclude that 'in addition to processing efficiency, enriched representations, and structural expansions, experts can flexibly use strategies, by recruiting the associated brain regions, to solve a range of problems, whereas novice performers can not'.

For Didau (2017), there are two hallmarks of expertise, and these are directly relevant to mathematics:

1. Automaticity of foundational knowledge

Have a look at the following problem:

Find 25% of £300

How did you get on? How long did it take you to arrive at the answer of £75?

I imagine you answered this question pretty quickly and without too much effort. The answer of £75 probably came to you almost instantly, without you really being conscious of what you had done. You immediately and unconsciously recognised the percentage sign and what it means, and you know that 25% is the same as a quarter. You have automated this knowledge, and hence it takes up no space whatsoever in your working memory. You have also automated the procedure for finding a quarter of something, so you can quickly work out a quarter of £300 by halving and halving again without imposing too much strain on your working memory. You can do all of this because you are an expert. You could also probably do this calculation if music was playing in the background, your partner was asking

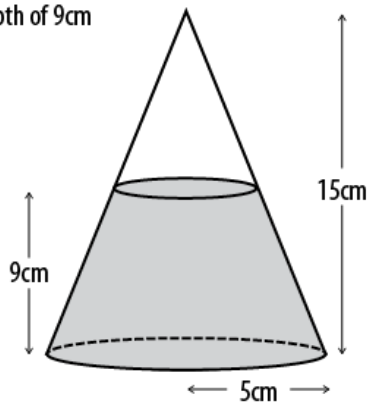
why you haven't done the washing up, and you were attempting to juggle. The fact you have automated much of the knowledge necessary to answer the maths question frees up space in your working memory to attend to other things.

A novice learner is also likely to be able to find 25% of £300, although it will be more cognitively demanding. Interpreting what the question is asking may not be instantaneous, they may not be aware that 25% is the same as a quarter and so may start by finding 10% and need to consider how to do this, and then again with 5%, and then they will need to add all these together. You both get the question right, but it is less cognitively demanding for you.

However, the real difference between domain-specific experts and novices when it comes to the automation of knowledge can be seen in more difficult problems. Consider the following higher tier, non-calculator GCSE question:

Volume of cone = $\frac{1}{3}\pi r^2 h$ where r is the radius and h is the perpendicular height.

A cone has a
horizontal base of radius 5cm
height of 15cm
The cone contains water to a depth of 9cm



Work out the volume of the water, in cm^3
Give your answer in terms of π. **[4 marks]**

Figure 1.2 – Source: AQA 2017 GCSE Maths Higher Paper 1

Does the answer arrive instantly this time?

Now, this problem is likely to be cognitively demanding for both you and a novice learner. But the difference is that you have automated all the basic knowledge involved. You do not need to expend any effort contemplating what each term in the formula means, how to square 5, how to find a third, or what to do with the pi symbol. This frees up cognitive capacity to attend to more global features of the problem, such as correctly interpreting the question and devising a strategy to get to the solution. Not only are you more likely to solve the problem, you are also more likely to learn from the experience. A novice learner, who does not have the relevant knowledge automated, is likely to get bogged down in the minutiae of the problem, experience cognitive overload, and potentially not learn anything transferable. More on this, and the consequences for us as teachers, in Chapter 9 on problem-solving.

2. Ability to see deep structure within domains of expertise

Understanding what problems are really about can be one of the toughest skills for students to master. But here's the thing – it is not a skill, at least not in the same way that adding together two fractions is a skill. It is a feature of being an expert. Consider the following question, this time taken from the GCSE foundation calculator paper, and ask yourself what the question is about:

The average age of teachers at a school is 36 years.

Mr Smith's age is $\frac{11}{9}$ of the average.

How old is Mr Smith?

[2 marks]

Answer _____ years

Figure 1.3 – Source: AQA 2017 GCSE Maths Foundation Paper 2

How long did it take you to discern that this question is essentially asking you to find $\frac{11}{9}$ of 36? How did you manage to avoid the irrelevant, potentially distracting, surface structures such as 'age', 'school', and possibly the toughest to avoid of all, 'average'?

You probably did this pretty quickly, and you did so because you are an expert. Your experience of answering lots of related questions, the significant volume of knowledge you have stored in long-term memory and – no less significant – the way that knowledge is organised allows you to avoid the surface features and identify the problem's deep structure.

Would all your students be able to do this – to get to the heart of what the question is asking? I know many of mine wouldn't. The ability to identify the deep structure of a problem, whilst avoiding the surface structure, is a key feature of expertise, and has huge implications for problem-solving, which will be discussed in Chapter 9.

The curse of knowledge

These two key differences between novice and expert learners can be difficult for us teachers to appreciate for the very reason that we ourselves are experts. We suffer from a '**curse of knowledge**' (*eg* Wieman, 2007), whereby when you know something it can be difficult to think about it from the perspective of someone who does not know it.

Much of our mathematical knowledge is automated so we carry out basic processes without imposing significant strain on our working memories. Likewise, we immediately recognise the deep structure in problems without really knowing how. This can make it – for me, at least – very difficult to see things from a student's perspective, empathise with the difficulties they have, and help them overcome them. I simply cannot remember what it was like to not be able to answer the questions in this section, and that is a problem.

The curse of knowledge is why subject knowledge alone is not enough to be a good teacher, and why asking students who have mastered a concept to help someone who is struggling sometimes does not always work. Novices are unable to think like experts, and experts are unable to think like novices,

and the resulting exchange can be a frustrating experience for both parties. Understanding how students think and struggle comes with experience, but an appreciation of the misconceptions they are likely to hold is of vital importance. We will discuss this in Chapter 11.

In the past, my approach to helping students think like an expert has been to simply model expertise. By showing students how I approach and solve problems, I thought I could help them become experts as well. This is not a completely fruitless exercise, but it neglects one incredibly important point – the only reason I can think like an expert is because I have vast amounts of domain-specific knowledge. That is why I can automate and recognise deep structure. Until this knowledge has been taught, practised, assessed and retrieved, attempts to enable our students to think like experts may well be in vain.

What I do now

As a general rule, I have assumed my students are experts before they are, or I have tried to fast-track them on the road to expertise by modelling how an expert thinks. I am now acutely aware of the fundamental importance of domain-specific knowledge. It is the distinguishing feature between expert and novice learners. Such knowledge helps our students think better, acquire new knowledge, self-explain, solve problems, and become the independent learners we want them to be. It also helps determine the optimal instructional techniques we should employ to help them learn effectively. Without such knowledge, we condemn our students to be forever novices, and no amount of showing them how to think differently will help.

This book is primarily about the best way to help students acquire and retain that knowledge.

OceanofPDF.com

1.3. What are they thinking about?

What I used to think

I used to think that lessons needed to be memorable to make them…well, memorable. This led to 'The Swiss Roll Incident': Bolton, Lancashire, 2014.

That year I taught a delightful Year 7 class, and we had just come to the end of our unit on fractions. We had covered all of the basics, and now it was time for them to apply their newly acquired skills to some different contexts. Hence, I set them the following problem to think about overnight:

Imagine you have 7 Swiss rolls, and you can stack them on top of each other. What is the fewest number of cuts you must make to ensure 12 people get the exact same amount of Swiss roll, and there is nothing left over?

That evening, whilst my students were pondering this tasty problem, I took a trip to Tesco. £12.96 later (these were deluxe Swiss rolls, I will have you know), I was tooled-up with all I needed for what was sure to be an amazing lesson.

And indeed it was.

Armed with a plastic knife and a bundle of paper towels, I set about presenting a real-life, practical solution to the Swiss Roll problem. Cream and jam went everywhere, and a particularly tasty looking quarter of a Swiss Roll was lost forever to the dusty classroom floor. But in the end we had done it. Maths in action, and the kids were loving it.

At the time of writing, these Year 7s are now in Year 11, and unfortunately I no longer teach them. However, I happened to see one of them in the corridor the other day, and after exchanging a few pleasantries, I asked her how she was getting on in maths:

‘Not bad, sir, but it’s a bit boring.’

‘What do you mean? Maths can never be boring!’

‘Well, it’s alright, I guess, but we don’t do anything fun like that Swiss roll lesson we did in Year 7.’

‘Veronica, just out of interest, do you remember what that lesson was about?’

‘Swiss rolls!’

And there we have the problem.

Sources of inspiration

- Barton, C. (2017) ‘Peps Mccrea’, *Mr Barton Maths Podcast*.
- Coe, R. (2013) *Improving Education: A Triumph of Hope ever Experience, CEM Inaugural Lecture*. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf
- *TES* (2015) ‘Experimental Probability – Bottle Flipping’. Available at: https://www.tes.com/teaching-resource/experimental-probability-bottle-flipping-11478338
- Willingham, D. T. (2003) ‘Ask the cognitive scientist. Students remember…what they think about’, *American Educator* 27 (2) pp. 37-41.
- Willingham, D. T. (2008) ‘What will improve a student’s memory?’, American Educator 32 (4) pp. 17-25.
- Willingham, D. T. (2009) *Why don’t students like school?: A cognitive scientist answers questions about how the mind works and what it means for the classroom.* Hoboken, NJ: John Wiley & Sons.

My takeaway

As soon as the ink has dried on my ‘experts and novices think differently’ tattoo, the next on my list is ‘**memory is the residue of thought**’. I do not

think there are two more important phrases for teachers to be aware of. Both will feature prominently throughout this book, but it is the second of these delights, coined by Willingham (2003), that is the cornerstone of this section.

Memory is the residue of thought, so students remember what they think about.

Consider the model of thinking introduced in Section 1.1. Before we start worrying about things like working memory capacity, relevant schemas, transfer and problem-solving, we need to ensure that the items being processed in working memory are the items that need to be processed. Learning is a change in long-term memory, but if the relevant ideas, skills and concepts have not found their way into working memory in the first place, then learning is not going to happen – or at least not the type of learning we intend. It is like having a top of the range juice-maker, but continually feeding it parsnips.

In the lesson I described above, I had hoped my students would have been considering fractions. Instead, they were undoubtedly thinking about jam. Whereas I wanted their working memories to be filled with common denominators, factors of 12 and families of equivalent fractions, instead it was filled up with thoughts of cream, how hungry they were, and the fact Mr Barton managed to get some jam on his glasses. Students remember what they think about. So, was it any wonder that, four years on, Veronica remembered Swiss rolls but not the content of the lesson, let alone the solution to the problem I posed (which is 3 cuts, by the way)?

Once I have the mantra 'memory is the residue of thought, so students remember what they think about' tumbling around in my head when I am weighing up lesson activities, I start to view things in a whole new light. Hence, when I came across across a five-star-rated resource on *TES* entitled 'Experimental Probability – Bottle Flipping', I shuddered.

This is the resource description provided by the author:

> I created this PowerPoint as an engaging end of term activity for my low attaining year ten class to investigate the relationship between

> amount of water in a bottle and the chances of it being flipped successfully. I used the resource to revise experimental probability but the wider investigation covers collecting data, a bit of FDP *etc*. I know bottle flipping is a bit of a sore point … for some but my class were fully engaged and it helped give them an appreciation of experimental probability.

I am sure the students were engaged, but engaged in what? The finer points of experimental probability, or the best technique to flip a bottle of water?

Which brings us nicely onto the subject of engagement. It is an oft-cited claim that students need to be engaged to be learning, and hence attempting to engage students can become an overriding aim of teaching. But does engagement actually lead to more learning? According to Coe (2013), not necessarily. In his much-discussed list of 'Poor Proxies for Learning', Coe includes occasions when 'students are engaged, interested, motivated'.

Now, Coe is not saying that engagement prohibits learning taking place, nor is he dismissing the possibility that engagement can support learning. He is saying that observing engagement alone does not imply learning. And the Swiss roll and Bottle Flipping lessons are classic examples of this. I am sure if someone walked into the room whilst those lessons were taking place they would have been met with a sea of engaged, interested and motivated students. The problem is, of course, what exactly were they engaged in? Unless we have further evidence (a test of retention, for example), we must be extremely careful in concluding that learning is taking place.

All of this still leaves us with a question: would Veronica have remembered that fractions lesson four years on had I left my Swiss rolls at home? Probably not. But would she have remembered the maths involved? I think so. I cannot remember the specific moment I learned to add fractions, or solve equations, or draw tree diagrams. But I can do them. And that is because of lessons where as much of my attention as possible was directed to thinking – and thinking hard – about the maths the teacher intended me to think about.

So, am I saying that maths lessons should be boring? No. I firmly believe – and as I hope to demonstrate in the chapters that follow – that all students can draw motivation and pleasure from learning and achieving in mathematics. But I am saying that there is a real danger that students can latch onto the surface structures of the experiences we provide, so much so that they fill up their working memories and leave little space for anything else. And if this is the case, then we had better be extremely careful when deciding what those surface structures are, for they are the things students are likely to remember.

What I do now

Willingham (2009) offers the following advice: 'review each lesson plan in terms of what the student is likely to think about. This sentence may represent the most general and useful idea that cognitive psychology can offer teachers'.

When I interviewed Peps Mccrea for my podcast, he suggested what I feel is a useful refinement to Willingham's original statement: **students remember what they *attend to***. Whilst the sentiment is the same, Peps made the point that as teachers we may have more success directing what students attend to rather than what they think about.

So, each time I plan a lesson I ask myself one question: *at this stage of the lesson, what will my students be attending to?* And if the answer is not what I need them to be attending to, I change my plan. Of course, I cannot control the focus of my students' attention, and even the most expertly planned lessons can fall victim to students' wandering minds. However, I can certainly boost my chances via the activities I choose, and also via the general principles of good instructional design discussed in Chapter 4.

I am also far less concerned about engagement – it comes a far distant second when planning lessons to considerations of learning, because I have realised that the path from engagement to learning is not as clear-cut as I previously thought. Of course I want my students to be interested in the work we are doing, but I feel there is a more reliable and sustainable way of achieving this, which will be covered in Chapter 2.

Hence, no more Swiss rolls or bottle flipping. No more ‘make a poster or a PowerPoint’ for revision, with students’ working memories more consumed with thoughts of colours, fonts and animations than the mathematical content. No more Tarsia jigsaws that the students need to cut out themselves, lest concerns about the neatness of the triangular pieces divert their attention away from the maths. And no more loop cards, with students possibly more concerned about when their turn is coming, suffering anxiety about speaking in public, or switching off when their turn has passed.

OceanofPDF.com

1.4. Expanding working memory capacity

What I used to think

My limited knowledge of working memory led me to what I considered to be an incredibly profound thought: if I could help boost the size of my students' working memories, then surely I could enable them to store more information and hence think better? Such a revelation led me to the inevitable conclusion that the key to successful thinking was *brain training*. After all, there is a multi-million pound industry built around the promise of improving people's working memories.

So, for an entire term, my Year 9 students would engage in their regular weekly dose of in-lesson brain training, aptly named as *The Mr Barton Brain Training Programme*.

A favourite website of mine was brainmetrix.com. Their worryingly addictive Memory Game involves remembering a sequence of flashing rectangles, and contains the following promising description: 'Here you will be able to test your memory, testing it will help you improve it, this game simulates many areas in your brain responsible for storing and retrieving information, you will be enhancing your memory while still having fun'.

Perfect, that's exactly the kind of scientific insight my students need, I thought. This was bound to work.

However, after persisting for a number of weeks, and encouraging my students to practise whenever they got a chance at home, something very interesting happened. The students did indeed appear to have successfully expanded the capacity of their working memories by playing the Memory Game – indeed, one student reached an incredibly impressive Level 21, with

me peaking at Level 12 – and yet their mathematical thinking did not seem to improve at all.

Something was not right.

Sources of inspiration

- Engle, R. W. and Kane, M. J. (2003) 'Executive attention, working memory capacity, and a two-factor theory of cognitive control' in Ross, B. H. (ed.) *Psychology of Learning and Motivation, Vol 44*. San Diego, CA: Elsevier Academic Press, pp. 145-199.
- Melby-Lervåg, M. and Hulme, C. (2013) 'Is working memory training effective? A meta-analytic review', *Developmental Psychology* 49 (2) p. 270.

My takeaway

The first point to make is that (for once) my intuitions were correct. A larger working memory capacity does appear to make thinking easier. As Engle and Kane (2003) put it, 'one of the most robust and, we believe, interesting, important findings in research on working memory is that working memory capacity span measures strongly predict a very broad range of higher-order cognitive capabilities, including language comprehension, reasoning, and even general intelligence'.

This makes perfect sense when we consider our model of thinking in Section 1.1. The larger the working memory capacity, the more ideas, information, facts and procedures can be held there and processed at any one time. Indeed, one could argue that if we could expand working memory capacity indefinitely, then the importance of the systematic acquisition of knowledge that will be a theme throughout this book flies out of the window. Why would you bother going to the hassle of learning and storing all that knowledge if you can just chuck it all into working memory whenever you need it, process it, and spit it back out?

Therefore, wasn't my Year 9 student who got to Level 21 effectively demonstrating that she could hold 21 units of information in her working

memory all at once?

Unfortunately not. As the meta analysis by Melby-Lervåg and Hulme (2013) concludes, any benefits from such exercises are short-term, and crucially do not transfer to other situations. So, whilst you may be able to train your brain to hold more of a specific type of written or oral information, this will not transfer across to enable you to solve more complex maths problems. My Year 9 had probably learned a very sophisticated way of chunking together those 21 separate units of information, but alas such a skill will not allow her to process a 21-step algebra problem with equal ease. Indeed, given that she was spending time at home getting better at counting purple flashing shapes instead of working on her algebra, it is fair to say the effect on her mathematical achievement was the opposite of what I wanted.

What I do now

Sadly, I have disbanded the *Mr Barton Brain Training Programme*. However, all is not lost. Instead of focusing on expanding the size of working memory, I ensure it is filled up with the right things, using the principles of Cognitive Load Theory and the Cognitive Theory of Multimedia Learning that will be discussed in Chapter 4. I also make use of the unlimited capacity of long-term memory by ensuring students have acquired and practised domain-specific knowledge. I achieve this using the principles of Explicit Instruction (Chapter 3) and carefully planned exercises (Chapter 7). This in turn helps students form the schemas that make thinking and problem-solving easier.

OceanofPDF.com

1.5. Methods that last

What I used to think

One way I used to approach my planning was to look at the kind of questions students needed to answer by the end of a given topic unit, and teach them the best method to enable them to answer these questions. What could possibly be wrong with that?

Well, let's take the topic of solving linear equations, and how I used to present it to my Year 7s. A quick glance at the scheme of work revealed that by the end of the *Introduction to Linear Equations* topic unit, students had to be comfortable with: *solving two-step linear equations with the variable appearing on one side, some of which involve negative numbers* (we will go to town on concepts such as interleaving and the order of schemes of work in Chapter 12).

Now, even back then I was against using the 'change side, change sign' approach, considering it a method prone to misuse (I lost count of the number of times a 'divide by negative 5' got magically transformed into a 'multiply by positive 5'), and one which was not built upon solid mathematical foundations. However, I was a huge fan of the approach known as 'I am thinking of a number'.

My lesson would start with something like this…

I am thinking of a number, I add 5 and I get 8. What was my number?

The students would happily tell me it was obviously 3. How did they know, I would ask them. Eventually we would tease out that you could subtract 5 from 8 to get the answer.

I then proceeded to explain that we could save time (students love saving time) by writing the problem as $n + 5 = 8$, alongside the presentation of the following diagram, which subtly introduced the concepts of a variable and inverse operations:

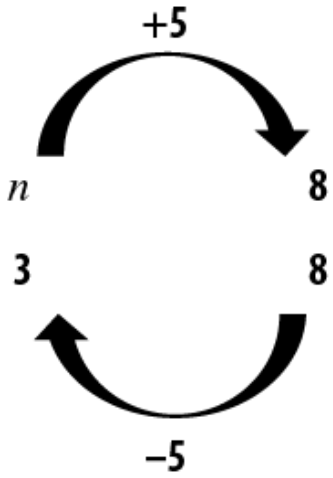


Figure 1.4 – Source: Craig Barton

The beauty of this approach (or at least, so I thought), was that we could tackle more complex problems in exactly the same way. So, *I am thinking of a number, I multiply by 2, then subtract 1, and I get 11*, became $2x - 1 = 11$, which we could solve using:

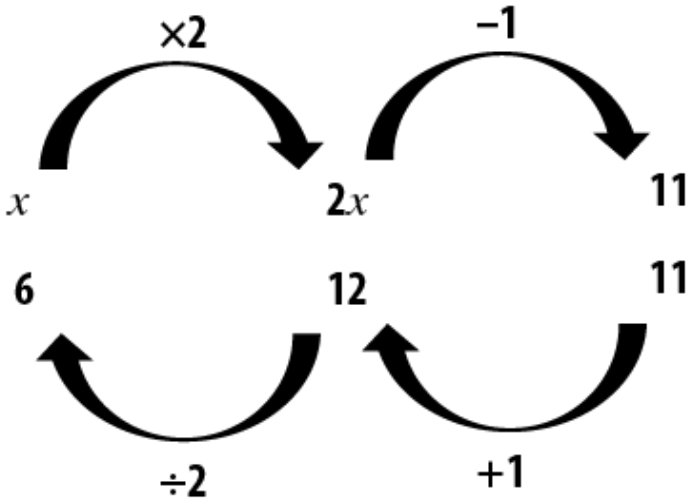


Figure 1.5 – Source: Craig Barton

Indeed, by the end of the lesson, my beaming Year 7s could even answer a question – 'a Grade A GCSE question!', as I proudly told them – like this:

$$\frac{4\,(2x - 3)}{5} + 6 = 10$$

My students had made rapid progress, algebra was no longer scary, and they even had a method that could be used to rearrange equations. This very lesson led to me receiving an Outstanding in my first Ofsted inspection as an NQT. High-fives all round. I had mastered this teaching lark in my first few months.

I guess the only fly in the ointment was the fact that my students would not have a clue how to solve equations such as $5 - x = 3$ or $3x - 2 = x + 4$. But these equations wouldn't come up until Year 8, so what was the point in worrying about them now?

Sources of inspiration

- Anderson, J. R. (1996) 'ACT: A simple theory of complex cognition', *American Psychologist* 51 (4) p. 355.
- Barton, C. (2017) 'Dani Quinn: Part 1', *Mr Barton Maths Podcast*.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Using the model of thinking introduced in Section 1.1, let's consider what happens when students who have been taught via 'I am thinking of a number' encounter an equation along the lines of $3x - 2 = x + 4$. The sight of the x and the equals sign allows them to immediately recognise this as an equation. This triggers a search for relevant schemas inside their long-term memory. Sure enough, they are likely to find a well-developed schema for solving equations. The problem is, none of the them look like this one.

So what does the student do? Well, they may attempt to solve it using an inefficient strategy – some variation of trial and improvement, perhaps – which is likely to place a great burden upon their fragile working memory. As we shall see in Chapter 9, such an approach may not be conducive to learning, even if students do get the right answer. Or they may try to apply their existing, deficient schema, which will inevitably lead to an incorrect solution and the possible development of a minefield of misconceptions that will need resolving. In fact, probably the most beneficial thing students could do for all concerned is to leave the question out.

Imagine that you are the Year 8 teacher who has inherited my class of students brought up on a diet of 'I am thinking of a number'. What do you do? Well, your only real option is to teach students another way of solving linear equations that works for these new types of questions – probably some variant of the balance method. Then you have two choices: either tell students to abandon their old way, which they will probably be reluctant to do as it is so easy, and miles better than the new, tricky way you are trying to teach them; or give them a complex set of rules which enables them to spot which questions require which methods, which effectively means they need a different schema for each variation, and maths quickly becomes the disparate set of meaningless rules that countless scores of children and adults view it as. Either way, you are likely to be fighting a losing battle, and it is undoubtedly all my fault.

Knowing when to make thinking easy and difficult for students will be a recurring theme throughout this book, and here is an example of when a more difficult formal method beats a quick-fire short-cut any day of the week.

It was through my podcast interview with Dani Quinn, the Head of Maths at Michaela Community School, that I first really considered this. When it comes to solving equations, Dani teaches her students the formal balance method from Day 1.

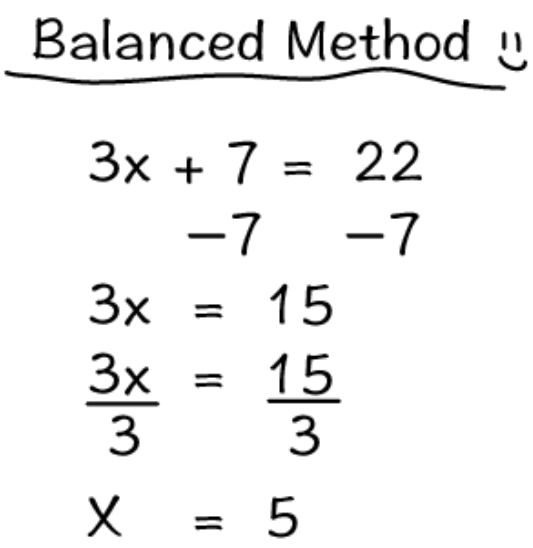


Figure 1.6 – Source: Craig Barton

Moreover, this is a department-wide approach to teaching solving linear equations which is non-negotiable. It ensures consistency among her teachers and a firm algebraic foundation for her students. Sure, this method is initially more difficult for students to grasp, and of course it looks like they are making far less 'progress' than the students racing through 'I am thinking of a number' problems, but the balance method for solving equations is one that will last.

What I do now

I think very carefully about the methods I teach my students, asking myself:

- How long will this method last them?
- How do I move them on from this method?
- Is it built upon solid mathematical foundations?

If necessary, I teach them in a way that is more difficult in the short-term, but which leads to greater long-term gains. And if I feel my students are not ready for this more difficult method, then I come back to it when they are.

OceanofPDF.com

1.6. Maths anxiety

What I used to think

I used to be of the opinion that there were two types of students who were anxious about maths:

1. The ones who struggled, who were anxious about maths precisely because they were no good at it.
2. The clever students, who claimed they were anxious, but secretly they were absolutely fine.

I found the first group easier to deal with than the second as I could understand their anxiety. When I was at school, I was absolutely terrible at CDT (it's a subject that doesn't really exist any more, but it basically involved me trying – and failing – to build a miniature set of shelves for an entire year, armed with wood, a hammer and some nails). The lesson was on a Thursday morning, and the thought of it would plague my entire week. I would genuinely lie awake all Wednesday night dreading it. During the lesson itself, I was a quivering wreck. Simple tasks, such as hammering in a nail or even speaking in coherent sentences, were suddenly beyond me. It even got to the stage that once the relief of having done the lesson subsided around Thursday lunchtime, the fear of next week's lesson would immediately kick in.

I was anxious about CDT because I was rubbish at CDT, and hence I could fully understand why some students in my lower set maths classes describe the same feelings about maths. What I found harder to explain, however, were high-flying students who claimed maths instilled the same set of fears despite notching up a string of top grades. In my mind, these students were Drama Kings and Queens.

Looking back, I am simply ashamed of my ignorance of the severity of maths anxiety.

Sources of inspiration

- Ashcraft, M. H. (2002) 'Math anxiety: Personal, educational, and cognitive consequences', *Current Directions in Psychological Science* 11 (5) pp. 181-185.
- Beilock, S. L. and Willingham, D. T. (2014) 'Math anxiety: can teachers help students reduce it? Ask the cognitive scientist', *American Educator* 38 (2) p. 28.
- Caviola, S., Carey, E., Mammarella, I. C. and Szucs, D. (2017) 'Stress, time pressure, strategy selection and math anxiety in mathematics: a review of the literature', *Frontiers in Psychology* 8, p. 1488.
- O'Leary, K., Fitzpatrick, C. L. and Hallett, D. (2017) 'Math anxiety is related to some, but not all, experiences with math', *Frontiers in Psychology* 8, p. 2067.
- Ramirez, G., Gunderson, E. A., Levine, S. C. and Beilock, S. L. (2013) 'Math anxiety, working memory, and math achievement in early elementary school', *Journal of Cognition and Development* 14 (2) pp. 187-202.
- Wong, M. and Evans, D. (2007) 'Improving basic multiplication fact recall for primary school students', *Mathematics Education Research Journal* 19 (1) pp. 89-106.

My takeaway

Ashcraft (2002) defines maths anxiety as '**a feeling of tension, apprehension, or fear that interferes with maths performance**'. Indeed, one of the really sad things about maths anxiety is that it stops students reaching their potential. According to figures quoted by Beilock and Willingham (2014), in the United States an estimated 25% of four-year college students and up to 80% of community college students suffer from a moderate to high degree of maths anxiety. Most students report having at

least one negative experience with maths at some point during their schooling.

Ashcraft summarises 30 years of work into the study of maths anxiety. There are a number of key points raised in the paper, but here is a selection that caught my eye:

1. Highly maths-anxious individuals avoid maths.
Now, of course our secondary school students cannot formally opt out of studying maths until A Level, but they can informally do so. They can study less at home, take part less in class, and talk themselves out of even so much as attempting a problem. And with practice in maths being so vital to success, we end up in a vicious cycle with avoidance leading to poorer performance, which inevitably adds to the existing anxiety.

2. Maths anxiety is only weakly related to overall intelligence.
This was a surprise to me. I had assumed a strong, negative relationship between intelligence and maths anxiety. But when I think about it, I can recall many students I have taught over the years who were highly able mathematically, and yet had (what seemed to that stupid, uninformed me) an irrational fear of the subject.

3. Timed tests seem to cause anxiety.
Ashcraft's third claim poses us with two major problems. First, every major exam that students encounter has a timed element to it, and so at some point students need to be able to cope with it. Second, as Wong and Evans (2007) find, without some time pressure students are unlikely to develop the kind of automated knowledge of key number facts that they need to free up capacity in working memory. For example, a student solving 6×8 by counting up in 6s on their fingers will be at a significant disadvantage when solving a complex problem that requires such knowledge compared to a student who can recall the answer of 48 instantly. Perhaps the key is to introduce the timed element slowly and carefully, in a supportive atmosphere, together with a policy of rarely collecting in or announcing students' scores. This is related to the low-stakes quiz approach that we will cover in Chapter 12.

It is also worth noting that the association between timed tests and maths anxiety is by no means universally accepted. Caviola *et al* (2017), in a literature overview of the past 30 years on the effect of stress and/or time pressure on math proficiency, conclude 'assuming causal relations between time pressure and inducing math anxiety currently does not have evidential support'.

4. Maths anxiety lowers performance because it takes up vital space in working memory.

On reflection, I experience a form of maths anxiety each Tuesday evening in our A level Maths Clinics. Despite having taught A level maths for 12 years, most weeks I am caught out by a tricky question posed by a struggling Year 13. Usually it is Part d) of a 13-mark epic, and as I frantically scan the question and try to piece together the answer on the spot, I can feel the expectant eyes of the student on me. And as the seconds pass and their confidence in me decreases, I physically feel something inside my head. There seems to be something blocking my thoughts. And the more I worry and focus on this feeling, the more dominant it becomes, and the less I am able to process the maths necessary to solve the problem. It is a vicious cycle that is often only resolved by me faking the need to nip to the toilet so I can go outside and compose myself. Ashcraft provides an explanation for what is happening to me, and what also happens to many students on a regular basis – maths anxiety takes up valuable space within working memory. As we saw in our model of how student think in Section 1.1 – and will develop further in Chapter 4 when we come to look at Cognitive Load Theory – students' working memories are fragile. In order to process thoughts successfully, students need all the capacity they can get. Hence maths anxiety can be considered an extraneous load that gets in the way and hence hinders thinking and learning.

5. Maths anxiety does not lower performance in all areas of maths – just the more cognitively demanding ones.

This can again be seen in my experiences in A level Maths Clinics. Basic arithmetic processes – particularly if they have been automated – impose relatively little strain on working memory, and hence are less susceptible to the debilitating effects of maths anxiety. However, more complex

mathematics places more of a burden upon working memory, which is thus affected by maths anxiety. It is a cruel irony that when students most need their working memories focused on the task in hand, maths anxiety is at its most disruptive.

6. When investigating high-level math topics, it is difficult to distinguish clearly between the effects of high math anxiety and low math competence.
This is a major problem for us teachers. Is the student who is unable to answer a complex question lacking knowledge or suffering from maths anxiety? It is important to know, because the remedies will likely differ.

Ramriez *et al* (2013) cite another key finding from research into maths anxiety that is both fascinating and worrying: 'students with the highest level of working memory capacity show the most pronounced negative relation between maths anxiety and maths achievement'. The authors suggest that students with higher working memory capacities tend to rely on more advanced problem-solving strategies as they have sufficient capacity to process them. The disruption to working memory caused by maths anxiety can result in such students switching to less successful problem-solving strategies as a means of circumventing the burden of maths anxiety. Ironically, something that usually helps students in maths – large working memory capacity – becomes most vulnerable to disruption when students are anxious.

What I do now

I take the issue of maths anxiety very seriously, realising it is a problem that could affect any student (or teacher) no matter what their maths achievement level. But what are we to do?

Well, a good start is to follow four practical strategies teachers can employ to help reduce maths anxiety in their students, outlined by Beilock and Willingham (2014):

1. *Ensure fundamental skills*. Domain-specific knowledge makes thinking and learning easier, and the automation of knowledge further reduces the strain on working memory. Thus the debilitating effects of maths anxiety are likely to be less as working memories have more capacity to cope.

2. *Focus on teacher training*. This is based on the finding that a teacher's anxiety about maths can be transferred to their students. As described above, this regularly happens to me in A level Maths Clinics, but also during lessons. I remember a few years back I was 'teaching' (in the loosest sense of the word) an A level mechanics lesson. I was not at all comfortable with the material, my brain felt like it froze up during a particular worked example, and the effect on my students was palpable. It is perhaps impossible to avoid this entirely, but preparation and experience are vital in minimising the occurrences. I would further argue that well-planned lessons following a model of explicit instruction that will be developed in this book are easier to deliver with confidence than lessons that follow a less guided approach.

3. *Change the assessment*. As discussed above, removing the time element of tests may be beneficial in that it reduces pressure and gives students extra time to consider their answers. However, we also need to consider the impact on the development of automated knowledge, as well as the logistics of giving every student a different amount of time. We will discuss these points further when we look at low-stakes quizzes in Chapter 12.

4. *Think carefully what to say when students struggle*. Often consolation in the form of, 'It's okay, not everyone can be good at these kind of problems', validates students' view that they are not good at maths. Better to focus on how you are convinced that hard work will help them get better, crucially following this up with concrete, effective study strategies like those outlined in Chapter 12.

A study by O'Leary *et al* (2017) also warrants a mention. They found that a greater perceived level of support from teachers by their students was associated with lower maths anxiety. More specifically, they found that 'there was a significant decrease in [maths anxiety] when participants reported that their teachers provided plenty of examples and practice items, and this remained after controlling for general and test anxiety'. A key argument of this book will be that a model of explicit instruction is likely to be more beneficial to novice learners than a less guided approach, and moreover that carefully chosen and presented worked examples are a key component of this model. For what it is worth, since I have adopted the model of explicit teaching that I will describe in the chapters that follow, I have noticed a

significant change in many of my students. They seem more confident, more willing to ask questions, more willing to embrace mistakes as the learning opportunities that they are, and just generally a lot happier. This has undoubtedly been one of the most important effects of the changes to my teaching that I have made.

OceanofPDF.com

1.7. If I only remember 3 things…

1. Expert and novice learners don't just know different things, they think in a fundamentally different way from each other.
2. Students remember what they attend to, so we should plan lessons and evaluate activities using this principle.
3. Anyone can suffer from maths anxiety. It can be debilitating, but there are ways we can help.

OceanofPDF.com

2. MOTIVATION

As a child, I had always been motivated to learn mathematics, and maths was always my favourite lesson. I was (and am) – for want of a better expression – a geek. So it was quite the shock, when I started teaching, to discover that not every teenager's idea of a good way to spend Friday night was solving quadratic equations.

For the vast majority of my career, I have tried all kinds of ways to motivate/beg/trick the disaffected adolescent into doing maths. It is knackering, slightly soul destroying and, when sweets were involved, rather expensive. So I sought to discover if there was a better, more sustainable way to motivate my students.

OceanofPDF.com

2.1. Models of Motivation

What I used to think

In much the same way as I had never really considered how students think, I had not really put much thought into what motivates my students. I had a sense that motivation was both important and desirable, but I had pretty much assumed that students either enjoyed maths or they didn't. Sure, there were little tricks I could pull to get students through tricky times, but ultimately their views on maths were fixed. Moreover, if asked why that was the case, I would probably have said that it was largely determined – with a few exceptions – by whether or not they were any good at it.

It turns out that it is not quite as simple as that.

Sources of inspiration

- American Psychological Association (2015) 'Top 20 Principles from Psychology for PreK-12 Teaching and Learning'. Available at http://www.apa.org/ed/schools/teaching-learning/top-twenty-principles.pdf.
- Deci, E. L., Koestner, R. and Ryan, R. M. (2001) 'Extrinsic rewards and intrinsic motivation in education: reconsidered once again', *Review of Educational Research* 71 (1) pp. 1-27. Lepper, M. R. (1988) 'Motivational considerations in the study of instruction', *Cognition and Instruction* 5 (4) pp. 289-309.
- Martin, A. J. (2016) *Using Load Reduction Instruction (LRI) to boost motivation and engagement*. Leicester: British Psychological Society.
- Middleton, J. A. and Spanias, P. A. (1999) 'Motivation for achievement in mathematics: Findings, generalizations, and criticisms of the

research', *Journal for Research in Mathematics Education* 30 (1) pp. 65-88.

- Pink, D. H. (2011) *Drive: the surprising truth about what motivates us*. London: Penguin.
- Tollefson, N. (2000) 'Classroom applications of cognitive theories of motivation', Educational Psychology Review 12 (1) pp. 63-83.
- Zimmerman, B. J. (2000) 'Self-efficacy: an essential motive to learn', *Contemporary Educational Psychology* 25 (1) pp. 82-91.

My takeaway

Whilst we have seen in Chapter 1 that student engagement is no guarantee of learning – as we cannot be sure exactly what students are engaged in – it seems sensible to assume that intrinsic motivation in mathematics, which can be thought of as wanting to learn for its own sake, is a desirable quality for our students to possess. Indeed, Lepper (1988) explains that when individuals engage in tasks in which they are motivated intrinsically, they tend to exhibit a number of pedagogically desirable behaviours including 'increased time on task, persistence in the face of failure, more elaborative processing and monitoring of comprehension, selection of more difficult tasks, greater creativity and risk taking, selection of deeper and more efficient performance and learning strategies, and choice of an activity in the absence of an extrinsic reward'.

Sounds great!

So, the natural questions arising from this are: what causes our students to be motivated? And what can we do as teachers to help?

Now, any book that attempts to answer these questions in a single chapter is asking for trouble. There are thousands of books on motivation, spanning the disciplines of psychology, cognitive science, behaviourism, economics and more. However, whilst I can do more than scratch the surface of this issue, it is something that cannot be ignored for two reasons. First, learning and motivation share an important, complex relationship. Second, accusations aimed at the explicit instruction model of maths teaching I will

introduce in Chapter 3 and develop throughout this book often revolve around the issue of motivation. Comments such as, 'sure that method may be effective, but if you bore the kids to tears, what good is it?', are all too common. Maths is often perceived, by students and parents alike, as a difficult, boring subject, and this can lead to misguided attempts to motivate students that can end up doing more harm than good.

So, before we crack open that can of worms, let us take a brief look at some of the general research findings related to motivation. Just as we found in our models of thinking, there is no universally accepted model of motivation, but there are common themes about what motivates us.

A sense of control

In his book *Drive*, Pink (2011) suggests a key determinant of motivation is *autonomy* – the desire to direct our own lives. The desire for control also looms large over the work of Deci *et al* (2001) in relation to Cognitive Evaluation Theory (CET), an interesting subset of their larger Self Determination Theory. According to the theory, external events such as the offering of rewards, the delivery of evaluations, the setting of deadlines, and other motivational inputs affect a person's intrinsic motivation to the extent that they influence that person's perception of self-determination. In short, if we feel we are not in control, or someone is trying to control us, we are less likely to be intrinsically motivated.

A belief that the work they are doing has value (intrinsic or extrinsic)

Pink (2011) identifies *purpose* – the yearning to do what we do in the service of something larger than ourselves – as a second key determinant of motivation. Martin (2016) agrees, including how much students believe what they learn at school is useful, relevant, meaningful, and important as a key component of motivation in his Load Reduction Instruction model. Tollefson (2000) discusses the concept of Expectancy × Value Theory, which proposes that the degree to which students will expend effort on a task is a function of: (a) their expectation they will be able to perform the task successfully and by so doing obtain the rewards associated with successful completion of the task; and (b) the value they place on the rewards associated with successful completion of the task. Specifically – and rather mathematically – the model assumes the amount of effort invested is a

product of the expectation of success and the value of the reward. Whilst this value can be intrinsic or extrinsic, both Middleton and Spanias (1999) and a report from the American Psychological Association (2015) argue that students tend to enjoy learning and perform better when they are more intrinsically than extrinsically motivated to achieve. Mastery

Pink's (2011) third element of motivation is *mastery* – the urge to get better and better at something that matters. The report from the American Psychological Association (2015) concurs, explaining that students persist in the face of challenging tasks and process information more deeply when they adopt mastery goals rather than performance goals. For Martin (2016), a 'mastery orientation' is a key component of motivation, explaining students' interest in and focus on learning, developing new skills, improving, understanding, and doing a good job for its own sake and not just for rewards or the marks they will get for their efforts.

A feeling that they are successful, or can be successful
Inherently related to mastery is the feeling of success. Middleton and Spanias (1999) explain that students' perceptions of success in mathematics are highly influential in forming their motivational attitudes. For Zimmerman (2000), self-efficacy (defined as one's belief in one's ability to succeed in specific situations or accomplish a task) is the key to motivation, explaining that self-efficacious students participate more readily, work harder, persist longer, and have fewer adverse emotional reactions when they encounter difficulties than do those who doubt their capabilities. Martin (2016) agrees, identifying students' belief and confidence in their ability to understand or to do well in schoolwork, to meet challenges they face, and to perform to the best of their ability, as key determinants of motivation.

For the remainder of this chapter we will assess strategies to impact upon student motivation in the light of these areas.

What I do now

Carefully consider what motivates my students, and specifically what actions I can take to influence this. As we shall see, some actions are easier than others, and some may actually do more harm than good.

OceanofPDF.com

2.2. Do students make good decisions?

What I used to think

Wherever possible, I wanted to give my students choices over their learning. This often involved allowing them to choose which questions, worksheet or task to do. This was all part of a grand plan both to help my students take more responsibility for their learning, and because I always assumed that students enjoyed choice.

The problem was, my students did not always make the choices I felt were best for them.

Sources of inspiration

- Clark, R. E. (1982) 'Antagonism between achievement and enjoyment in ATI studies', *Educational Psychologist* 17 (2) pp. 92-101.
- Deci, E. L., Koestner, R. and Ryan, R. M. (2001) 'Extrinsic rewards and intrinsic motivation in education: reconsidered once again', *Review of Educational Research* 71 (1) pp. 1-27.
- Kirschner, P. A. and Van Merriënboer, J. J. G. (2013) 'Do learners really know best? Urban legends in education', *Educational Psychologist* 48 (3) pp. 169-183.
- Nuthall, G. (2007) *The hidden lives of learners*. Wellington: NZCER Press.
- Pink, D. H. (2011) *Drive: the surprising truth about what motivates us*. London: Penguin.

My takeaway

The models of motivation we looked at in Section 2.1 suggest that choice is indeed important. Pink (2011) lists 'autonomy' as one of his three essential elements of motivation, and the work of Deci *et al* (2001) has the concept of self-determination at its core. This implies that giving students a sense of control over their learning is intrinsically motivating.

But when given this control, do students make the correct decisions?

According to the work of Kirschner and Van Merriënboer (2013), apparently not. In their paper, intended to debunk common misconceptions about learning, they identify three major problems with giving the learner more control:

1. *Learners may not have the capacity to weigh up both the demands of the task and their own learning needs in relation to that task*. A student with the very best of intentions may not be able to select the most appropriate task for their learning needs because doing so requires them to process both what the task is about, and where their own strengths and weaknesses lie. When presented with a choice of unfamiliar tasks on a topic they are relatively new to, it is not surprising that students would struggle to choose optimally.

2. *Learners often choose what they prefer, but what they prefer is not always what is best for them*. In particular, students continue to practise tasks they like or are already proficient in, but are reluctant to start with new, yet unfamiliar tasks. Nuthall (2007) explains: *our research shows that students can be busiest and most involved with material they already know*. I have seen this countless times with my students. I remember one Year 11 girl who would happily add fractions together all day long, but would not touch an equation with a barge pole. When asked about her fondness for fractions, she replied 'I am good at these'.

3. *The paradox of choice*. People appreciate having the opportunity to make some choices, but the more options that they have to choose from, the more frustrating it is to make the choice. More choice almost inevitably means more regret over what has not been chosen. This is just as true in a restaurant with an abundant menu as it is in a maths lesson. When giving

my students freedom over the questions they answer or the task they attempt, many would say, 'Just tell me which ones to do, sir'.

An important finding comes from Clark (1982). In a review of relevant research he came to the conclusion that 'students often report enjoying the method from which they learn the least'. Specifically, 'low-ability' (his phrase, not mine) students typically report preferring less teacher-led instructional methods (such as discovery-based learning), apparently because they allow them to maintain a 'low profile' so that their failures are not as visible. However, as we will see in Chapter 3, such students benefit more from more teacher-led methods (such as explicit instruction) which lower the information processing load on them. Conversely, high-ability students prefer more structured methods (such as worked examples) which they believe will make their efforts more efficient, and yet these lower load methods seem often to interfere with their learning, as we will see when we consider the Expertise-Reversal Effect in Section 6.7. High-ability students actually seem to learn more from less teacher-led approaches (such as solving problems independently) which allow them to apply their existing knowledge and hence further develop their schemas. Therefore, students' preferences are likely to be influenced by their levels of enjoyment, which are poor indicators of how much they are learning

But does it really matter if students do not make optimal choices? After all, they are likely to be more motivated as a result of this feeling of control, and motivation seems important for learning.

Well, yes it does, and for two reasons. First, as we have seen, the decisions students make are likely to negatively and directly impact on their learning. Second, if – as I will argue later in this chapter – the feeling and experience of success is also a key motivating factor, then these misguided choices could wipe out any motivational gains via their negative impact on learning. So, by allowing students to make too many decisions, we may end up with the worst of both worlds: they learn less and their motivation does not improve.

What I do now

In general, I now give students far less choice in the work they do in class than I used to. If I have put the thought into planning a carefully designed sequence of examples and practice questions (Chapter 7), or a Purposeful Practice activity to enable students to gain vital practice of key skills whilst also developing their problem-solving capabilities (Chapter 10), then I want to ensure they do it. This supports my view that the most effective form of differentiation is by time not task, which will be a recurring theme throughout this book.

In order to tap into some of the motivational benefits of choice, as well as helping to develop students' metacognitive skills, when I do give students a choice over a task or a set of questions, it is either:

- from a list that I have chosen
- following a low-stakes quiz (Chapter 12) so they are better informed as to their current understanding

OceanofPDF.com

2.3. Real-life Maths

What I used to think

'Sir, when will I ever use this in real life?'

If there is a more painful set of ten words to intrude upon the ears of a maths teacher, I am yet to hear it.

When asked how to make maths more interesting to students, the answer is a simple one: you make it more relevant to their lives, of course. Next question, please.

So, for many, many years, I would spend ages thinking of ways to shoehorn real-life contexts and 'helpful' analogies into the classroom. Could I use David Beckham's last-minute free kick against Greece in 2001 to help convey the principles of quadratic equations? Of course! Something happened on the news tonight – let's find a way to work it into tomorrow's lesson. And based on the number of maths topics I managed to extract from the 2012 London Olympics, I am probably due a gold medal myself.

The problem was, very rarely did these contexts lend themselves perfectly to the maths I wanted to teach. Often I would need to present a modified, simplified version that bore little resemblance to the original context. Moreover, not all my students found these contexts useful or even that interesting.

Sources of inspiration

- Barton, C. (2017) 'AQA Chief Examiner Trevor Senior', *Mr Barton Maths Podcast*.Boaler, J. (1994) 'When do girls prefer football to fashion? An analysis of female underachievement in relation to

"realistic" mathematics contexts', *British Educational Research Journal* 20 (5) pp. 551-564.

- De Bock, D., Verschaffel, L., Janssens, D., Van Dooren, W. and Claes, K. (2003) 'Do realistic contexts and graphical representations always have a beneficial impact on students' performance? Negative evidence from a study on modelling non-linear geometry problems', *Learning and Instruction* 13 (4) pp. 441-463.
- Little, C. and Jones, K. (2010) 'The effect of using real world contexts in post-16 mathematics questions', *Proceedings of the British Congress for Mathematics Education* 30 (1) no pagination.
- Nuthall, G. (2007) *The hidden lives of learners*. Wellington: NZCER Press.
- Wiliam, D. (1997) 'Relevance as MacGuffin in mathematics education', British Educational Research Association Conference, York, September 1997.
- Willingham, D. T. (2003) 'Ask the cognitive scientist. Students remember…what they think about', *American Educator* 27 (2) pp. 37-41.

My takeaway

We have seen in Section 2.1 that giving students a sense that their work has a wider value is likely to contribute positively to their levels of motivation. So, it would seem sensible to conclude that attempting to make the maths that students study relevant to their lives would be a good thing. However, it is not that straightforward.

Little and Jones (2010) summarise the dilemma perfectly:

> On the one hand, by making a connection between the abstract world of mathematics and everyday contexts, we are reinforcing the utility of mathematics as a language for explaining the patterns and symmetries of the 'real' world. On the other hand, if we manipulate and 'sanitise' real-world experiences to enable them to be modeled by a preordained

> set of mathematical techniques, then the result can appear to be artificial and contrived.

Nuthall (2007) puts it more succinctly: 'Students are constantly on their guard against being conned into being interested'.

Indeed, that has been my experience. Students are not stupid – and when I am standing there claiming that David Beckham considers the properties of the resultant quadratic curve when lining up his free-kick, I am fooling no one apart from myself.

Wiliam (1997) identities three different types of problematic real-life contexts that plague school mathematics:

1. Pointless Contexts
Consider the following question from a maths textbook:

Alan drank $\frac{5}{8}$ *of his pint of beer. What fraction was left?*

(next to a picture of a middle-aged man drinking a pint of beer)

When evaluating the use of questions like this, we need to consider what they are intended to achieve.

If we want our students to develop automaticity in subtracting fractions from 1, then let's cut out the context altogether and just ask our students to calculate $1 - \frac{5}{8}$. After all, the information surrounding the context will still need to be processed by the student which, as we shall see in Chapter 4 via the Redundancy Effect, risks overloading their fragile working memories and inhibiting learning.

If we want our students to practise overcoming the surface features of a problem and recognising the deep structure, then an argument can be made for the use of questions like this – after all, questions where the context disguises an otherwise straightforward mathematical procedure come up all the time in high-stakes exams and catch students out. However, as shall be

discussed in Chapter 7, batching contextual problems with the same deep structure makes the surface structure largely irrelevant as students can figure out what they need to do without thinking hard.

Finally, if we think questions like this are in any way motivating to our students in so much as they relate to their wider lives, then I think we might be a little deluded. After all, as Wiliam points out, beer is not measured in eighths of pints, and even were it to be so measured, it is unlikely that Alan would be overly concerned about the precise measurement while enjoying his drink.

2. Confusing Contexts
Consider the following example:

This is the sign in a lift at an office block:

This lift can carry up to 14 people.

In the morning rush, 269 people want to go up in the lift.

How many times must it go up?

The calculator gives an answer of 19.21428571, and we would hope our students would deduce that '20 times' is thus the correct answer. However, is it definitely that clear-cut? As Wiliam explains, a number of assumptions must be made for 20 to indeed be correct:

- that the lift is full apart from the last trip
- that no one gets fed up and uses the stairs
- that the restriction is based strictly on number rather than mass or volume
- none of the people using the lift are wheelchair users

Similar ambiguities are often noted by teachers reviewing questions in high-stakes maths exams, where equally valid cases can be made for two or more answers. The fact that these exams are prepared years in advance and checked by several people just goes to show how difficult it can be to create

non-confusing, unambiguous questions. Indeed, when I interviewed AQA's Chief Examiner, Trevor Senior, he described the difficulty of using contextual questions in maths GCSEs. With this in mind, I am far more wary of any contexts I use in class, in case they cause unnecessary confusion.

3. Dangerous Contexts
This is perhaps the most common and most problematic. By attempting to appeal to students' interests we risk excluding those students who do not share those interests. Hence the (completely fair) groan that would resonate around the room, from both girls and boys, when I lazily fell back on yet another football-related example.

However, the very opposite can be true. In attempting to appeal to students' interests we are, by definition, moving the content of the learning into a domain which they have more knowledge of. This sounds like it must be a good thing. However, as I have already indicated, often when using a real-life context we need to modify and simplify it. Hence, our representation of the context may be at odds with that of our more knowledgeable students. Furthermore, they may bring their existing knowledge and experiences to bear in a way that we had not intended.

A fascinating example of this is provided in a paper by Boaler (1994). The title says it all: 'When do girls prefer football to fashion? An analysis of female underachievement in relation to "realistic" mathematical contexts'. Boaler found that girls may seek to relate in-context problems to their existing knowledge, while boys are often content to tackle the problem in isolation from their previous experience. This results in boys tending to be more successful on questions set in a context than girls. So, by trying to appeal to a subset of students' special interests in order to motivate and engage them, we could well end up doing them more harm than good.

Videos
A related means of motivating students is to use different mediums of presentation within lessons. Now, we will further consider the use of images, animations and dynamic geometry software in Chapter 4, but here I just want to briefly mention the use of video. Often video can be seen as an easy way to funnel real-life maths and relevance into the maths classroom. I have

been guilty myself, making use of the classic *Maths4Real* videos (one of TV star Ben Shephard's first presenting jobs, I'll have you know), and even things like the maths feature on Chris Moyles's short-lived TV quiz show, *Quiz Night*.

Two findings need bearing in mind when considering using such things as a means to motivate students. The first is the (hopefully) familiar mantra from Willingham (2003): 'students remember what they think about'. If students' attention is diverted away from the maths to another feature of the video (such as how young Ben Shephard looks), then less learning is likely to take place.

The second is more subtle. De Bock *et al* (2003) conducted a study involving 152 eighth graders (13- to 14-year-olds) and 161 tenth graders (15- to 16-year-olds). They were given a paper-and-pencil test about the relationships among the lengths, areas and volumes of different types of rectilinear and non-rectilinear figures. One group was then shown instructional videos set in a real-life context, whilst the other was given a set of written instructions. The researchers found that students who watched the videos performed significantly worse than the students from the other groups. They speculated that students perceive video as a less difficult medium than written materials, and are therefore inclined to invest less mental effort compared to an instructional setting that was perceived as more difficult.

What I do now

On reflection, the real-life contexts I have used in the classroom have not been all that motivating. By modifying and simplifying the presentation of the context, I have usually stripped away its interest and relevance. Perhaps more importantly, we have seen how such contexts can be detrimental to students' learning, confusing them or diverting their attention away from the things that really matter. In short, I have concluded real-life maths is often more trouble than it is worth.

As I hope to show in the next two sections, there are better ways we can help students value the work they do in mathematics.

For a final point, I have a theory about the question, 'When will we ever use this in real-life?'. What I really think students are saying is, 'I don't understand this'. This is based solely on my observation that a student who has just got a load of questions correct – regardless of the topic and context – has never once asked me that question. And yet I have heard it many a time following a ropey explanation by me, or an ill-planned activity. If I teach my students well so that they can achieve success, that rather annoying question seems to disappear.

OceanofPDF.com

2.4. Teacher influence

What I used to think

How much influence do I have over my students' feelings and motivations about mathematics?

To be honest, this is a question I had never really considered all that deeply. I saw it as my responsibility to try to teach my students to the best of my abilities in a safe learning environment, all whilst being a positive role model. I think I even said this in a job interview once.

The problem was, I did not have the slightest clue what being a positive role model actually meant.

Sources of inspiration

- Boaler, J. (2015) *The elephant in the classroom: helping children learn and love maths*. London: Souvenir Press. Greany, T., Barnes, I., Mostafa, T., Pensiero, N. and Swensson, C. (2016) 'Trends in maths and science study (TIMSS): national report for England'. Available at: http://dera.ioe.ac.uk/28040/1/TIMSS_2015_England_Report_FINAL_for_govuk_-_reformatted.pdf.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Middleton, J. A., and Spanias, P. A. (1999) 'Motivation for achievement in mathematics: findings, generalizations, and criticisms of the research', *Journal for Research in Mathematics Education* 30 (1) pp. 65-88.

My takeaway

We have seen how the careless use of real-life maths may not cause students to value a topic or concept any higher, and hence is unlikely to contribute positively to their levels of motivation. However, there may be a way we can help students value the subject of mathematics as a whole.

A conclusion Middleton and Spanias (1999) reach in their comprehensive review of research is that 'motivations toward mathematics are developed early, are highly stable over time, and are influenced greatly by teacher actions and attitudes'. In particular they find that:

1. Students will feel more comfortable taking risks if they know that they will not be criticised or humiliated for making mistakes.

2. Students tend to attribute their feelings about mathematics to their identification with influential teachers or to their reactions to bad experiences, for which they blame teachers.

I have come to the realisation that, as their maths teacher, I am possibly the only positive mathematical role model in many of my students' lives. Sadly, students are likely to be surrounded by maths-haters or maths-avoiders. This is unsurprising, given what we have learned about maths anxiety in the previous chapter, together with the portrayal of maths in much of the media as a challenging, abstract subject accessible only to geniuses and weirdos. An overwhelming number of parents will openly announce in front of me and their children at Parents' Evening that they were never good at maths themselves. A significant proportion of a student's friends may tell them maths is the subject they hate the most. Indeed, in a 2016 survey, Greany *et al* found that the proportion of pupils who reported that they do not like learning maths was 17% at Year 5, and 48% at Year 9. As both Middleton and Spanias (1999) and the work of Boaler (2015) suggest, students may even encounter negative messages from the maths teachers they have had over the years, which has the potential to cause long-lasting damage. In short, the sole witness for the defence in the trial against maths is me, their current teacher, and I need to up my game.

I need to portray maths in the best possible light. But that alone is not enough. I also need to create a learning environment that is conducive to the

enjoyment of a subject that many students do find a challenge.

What I do now

I am far more aware of the power I have to influence my students' perceptions of mathematics, and as the great philosopher Ben Parker (Spiderman's uncle) once said, with great power comes great responsibility. Therefore, I see it as my responsibility, not just to teach my students maths to the best of my abilities, but to be an *actively* positive role model in their learning of the subject. This is likely to help them value mathematics more, and hence be more motivated to succeed in it.

So, I have created a set of Golden Rules that I try my best to stick to with every single class I teach, no matter their age or set, or what topic I am teaching them:

1. I am empathetic, not apologetic

I empathise with students when they are struggling with a particular concept or topic. I explain that I often struggle myself – when I make a mistake in a lesson (which only happens about 78% of the time these days), I make sure all the students are aware of this, so they see that everyone struggles and it is perfectly normal to make mistakes. But I never, ever, ever apologise for asking them do mathematics, no matter how challenging or dull (in their opinion) it may be. Apologising is neither right nor sustainable. It leads to an unhealthy dynamic whereby my students have the upper hand, and I need to keep providing external rewards in an attempt to buy their effort – and as we shall see later in this chapter, this is unlikely to work in the long-term. Furthermore, it is simply draining. As I know from losing many arguments with my wife, saying sorry is tiring, and teaching is tiring enough without this extra burden. Lemov (2015) sums all this up in his strategy Without Apology: 'embrace – rather than apologise for – rigorous content, academic challenge, and the hard work necessary to scholarship'.

2. I never hide my love of maths…from anyone

I am a geek. I flipping love maths. I did not have the guts to say that for much of my adolescence, desperate as I was to fit in with my peers. Likewise, for the first few years of my teaching career, I tended to hide this love away.

Perhaps somewhat tellingly, I did not do this from all students. When teaching top sets, I would revel in my geekiness, surrounded as I often was by enough like-minded people to feel like I was in the majority with my passions. But give me a middle set of Year 11s – one populated by stereotypical football-loving lads and girls whose top priority from Monday morning onwards was to start planning where they were going out on Friday night – and I was more likely to play down my love of maths, admitting that it can be dull, and a struggle, but it is just something they have to do, especially if they want to pass the exam. But what good is that really doing them? Sure, it might make them like me a little more for that lesson and get me (at best) an extra 10 minutes of effort from them. But teenagers see through dishonesty more easily than anyone else, and in the long run I feel it is better for everyone if I admit I love maths and try to do my best to show my students, no matter who they are, exactly why. Not every student will come around to my way of thinking, many will think me mad, others may pity me. However, I feel more will trust and respect me because of it, which can only be beneficial for our long-term relations and their subsequent achievement.

Now, there will be maths teachers reading this who did not like maths when they were a school, and for whom extolling the wonders of the subject would be just as false as me claiming I do not get quite the thrill from simplifying a particularly challenging algebraic fraction. And such teachers will probably always have an advantage over the likes of me as they can understand students' frustrations and empathise with their struggles in a way that I will probably never be able to do, no matter how many mathematically reluctant students I work with. But I still believe that for these teachers the message they communicate needs to be a clear one and a positive one: okay, maths can be difficult, I found it difficult, but look at me now, I teach it, I enjoy it, I can do it, and you can too.

3. If I dislike a topic, I don't show it

Mathematics is a huge topic, few people are good at all aspects of it, and fewer people – even self-confessed geeks like me – enjoy all aspects of it. Keep this quiet, but I flipping hate 3D trigonometry. Any time the question asks for the angle between the line and the plane, I am pretty much reduced

to guessing, unable to spot the relevant 2D right-angled triangle. I also find compound measures a little bit on the dull side, and do not get me started on the torture that is A Level mechanics. But I try my very best not to convey these negative feelings to my students. Sure, I will admit that I find certain parts of maths tough, but I try to never let my enthusiasm for the subject falter. For if I did – given the influence the research suggests teachers have – it would unfairly taint the topic or concept in the minds of my students. I am sure that there have been several students for whom I have permanently diminished their enjoyment of friction and kinematics, but no more!

4. If someone laughs or is in anyway nasty about another student's mistake, I go mental

I rarely shout in lessons. If a student messes around, then I prefer a stern look and (in my eyes, at least) a spine-chilling whisper of 'see me at the end'. However, if a child dares to snigger, snort, sniff or in any other way indicate their derision towards another child's answer, then they are for it. And this will not be a quiet word. Oh no, I will ensure every other student knows exactly why Mr Barton is currently going red in the face. I will argue later in this chapter that I feel that I have, in the past, encouraged mistakes too much, but that does not mean I do not value the identification and resolution of mistakes as a key part of the learning process, and it certainly does not mean that I want to discourage students from voicing an answer for fear of ridicule if it is wrong. For Lemov (2015), such teacher actions are a fundamental part of what he terms the 'Culture of Error'. A sharing of and respect for answers, whether they be right or wrong, is a fundamental feature of any well-functioning lesson and safe learning environment, and woe betide anyone who tries to stop that.

And breathe…

OceanofPDF.com

2.5. Providing a Purpose

What I used to think

As I have already said, I used to believe that the only way I could help my students see the value in maths was to make it relevant to their lives. As such, I would search desperately to find relevance in topics such as factorising quadratics, laws of indices, and so on. These seemed so distant from anything that would ever matter to my students that I would end up coming up with the most contrived, pathetic real-life scenario in the world, or simply say to them: look, this is just something we need to do.

The only purpose that served was to cause some students to dislike maths even more.

Sources of inspiration

- Barton, C. (2016) 'Dan Meyer', *Mr Barton Maths Podcast*.
- Fuller, E., Rabin, J. M. and Harel, G. (2011) 'Intellectual need and problem-free activity in the mathematics classroom', *Jornal Internacional de Estudos em Educação Matemática* 4 (1) pp. 80-114.
  Meyer, D. (2015) 'If math is the aspirin, then how do you create the headache?', *dy/dan blog*, 17 June. Available at: http://blog.mrmeyer.com/2015/if-math-is-the-aspirin-then-how-do-you-create-the-headache/
- Wu, H. (1999) 'Basic skills versus conceptual understanding', *American Educator* 23 (3) pp. 14-19.

My takeaway

Despite our best efforts, and following the Golden Rules suggested in the previous section, not every student is going to value mathematics as highly as we would like, and hence enjoy the motivation that such value brings. Moreover, there are certain topics that even the keenest mathematicians struggle to see the wider value in.

Whilst shoehorning dodgy real-life contexts into lessons is unlikely to provide the answer, there is one other technique I have learned, courtesy of the US maths teaching expert, Dan Meyer. When I interviewed Dan for my podcast, we talked about his 'Headache-Aspirin' series of blog posts (see Meyer, 2015). Dan asks the question: 'If Math Is The Aspirin, Then How Do You Create The Headache?'. He argues that real-world applications of many maths skills we teach in school are a lie, going on to say that if our theory is 'maths is interesting only if it's real-world', then we will struggle to find interest in many of the things we teach. Instead, we should ask ourselves, '**Why did mathematicians think this skill was worth even a little bit of our time?**'.

Meyer's work is influenced by Fuller *et al* (2011) who propose the Need for Computation. This is defined as 'the need to find more efficient computational methods, such as one might need to extend computations to larger numbers in a reasonable "running time"'. The key here is that in order for the need for computation to be met, students need to first experience the 'longcut' before they learn the shortcut, for only then will they appreciate the power of that shortcut.

Let's take an example that I have used myself many times over the last two years: factorising a quadratic expression. Now, you would be hard pushed to find a real-life application of factorising quadratics that is likely to be relevant and appeal to your average group of 14- and 15-year-olds. However, that does not mean that we cannot create a purpose and hence a value in them learning how to do it.

The start of the lesson may go something like this:

- Ask students to pick a number between 1 and 10 and write it down.
- Now write the following expression on the board: $x^2 - 2x - 35$

- Ask students to substitute their number into the expression.
- After a minute or so, ask, 'Has anyone got an answer of zero? If you haven't, try another number'.
- Pretty soon, it is likely that someone will try '7' and get 0. Write this on the board and use the example as a means of checking students are substituting numbers correctly.
- Then announce: 'There is actually another number that also gives you 0. Can you find it?'.
- After five minutes of struggle, with no clues as to the nature of this mystery number, it is likely students will be getting frustrated.
- At this point, you announce: 'Wouldn't it be great if there was an easier way to find this mystery number?'.
- Hence, you now have a purpose for introducing the technique for factorising quadratic expressions, and the technique has value to the students.

The key principle is that students should experience inefficient computation before we help them develop efficient computation. Notice also, however, that students experience early (and partial) success in finding the first number. This may be just as important in maintaining their interest and effort in the task.

Wu (1999) describes a similar approach to show students why the standard algorithms of written addition, subtraction, multiplication and division are useful. By showing students how difficult life would be without them, whilst at the same time relating the long-form arithmetic calculations to the formation of the algorithms, Wu is able to create both a purpose for their use together with the development of conceptual understanding.

What I do now

Wherever possible (and it is more possible than you might think) I show my students that the skills I teach them have a use and a purpose, even if that

purpose seems completely apart from their 'real lives'. I do this by first asking them to solve problems without these 'shortcuts', and by using the principle of Purposeful Practice we will meet in Chapter 10.

I would strongly recommend checking out Dan Meyer's 'Headache and Aspirin' blog series. Other areas covered are:

- Simplifying Rational Expressions
- Proof
- Graphing Linear Inequalities
- Functions
- Laws of Indices

Paul Smith (@drpaas1001) posted this wonderful tweet that sums this whole concept up nicely:

$$\frac{6}{55} \times \frac{35}{3} \times \frac{11}{14}$$

I loved the reaction of my S1s when I showed them the easy way...after they'd done it the hard way. They called me evil. #mathsrockedtoday

And once you get into this way of thinking, it is quite easy to come up with similar ways to introduce other topics.

OceanofPDF.com

2.6. Rewards and Sanctions

What I used to think

Have you ever bribed a student to do some maths? I know I have. Whilst I have not formally handed over cold, hard cash, I have certainly plied my Year 11s with costly confectionery in the weeks leading up to an exam, and partaken in senior-leadership sponsored Pizza Revision Nights for the old C/D borderline GCSE students. Likewise, I have cajoled many a flagging student with praise for their efforts and the promise of a 'fun lesson' at the end of term. At the other end of the spectrum, I have placed students into detention for poor homework, threatened to call parents at break unless they answer at least ten more questions this lesson, and invoked the ultimate sanction for any Year 11 student – a suggestion that they might be banned from Prom.

In essence, I was providing sources of external motivation, both positive and negative. I never felt entirely right doing a lot of them – especially those that cost me money (I am a tight northerner, after all) – but they did buy me some short-term results.

Could I have done any better?

Sources of inspiration

- Bennett, T. (no date) 'Punishing with rewards: when praise becomes a sanction', TES *Behaviour and classroom management collections*. Available at: https://www.tes.com/articles/behaviour-punishing-rewards-when-praise-becomes-a-sanction
- Bennett, T. (2010) *The behaviour guru: behaviour management solutions for teachers*. London: A&C Black.

- Deci, E. L., Koestner, R. and Ryan, R. M. (2001) 'Extrinsic rewards and intrinsic motivation in education: reconsidered once again', *Review of Educational Research* 71 (1) pp. 1-27.
- Didau, D. and Rose, N. (2016) *What every teacher needs to know about … psychology*. Woodbridge: John Catt Educational Limited.
- Dweck, C. (2014) *Mindset, the new psychology of success*. New York, NY: Ballantine Books.
- Dweck, C. (2016) 'Praise the effort, not the outcome? Think again', *TES Magazine*. Available at: https://www.tes.com/news/school-news/breaking-views/praise-effort-not-outcome-think-again
- Middleton, J. A. and Spanias, P. A. (1999) 'Motivation for achievement in mathematics: Findings, generalizations, and criticisms of the research', *Journal for Research in Mathematics Education* 30 (1) pp. 65-88.
- Rose, N. (2017) 'Why punishments and rewards don't work', *TES Magazine*, 8th September, 2017.
- Tversky, A. and Kahneman, D. (1991) 'Loss aversion in riskless choice: a reference-dependent model', *The Quarterly Journal of Economics* 106 (4) pp. 1039-1061.
- Willingham, D. T. (2006) 'How praise can motivate – or stifle', *American Educator* 29 (4) pp. 23-27.

My takeaway

If attempting to tackle the subject of motivation in a chapter was a little ambitious, any hope of doing the issue of rewards and sanctions any sort of justice in a single section is bordering on delusional. For more comprehensive coverage of the key psychological principles of rewards and sanctions, I would recommend David Didau and Nick Rose's *What every teacher needs to know about … psychology*, and for practical strategies there is Tom Bennett's *Behaviour Management Solutions for Teachers*. Likewise, much of the work in making rewards and sanctions effective needs to be

done at a whole-school level, ensuring students know the expectations placed on them in every class, along with the consequences of failing to meet them. As individual teachers, we are often slaves to the culture created by our senior leaders. However, there are actions we can take in our own classrooms. The following provides a brief overview of the things I wish I had known about the rewards and sanctions far earlier on in my career.

Value

In Section 2.1 we saw that a key determinant of students' motivation is the belief that the work they are doing has value. Ideally this value will come intrinsically – from the task itself or the culture the teacher has created – whereby learning itself is regarded as being valuable. However, offering a reward for the completion of a task, or threatening a sanction for non-completion will have an immediate impact on the value of that task to the student, depending on how they value that reward or sanction.

Sanctions may be more successful at ascribing actions a value than external rewards such as merits, chocolates and praise due to Tversky and Kahneman's (1991) concept of 'loss aversion'. Humans often differ in our response to positive reinforcement and negative punishment, such that when people weigh up similar gains and losses, people tend to prefer avoiding losses to making gains. Hence, sanctions like the confiscation of a mobile phone, losing some of your break time, or exclusion from Prom, may be more effective than giving merits or rewards.

Deci *et al* (2001) warn that engagement-contingent rewards (those offered explicitly for engaging in an activity, regardless of the outcome of that activity) and completion-contingent rewards (those given for completion of an activity, again regardless of the outcome) may boost short-term value, but are likely to significantly diminish intrinsic motivation. Performance-contingent rewards (defined as rewards given explicitly for doing well at a task or for performing up to a specified standard) can maintain or enhance intrinsic motivation if the receiver of the reward interprets it informationally, as an affirmation of competence. Yet, because performance-contingent rewards are often used as a vehicle to control not only what the person does but how well he or she does it, such rewards can easily be experienced as very controlling, thus undermining intrinsic motivation.

Control

In Section 2.1 we also saw how a student's sense of control was an important factor for their motivation. Whilst it is clear that sanctions are likely to be considered controlling due to the direct influence they are intended to have on behaviour (*eg* the completion of a homework, or stopping talking), Deci *et al* (2001) argue that rewards can also be seen as controlling and hence have a negative effect on motivation. For example, tangible rewards, such as merits and prizes, decrease intrinsic motivation because such rewards are frequently used to persuade people to do things they would not otherwise do. Tangible rewards may control immediate behaviours, but they have negative consequences for subsequent interest, persistence, and preference for challenge, especially for younger children. Perhaps unsurprisingly, unexpected rewards are not found to be detrimental to intrinsic motivation, whereas expected rewards are. The reasoning is that if people are not doing a task in order to get a reward, they are not likely to experience their task behaviour as being controlled by the reward.

Intrinsic v External Motivation

Clearly we would like all our students to be intrinsically motivated – to do maths for the love of doing maths – instead of needing to rely on external rewards or the threat of sanctions to motivate them. But it is not quite as simple as that. After all, it can be very difficult to separate intrinsic motivation from extrinsic motivation. When I was at school, I found genuine pleasure doing maths. But if there were no good grades or teacher praise to reward my efforts along the way, or a final exam to aim towards, I am not sure I would have been quite as keen. Likewise, as much as I love my job, one thing that certainly compels me to pick up my red pen on a Sunday afternoon is an impending book scrutiny, and the consequences if my marking is not up to scratch.

In their comprehensive review of research into motivation in mathematics, Middleton and Spanias (1999) conclude that 'providing opportunities for students to develop intrinsic motivation in mathematics is generally superior to providing extrinsic incentives for achievement'. They further explain that to facilitate the development of students' intrinsic motivation, teachers must teach knowledge and skills that are worth learning. Students

who come to value and enjoy mathematics increase their achievement, their persistence in the face of failure, and their confidence. However, they also concede that providing external incentives for success can and does encourage students to achieve.

There is also the possibility of a virtuous cycle. External motivation to put in extra effort in the form of a reward or a sanction may lead to the development of intrinsic motivation. The student who is cajoled into revising properly for a test may perform well, feel good about themselves, realise they can be successful, and be motivated themselves to try harder again in the future. The important relationship between success and motivation will be discussed in the final sections of this chapter.

Praise

Perhaps the most immediate external reward that us classroom teachers have available to us is praise. Willingham (2006) offers the following guiding principles to make praise as effective as possible:

1. Praise should be sincere.

If praise is dishonest, controlling or unearned, it is likely to have negative consequences. I have struggled in particular with the last one of these. A student who never hands in their homework finally hands something in, and it is a load of rubbish. Do I praise the fact that at least he handed it in, or does that send a signal that such an effort is acceptable? For Willingham, the answer is to say/write something like 'It's great that you finished the assignment, but I'm a little disappointed in the quality of this work because I know you can do better'. Chair of UK Education Behaviour Group Tom Bennett summarises his thoughts on the matter as follows: 'praise and rewards have to be earned to be meaningful, or the class will see you as a kind of Santa Claus in an agreeable cardigan, good for sweets and treats, but not to be treated as anything other than a biscuit barrel'.

2. Praise should emphasise process not ability.

Praising ability may lead students to have a fixed view of ability, which may be detrimental to their long-term development – this is related to Dweck's work on growth mindsets (*eg* Dweck, 2014) that we will dive into further in the next section. However, simply praising effort has complications as well –

often it is socially more acceptable for students to put as little effort in as possible, and being told 'you tried really hard' may be interpreted by a student as 'you are thick, but nice try'. Willingham's suggestion is to praise the product of the process – 'that is a brilliant solution', as opposed to 'you worked really hard on that solution'. Indeed, in a 2016 article for *TES*, Dweck herself explains the potential pitfalls with praising effort and ignoring achievement, concluding: 'so, praise the effort not the outcome? Let's change that to: praise the effort (as well as the strategies, focus, perseverance and information-seeking) in relation to the outcome – with particular emphasis on learning and progress. True, it may not roll off of the tongue quite as easily, but it will certainly help our students more'.

3. Praise should be immediate and unexpected.
Praise loses much of its informational and motivational impact if the teacher praises a child for having shown good effort two weeks ago. Making praise unpredictable is hard to do, but can be of huge benefit. The goal is not simply to get the child to stop asking for praise; it is to help the child to think of their work differently – as something that is done for the student's own satisfaction, not to garner praise from the teacher. This might be thought of as being part of a longer-term strategy to help students see the value in the mathematics they study, by being a positive role model (Section 2.4), and providing a purpose to their work (Section 2.5).

What I do now

I am much more aware of the dangers of a reliance on external rewards and sanctions. Simply rewarding students for turning up to revision or completing a homework regardless of their efforts, or dishing out praise left, right and centre, are likely to diminish levels of intrinsic motivation, and cause me to be locked in a potentially futile battle to get that intrinsic motivation back. It can also lead to what I call 'reward inflation', whereby I keep needing to up the value of a reward to get the same effect. This ends up being very costly to a tight northerner like me. Likewise, if students learn to expect rewards any time they find something difficult or boring, then suddenly everything becomes difficult and boring, and all hell could break loose if those rewards suddenly disappear.

However, I am not dismissive of external rewards and sanctions. Their effects can be immediate and significant.

So, when I find myself, three weeks before a GCSE exam, faced with a Year 11 lad, head on the desk, with the question on straight line graphs apparently having drained his will to live, external motivation may well be necessary, in whatever form I predict is likely to work best for that given student and that given situation. It may be a stern warning, a word of comfort, or a Polo mint. Likewise, if a class is struggling with a difficult concept, praise or the threat that we will keep going into break may be exactly what they need to give me their full attention and effort in the short-term. Both of these scenarios may kick-start a virtuous cycle, whereby the initial (external) motivation leads to success, which breeds more sustainable motivation in the long run.

OceanofPDF.com

2.7. Why struggle and failure aren't always good – Part 1

What I used to think

For many years, I have been a firm believer in the key role of mistakes in the learning process. In a video on her YouCubed.orgwebsite, Jo Boaler explains how 'a research study found that when people make mistakes their brains grew more than when they got work right', going on to explain that this finding is down to the dual-firing of synapses.

It was this sort of thinking – which was reinforced on pretty much every training course I went on in the 2000s – that led me to conclude that mistakes were the key to learning, and thus struggle and failure must be good.

Indeed, this thinking also prompted a school I once worked at to run a twilight session entitled 'How to be a Bastard Teacher'. The idea was simple – we were to set students difficult challenges, and leave them to struggle, safe in the knowledge that their brains were growing so we were doing them good.

Hence, I would set my students 'extreme challenges' throughout my lessons. Here is an example I gave to my middle-set Year 10s, just in case you would like to play along at home:

Extreme Sequences Challenge
Here is a sequence:

2, 2√7, 14, 14√7

What is the value of the 21st term divided by the 17th term?

15 minutes.

Good luck!

I would then take a seat, be a bastard, and enjoy the struggle that followed. The strange thing was, many of my students didn't seem to really enjoy this approach, and quite a few simply gave up on the problems I presented them with.

Sources of inspiration

- Bahník, Š. and Vranka, M. A. (2017) 'Growth mindset is not associated with scholastic aptitude in a large sample of university applicants', *Personality and Individual Differences* 117, pp. 139-143.
- Barton, C. (2017) 'Nick Rose', *Mr Barton Maths Podcast*.
- Boaler, J. (2015) *Mathematical mindsets: unleashing students' potential through creative math, inspiring messages and innovative teaching*. Hoboken, NJ: John Wiley & Sons.
- Chivers, T. (2017) 'A mindset "revolution" sweeping Britain's classrooms may be based on shaky science', Available at: https://www.buzzfeed.com/tomchivers/what-is-your-mindset?utm_term=.loBQ1po0k#.jkPZbL0y8
- Cockcroft, W. H. (1982) *Mathematics counts*. London: HM Stationery Office.
- Coe, R. (2013) *Improving education: a triumph of hope over experience, CEM Inaugural Lecture*. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf
- Dweck, C. (2007) 'Boosting achievement with messages that motivate', *Education Canada* 47 (2) pp. 6-10.
- Dweck, C. (2014) *Mindset, the new psychology of success*. New York, NY: Ballantine Books.

- Hattie, J. (2017) 'Misinterpreting the growth mindset: why we're doing students a disservice', *Education Week blog*. Available at: http://blogs.edweek.org/edweek/finding_common_ground/2017/06/misinterpreting_the_growth_mindset_why_were_doing_students_a_disservice.html
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Li, Y. and Bates, T. (2017) 'Does mindset affect children's ability, school achievement, or response to challenge? Three failures to replicate'. Available at: https://osf.io/preprints/socarxiv/tsdwy/
- Middleton, J. A. and Spanias, P. A. (1999) 'Motivation for achievement in mathematics: findings, generalizations, and criticisms of the research', *Journal for Research in Mathematics Education* 30 (1) pp. 65-88.
- *TES Podcast* (2017) 'Growth mindset, cognitive load and the role of the researcher', 13th September, 2017.
- *TES Podcast* (2017) 'Professor Carol Dweck on growth mindset', 18th October, 2017.

My takeaway

I no longer believe that mistakes, struggle and failure are the key to learning, and indeed I think they can have quite the opposite effect. Let me try to explain what I mean.

A key element of motivation identified in Section 2.1 was some feeling of success, or belief that success was in reach. This seems logical – we tend to enjoy doing things we are either good at or believe we can become good at.

So, how would students' expectations of their likely success be affected by the presentation of complex problems combined with my bastard teacher approach? Telling them that the problem is a difficult one, offering no hints, and making it very clear that I was not going to give them support is hardly the kind of recipe that is likely to raise their expectations.

In their comprehensive review of research related to motivation in mathematics, Middleton and Spanias (1999) reach a similar conclusion. They argue that 'students' perceptions of success in mathematics are highly influential in forming their motivational attitudes'. They explain that the effort a person is willing to expend on a task is determined by the expectation that participation in the task will result in a successful outcome, and in order to allow students to feel successful in mathematics without undermining either the value of success or a healthy attitude toward failure, teachers must structure tasks such that they present an appropriate level of challenge and difficulty for students. Cockcroft (1982) adds support to this view: 'whatever their level of attainment, pupils should not be allowed to experience repeated failure. If this shows signs of occurring, it is an indication that the advance has continued too far and that a change of topic is needed'. Thus, mathematics activities must be difficult enough that students are not bored, yet must also allow for an experience of success given appropriate effort by the student.

So, my choice and presentation of task was clearly a poor one for many of my students. It was no surprise that so many of them lost motivation, gave up, and hence learned very little.

But I believe there is a bigger issue at play here. I think there is a widespread belief across education that we can somehow magically instil in our students that gritty determination to cope with anything that may come their way, and hence thrive in times of struggle and failure. I don't think it is that simple.

Let's take Dweck's (2014) concept of a growth mindset – something that much of Boaler's work is based on. Students with a growth mindset don't give up. They take on challenges, work hard, and do not fear struggle, failure or mistakes because they view them as learning opportunities. Many schools across the UK have gone growth-mindset mad over the last few years, with posters adorning the walls and motivational assemblies ringing in students' ears. Surely if we can instil a growth mindset into our students, then my concerns with the dangers of struggle and failure disappear into insignificance, because students with a growth mindset embrace such difficulties?

There are two problems with this.

First, several studies have been unable to replicate Dweck's original findings. For example, Li and Bates (2017) cite three different attempts to replicate Dweck's methods, all of which failed to produce results suggesting a significant effect of a growth mindset. Likewise, a study by Bahník and Vranka (2017) measured the mindset of university applicants taking a scholastic aptitude test. They found growth mindset was not positively associated with results of the test, did not predict change of the results for those who retook the test, did not predict participation in a future administration of the test, and did not predict the total number of tests taken.

But failure to replicate findings is common in research, and does not, on its own, falsify a theory. Indeed, Dweck herself responded (see Chivers, 2017) to the findings of the first study by saying: 'Not anyone can do a replication. We put so much thought into creating an environment; we spend hours and days on each question, on creating a context in which the phenomenon could plausibly emerge'. But if researchers cannot replicate the ideal conditions to foster growth mindsets, what hope is there that schools and teachers can do so? A few posters and assemblies at the start of the year are hardly likely to do the trick. Indeed, in a 2017 interview with John Hattie, Dweck herself concedes that much of her work on mindsets has been misinterpreted and misapplied in schools. And in an interview with *TES* later that year, Dweck explained: 'some educators have seen [growth mindset] as reducing to effort. They are just telling kids to try hard. Far from a growth mindset, that's called nagging … If you just go in there and explain growth mindset you can't expect to create it … there are many complexities and subtleties that need to be transmitted'. So the problem may not be the concept of a growth mindset itself, but the difficulty of instilling it in our students, especially in the short-term. It is a long-term belief from teachers that every child can improve that is reinforced to students every day. It is not enough to put posters up or have a few motivational assemblies, nor is it a case of simply saying *try harder*. In the same *TES* interview, Dweck explained that a growth mindset is 'developed over time through learning,

mentoring, hard work, good strategies'. And I fear many schools have failed to realise this.

But there is a bigger problem. I feel that a key component of a growth mindset that is often overlooked is the importance of success. Past success is where the belief that students will succeed again comes from. As one of my Year 9s, having sat through an assembly on mindsets, so eloquently put it: 'it's kind of hard to have a growth mindset when I keep doing shit on tests, sir'.

All of this is not to say that students should never be allowed to struggle, fail, or make mistakes. Far from it. Struggling, failing, making mistakes and learning from the experience is clearly a key part of mathematical development. If students down tools at the first sign of difficulty, how will they ever learn? No one is born with fully formed schemas, everyone holds misconceptions, and even the most carefully planned-out explanations are prone to misinterpretation. Furthermore, without the kind of classroom atmosphere where students are not afraid to try things, make mistakes, and admit to those mistakes, we as teachers are blind to their weaknesses and powerless to help them. When mistakes do occur they should be embraced and treated as the learning opportunities that they truly are. Lemov (2015) calls this a Culture of Error, emphasising the need to 'create an environment where your students feel safe making and discussing mistakes, so you can spend less time hunting for errors and more time fixing them'. This is 100% true, and I will dig deeper into how to identify and resolve errors in Chapter 11.

But my fear is that in the past I have gone too far. I have prioritised struggle over success. Students cannot struggle and fail all the time. Without the belief that they may succeed that only comes with the experience of past success, it will take a special kind of student to keep going, and going, and going.

What I do now

For Coe (2013), learning happens when students think hard. But I now realise that students may only be willing to think hard if they believe that

effort will pay off. Too much experience of past struggle and failure will only dampen that belief.

The work of Boaler (*eg* Boaler, 2015) and Dweck (*eg* Dweck, 2014) certainly does not say students should struggle and fail all the time. But I believe this is how it has been interpreted in many schools. By emphasising the importance of mistakes in the mistaken belief that we can magically produce gritty, determined students with growth mindsets, we are in danger of overlooking the importance of success. Success is the foundation upon which grit and a growth mindset must be built.

Indeed, I don't think mistakes are the key to learning – if repeated mistakes, struggle and failure lead to students putting in less effort, and eventually giving up, how can they be? I think identifying, understanding and resolving mistakes is one of the most important things we as teachers can do, but I believe success is the key to learning. It is success that motivates students to try harder, and helps them cope and thrive when things get tough.

So, how can I help my students avoid too much struggle and failure, whilst still ensuring they are thinking hard?

I can start by improving the tasks I give my students. The sequences challenge in the form I presented it above may be suitable for some students – those with strong domain-specific knowledge and experience of past success who will revel in the challenge – but for many other students it caused them to down tools. Hence, it can be restructured as:

Sequences Challenge
Here is a sequence:

2, 2√7, 14, 14√7

a) What do you multiply by to get from one term to the next?

b) What is the next term?

c) What is the nth term rule?

d) What is the 17th term?

e) What is the value of the 21st term divided by the 17th term? Try on your own, and ask your partner if you are stuck.

This is better, but still not great. The choice of activities we give our students is crucial. Ideally we should give them a taste of immediate success (within 20 seconds as a general rule), but with enough challenge to keep them thinking hard. In Chapter 10 we will look at a collection of tasks that I feel do exactly this, aiding the development of both procedural fluency and conceptual understanding.

I can also be more aware of the language that I use. Being told something is really difficult may well motivate some students, but for others it will only reduce their expectancy of success, and hence have the opposite effect on motivation than what we intend. Of course, this all comes down to our knowledge of the students we teach – there is no black-and-white rule that works for everyone. But I certainly have toned down the adjectives I use to describe the tasks I give my students. And following the advice of Dylan Wiliam in a 2017 podcast interview for *TES*, when a student says they can't do something, I correct them by saying: 'you can't do it yet'.

But this alone is not enough to help students feel they can be successful. Fortunately we have a whole book to try to figure out how to do it better.

OceanofPDF.com

2.8. Achievement and Motivation

What I used to think

I used to believe that motivation led to achievement. Hence, a lot of my time was spent planning for motivation.

I would try every trick in the book to motivate my students. Often this would fall under the category of 'external motivation' as described in Section 2.6, and not surprisingly has at best limited, short-term success. I found developing intrinsic motivation difficult. It was all well and good extolling my love of mathematics, creating a safe environment in which to make mistakes, banging on about the benefits of a growth mindset, and designing tasks that were accessible and gave students a sense of purpose, but if my students lacked the underlying mathematical knowledge, I often found myself fighting a losing battle.

But I carried on regardless, never for one moment considering that the direction of causation between motivation and achievement could in fact be the other way around.

Sources of inspiration

- American Psychological Association (2015) 'Top 20 Principles from Psychology for PreK-12 Teaching and Learning'. Available at http://www.apa.org/ed/schools/teaching-learning/top-twenty-principles.pdf.
- Ashman, G. (2015) 'Motivating students about maths', *Filling the Pail blog*. Available at: https://gregashman.wordpress.com/2015/11/12/motivating-students-about-maths/

- Didau, D. and Rose, N. (2016) *What every teacher needs to know about … psychology*. Woodbridge: John Catt Educational Limited.
- Garon-Carrier, G., Boivin, M., Guay, F., Kovas, Y., Dionne, G., Lemelin, J., Séguin, J. R., Vitaro, F. and Tremblay, R. E. (2016) 'Intrinsic motivation and achievement in mathematics in elementary school: A longitudinal investigation of their association', *Child Development* 87 (1) pp. 165-175.
- Guay, F., Marsh, H. W. and Boivin, M. (2003) 'Academic self-concept and academic achievement: developmental perspectives on their causal ordering', *Journal of Educational Psychology* 95 (1) p. 124.
- Martin, A. J. (2016) *Using Load Reduction Instruction (LRI) to boost motivation and engagement*. Leicester: British Psychological Society.
- Pink, D. H. (2011) *Drive: the surprising truth about what motivates us*. London: Penguin.
- Zimmerman, B. J. (2000) 'Self-efficacy: an essential motive to learn', *Contemporary Educational Psychology* 25 (1) pp. 82-91.

My takeaway

I apologise if the following sentence is the most obvious thing you have ever read in your life, but it certainly was not to me: *imagine if it was not the case that motivation in maths led to achievement, but that achievement led to motivation*.

Whilst it seems clear that there is a reciprocal relationship between the two – hence opening up the wonderful possibility of a virtuous cycle, whereby motivation leads to achievement, which leads to more motivation, and so on – evidence suggests that the direction of causation runs more strongly in the opposite way than you may think. Specifically, motivation may not predict achievement, but achievement does predict motivation. There are numerous studies and pieces of work that reach such a conclusion.

Garon-Carrier *et al* (2016), in their study on motivation and achievement in maths in elementary schools in the US, concluded: 'Contrary to the

hypothesis that motivation and achievement are reciprocally associated over time, our results point to a directional association from prior achievement to subsequent intrinsic motivation'. In other words, intrinsic motivation does not predict achievement in maths for primary school students, but achievement does predict intrinsic motivation.

Guay *et al* (2003) tested theoretical and developmental models of the causal ordering between students' self-perception and academic achievement. The authors studied students across three age groups (Grades 2, 3, and 4 from ten elementary schools) in an attempt to identify the direction of the relationship between self-perception and achievement. Across all three age groups they found support for a reciprocal effects model, giving us a virtuous cycle. Perhaps the most interesting finding was that while there is a strong correlation between self-perception and achievement, the actual effect of achievement on self-perception is stronger than the other way round.

Zimmerman (2000) argues that self-efficacy (defined as one's belief in one's ability to succeed in specific situations or accomplish a task) plays a major role in determining a student's motivation for their subject. He explains that self-efficacy beliefs are predictive of two indicators of students' motivation: rate of performance and expenditure of energy. Zimmerman concludes: 'this empirical evidence of its role as a potent mediator of students' learning and motivation confirms the historic wisdom of educators that students' self-beliefs about academic capabilities do play an essential role in their motivation to achieve'.

In his Load Reduction Instruction Model, Martin (2016) outlines practical strategies that teachers can follow to increase motivation. In early skill acquisition, these follow an explicit instruction model (Chapter 3), making use of many of the key principles of Cognitive Load Theory (Chapter 4). Martin concludes: 'Research shows there is a cycle that operates such that learning ("skill") fosters subsequent motivation and engagement ("will"). For example, self-efficacy is likely to be enhanced (or sustained) through the academic knowledge and skill that explicit instruction is shown to develop'.

The American Psychological Association's 'Top 20 Principles from Psychology for PreK-12 Teaching and Learning' (2015) states that as students develop increasing competence, the knowledge and skills that have been developed provide a foundation to support the more complex tasks, which become less effortful and more enjoyable. When students have reached this point, learning often becomes its own intrinsic reward.

In *Drive*, Pink (2015) argues that a key source of intrinsic motivation is mastery of a given topic or concept. Only when mastery has been achieved can something truly become enjoyable. As we become proficient in the basics, tasks require less effort and become more enjoyable.

Prominent maths blogger Greg Ashman (2015) states:

> The problem is that achievement and motivation are correlated with each other and so many commentators then assume that we can increase achievement by taking steps to increase motivation; motivation causes achievement. This is often the rationale for inquiry learning initiatives. Yet we can't be sure that it works this way around. It could just as well be true that achievement causes intrinsic motivation. This is what I tend to think. Or it could be the case that the cause acts in both directions; a virtuous circle.

Finally, in their book *What every teacher needs to know about … psychology*, Didau and Rose (2016) state: 'the more success students have experienced, the more likely they are to be motivated to work harder. Rather than motivation resulting in improved performance, it seems that improved performance leads to increased motivation'.

Assuming that it is in fact the case that achievement – and crucially, students' perception of their own ability to succeed – leads to motivation, then what are the implications?

Well, we can start by creating conditions favourable to motivation as described so far in this chapter. Control, purpose, the environment and the example set by the teacher are all important. But above all else, we need to plan for achievement, not for motivation. We need to teach our students in such a way as to maximise the chance of them learning and understanding.

We need to choose explanations, examples and activities that maximise learning, not ones which we believe will be motivating due to some misguided notion of relevance. For them to truly believe they can be successful, students need to taste success for themselves, not just be told they can be successful.

What I do now

I now believe that the most significant influences on students' levels of motivation are the interconnected feelings that they are successful, or that they can be successful, and that these feelings come directly from achievement. This can be either a virtuous or a vicious cycle. If students enjoy success, then they are motivated, which can lead to more success via the extra effort they put in. However, if students experience too much failure, then they may lose motivation, which leads to less effort and increases the chances of subsequent failure. Perhaps more importantly, I can have a positive influence on this driver of motivation, not through tricks and gimmicks, but through good teaching.

What does that good teaching look like? Well, we have the remainder of the book to tackle that. But here is a sneak preview:

1. I take a far more active role in lessons than previously, no longer afraid to use my expertise to teach explicitly (Chapter 3).
2. I think very carefully about the presentation of information so that my students' fragile working memories are focused upon the things that matter (Chapter 4).
3. I provide carefully planned explanations, examples and exercises designed to give students the very best chance of understanding a given concept (Chapters 6 and 7).
4. I plan activities that give my students an almost immediate sense of success, whilst having enough challenge that they are still thinking hard (Chapters 7 and 10).
5. If I am introducing a topic, I use the principles of Deliberate Practice to isolate, develop and assess a specific skill (Chapter 8).

6. If I am reviewing or revisiting a topic, I use the principles of Purposeful Practice, allowing students to develop procedural fluency whilst providing opportunities for deeper conceptual understanding (Chapter 10).

7. During all of this I ensure the environment in my classroom is open, friendly and positive – a place where students are not afraid to ask questions or admit mistakes.

8. I give myself the best chance of identifying, understanding and resolving any misconceptions revealed in such an environment using the principles of formative assessment multiple times every single lesson (Chapter 11).

9. I quiz repeatedly – every lesson, in fact – not for assessment purposes, but for learning purposes (Chapter 12).

10. I carefully introduce desirable difficulties to improve the storage and retrieval strengths of the knowledge contained in my students' long-term memories (Chapter 12).

11. Once students have reached a sufficient level of domain-specific expertise, I develop their ability to solve problems, enabling them to transfer their knowledge to different situations (Chapter 9).

12. When times get tough – as they inevitably will – I am careful with my use of external rewards, using them as a last resort to bump students into this virtuous cycle.

This is by no means a revolutionary approach, and to the casual observer it may appear a little bit, well, dull. But consider this: what is more motivating than being taught a subject well, in a supportive environment, by someone who cares, and seeing that you can achieve? I'd take that over a convoluted lesson about a David Beckham free kick any day of the week.

Such a realisation has been liberating for me. It has meant that I can focus all my energy and attention on helping my students both to be successful and to feel successful, instead of worrying too much if individual activities were

likely to be motivating or not. I can plan for achievement, not for motivation. In short, my job now boils down to this: *teach my students well*. If I can do that, motivation will take care of itself.

OceanofPDF.com

2.9. If I only remember 3 things…

1. It is important that maths matters to students, but real-life contexts are rarely the best way to go about achieving this.
2. Intrinsic motivation is better for the long run, but external rewards and sanctions have a key role to play, and may help nudge students into a virtuous cycle.
3. Motivation is directly influenced by achievement. If students are successful and believe they can be successful, they will be motivated.

OceanofPDF.com

3. EXPLICIT INSTRUCTION

Up until relatively recently, I did not know there was a huge debate in education – one that divides people into two distinct camps. In fact, if you had asked me to name the most contentious issue that I had encountered in my 12 years of teaching – one where I had to put my neck on the line and choose a side – it would have been Promethean *v* Smart Boards. I've had some heated debates at maths conferences over that one, I can tell you.

But these days I cannot wander through my Twitter timeline without coming across a back-and-forth – mostly friendly in nature, sometimes distinctly not – between a so-called Progressive and a so-called Traditionalist.

The concepts of traditional and progressive education are rooted in the ideas of John Dewey, an American philosopher and educational reformer, who identified two distinct types of education. In its simplest form, traditional teaching favours explicit instruction with a specific end-goal (usually an assessment), whereas progressive teaching sees schooling as part of a much wider approach to education. However, as with all labels, things are not quite that simple.

Just to complicate matters further, there is a related debate which pits two distinct approaches to teaching against each other. On one side there are those who suggest students learn best in an unguided or less guided environment, where students, rather than being presented with essential information, must discover, request or construct this information for themselves. Such an approach is given a variety of names including inquiry, problem, project or discovery-based learning – although as we shall see

these approaches have key differences from each other. On the other side are those who suggest that students, particularly in the early stages of knowledge acquisition, should be provided with direct instructional guidance on the knowledge required to understand a given concept.

Whilst I was unfamiliar with these debates, looking back over my career I am in no doubt at all which camp I sat in. I was a progressive teacher who felt it was clearly better for students to discover relationships for themselves. I liked nothing more than a rich task, investigation or problem-solving lesson. My job was not to impart my wisdom and insights on my students, but to provide the conditions under which they could generate their own wisdom and insights. I was the guide on the side, not the sage on the stage.

This chapter is an attempt to explain why I have completely changed my mind.

OceanofPDF.com

3.1. What makes great teaching?

What I used to think

I used to believe that there was no 'right' way to teach. At a number of presentations I have delivered to teachers across the country, I have uttered the phrase, 'Just as every child learns differently and we must cater for that, so every teacher should be allowed to teach differently'. This once prompted a standing ovation.

I am now no longer convinced I agree with either part of that sentiment.

Sources of inspiration

- Barton, C. (2017a) 'Dani Quinn – Part 1', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Greg Ashman', *Mr Barton Maths Podcast*.
- Centre for Education Statistics and Evaluation (2014) *What works best: evidence-based practices to help improve NSW student performance*. Available at: https://www.cese.nsw.gov.au/images/stories/PDF/What-works-best_FA-2015_AA.pdf.
- Coe, R., Aloisi, C., Higgins, S. and Major, L. E. (2014) *What makes great teaching? Review of the underpinning research*. Available at: https://www.suttontrust.com/wp-content/uploads/2014/10/What-Makes-Great-Teaching-REPORT.pdf.
- Luke, A. (2014) *On explicit and direct instruction*. Available at: https://www.alea.edu.au/documents/item/861.
- Rosenshine, B. (2008) *Five meanings of direct instruction*. Lincoln, IL: Center on Innovation & Improvement. Available at: http://www.centerii.org/search/Resources/FiveDirectInstruct.pdf.

- Rosenshine, B. (2012) 'Principles of instruction: research-based strategies that all teachers should know', *American Educator* 36 (1) pp. 12-39.

My takeaway

Now, before you throw this book on the floor in disgust, I am by no means suggesting that everyone needs to teach the exact same way. In my podcast interviews with Dani Quinn and Greg Ashman you can hear me balk at the idea of centrally planned lessons, as I fear that exciting new ideas for conveying a given concept may never see the light of day under such a system. Likewise, all teachers have unique sets of strengths and should play to those.

However, I have come to the realisation that there is an approach to teaching that is more effective than any other in what I will term the early knowledge acquisition phase of learning – that is, when students are being introduced to a concept or skill for the first time.

Now, you would be mad to take my word for this. After all, this is the man who tried to teach his Year 7s fractions via the medium of Swiss rolls (see Section 1.3). So, it is fortunate that there have been several comprehensive, evidence-based summaries, written by authors in different parts of the world, that delve into the question of what makes great teaching.

In his 'Principles of Instruction: Research-Based Strategies That All Teachers Should Know', Rosenshine (2012) presents 10 research-based principles from cognitive science and studies of master teachers, together with practical strategies for classroom implementation. The principles are:

1. Begin a lesson with a short review of previous learning.
2. Present new material in small steps with student practice after each step
3. Ask a large number of questions and check the responses of all students.
4. Provide models.

5. Guide students' practice.
6. Check for student understanding.
7. Obtain a high success rate.
8. Provide scaffolds for difficult tasks.
9. Require and monitor independent practice.
10. Engage students in weekly and monthly review.

We will return to several of these principles throughout this book, but for now I want to draw your attention to numbers 2, 3, 4, 5 and 8. These are the hallmarks of well-planned, teacher-led, fully guided instruction.

In their report for CEM entitled *What makes great teaching?*, Coe *et al* (2014) list the following together with their view on the evidence of impact on student outcomes:

1. Content knowledge (Strong evidence)
2. Quality of instruction (Strong evidence)
3. Classroom climate (Moderate evidence)
4. Classroom management (Moderate evidence)
5. Teacher beliefs (Some evidence)
6. Professional behaviours (Some evidence)

We will return to the first of these when we look at formative assessment (Chapter 11), distinguishing between the importance of knowing your subject and knowing where students are likely to struggle. However, it is the second that is relevant here. The authors give more detail:

> Quality of instruction is at the heart of all frameworks of teaching effectiveness. Key elements such as effective questioning and use of assessment are found in all of them. Specific practices like the need to review previous learning, provide models for the kinds of responses

> students are required to produce, provide adequate time for practice to embed skills securely and scaffold new learning are also elements of high quality instruction.

Once again, we have the hallmarks of teacher-led, guided instruction.

Finally, we travel to the southern hemisphere for the 2014 Centre for Education and Statistics report, *What works best: evidence-based practices to help improve NSW student performance*. Again, areas are identified that are likely to improve student learning, each presented with relevant data and links to research. These areas are:

1. High expectations
2. Explicit teaching
3. Effective feedback
4. Use of data to inform practice
5. Classroom management
6. Wellbeing
7. Collaboration

Explicit teaching once again finds a place high up the list.

Delving into why explicit instruction is so effective – and crucially identifying the practical strategies and techniques that have such a high impact on student learning – is a key aim of this book. However, at the outset I want to make clear that explicit instruction – as I will use the term – is neither 'chalk and talk', nor is it lecturing. Students are not passive recipients of information – they are fully involved in the learning process, indeed arguably more so than with less guided approaches where there may be opportunities for students to hide away. Neither is it necessarily boring, not least because of the conclusions from Chapter 2 on the positive effect of achievement on motivation, but also because teachers have the opportunity to take topics and concepts further and deeper, arming students with the

knowledge needed to excel. Neither does this approach preclude the use of high-quality rich tasks, problem-solving and inquiry-based activities – indeed I will argue that it allows students to get the very best out of these types of activities at a time when they are most ready for them. However – and this needs to be made very clear – explicit instruction certainly is teacher-led instruction.

One final point – the difference between explicit instruction and direct instruction is a little murky, and the two terms are often used interchangeably in the literature. For those wanting to know more, Rosenshine (2008) discusses 'five meanings of direct instruction', and Luke (2014) provides a detailed discussion of the similarities and differences between direct and explicit instruction. However, the key principles behind both models of instruction are the same, and hence for consistency and simplicity I will refer to 'explicit instruction' throughout this book.

What I do now

I strongly favour an explicit instruction model of teaching, especially in the early knowledge acquisition phase of learning. So, when I am introducing a topic for the first time, regardless of the age or prior achievement of the class, I will use an explicit instruction approach. I will use the strategies outlined in this book to make this approach to teaching as effective as possible.

'But why exactly are the less guided approaches that I used to swear by not as effective as explicit instruction?', I hear you say. That is the focus of the next two sections.

OceanofPDF.com

3.2. Are some students natural mathematicians?

What I used to think

I used to believe that maths came naturally to some students. Indeed, many a time I have adorned student reports with phrases such as 'Josie is a natural mathematician'.

So, is it true that some students are natural mathematicians?

Sources of inspiration

- Geary, D. C. (2007) *Educating the evolved mind: conceptual foundations for an evolutionary educational psychology*. Charlotte, NC: Information Age Publishing.
- Lin, J. and Geary, D. C. (1998) 'Numerical cognition: age-related differences in the speed of executing biologically primary and biologically secondary processes', *Experimental Aging Research* 24 (2) pp. 101-137.
- Paas, F. and Sweller, J. (2012) 'An evolutionary upgrade of cognitive load theory: using the human motor system and collaboration to support the learning of complex cognitive tasks', *Educational Psychology Review* 24 (1) pp. 27-45.

My takeaway

Geary (2007) argues that there are two types of knowledge and ability: those that are biologically primary and emerge instinctively by virtue of our evolved cognitive structures; and those that are biologically secondary and exclusively cultural, acquired through formal or informal instruction or training. Evolution over millions of years has led to us humans developing

brains that eagerly and rapidly acquire those things that are biologically primary, whereas the brain has simply not had enough time to adapt to make biologically secondary knowledge and ability as easy to acquire. Being able to communicate and being social have provided an evolutionary advantage to humans over millions of years in a way that being able to plot the graph of y = sin(x) has not – as my long, painful spell at internet dating clearly revealed. The key point is, much of the mathematics students learn in school is biologically secondary by its very nature, and hence does not come naturally to them.

Geary and Lin (1998) analyse three key areas of mathematics in an attempt to distinguish between primary and secondary knowledge:

1. *Enumeration:* Human infants and animals from many other species are able to enumerate or quantify the number of objects in sets of three to four items. The authors call this process 'subitizing', defined as 'the ability to quickly and automatically quantify small sets of items without counting'. But what about for larger sets? Well, adults typically use a combination of counting and estimation. This likely reflects the use of both primary and secondary competencies. Primary features include an implicit understanding of counting, and secondary features include learning the quantities associated with numbers beyond the subitizing range. Thus, the authors conclude that subitizing appears to represent a more pure primary enumeration process, whereas actual counting is likely to be biologically secondary.

2. *Magnitude Comparison:* The speed of determining which of two numbers or quantities is smaller or larger becomes slower as the magnitude of the numbers increases, but becomes faster and less error prone as the distance between the two numbers increases. No surprise there. The authors explain that the most conservative approach is to assume that the limit of these innate, primary representations is quantities associated with 1 to 3, inclusive, with the rest being biologically secondary.

3. *Addition and Subtraction:* It appears that human infants, preverbal children, and even the common chimpanzee are able to add and subtract items from sets of up to three (sometimes four) items. Although this

primary knowledge almost certainly provides the initial framework for the school-based learning of simple addition and subtraction, most of the formal arithmetic skills learned in school appear to be biologically secondary. For example, effective borrowing is dependent on a conceptual understanding of the base-10 structure of the Arabic number system and on school-taught procedures.

Hence, whilst humans have evolved to naturally acquire certain basic mathematical skills that form the foundation for more complex skills, these are not likely to get us very far through a school's curriculum, nor indeed through much of the real world. No one has a natural predisposition for learning trigonometry or even the times tables in the way that they are predisposed to speak or play. To be able to do such things we need to be taught, followed by practice. Simply immersing students in a series of unstructured maths tasks is not likely to enable them to learn the fundamentals of the topic in a way that immersion helps them to develop the ability to speak and communicate.

So why do some students appear to be more 'natural' mathematicians than others? This is likely the result of a complex concoction of prior knowledge, intelligence, motivation and pedagogical approach, with prior knowledge perhaps dominating in the way it aids the acquisition of new knowledge as discussed in Section 1.1. The key point is, however, that the majority of maths that students encounter at school does not come naturally to anyone. By assuming it does, we are in danger of placing too much of a burden upon our students. Most of mathematics is biologically secondary and hence, in the words of Geary (2007), 'is predicted to be heavily dependent on the teacher'.

Moreover, enabling students to acquire such biologically secondary content is difficult. Socialising with other students is likely to appeal more to students' biologically primary predispositions than the lovely algebra task you have laid on for them. As Geary (2007) puts it: 'children's inherent motivational biases and conative preferences are linked to biologically primary folk domains and function to guide children's fleshing out of the corresponding primary abilities. In many cases, these biases and preferences are likely to conflict with the activities needed for secondary learning'.

What I do now

I have reached the conclusion that no one is a natural mathematician – never shall that phrase appear on a report again. All students need help learning the biologically secondary mathematical knowledge that will get them through secondary school and beyond. The question is, how best to help them acquire that knowledge? Given that such knowledge does not come naturally, and the classroom is awash with biologically primary distractions, it is my view that an approach involving less guidance during instruction is not optimal.

When something does not come naturally, students need all the help they can get to understand the content and to remain focused. The assumption that students will acquire this knowledge themselves, even under a model of guided discovery that we will discuss in Section 3.4, is likely to lead to difficulties. It places too much of a burden on students. As teachers – experts in our field – we are best placed to help students acquire that fundamental knowledge, presenting it in a way that is meaningful and accessible, accounting for the limits of their fragile working memories. Indeed, it is our duty to do so.

Geary (2007) is very clear on both his appreciation of the challenge of enabling students to acquire biologically secondary knowledge, and his solution:

> The gist is that the cognitive and motivational complexities of the processes involved in the generation of secondary knowledge and the ever widening gap between this knowledge and folk knowledge leads me to conclude that most children will not be sufficiently motivated nor cognitively able to learn all of [the] secondary knowledge needed for functioning in modern societies without well organized, explicit and direct teacher instruction.

A model of carefully planned explicit instruction is likely to be more successful in terms of learning than leaving students to their own devices under a model of partial guidance, whether we perceive our students to be 'naturals' or not.

OceanofPDF.com

3.3. When and why less guidance does not work

What I used to think

I flipping love an investigation, inquiry or strange problem to solve. These were the kinds of activities that got me interested in maths as a child. As such, I wanted to use as many as possible with my own students – indeed, it was one of the main reasons I became a teacher (closely followed by the holidays). I saw it as my duty to present students with tasks, problems and activities that would allow their creativity and imagination to flourish, and then make sure I got out of their way.

And so the schemes of work that I co-wrote in my current school contained a 'compulsory rich task' for every topic we covered. These were high-quality activities that all classes would do, and which called for less teacher guidance. Here are two of my favourites.

Diagonals of a Rectangle

Draw any size rectangle on square paper (whole number width and height)

Draw a line from the bottom-left corner to the top-right corner

Count how many squares the line passes through

Investigate!

I adapted this from an old GCSE coursework task, and all our Year 8 classes did it during the second term before a unit on Factors and Multiples.

Surds Inquiry

$$\sqrt{2\tfrac{2}{3}} = 2\sqrt{\tfrac{2}{3}}$$

Figure 3.1 – Source: Rachael Reed, via Inquiry Maths, available at http://www.inquirymaths.org/home/number-prompts/surds

This beauty is taken from Andrew Blair's Inquiry Maths website (http://www.inquirymaths.org). As explained on the homepage: 'Inquiry Maths is a model of teaching that encourages students to regulate their own activity while exploring a mathematical statement (called a prompt)'. We used this particular prompt with our higher tier Year 11 students at the end of a unit on surds.

I looked forward to delivering these lessons so much, imagining all the discoveries my students would make, the lines of inquiry they would pursue, and the love of maths that would develop. But they rarely went as well as I hoped. Not once did all my students end up where I wanted them to be, both in terms of their enjoyment and achievement. The insights and conclusions they reached were rarely complete, and often erroneous. Precious time would be spent discussing and debating, correcting and re-explaining – time that could have been spent practising and applying our new-found knowledge. It broke my heart.

At the time, I blamed myself and my deficiencies as a teacher. And to a large extent, I still do. But until now, a question I never really asked myself was: are these approaches that favour less teacher guidance really the best thing for all my students all of the time?

Sources of inspiration

- Barton, C. (2016) 'Dylan Wiliam', *Mr Barton Maths Podcast*.
- Barton, C. (2017) 'Andrew Blair', *Mr Barton Maths Podcast*.
- Blair, A. (2017) *Inquiry Maths*. Available at: www.inquirymaths.co.uk/
- Clark, R., Kirschner, P. A. and Sweller, J. (2012) 'Putting students on the path to learning: The case for fully guided instruction', *American Educator* 36 (1) pp. 6-11.
- Engelmann, S., Becker, W. C., Carnine, D. and Gersten, R. (1988) 'The direct instruction follow through model: design and outcomes',

Education and Treatment of Children 11 (4) pp. 303-317.

- Geary, D. C. (2007) *Educating the evolved mind: conceptual foundations for an evolutionary educational psychology*. Charlotte, NC: Information Age Publishing.
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.
- Kirschner, P. A. and Van Merriënboer, J. J. G. (2013) 'Do learners really know best? Urban legends in education', *Educational Psychologist* 48 (3) pp. 169-183.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Martin, A. J. (2016) *Using Load Reduction Instruction (LRI) to boost motivation and engagement*. Leicester: British Psychological Society.
- Mourshed, M., Krawitz, M. and Dorn, E. (2017) *How to improve student educational outcomes: new insights from data analytics*. Available at: www.mckinsey.com/industries/social-sector/our-insights/how-to-improve-student-educational-outcomes-new-insights-from-data-analytics
- Schoenfeld, A. (1983) 'Episodes and executive decisions in mathematical problem solving' in Lesh, R. and Landau, M. (eds) *Acquisition of mathematics concepts and processes*. New York, NY: Academic Press, pp. 345-395.
- Schoenfeld, A. (1985) *Mathematical problem solving*. Cambridge, MA: Academic Press.
- Schoenfeld, A. (1987) 'What's all the fuss about metacognition?' in Schoenfeld, A. (ed.) *Cognitive science and mathematics education*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 189-215.

- Schoenfeld, A. (2009) 'Learning to think mathematically: problem solving, metacognition, and sense-making in mathematics' in Grouws, D. (ed.) *Handbook for research on mathematics teaching and learning*. New York, NY: MacMillan, pp. 334-370.
- Sweller, J., Mawer, R. F. and Howe, W. (1982) 'Consequences of history-cued and means-end strategies in problem solving', *The American Journal of Psychology* 95 (3) pp. 455-483.
- Wu, H. (1999) 'Basic skills versus conceptual understanding', *American Educator* 23 (3) pp. 14-19.

My takeaway

What is less guidance?
At the outset, it is important to clarify what I mean by 'less guidance'. In the literature – in particular in the influential 2006 paper by Kirschner *et al* entitled 'Why minimal guidance during instruction does not work' – many approaches are lumped together under a single phrase, including inquiry-, problem-, project- and discovery-based learning. In addition, such minimal or partial guidance can also refer to periods of lessons where students are given open-ended problems to tackle, puzzles to solve, or investigations to work through.

Such all-encompassing labels are dangerous. Indeed, in a 2017 blog post entitled 'Inquiry is NOT discovery learning', the creator of the Inquiry Maths website, Andrew Blair, is keen to point out the differences in his approach. He concludes, 'in discovery learning, the teacher attempts to preserve the pretence of discovery, even to the extent of withholding knowledge; in inquiry, the teacher, as a participant in the classroom activity, aims to introduce subject-specific knowledge when it is most relevant and meaningful to her students'. Furthermore, when I interviewed Andrew for my podcast, he explained that there were different types of inquiry, ranging from structured, to guided, right through to open – and indeed, when students request information via use of their Regulatory Cards, spells of teaching that closely resemble explicit instruction can occur.

However, whilst there are important differences between the approaches that call for less teacher guidance, they do have one thing in common. This is summarised nicely by Andrew in his description of Inquiry Maths on the website's homepage: 'students take responsibility for directing the lesson with the teacher acting as the arbiter of legitimate mathematical activity'. It is this point – the fact that greater responsibility is placed in the hands of students – that distinguishes less guided approaches to teaching from a model of explicit instruction.

Take the two activities above – Diagonals of a Rectangle and Surds Inquiry. There is no worked example, or Question 1 that all students start on. Sure, the teacher may give hints, but control and responsibility for what happens next and throughout the tasks rests firmly in the hands of the students.

Is this necessarily a bad thing? Well, I fear it might just be.

Why less guidance may not be good for learning

Students may not have the capacity to weigh up both the demands of the task and their own learning needs in relation to the task

This was one of Kirschner and van Merriënboer's (2013) findings discussed in Section 2.2. The kinds of tasks and activities that favour less teacher guidance are brilliant in the sense that students do not know where they will lead, as is the case in both the Rectangles and Surds tasks. But this causes us a problem when we are also expecting students to work at an appropriate level of challenge throughout them. With students free to pursue their own paths and lines of inquiry, how do we know that what they are working on is best for their learning?

Students continue to practise tasks they like or are already proficient in, but are reluctant to start with new, yet unfamiliar tasks

Another one from Kirschner and van Merriënboer's (2013) that I've observed many times. Take the Rectangles task. Students will happily draw rectangles all day long, neatly recording their results in a table, but are less willing to look for patterns, make predictions, and test those predictions out – something that is needed to develop their understanding. Schoenfeld (*eg* 1983, 1985, 1987) demonstrates that expert problem-solvers frequently

engage in metacognitive acts in which they pause and reflect on the approaches they are using, whereas novice problem-solvers are often observed to become fixated on an approach and pursue it relentlessly, however unprofitably. But, as I will argue in Section 9.2, metacognition is not something we can easily teach as it is inherently tied to domain-specific knowledge.

Some students do not know where to start
How would your students react to the Rectangles and Surds tasks? Some of my students would dive right in, but I know many others who would be completely stumped. Schoenfeld (2009) presents an interesting analysis of what happens when novice learners attempt to solve problems without the guidance of a teacher. In a 20-minute period, around 18 minutes are typically spent exploring. In contrast, an expert in the subject will cycle through stages of analysing, planning, implementing and verifying. Now, of course, we can offer our students support and guidance to help them approach these tasks in a structured and systematic way, but by doing so are we not moving further towards a more teacher-led form of explicit instruction?

The activities and lessons harder to plan
This may well be my shortcomings as a teacher, but I find such activities far more difficult to plan than the teacher-led approach I will develop throughout this book, and as such I feel they can have a negative effect on my students' learning. Lemov (2015) discusses the need for teachers to Plan for Error. With students pursuing many different paths through these activities, it becomes far harder to anticipate where they may go wrong, and hence plan appropriate support.

Mistakes and misconceptions are harder to identify and act upon
Kirschner *et al* (2006) argue that students are more likely to acquire misconceptions, or incomplete or disorganised knowledge via less teacher-led approaches due to the lack of guidance. Indeed, the problems of over- and under-generalisation of rules that we will discuss in Section 7.2 are tricky enough to overcome when we as teachers have control over the examples we present, so it is little wonder they develop when students are left to their own devices. Wu (1999) makes a related argument when

questioning the practice of letting students devise their own methods to solve problems: 'What is left unsaid is that when a child makes up an algorithm, the act raises two immediate concerns: 1) whether the algorithm is correct 2) whether it is applicable under all circumstances'. But is this a problem? After all, students are bound to make mistakes and form misconceptions, no matter how carefully the lesson is planned. However, the key point is that under a less guided approach these mistakes and misconceptions are harder to identify and hence deal with appropriately, with students working on many different things. Indeed, issues of correctness and generality are tricky enough to discern and resolve when working one-to-one with a student, but how about in a class of 30, with several students enthusiastically asking you to check their work at once? Consider the surds inquiry versus a more structured lesson involving worked examples and carefully varied practice – I have to concede that mistakes and misconceptions are likely to be fewer and easier to identify in the second case. As we have seen with Anderson's model of thinking in Section 1.1, practice makes permanent, and hence once developed, this incomplete or erroneous knowledge may be hard to rectify.

Students may struggle and not learn

This is the big (and most controversial) one. Sweller *et al* (1982) argue that '**learners can engage in problem-solving activities for extended periods and learn almost nothing**'. To fully understand the rationale behind this bold claim, we need to understand the concept of cognitive overload (Chapter 4) and the different ways expert and novice learners solve problems (Chapter 9). Much more of this to come, but it is worth considering at this stage whether all the times you see your students struggling on problems, tasks and activities they are in fact learning as effectively as they could be with more teacher guidance. I know my students were not.

Advantages of less guidance

Let us now turn our attention to the potential benefits of the less guided forms of instruction that I believed so passionately in for most of my career.

They are more motivating, providing students with a sense of purpose and control

We have seen in Chapter 2 how a sense of purpose and control are key components of motivation, and less teacher-led forms of instruction certainly provide both of these. Hearing Andrew Blair talk on my podcast about the way students are hooked in by the prompts in an inquiry lesson, propelling them to request knowledge from the teacher in order to take their inquiry further, is inspiring. It certainly provides a sense of purpose that I would have found motivating as a student. But we need to be careful on a number of points. First, I was a bit of a weirdo at school – I loved all maths, but especially when I was encouraged to investigate, experiment and inquire. Looking back some 20 years later – and indeed, as I will argue in this book – I think that was because I had the domain-specific knowledge to investigate, experiment and inquire. Without such knowledge, these activities would probably have been the frustrating experience they were for many of my classmates and now are for many of my students. We have already discussed how handing over too much control can be a bad thing in terms of the poor decisions students may make, but it is also worth mentioning that many students – particularly novice learners – appreciate and enjoy more guidance from the teacher, reducing feelings of anxiety and vulnerability. Second, it is not fair to dismiss explicit instruction as dull and boring where students are passive in the learning experience. As I hope to show, the model I propose is both interactive, challenging and ultimately rewarding. Finally, we need to remember from our work in Chapter 2 that one of the key drivers of motivation is success and a feeling that students can be successful. I believe that during the initial stages of learning a new skill or concept, students are more likely to be successful under the model of explicit instruction that I will develop throughout this book.

Students can do the activities in lots of different ways

This is certainly true for many such activities – the Rectangles task is a classic example – but is this necessarily a good thing? We have seen above why such freedom may not be all that motivating for all students. But what is the potential impact on their learning? Students remember what they think about or attend to. By giving students such freedom we reduce the chances of them attending to the things that really matter. During an inquiry students may go off on a mathematical tangent, or during an investigation students are likely to choose examples in a non-systematic way, with too

much variation to enable them to spot key patterns or relationships. In Section 7.8, I will argue that a key strength of the explicit model of instruction I propose is the use of Variation Theory, controlling the sequence of questions students answer to direct their limited attention to the most important elements that are changing in order to develop their conceptual understanding. Likewise, one of the strengths of Purposeful Practice (Chapter 10) is that students have to practise the relevant procedure.

They allow students to understand concepts, not just carry out meaningless skills

This all comes down to what we mean by the term 'understand'. In section 3.9 I will argue that understanding is a combination of procedural fluency and conceptual understanding, and that it does not necessarily follow that conceptual understanding should come first. Taking this further, I believe that a model of explicit instruction is the best way for students to gain the procedural fluency that is necessary to develop understanding in many of the concepts in mathematics.

They allow students to develop wider skills, such as independence and problem-solving

When I interviewed Dylan Wiliam for my podcast, he described Project Follow Through as the most important piece of educational research ever conducted – and yet I had never heard of it. Project Follow Through is the most expensive and extensive educational experiment in history. Beginning in 1968 under the sponsorship of the US government, it was charged with determining the best way of teaching at-risk children from kindergarten through to Grade 3. Over 200,000 children in 178 communities were included in the study. Nine models, grouped into three broad teaching approaches were used: academic focus, problem-solving and self-esteem focus. The results were startling. Siegfried Englemann's model of Direct Instruction outperformed all other models in basic skills. However, what is perhaps most interesting is that Direct Instruction also outperformed all other models on measures of problem-solving and self-esteem. As I will argue in Chapter 9, domain-specific knowledge is the key to the

development of both problem-solving and independence, and explicit instruction is the best way to acquire this knowledge.

Do we need a balance?

One conclusion to draw from this is that a balance is needed – a solid dose of explicit instruction mixed in with a sprinkle of partial guidance. Indeed, a 2017 report from McKinsey concluded that 'students who receive a blend of teacher-directed and inquiry-based instruction have the best outcomes'. So, in the past, I might choose to introduce one concept via formal, teacher-led guided instruction, and the next via a discovery or inquiry-based task.

But that is not quite right.

A balance certainly is needed between these two type of instruction, but the appropriateness of each is determined by one thing and one thing only – the domain-specific knowledge of the students.

Clark *et al* (2012) do not advocate using all aspects of explicit instruction all the time. Indeed, they recognise the need for learners to be given the opportunity to solve problems independently, but assert this should be used as a means for practising newly learnt content and skills, not to discover information themselves. Likewise, Kirschner *et al* (2006) are careful to point out that some students may well benefit from a more partial guided form of instruction. Such students – experts in a domain-specific area – reach a stage where fully guided instruction in the form of the teacher-led worked examples we will consider in Chapters 6 and 7, followed by Chapter 8's Deliberate Practice, stops being beneficial, and a more independent form of problem-solving is required to further increase their understanding. Such a phenomenon is known as the Expertise-Reversal Effect and will be considered further in Section 6.7. When students reach this stage, carefully chosen inquiries, investigations and more open-ended tasks are likely to be the most appropriate form of activity.

But such less guided instruction must come at the end. Martin (2016) advocates the following five principles as key points in the learning process:

1. Reducing the difficulty of a task during initial learning

2. Instructional support and scaffolding through the task
3. Ample structured practice
4. Appropriate provision of instructional feedback
5. Independent practice, supported autonomy, and guided discovery learning.

The key point is that this is a progression. In the past I have either started with the final stage, or been too quick to label my students as 'experts' and move them onto a less teacher-led approach. I have been very much of the mindset that once I have shown students the basics, away they go to figure out the rest, and we can fill in any gaps in knowledge alongside the development of problem-solving skills. I now believe this was wrong. As I discussed in Section 1.2, the distinction between novice and expert learners is a blurry one, with many shades in between. Nevertheless, until students have acquired and automated key knowledge, more teacher-led instruction is needed.

We must also be careful in concluding that higher-achieving students need less guided instruction than lower-achieving students. There are very few students who are either an expert or a novice in all areas of maths. The key is to determine the domain-specific level of knowledge – regular low-stakes testing (Chapter 12) and effective formative assessment (Chapter 11) are ways to identify this – and decide the appropriate style of instruction from there.

I would even go one stage further. The first time anyone – top-set, straight-A student, or someone who always struggles – meets a skill or concept for the first time, they are all novice learners together. Sure, the former is likely to have an advantage with more related knowledge stored, organised and automated in long-term memory, which will make the acquisition of new knowledge that much quicker and easier. But they are novices nonetheless. Hence, approaches and activities that require less guidance during instruction are never appropriate in the early knowledge acquisition phase. I know that is a big claim, but I am sticking with it. Likewise, in the next

section I will go on to argue why halfway-houses such as guided discovery are not appropriate in this phase of learning either.

What I do now

My teaching used to be filled with less guided activities such as the Rectangles and Surds tasks. Now it is not – at least not in the initial phase of acquiring knowledge. That makes me a little bit sad, but I believe ultimately it is better for my students, both in terms of their learning and their overall enjoyment of mathematics.

When students first encounter a concept, I believe they need clear explanations, well chosen examples and what I will come to describe as Intelligent Practice. These are the hallmarks of explicit instruction. When first introducing a skill or concept to students, the responsibility for directing the lesson has to belong to the teacher. We are the experts, and it is at this stage in their learning that students need us most. Resting this responsibility on the shoulders of novice learners risks jeopardising their learning, and that is simply not fair. It is our job to teach – and teach well – not to facilitate.

In our podcast interview, Andrew Blair explains that he uses inquiries to introduce key skills and concepts. The idea is that students reach a point in the process where they realise they are stuck and request help (via a Regulatory Card) which Andrew then provides via a model that closely resembles explicit instruction. Students realising they are stuck provides the purpose and motivation for them to want to acquire the new skill. I believe students can and do learn basic skills from these activities, but it will take a lot longer, and with the likelihood of many bumps and bruises along the way.

I prefer to do it the other way around. I will assess students' initial understanding (see Chapter 11), and then provide the examples and Intelligent Practice (Chapters 6, 7 and 8) needed to help students on the journey from novice to expert. This does not mean my students will be deprived of rich, challenging, interesting maths. Far from it. The principle of Purposeful Practice (Chapter 10) will ensure they develop procedural fluency in an interesting way, and then the collection of less guided activities

that I love so much will be ready and waiting, at a time when my students are best placed to benefit from them. 3.4. The problem with guided discovery

What I used to think

There is an approach to teaching that sits somewhere between explicit instruction and models that require less teacher guidance, and it is often used to introduce and develop knowledge during the initial phase of learning. It is an approach that I believed had all the advantages of the latter in terms of opportunities for purpose, independence and students taking ownership over their learning, combined with the predictability and control that I felt explicit instruction delivered.

Such an approach goes by several names – *guided or structured discovery* and *pattern sniffing* are some of the most common – and for many years I felt it was the Holy Grail.

Lessons and activities in the mould of guided discovery usually allow students to experiment and theorise, but in a controlled way. The idea is not to prove something, but to spot a pattern or a relationship. Humans are naturally good pattern-spotters – indeed Mattson (2014) argues that 'superior pattern processing' is the fundamental basis of most, if not all, unique features of the human brain. And so I convinced myself that once students identify such a pattern or a relationship, they are then more invested in the concept and eager to learn more. Indeed, if students do spot the relationship, then my job may well be done, and I can get a cup of tea.

Guided discovery lesson plans can be found on most topics in maths. Geometry is a particularly fertile breeding ground. Take something like circle theorems. Instead of simply explaining to students the Angle at the Centre relationship, why not have them discover it for themselves? Give them a set of blank circles, instructions to construct several formulations of the theorem, each time giving them complete freedom as to where they place their three points on the circumference, challenge them to measure the two relevant angles and then see what they notice. Students get important practice of measuring angles, a feeling of involvement in their

own learning, and may even teach themselves a key GCSE topic without me needing to say a word. What could possibly go wrong?

Similar approaches are not hard to imagine on topics such as angles in polygons, Pythagoras, trigonometric ratios, the formula for the area of a trapezium, and so on. Indeed, examples of all of these can be found on the *TES* website within seconds.

I was particularly proud of a guided discovery task I came up with for introducing some of the more complex laws of indices to my Year 11 class two years ago. The worksheet looked like this:

Discovering Laws of Indices Investigation

By experimenting with your calculator, see if you can discover the rules for each of the following type of indices.

Index Law	Examples
$(\frac{p}{q})^a$	$(\frac{1}{2})^2$, $(\frac{1}{4})^3$, $(\frac{2}{3})^2$, $(\frac{3}{5})^4$
n^{-1}	5^{-1}, 2^{-1}, 0.25^{-1}, $(\frac{2}{5})^{-1}$
n^{-a}	3^{-2}, 4^{-3}, 0.1^{-5}, $(\frac{4}{5})^{-2}$
$(n)^{\frac{1}{a}}$	$(4)^{\frac{1}{2}}$, $(25)^{\frac{1}{2}}$, $(\frac{9}{16})^{\frac{1}{2}}$, $(64)^{\frac{1}{3}}$
$(n)^{\frac{b}{a}}$	$(9)^{\frac{3}{2}}$, $(8)^{\frac{5}{3}}$, $(\frac{16}{25})^{\frac{3}{2}}$, $(64)^{\frac{2}{3}}$
$(n)^{-\frac{b}{a}}$	$(25)^{\frac{1}{2}}$, $(16)^{\frac{3}{2}}$, $(\frac{27}{125})^{\frac{1}{3}}$, $(\frac{8}{27})^{-\frac{2}{3}}$

Each time you think you have it, test it out with a few more numbers and then try to explain it in your own words.

Figure 3.2 – Source: Craig Barton

Nice, eh? Again, I ask the question: what could possibly go wrong?

Well, quite a lot, as it turns out.

Sources of inspiration

- Coe, R. (2013) *Improving education: a triumph of hope over experience, CEM Inaugural Lecture*. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.
- Mattson, M. P. (2014) 'Superior pattern processing is the essence of the evolved human brain', *Frontiers in Neuroscience* 6 (8) pp. 1-17.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

When considering a guided discovery task, the question I should have asked myself is: *what is the best that can happen?*

Take the laws of indices lesson. The best that can happen is that all students discover the laws of indices for themselves, leaving no gaps in their knowledge, nor developing any misconceptions, in a reasonable time frame. We can then proceed with the rest of the lesson, maybe moving on to application questions, or interleaving other topics into the examples (see Chapter 12), such as indices involving surds or fractions.

How often does that actually happen?

In my experience, literally never.

What actually happens is that one or two students discover exactly what I wanted them to discover. They are feeling great about themselves, and rightly so – as we have seen in Chapter 2, success is motivating. A handful of students have some kind of idea what is going on, but with an eclectic mix of gaps in their knowledge and newly formed misconceptions. Some of these students are aware they have gaps and misconceptions, others are blissfully

ignorant. And the rest of the students do not have a flipping clue what is going on. They are feeling confused and pretty down about themselves when they see their fellow classmates have figured it out.

Any form of decent formative assessment strategy (Chapter 11) quickly reveals this disparity between levels of understanding, and as such I cannot move on with the lesson. So what do I inevitably end up doing? Teaching the laws of indices, of course! Maybe I will set those students who seem to have understood it off on the work I hoped everyone else would be moving on to – mind you, I would really like them to hear my explanation and do the worked examples, but how can I justify doing so when they have demonstrated their understanding?

Hmmm…

Anyway, back to the rest of the class.

By this stage, I am 30 minutes into a 50-minute lesson, rattling through a series of worked examples on the laws of indices far quicker and with much less care than I should. There is zero time for the students to practise their newly acquired skills and hence consolidate their knowledge, nor sufficient time for me to do any kind of application questions which would show them the full breadth of the topic.

But it is even worse than that. Even if I could somehow freeze time and spend those lost 30 minutes going through carefully structured and well-chosen worked examples (Chapters 6 and 7), I am not back at square one. I am behind square one, because my students are no longer coming at the topic with fresh eyes. Many of those who failed to 'discover' the key relationships have already decided that indices are difficult, and yet another area of maths that they don't understand. It's going to take more than my magically retrieved 30 minutes to turn that one around.

And so I wave goodbye to a group of confused students trundling out of the door, promising that we will pick this up again tomorrow, assuring them it will all be fine. I am already dreading the lesson, wanting to open up proceedings by saying 'okay, everyone, forget what happened yesterday'. In the past, I blamed my students for this – if only more of them could have

figured it out. Now, I know the blame rests squarely at my feet. I never gave them a chance.

It is even easier to imagine things going wrong in the Circle Theorem discovery task. A misdrawn line here, a poorly measured angle there, and before you know it students have invented a whole new set of circle theorems, and we are left to pick up the pieces.

Now, you could argue that I am painting an unrealistic picture here. And, maybe I am. But I tell you what: I think this picture is a damn sight more realistic than one in which all students have discovered the thing we want them to discover.

What I do now

Now, I know what you are thinking – 'God, your lessons must be so boring. What do you do – simply tell students how to do things, and that is that?'.

Sometimes, yes, but not always. Furthermore, I don't think it is boring, and I certainly believe it is more effective. Let me try to show you what I mean.

Introducing Circle Theorems

How would I introduce this lesson now?

I would begin by asking myself what I hope to achieve. Is it a complete understanding of why the Angle at the Centre theorem works? No, it is not. In Section 3.9 we will consider when the How should come before the Why, and this is a prime example. The proof is too conceptually complicated, and can be covered later on, perhaps in a unit on geometrical proof as a way of recapping circle theorems. No, my aim here is to develop *procedural fluency*. We will discuss the true meaning of this phrase and its relation to conceptual understanding in that same section, but for now it is enough to say that I want my students to be able to recognise the relationship and answer a wide range of questions on it. For Willingham (2002) the development of such inflexible knowledge is a necessary step towards expertise.

I would ask students to look at the board in silence. Starting with a blank screen in Geogebra, and without saying a word, I would construct the Angle

at the Centre theorem with both angles clearly marked on.

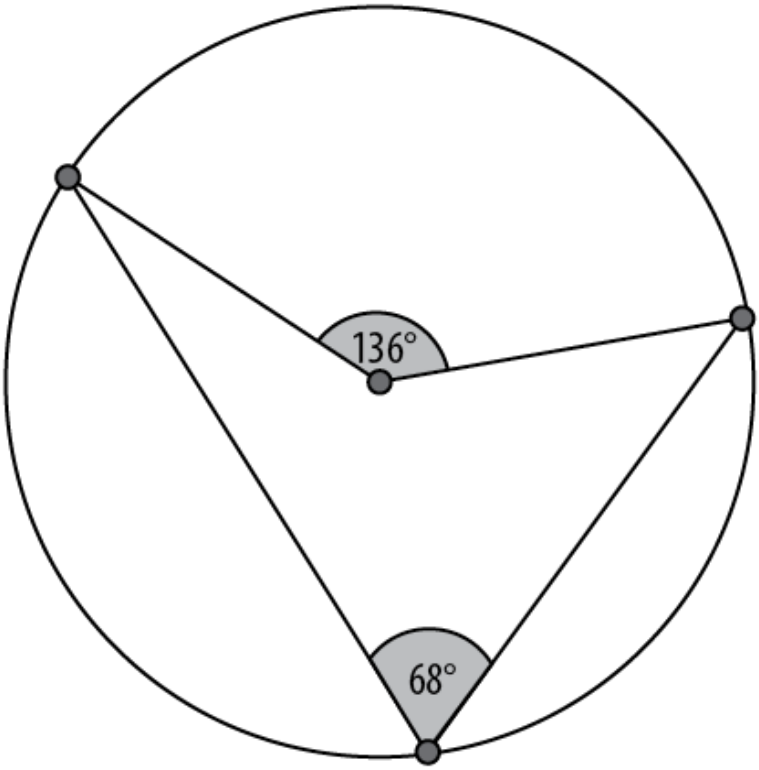


Figure 3.3 – Source: Craig Barton, created using Geogebra

I would then move one of the points on the circumference so the sizes of the angles changed. Again, I remain silent and so too are my students – no questions or discussions. I would then move the other point to a new a position (note: this approach makes use of the Modality Effect that will be covered in Section 4.3, as well as Section 4.8's Silent Teacher). At this stage I will include a non-example. The importance of non-examples in the context of mathematical rules will be discussed in Chapter 7, but just to whet your appetite, I would modify my Geogebra animation so it looked like this:

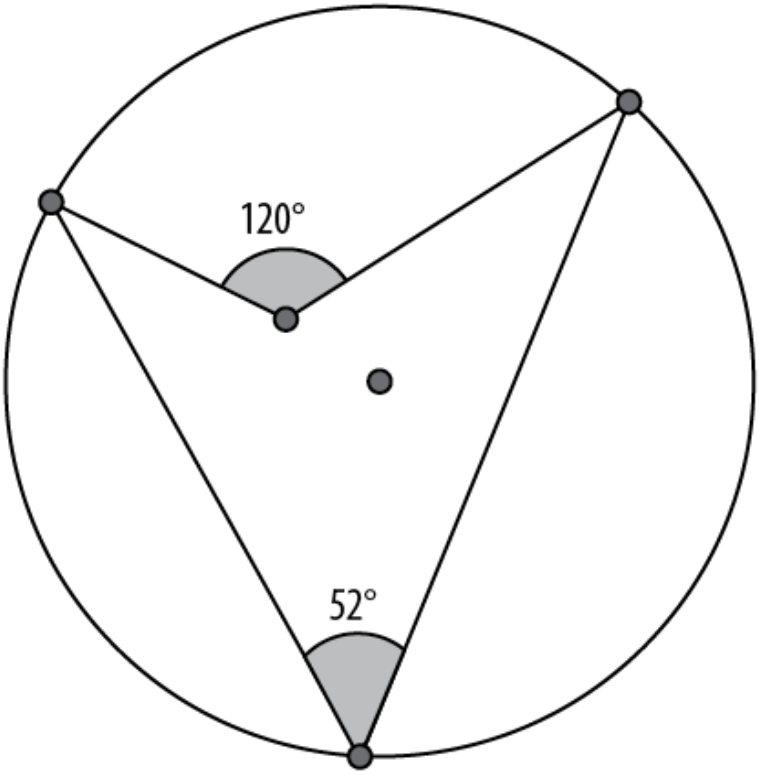


Figure 3.4 – Source: Craig Barton, created using Geogebra

Having gone through a series of silent examples and non-examples, I would then say, 'Have a think for 30 seconds what you reckon is going on, and then when I say, try to describe it to the person next to you'. The importance of this moment of pausing, reflecting and self-explanation before discussion with a peer is covered in Chapter 5.

Given that this is a relatively straightforward relationship to spot when presented in the controlled way I have described, I will choose someone and ask them to offer their interpretation of what is happening. If the

relationship was more subtle, then I would have no fear in offering my explanation. My priority is to prevent confusion spreading around the class by making the conclusion of the demonstration as clear and unambiguous as possible.

A student may say something like: 'halve the angle in the centre, you get that one on the edge'.

The use of technical language in mathematics is both a fascinating and problematic area. Of course we want all students to be able to describe the relationship using the proper terminology – something along the lines of 'the angle subtended at the centre of a circle is double the size of the angle subtended at the circumference from the same two points'. Indeed, in order to get full marks at GCSE, students will need to use such language. However, how important is it in the context of this demonstration? Remember, I am attempting to help my students develop procedural fluency. The careful choice of examples and non-examples has enabled them to build a good understanding of the Angle at the Centre relationship, and me then insisting they use an abstract, confusing phrase almost seems like a back step. Furthermore, if students' fragile working memories are consumed trying to recognise the relationship, process the calculation and remember the correct language, the development of procedural fluency that is my overarching aim may be jeopardised – more on this in Chapter 4's Cognitive Load Theory. Matching up examples of such rules to their proper technical definitions is a separate skill that can – and should – be taught apart from the development of an understanding of the rule itself.

So, at this stage I would probably insist that we use 'circumference' in place of 'edge', as this is a term students are (or at least should be) familiar with, but leave 'subtended' for much later on. I may get students to make a note of this rule in their books, but this is by far the least important part of the process. As we shall see in Chapter 7, rules are prone to over- and under-generalisation, and it is through the careful selection and presentation of examples and non-examples that true understanding develops.

Likewise, when students first come to answer questions on the Angle at the Centre, I no longer insist they provide reasons why as I always did in the

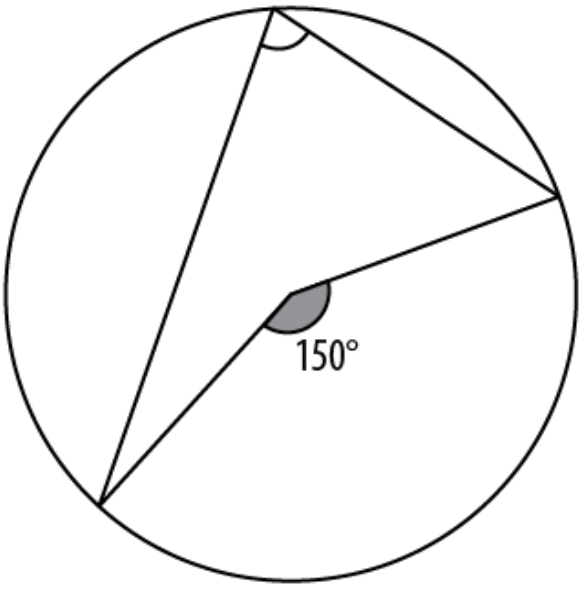
past. Again, the recalling and articulating of such rules – whilst undoubtedly important in the long-run – is not needed at this stage, and the insistence students provide such rules could be cognitively overwhelming and hinder learning. I go back to my top priority – I want my students comfortable recognising and using the rule, and tasting success. I want this knowledge transferring from working to long-term memory so it is accessible and as much of it as possible is automated. Then, and only then, will I introduce an insistence on the writing of rules.

With the relationship established, I will quickly assess understanding. This will take the form of multiple-choice diagnostic questions, with students voting with their fingers so I get a whole-class picture within seconds. This technique will be covered in depth in Chapter 11, but as a teaser, on the next page is a question I might project onto the board.

After voting and a discussion (again described in Chapter 11), we will check the answer on Geogebra.

I am now ready to move on with the rest of the lesson, which will probably involve a period of Intelligent Practice (Section 7.8), before introducing the next circle theorem in a similar way.

What is the size of the marked angle?



A 30°

B 75°

C 210°

D 300°

Figure 3.5 – Source: Craig Barton, created for Diagnostic Questions

This entire demonstration takes maybe five minutes, as opposed to the 20 or so that could be taken up handing out sheets, drawing lines, measuring angles, and hoping the right answer materialises. I can use this gained time to get the students practising and honing their skills, or ask more interesting questions such as, 'What do you think happens when the angle at the centre goes above 180°?' – something I can again illustrate using the dynamic geometry package.

I have introduced the concept in a way that I had more control over. I have incorporated non-examples, the potential for misconceptions or incomplete knowledge has been reduced, and my students have had more opportunity to practise and develop procedural fluency. At the end of this I believe they will feel successful, and as we saw in Chapter 2, success and the perception of success is a key factor in motivation.

An apparent advantage of the guided discovery approach in this particular instance is that students got to practise measuring angles – something that is lacking in my teacher-led demonstration. However, we must remember that the aim of this activity is to develop procedural fluency in the Angle at the Centre theorem. Any gaps in knowledge or misconceptions students have when it comes to measuring angles risks harming the development of this procedural fluency – if students cannot accurately measure angles, they are not likely to spot the relationship. So why take the risk? We will examine the concept of interleaving in Chapter 12, but it is worth pointing out here that interleaving of a previous topic should never risk harming understanding of a new topic. Procedural fluency of the new topic must be secure before other concepts are carefully woven in. Indeed, once I am satisfied that students can answer a sufficiently wide variety of basic questions on the Angle at the Centre, I may well set them the task of measuring angles to check the theorems, just as I am likely to include algebraic or fractional angles. But procedural fluency in the new topic must come first.

Introducing the Laws of Indices

Once again let us start with the aim of developing procedural fluency.

I want to tap into students' natural propensities to spot patterns, but in a more controlled way. The first thing to point out is that the start of the lesson would involve three multiple-choice diagnostic questions (Chapter 11) assessing students' ability to convert between fractions and decimals. Why these particular topics? Because they are fundamental to developing procedural fluency in the more complex laws of indices I want students to learn, in a way that being able to measure angles was not for circle theorems. If my students' fragile working memories are filled up trying to figure out what 0.4 is as a fraction, then the chance of them developing procedural fluency in the laws of indices is drastically reduced. The process of assessing baseline knowledge is tackled in depth in Section 11.6, but for now it is enough to say that if I discovered that students were not fluent in their ability to convert between fractions and decimals, I would abandon any notion of continuing with the laws of indices and get this sorted.

Second, I would tackle each law in isolation, get students comfortable with it, then give them time to practise on their own before moving on to the next. The batching of worked examples without sufficient student practice in between is a mistake I have been making for years, and will be covered in Chapter 6.

So, let's take the index law n^{-1}. I would have prepared a simple Excel spreadsheet. It allows me to type any number into a cell and then automatically calculates the decimal and fractional representations of the reciprocal.

	A	b	c	D	E	F	G
1	n	$(n)^{-1}$ decimal	$(n)^{-1}$ fraction		n	$(n)^{-1}$ decimal	$(n)^{-1}$ fraction
2	8	0.125	⅛		½	2	2
3	5	0.2	⅕		⅕	5	5
4	-8	-0.125	-⅛		⅜	1.666666667	1⅔
5	-5	-0.2	-⅕		⅘	1.25	1¼
6	0.8	1.25	1¼		-½	-2	-2
7	0.25	4	4		⅜	2.666666667	2⅔
8	-0.1	-10	-10		4½	0.222222222	⅔

Figure 3.6 – Source: Craig Barton, created using Microsoft Excel

Two questions may emerge from this.

1. Why don't I type these numbers into a calculator instead of Excel, as a calculator is the main instrument students use in maths? Surely it would be better to use a visualiser or a calculator emulator so students gained experience of using a calculator efficiently? This is all true, but Excel has two significant advantages for this particular demonstration. First, it is quick. I can type a number in and instantly see it displayed in two different formats. Second, all the results remain on the screen, thereby better enabling students to spot patterns.

2. Why don't I get students to do the work here by asking them to enter the examples into their calculators? Because, just like measuring angles with circle theorems, I do not want students' lack of competence in a separate skill to threaten their learning of a new concept. They probably would be okay – and of course I could assess their competency at the start of the lesson – but again why take the risk? This lesson is about them learning the laws of indices, not how to use their calculators. Competency on calculators is not necessary for learning the laws of indices in the way the ability to convert between fractions and decimals is, so I will leave it out of the early knowledge acquisition phase. Students can use their calculators to check their answers later, once the laws of indices are secure, and hence still benefit from interleaving, but in an order that makes sense to maximise the chance of them learning.

What follows is a structure very similar to the circle theorems demonstration. I may ask: 'If I type 8 into the next empty cell, write down on your mini whiteboards what you think will appear in the next 2 cells. Then, when I tell you, compare answers with your neighbour'. I will tackle positive integers first, before moving on to negative integers, positive decimals and negative decimals. Then I will move onto fractions. As we will discuss in Chapter 7, I will have chosen all the content and order of these examples in advance and with great care – they cover the full domain of content and avoid ambiguities that could lead to erroneous generalisations.

Unlike with the circle theorem demonstration, which was relatively straightforward to spot, I need some information on how my students are coping with these examples so I can get a feel of when to move on. So, I may ask them to hold up their mini whiteboards at critical points, drawing students' attention to previous examples where needed. Finally, and at a point that I sense it appropriate, I will give students an opportunity to self-explain, then discuss with their neighbours, before we once again settle upon a description of this particular index law. This will be followed up by a multiple-choice diagnostic question, a quick example-problem pair (Chapter 6), a carefully chosen group of questions for them to try (Chapter 7's Intelligent Practice), and then on to the next index law.

Summary

None of this is revolutionary, and on the surface it does not look all that different from having the students try to discover the laws themselves and then having a discussion about it. But it is different. Whilst features like me taking control over the examples and non-examples I present, insisting on periods of silence, steering the discussion, and removing unnecessary content that may hinder learning may seem minor, they make a world of difference in this crucial early knowledge acquisition phase. It is also so much quicker, leaving time for the all-important practice.

A common complaint I hear from teachers when I describe this approach is that students are not as actively involved as they would be during guided discovery. My response is that it depends on what you mean by *active*, and its anthesis, *passive*. If active students are ones making noise, working in groups, moving around the classroom, going about the task several different ways, getting some things right but plenty of things wrong, whereas passive students are sitting there quietly, thinking hard about the mathematics I am presenting, then I know which one I would prefer, especially at this early knowledge acquisition stage. For me, such 'activity' is exactly the poor proxy for learning that Coe (2013) warns us about. Students may well be active, but active doing what? What are they thinking about? What are they expanding their precious, limited working memory reserves on?

During these demonstrations, my students are active in another sense. They are actively thinking hard about the matter in hand – or at least I am

creating conditions to give them the very best chance of thinking hard about the matter in hand, and nothing else. Such activity is impossible to see, hence it is often dismissed as passivity and a lack of engagement. But periods of quiet contemplation like this are the key to learning, especially when we consider in greater depth the limits of working memory in Chapter 4. This will be even more important when we discuss the presentation of worked examples in Chapter 6.

In summary, guided discovery in the way I used to do it is not the Holy Grail I once thought it to be. There is no doubt that some students get a lot out of such an approach, and through careful instructional design this number of students can rise. Likewise, we need to keep in mind that once students reach a level of expertise in a given field, then less guidance in instruction is a good thing. But here is the key point – guided discovery, almost by definition, is used with novice learners. If they were experts, there would be nothing for them to discover. And I am very firm on my belief that in the early knowledge acquisition phase, fully-guided, explicit instruction is the way forward. It is time efficient, minimises the chance of incomplete knowledge or the development of misconceptions, and via its positive impact on student achievement it is motivating.

OceanofPDF.com

3.5. Teaching lower achieving students

What I used to think

In the two secondary schools where I have spent the majority of my career, students in maths were placed into sets by some measure of achievement (usually a combination of Key Stage 2 SATs score, GCSE target, and internal assessments throughout the year). Both of my heads of department were very keen to share classes evenly and fairly across staff, and so most years I would be given either a bottom-set or a next-to-bottom-set class in Key Stage 3.

I'll be the first to admit that I do not consider myself particularly effective at teaching *lower-achieving students* (for want of a better phrase, although all options seem either patronising or limiting). Perhaps it was because I felt I could relate more to the high-flyers, or maybe because I enjoyed the more complex areas of mathematics. Either way, I found it difficult. The main issue was that my instincts told me that I needed to teach these students differently to how I would teach a top set. The idea of explaining a concept, modelling a worked example, and then getting students to complete a series of related questions seemed, well, too formal somehow. Surely such an approach would be daunting and off-putting for the students, many of whom had endured a negative relationship with mathematics for many years? Surely it was better to adopt a more informal approach? Let them play with numbers, let them experiment, let them discover.

The problem was, they never really seemed to learn much this way.

Sources of inspiration

- Gersten, R., Beckmann, S., Clarke, B., Foegen, A., Marsh, L., Star, J. R. and Witzel, B. (2009) *Assisting students struggling with mathematics:*

Response to Intervention (RtI) for elementary and middle schools. Princeton, NJ: What Works Clearinghouse.

- Rowe, K. (2007) 'Educational effectiveness: the importance of evidence-based teaching practices for the provision of quality teaching and learning standards' in McInerney, D. M., Van Etten, S. and Dowson, M. (eds) *Standards in education*. Charlotte, NC: Information Age Publishing, pp. 59-92.

My takeaway

We have seen in the previous sections that less guidance during instruction is simply not suitable for novice learners. More often than not, novices' lack of domain-specific knowledge leads to a frustrating, demotivating experience. Given that the students in the bottom sets are novices in many areas of mathematics, it follows that a more teacher-led, explicit form of instruction would be more effective.

Indeed, in a review of relevant studies into students who struggle with mathematics, Gersten *et al* (2009) provide eight recommendations, some of which are aimed more at a senior leadership or governmental level. However, one recommendation is directly relevant to our discussion in this section:

> Recommendation 3: Instruction during the intervention should be explicit and systematic. This includes providing models of proficient problem solving, verbalization of thought processes, guided practice, corrective feedback, and frequent cumulative review.

The authors conclude that the evidence supporting this recommendation is 'strong', and in the subsequent discussion they state:

> Explicit instruction typically begins with a clear unambiguous exposition of concepts and step-by-step models of how to perform operations and reasons for the procedures. Interventionists should think aloud (make their thinking processes public) as they model each step of the process. They should not only tell students about the steps

> and procedures they are performing, but also allude to the reasoning behind them (link to the underlying mathematics).

There is no reason why the model of teaching suggested by this recommendation – and which is developed and expanded upon throughout this book, complete with carefully planned explanations, example-problem pairs, Intelligent Practice, tests of retention, formative assessment, and carefully structured desirable difficulties – cannot be used with students who struggle mathematically. Indeed, you could very well make the argument that these are the first group of students such an approach should be used with as they are the ones who would benefit from it the most.

In a chapter in *Standards of Education* summarising research into teacher effectiveness, Rowe (2007) argues:

> The problem arises when student-centered constructivist learning activities precede explicit teaching, or replaces it, with the assumption that students have adequate knowledge and skills to effectively engage with constructivist learning activities designed to generate new learning. In many instances, this assumption is not tenable, particularly for those students experiencing learning difficulties, resulting in low self-esteem, dysfunctional attitudes and motivations, disengagement, and externalizing behavior problems at school and at home.

It may feel unfair to teach students in this way – but I would argue that it is unfair not to.

What I do now

I consider myself pretty poor at teaching students who struggle mathematically, and I have no large sample of reliable data to suggest that recent changes I have made have been effective. All I can offer is the subjective observation that my students seem happier and are retaining more now that I am following the principles of teacher-led explicit instruction.

OceanofPDF.com

3.6. Story structure

What I used to think

Once upon a time, when I first started teaching, it was all about the 3-part lesson: starter, main and plenary. Within a couple of years, a new character had appeared on the scene: the 5-part lesson, complete with a few mini-plenaries thrown in for good measure. Soon, *lesson-part inflation* had gripped the teaching world, and 7- and even 11-part lessons were not unheard of.

I often wondered if there was ever any thought given as to what makes an effective structure to a lesson. But when you have five 23-part lessons to plan for tomorrow, you don't tend to have too much time to dwell upon such matters.

Sources of inspiration

- Anderson, C. (2016) *TED Talks: the official TED guide to public speaking*. Boston, MA: Houghton Mifflin Harcourt.
- Geary, D. C. (2007) *Educating the evolved mind: conceptual foundations for an evolutionary educational psychology*. Charlotte, NC: Information Age Publishing.
- Meyer, D. (2011) 'The Three Acts Of A Mathematical Story', *dy/dan blog*. Available at: http://blog.mrmeyer.com/2011/the-three-acts-of-a-mathematical-story/
- Meyer, D. (no date) *Pizza Doubler*. Available at: http://threeacts.mrmeyer.com/pizzadoubler/

- Meyer, D. (no date) *Playing Catch-up*. Available at: http://threeacts.mrmeyer.com/playingcatchup/
- Willingham, D. T. (2004) 'Ask the cognitive scientist. The privileged status of story', *American Educator* 28 (2) pp. 43-45.

My takeaway

If we are going to adopt a model of explicit instruction, we need to ensure it is as effective for students' learning as possible. One way to do this may be to harness the power of stories.

Willingham (2004) argues that planning a lesson around the structure of a story is a good idea, not simply to engage students, but to help them learn and retain the content. He argues that stories are both easier to comprehend and easier to remember than if that information is presented in a non-story structure. Indeed, in *TED Talks: The Official Guide to Public Speaking*, Anderson (2016) devotes an entire chapter – entitled 'The Irresistible Allure of Stories' – to the power of a story-driven narrative, explaining that 'many of the best (TED) talks are anchored in storytelling. Unlike challenging explanations or complex arguments, everyone can relate to stories. They typically have a simple linear structure that makes them easy to follow. You just let the speaker take you on a journey, one step at a time. Thanks to our long history around campfires, our minds are really good at tracking along'. Returning to the work of Geary (2007), storytelling has existed in human societies since the evolution of language, whereas communicating in written form is a relatively new development. Hence, humans are likely to find learning from stories easier than from other forms of instruction.

With this in mind, Willingham (2004) suggests the following key components of a story: causality, conflict, complications and character. But how easy is it to structure a maths lesson around these components without it feeling forced, phony and downright unbelievable to the students? Well, not as difficult as you might think.

Causality can be as simple as showing one concept follows directly from another. There is a danger that many students come to view maths as a series of disconnected rules and concepts that bear no relation to each other.

Interleaving topics – for example, when studying angle facts, including questions with algebra, fractions and ratio – is one way to achieve this whilst also tapping into the benefits of desirable difficulties (Chapter 12). Another way is to present maths as a journey. So, because you have learned to find factors of numbers and to expand double brackets, you can now factorise quadratics, which will means soon you will be able to sketch quadratic graphs, which will eventually lead you to coordinate geometry, calculus, and a whole host of other wonderful things.

Conflict can be as straightforward as presenting students with a problem that cannot be solved, and indeed this can be a good hook to start a lesson. The most simple version of this is a worked example – students cannot solve it with their existing mathematical knowledge, but following the methods discussed in Chapters 6 and 7, the teacher can guide them through this conflict. Later in this chapter (Section 3.8), we will meet *cognitive conflict*, where something students thought they knew no longer fits in with what they are now observing. Such a technique is fraught with danger, but in the right circumstances it can be incredibly powerful.

Complications exist when our current mathematical tools let us down, or we encounter a surprising result. In Section 2.5 we discussed the importance of providing students with a purpose. If students come to realise that their time-consuming, inefficient way of approaching a problem is not going to work, they are likely to be more receptive to the new approach you present them with. This provides both a purpose and a resolution to the complications.

Character is perhaps a little more difficult, but searching the rich history of maths to see where the great ideas started from can breathe humanity into otherwise abstract ideas. I always introduce Pythagoras with a few tales about the man himself, and I am a big fan of Dr Frost's PowerPoints (www.drfrostmaths.com), as they often contain a brief history of the origins of concepts.

Dan Meyer's Three Act Math structure (*eg* Meyer, 2011) is a wonderful way to include more than one of these components:

Act 1: Introduce the central conflict of your story/task clearly, visually, viscerally, using as few words as possible.

Act 2: The protagonist/student overcomes obstacles, looks for resources, and develops new tools.

Act 3: Resolve the conflict and set up a sequel/extension.

Playing Catch-up and *Pizza Doubler* are two of my favourite Three Act Math activities. However, as I hope was made clear in the previous two sections, I would only be using such activities when I am confident students had sufficient domain-specific knowledge. For example, I would not use *Playing Catch-up* to teach speed, distance and time, but as an effective, interesting means of applying students' existing knowledge.

Hooks can be a great way to establish conflict, complication and character, but Willingham (2004) suggests: 'don't always have the hook at the beginning of the lesson'. Sometimes a hook mid-lesson is just what is needed, especially if students are flagging a little bit. This could be as simple as suddenly presenting students with a mixed number fraction to multiply, when they have been used to proper fractions, or a boundary example (Chapter 7.3) for students to ponder.

What I do now

Not all lessons lend themselves well to the components of causality, conflict, complications and character, and indeed core content should never be sacrificed in order to shoehorn one of these in. Perhaps more importantly we should keep in mind Willingham's other piece of advice: 'students remember what they think about'. Dressing up as Pythagoras with full blown robe, wig and beard – it is a moment I wish I could erase from my memory – unsurprisingly led to my students focusing more on the possible mental derailment of their teacher than the fundamentals of the theorem.

However, these four components are certainly useful to bear in mind when planning a lesson. Subtly harnessing the power of stories can help make explicit instruction the captivating, effective model of teaching that it has the potential to be, and we can all live happily ever after.

OceanofPDF.com

3.7. Analogies

What I used to think

Given the abstract nature of many of the concepts students encounter in mathematics, often a good old-fashioned analogy is called for. And I've used plenty in my time – from bowls of soup to unwrapping presents.

A question I never really asked myself was: are these analogies good for learning?

Sources of inspiration

- Reddy, B. (2014) 'How we teach addition & subtraction of negative numbers', *Mr Reddy Maths Blog*. Available at: http://mrreddy.com/blog/2014/07/how-we-teach-addition-subtraction-of-negative-numbers/
- Richland, L. E., Zur, O. and Holyoak, K. J. (2007) 'Cognitive supports for analogies in the mathematics classroom', *Science* 316 (5828) pp. 1128-1129.
- *TES* Editorial (2017) *Listen: What every teacher needs to know about memory – professors Robert and Elizabeth Bjork talk to* TES *Podagogy*. Available at: https://www.tes.com/news/school-news/breaking-views/listen-what-every-teacher-needs-know-about-memory-professors-robert
- Willingham, D. T. (2003) 'Ask the cognitive scientist. Students remember…what they think about', *American Educator* 27 (2) pp. 37-41.

- Willingham, D. T. (2009) 'Is it true that some people just can't do math?', *American Educator* 33 (4) pp. 14-19.
- Wiliam, D. (1997) 'Relevance as MacGuffin in mathematics education', British Educational Research Association Conference, York, September 1997.

My takeaway

In addition to story structure, another way to help boost the effectiveness of explicit instruction is to utilise the power of appropriate analogies.

Thinking back to our model of how students think and learn introduced in Chapter 1, it is clear that when an analogy is successful, it enables students to connect new information, concepts and skills to existing knowledge stored in their long-term memory. Indeed, in a 2017 interview for TES, Robert Bjork argued that after a relatively early point in our lives, all new learning relies on making connections to existing knowledge.

But not all analogies are created equally.

Familiarity and Visibility

According to Willingham (2009), there are a number of things a teacher should keep in mind when considering using an analogy in the classroom. We can look at these in the context of invoking the analogy of a balance scale to teach solving equations:

- *Familiarity.* Do students actually know what a balance scale is? These days, possibly not. Indeed, how many of the common analogies used by maths teachers do students actually encounter outside of their maths lessons? Thermometers and spinners, to name but two. Recently, one of my Year 7s claimed to have never seen dice before, which broke my heart a little.
- *Vividness*. Is there a physical balance scale for students to see? If not, is there an interactive version, or a clearly drawn blank template that can be used throughout the lesson?

- *Making the alignment plain*. Writing the two sides of the equation over the two sides of a drawing of a balance scale makes crystal clear the connection between the sides of the scales and the sides of the equation.
- *Continuing to reinforce the analogy*. By referring to the balance scale at appropriate times as the equation is solved, the concepts of balance and equality are made stronger.

Richland *et al* (2007) reinforce the importance of familiarity and visibility. When observing maths lessons in the US, Hong Kong and Japan, they found familiar, visible analogies were used less by teachers in the US than by teachers in the other two (higher-performing) regions. The authors conclude that 'if the source analog is not familiar and not visible, then students may struggle with processing. First, students will need to perform a taxing memory search to understand the source. Then, assuming that memory retrieval is successful, lack of visual availability will place further burdens on working memory during production of the relational comparison'.

However, for Wiliam (1997), visibility and familiarity are not enough to make an analogy suitable for use with students. He introduces two other criteria: Match and Range.

Match
Does the task (and possible interpretations of the task) match the core mathematical activities you want to convey, and/or will students need to suppress/ignore attributes to engage in the mathematical activity you intended?

We have already learned from Willingham (2003) that students remember what they think about. Similarly, in Section 2.3 we saw the potential pitfalls of using 'real-life' situations in an attempt to appeal to students' interests. The exact same problems apply to analogies. Analogies involving social media, sport, and technology in general may suffer from the dual burden of needing to be simplified so much that they no longer resemble their original form, together with the likelihood of students bringing their own

conceptions with them that may distract from the core principle we are trying to convey.

Range
How far does the model take you along your journey to understanding a topic?

This is closely related to the concept of teaching methods that last which we discussed in Section 1.6. Think about simplifying expressions. Many maths teachers will have been warned about the dangers of attempting to explain something like 3a + 2b using the analogy of *three apples and two bananas* – 'You cannot add apples to bananas!' – as it may lead students to believe that variables can only represent objects. But in fact any attempt to explain that you cannot combine unlike 'things' soon falls apart when students encounter 3a × 2b.

Or how about rearranging equations? My favourite analogy here is present wrapping. So, when faced with the challenge of making p the subject of $r = 4p^2 + q$, I may explain that we start with the precious gift of p, it is first wrapped in a layer of 'squared', then in a layer of 'multiply by 4', and finally in a layer of 'plus q'. To get to our gift, we need to unwrap the outer layer first by doing the inverse. So we start by subtracting q, then dividing by 4, and so on until we get to the centre. This is all well and good (actually, now I'm seeing this written down, 'good' may be an overstatement), but this analogy is rendered completely useless by equations such as: $r = \frac{a}{p}$ and $pr = pq + s$.

Finally, let's consider *range* in the mathematical context that possibly contains the widest variety of analogies – negative numbers.

One familiar approach is to use a thermometer. An issue already touched upon is: do students actually know what a thermometer is? Why on earth would they use something so antiquated when they can just look up the temperature on their smart phones? But aside from that, how well does the thermometer analogy hold up when it comes to key concepts related to negative numbers?

- Ordering negative numbers – no problem! The temperature was -40C in Blackburn and -20C in Preston. Which was colder and by how much?
- Adding or subtracting positive integers from negative numbers – yes! The temperature in Bolton is -30C. It falls by 50C. What is the new temperature?
- Answering questions like this: -3 + -2, or 4 – -2? Not so good.

For a long time, my go-to negative number analogy has been soup. Students are presented with a picture of a bowl of soup, and told that they can either put in or take out chillies or ice cubes (see Figure 3.7). Each chilli adds one to the temperature, and each ice cube reduces the temperature by one. This means that when confronted by a reasonably difficult question such as -5 – -2, students can say something like: *the soup was -5*0C*, then I took out 2 ice cubes, so the temperature must have gone up 2*0C*, so the answer is -3*.

Apart from the questionable chemistry involved in claiming the addition of a chilli and the removal of an ice cube has the same effect on temperature, there is a bigger problem. This analogy works for all addition and subtraction of negative number questions, but how about multiplication and division? Alas, not so great.

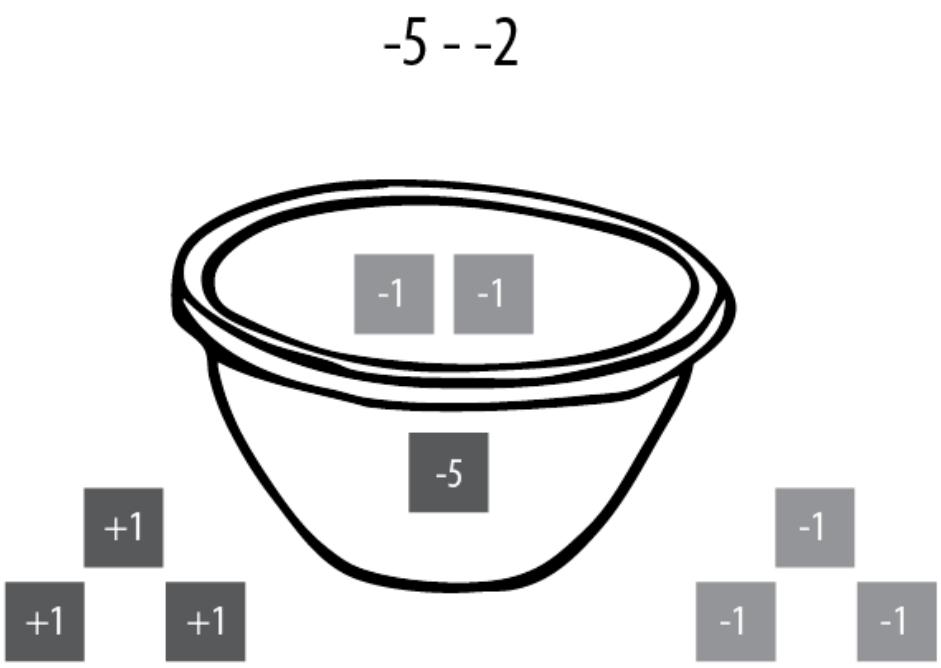


Figure 3.7 – Source: Craig Barton

I have found similar issues with related negative number analogies including sandcastles, debts and even trains going into and out of a tunnel. And that is a problem, because it means you need a new analogy for each concept, and then students may have difficulty remembering which one goes with which.

Perhaps for this reason, Bruno Reddy adopts a non-analogy approach. As he explains in a 2014 blog post:

> No analogies – When teaching addition and subtraction, we NEVER talk about 'two negatives make a positive' or use analogies about ice cubes, good/bad people, or use negative/positive tiles.
>
> Having reinforced that adding goes 'up' and subtracting goes 'down', we look at the effect of adding a negative and subtracting a negative. We explicitly teach them, without an analogy, that when adding a negative you go down and when subtracting a negative you go up. Again, lots of whole class practice with the air number line. It's slow and deliberate to start with but becomes high energy and high stakes as they get more proficient. By high energy, I mean doing lots of questions quickly as a

class, call-and-response style and by high stakes I mean playing Simon Says.

We will discuss the introduction of negative numbers more when we look at the choice of examples in Chapter 7.

What I do now

I am much more aware of the power of analogies as a tool for explicit instruction. When they work well, they enable students to connect new knowledge with existing knowledge. However, not all analogies are as effective as each other. When considering whether or not to use one, I ask myself:

- Are the students familiar with it?
- Can they see it?
- Will it lead to any unnecessary confusion?
- Does it have range to cover a sufficient amount of the content?

If not, then I will introduce the concept without an analogy, focusing on the development of procedural fluency and conceptual understanding through my choice of examples and practice questions. A bad analogy is worse than no analogy at all.

OceanofPDF.com

3.8. Cognitive conflict

What I used to think

Hands-up if the following rings any bells:

Me: What is

$\frac{1}{3} + \frac{1}{5}$?

Year 11 (complete with 11 years of schooling, at least 6 of which have covered the concept of adding fractions, 3 of which have been taught be me):

$\frac{2}{8}$

After realising that they are not joking, I die a little inside.

After composing myself, I do what I always do. I politely explain that they are not quite right, and then I present them with a beautiful explanation of how to add fractions correctly, offering both a pictorial representation and a clear presentation of the algorithm, followed by lots of examples and opportunities to practise, all the time convincing myself that this time they will get it.

The problem is, I have been saying and doing the same thing since I started teaching this class in Year 9.

Sources of inspiration

- Anderson, J. R. (1996) 'ACT: a simple theory of complex cognition', *American Psychologist* 51 (4) pp. 355-365.
- Barton, C. (2017a) 'Nick Rose', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Cook, J. and Lewandowsky, S. (2011) *The debunking handbook*. St Lucia, Australia: University of Queensland.
- Geary, D. C. (2007) *Educating the evolved mind: conceptual foundations for an evolutionary educational psychology*. Charlotte, NC: Information Age Publishing.
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.
- Limón, M. (2001) 'On the cognitive conflict as an instructional strategy for conceptual change: a critical appraisal', *Learning and Instruction* 11 (4) pp. 357-380.
- Lin, J. and Geary, D. C. (1998) 'Numerical cognition: age-related differences in the speed of executing biologically primary and biologically secondary processes', *Experimental Aging Research* 24 (2) pp. 101-137.
- Lucariello, J., and Naff, D. (2013) *How do I get my students over their alternative conceptions (misconceptions) for learning?*. American Psychological Association. Available at: http://www.apa.org/education/k12/misconceptions.aspx
- Mayer, D. (no date) *Popcorn Picker*. Available at: http://threeacts.mrmeyer.com/popcornpicker/
- Siegler, R. S. (2002) 'Microgenetic studies of self-explanation' in Granott, N. and Parziale, J. (eds) *Microdevelopment: transition processes*

in development and learning. Cambridge: Cambridge University Press, pp. 31-58.

- Steward, D. (2012) 'Subtraction misconception', *Median Maths Blog*. Available at: http://donsteward.blogspot.co.uk/2012/06/subtraction-misconception.html
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.
- Zohar, A. and Aharon-Kravetsky, S. (2005) 'Exploring the effects of cognitive conflict and direct teaching for students of different academic levels', *Journal of Research in Science Teaching* 42 (7) pp. 829-855.

My takeaway

As we have seen, Kirschner *et al* (2006) define learning as 'a change in long-term memory'. The problem is that students do not arrive in front of us with empty compartments within their long-term memories which are ready to be filled with lots of lovely correct knowledge. Instead, students arrive with cognitive baggage. They are likely to have erroneous beliefs and incomplete knowledge about all manner of concepts. These erroneous beliefs and incomplete knowledge form *misconceptions*.

Now, I am more than a little obsessed about misconceptions in mathematics. Most weekends I spend a good couple of hours prowling my Diagnostic Questions website to uncover yet another student misconception that I could not have predicted. We will discuss identifying misconceptions in Chapter 11, but here it is important to establish exactly what we mean by the term. Perhaps the most useful way to think of misconceptions is in contrast to something else: mistakes.

Mistakes tend to be one-off events – the pupil understands the concept or the algorithm, but may make a computational error due to carelessness or, as we will see in Section 4.2, cognitive overload. Give the student the same question again and they are unlikely to make the same mistake. Inform the student that they have made a mistake somewhere in their work, and they

are likely to be able to find it. Consequently, mistakes are arguably less serious than misconceptions, but also frustratingly less predictable.

In the addition of fractions example above, a mistake would be if the student attempted to express the fractions with the same denominator, but made an arithmetic error when calculating one of the numerators.

Misconceptions are a whole other kettle of fish. Misconceptions are the result of erroneous beliefs or incomplete knowledge. The same misconception is likely to occur time and time again. Informing the student who has made an error due to a misconception is likely to be a waste of time as, by definition, they do not know they are wrong. The student in my example above has not made a mistake when it comes to adding fractions, they have a misconception about adding fractions. Misconceptions are undoubtedly more serious than mistakes – not just because of the frequency of which they occur, but because they can significantly interfere with the acquisition of new knowledge. For example, the student who has misconceptions when it comes to adding fractions will likely suffer with algebraic fractions, equations with fractions, and so on. On the bright side, I will argue in Chapter 11 that these misconceptions can be predicted.

It is my belief that one of the greatest challenges facing a teacher, no matter what approach to teaching they choose, is how to identify, understand and resolve the misconceptions that our students hold. Indeed, when I interviewed Professors Robert and Elizabeth Bjork for my podcast, they made the point that misconceptions never disappear from long-term memory, and the best we can hope to achieve is to make the correct knowledge more accessible. Nick Rose made a similar point in our interview, explaining that misconceptions are likely to lie close to the surface of our consciousness, fighting for our attention, and it can be rather effortful to avoid them in order to select the correct piece of knowledge.

Pinpointing exactly where these misconceptions come from is very difficult. Some may have their origins in our evolution. We have already seen when looking at the work of Lin and Geary (1998) that the vast majority of the mathematics that students meet at school is biologically secondary knowledge, and the human brain has simply not had enough time to evolve

in a way to make this knowledge easy to acquire. Hence, if our primary instincts interfere with this acquisition, then incomplete knowledge, erroneous beliefs, and therefore misconceptions, may be the result. Wiliam (2011) suggests that misconceptions may arise from over-generalisation. For example, students are unlikely to have ever been taught that $2.3 \times 10 = 2.30$, but have instead (quite logically) inferred that result having observed that $23 \times 10 = 230$. Likewise, why would you not add the numerators and denominators when adding two fractions? It would certainly make life a lot easier.

Misconceptions may remain prominent and accessible in long-term memory as a result of repeatedly practising the wrong thing without corrective feedback. The model of thinking proposed by Anderson (1996), discussed in Section 1.1, proposes that after a certain amount of practice, the path of production becomes stable and procedural learning has occurred, and that once this path has been created it becomes difficult to 'rewire' it. This is why it seems difficult to unlearn something and learn new conflicting information or skills. This is all great, assuming of course that the knowledge we are learning is correct, but erroneous knowledge is just as hard to unlearn.

Hence, many of the misconceptions that our students hold are things that feel intuitively right, and have been allowed to linger. Examples can be found in every single area of mathematics, but I will focus on the following four in this section:

- Addition of fractions – adding the numerators and denominators
- Probability of combined events – when rolling two dice, all scores are equally likely
- Written subtraction – take the smallest number from the biggest number
- Volume – tall thin shapes have the same volume as short wide ones

So, how do we resolve such misconceptions?

If learning is indeed a change in long-term memory, then we require something rather significant to happen. Successfully transferring knowledge from working memory to long-term memory is tricky enough, but in these cases we also need this new knowledge to battle against existing erroneous knowledge stored in long-term memory and come out victorious. If this existing knowledge is deeply embedded, then this battle may be an incredibly tough one, and simply telling students the right way of doing things as I suggested with adding fractions may not be enough.

One strategy to instigate the necessary change in long-term memory is to induce *cognitive conflict*, but it comes with a big fat word of warning.

Cognitive conflict is based around the premise that students are more likely to both accept new information and make a conscious effort to adjust their existing beliefs and understanding if they are first shown clearly why their existing beliefs and understanding are incorrect.

This is related to Siegler's (2002) Overlapping Waves theory of learning. In short, the theory states that individuals know and use a variety of strategies which compete with each other for use in any given situation. With improved or increased knowledge, good strategies gradually replace ineffective ones. However, for more efficient change to occur, learners must reject their ineffective strategies, which can only happen if they understand both that the knowledge is wrong and why it is wrong.

Let's see cognitive conflict in action in terms of the addition of fractions misconception.

So, I may start the lesson with two diagnostic multiple-choice questions (Chapter 11). The first would be aimed at establishing whether my students are aware that $\frac{1}{3}$ is bigger than $\frac{1}{4}$, and the second to check that students know $\frac{2}{8}$ simplifies to $\frac{1}{4}$. With this assessed, I would then present the following:

$$\frac{1}{3} + \frac{1}{5} = \frac{2}{8}$$

$$\frac{2}{8} = \frac{1}{4}$$

I would then pause, and ask, 'Does anyone see a problem with this?'.

My hope is that students realise that they have started with something relatively large (⅓), added something positive on to it (1/5) and ended up with something smaller (¼). This could be presented in a more concrete form with visual images of fractions of shapes instead of just numbers. Hence, students would see that their method for adding fractions could not be correct, thus cognitive conflict would be induced, which would make them more willing to accept the correct method when I subsequently present it. This should be more likely to result in a change in their long-term memory than me simply presenting them with the correct way to add fractions without such conflict.

This approach certainly makes logical sense, and has been supported by many studies. For example, Lucariello and Naff (2013) found evidence in support of cognitive conflict, and recommend two strategies to induce it:

1. *Present students with anomalous data*. Just like we saw in the addition of fractions example above, this is data that does not accord with their misconception.
2. *Present students with refutational texts*. A refutational text introduces a common misconception, refutes it, and offers a new (alternative) theory that proves to be more satisfactory. This suggests that confronting such misconceptions head-on in the notes we give students in their books might even be a good thing. So, have them copy down the incorrect method for adding fractions, explain why it is wrong, and then copy down the correct method.

However, not all studies support the inducement of cognitive conflict as a tool for overcoming misconceptions, with mixed or even negative results on

learning. For example, Zohar and Kravetsky (2003) sought to compare the effectiveness of two teaching methods – Inducing a Cognitive Conflict (ICC) versus Direct Teaching (DT) – for students of two academic levels (low versus high). The key finding was that the ICC teaching method was more effective for high-level students following a test of retention 5 months later, while the DT method was more effective for low-level students. Moreover, inducing cognitive conflict actually delayed the progress of students in the low academic group.

Why would this be the case? Well two common reasons given for the apparent failure to demonstrate the effectiveness of cognitive conflict in the classroom are:

1. Students often fail to reach a stage of meaningful conflict as it requires a certain degree of both prior knowledge and reasoning abilities.
Limon (2001) states: 'to start the process of change, this conflict has to be meaningful for the individual'. We can see how this would be a barrier to learning if we consider the fractions example. In order to experience cognitive conflict, students must understand the significance of the incorrect demonstration. Hence, they must first understand that $\frac{1}{3}$ is bigger than $\frac{1}{4}$, then that $\frac{2}{8}$ and $\frac{1}{4}$ are equivalent, and finally that when you add two positive numbers together – indeed, not just numbers, but *fractions* – the result must be bigger than the number you started with. Understanding these three elements and then piecing them together into a coherent argument that refutes their previously held belief is likely to be cognitively challenging for students with low prior knowledge levels. Working memory is overloaded, or students cannot recognise the importance of the disconfirming evidence being presented as all their effort is taken up understanding the individual elements.

2. Students may not have an appropriate degree of motivation (goals, values and self-efficacy) that are potential mediators in the process of conceptual change.

Of course, it is one thing understanding an argument, but it is another acting on it. I have wheeled out my favourite Willingham quote many times already in this book, but here it comes again – memory is the residue of thought, so students remember what they think about. If students are not thinking about (or attending to) the argument, then no cognitive conflict can occur. Motivation is likely to play a significant role in this, and, as we have seen in Chapter 2, that is also intimately tied to knowledge. Zohar and Kravetsky (2003) speculate that students with low academic aptitudes and achievements tend to have a lower degree of prior knowledge, less advanced reasoning abilities and a lower degree of motivation than students with high academic aptitudes and achievements.

Taken together, these arguments suggest that inducing cognitive conflict is likely to be less effective for students with lower levels of domain-specific knowledge.

We may even go further and suggest the inducement of cognitive conflict may hinder the learning of novice learners by exposing them to incorrect methods that may offer support to existing misconceptions, or even create new ones. The student who does not fully understand a procedure is arguably just as likely to remember the wrong method as the right one. Indeed, Cook and Lewandowsky (2011) report a 'familiarity backfire effect' in which students' belief in the misconception strengthens after exposure to it. Students remember what they heard first (the misconception) and attention wanders for what follows.

What I do now

Sometimes good quality examples and explanations are not enough, and students keep making the same mistakes again and again. Inducing cognitive conflict, so that students see why they are wrong, seems like the perfect solution. However, for it to be successful, students need the ability to understand the conflict and the will to change their existing beliefs. Both of these requirements are intimately linked to domain-specific knowledge.

I am reluctant to induce cognitive conflict when I first introduce a topic. At this point of learning, everyone is a novice. Misconceptions surrounding key baseline knowledge should be identified and acted upon using the

techniques of Chapter 11, but once students are at the point where I am teaching them a new skill or concept, I will teach them the correct way to do it. As I will discuss in more detail when we cover worked examples in Chapter 6, I am also careful how much whole-class student involvement there is at this crucial stage of learning. I want the things students hear and see to be as clear, unambiguous and correct as possible, and I have come to realise that that is often more likely to happen if the explanation comes from me. If misconceptions later become apparent via formative assessment strategies (Chapter 11), I may attempt to reteach the correct method again, using a different analogy or choice of examples (Chapter 7) where appropriate. If the misconception persists, that is when I may turn towards cognitive conflict. Cognitive conflict is a tool of last resort for tackling misconceptions that have developed over time when they become apparent.

Returning to the other three examples of misconceptions listed above, I have used the following techniques to help overcome students' faulty thinking when teaching the correct method alone has failed:

Probability of combined events – when rolling two dice, all scores are equally likely

Probability is a minefield of counterintuitive results that feed on students' misconceptions. I have always been a fan of Probability Bingo (there is a horse race version that works exactly the same way), both as a means to introduce sample spaces and to address the common misconception that all events are equally likely. You will no doubt be familiar with how it works, but just in case: students are told that two dice will be rolled, the scores added together, and the total called out. This will be done 36 times. Their job is to choose the 36 numbers that they think will be called out, crossing off one occurrence of their number each time. When they see the completed grid, they are confronted with the unshakable truth that not all outcomes are equally likely (this is especially apparent to those students who were – and presumably still are – waiting for a total of 1 to be called). I often take a screenshot of the final result and return to this after we have covered experimental probability as a way of tackling the additional misconception that experimental outcomes match theoretical outcomes. As an aside, one of my favourite questions to ask Year 13 maths students attending university

interviews is: *given that 7 is the most likely outcome on any one roll, and that each roll is independent, why is the Bingo grid with the highest chance of winning not one comprised of thirty-six 7s?*

Written subtraction – just take the smallest number from the biggest number
Not for the last time in this book, I will bow to the genius of Don Steward. Imagine you are a child who believes that when performing the written subtracting algorithm you should always subtract the smallest number from the biggest. Would the following two examples be enough to induce cognitive conflict?

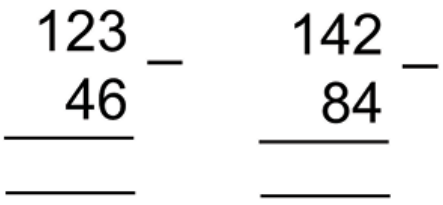


Figure 3.8 – Source: Don Steward, available at http://donsteward.blogspot.co.uk/2012/06/subtraction-misconception.html

Volume – tall thin shapes have the same volume as short wide ones
Finally, we have the king of the visual, Dan Meyer, who has produced a wonderful video to accompany his *Popcorn Picker* Three Act Math task (available at http://threeacts.mrmeyer.com/popcornpicker/). In it, Dan takes two identical sheets of A4 paper, folding one into an open-ended cylinder one way (tall and thin), and the other into an open-ended cylinder the other way (short and wide). Resting them on the table, he then begins to pour popcorn into each. Surely they will each hold the same amount of popcorn?

I think all three of these require relatively little background knowledge to appreciate the conflict, and yet can be profound in the effect they have on students. But this all comes down to the domain-specific knowledge of the students in question, and that can only be determined by formative assessment, testing, and a teacher's intimate knowledge of their students.

OceanofPDF.com

3.9. How before Why

What I used to think

For me, the following was a no-brainer: stuents should be taught why they were doing something before they were shown how to do it. Otherwise, surely maths just becomes a collection of meaningless rules? Few things in the teaching world were more obvious than that. Hence, I would always teach the Why before teaching the How.

The problem was, often the Why was rather complicated, and by the time we got to the How, my students were not exactly flying high with confidence and ready to embrace this lovely new concept.

Sources of inspiration

- Barton, C. (2017) 'Dani Quinn – Part 1', *Mr Barton Maths Podcast*.
- Boulton, K. (2017) 'Maths: conceptual understanding first, or procedural fluency?', *To the real maths blog*. Available at: https://tothereal.wordpress.com/2017/06/11/maths-conceptual-understanding-first-or-procedural-fluency/
- Donovan, M. S., Bransford, J. D. and Pellegrino, J. W. (eds) (1999) *How people learn: bridging research and practice*. Washington, DC: National Academy Press.
- NCTM (no date) *Procedural Fluency in Mathematics*. Available at: http://www.nctm.org/Standards-and-Positions/Position-Statements/Procedural-Fluency-in-Mathematics/
- Rittle-Johnson, B. and Alibali, M. W. (1999) 'Conceptual and procedural knowledge of mathematics: Does one lead to the other?',

Journal of Educational Psychology 91 (1) pp. 175-189.

- Rittle-Johnson, B. (2016) 'Ask a scientist: what should preschool math look like?', *Education Week*. Available at: http://blogs.edweek.org/edweek/early_years/2016/02/ask_a_scientist_what_should_preschool_math_look_like.html
- Rittle-Johnson, B., Schneider, M., and Star, J. R. (2015) 'Not a one-way street: bidirectional relations between procedural and conceptual knowledge of mathematics', *Educational Psychology Review* 27 (4) pp. 587-597.
- Skemp, R. R. (1976) 'Relational understanding and instrumental understanding', *Mathematics Teaching* 77 (1) pp. 20-26.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.
- Zukav, G. (2012) *The dancing Wu Li masters: an overview of the new physics*. New York, NY: Random House.

My takeaway

It was speaking to Dani Quinn on my podcast and reading Kris Boulton's 2017 blog post that really got me thinking about this issue, and it is by no means straightforward.

Before we dive in too deeply, we need to define exactly what we mean by How and Why.

'How' is relatively straightforward. When I say I want my students to know how to do something, what I really mean is that I want them to develop **procedural fluency**. The National Council of Teachers of Mathematics (NCTM) defines procedural fluency as 'the ability to apply procedures accurately, efficiently, and flexibly; to transfer procedures to different problems and contexts; to build or modify procedures from other procedures; and to recognize when one strategy or procedure is more appropriate to apply than another'.

‘Why’ is a bit more troublesome.

You could argue that in order to truly understand why something is the case, you need to prove it. But this immediately leads us into a whole world of trouble with much of the maths that students study at school. For example, in the UK, students need to be able to use and apply the formula for the volume of a sphere by the time they sit their GCSEs, and yet the proof of where that formula comes from requires calculus and an application of the principle of volume of revolution, which is usually covered towards the end of an A level course two years later. Or how about proving (not demonstrating) that the angles in a triangle equal 180° or why the process for dividing fractions works?

So, if ‘Why’ is not proof, then what exactly is it?

This is where the concept of **conceptual understanding** comes into play. Donovan *et al* (1999) define this as the student’s ability to connect new mathematics ideas with ideas she/he already knows, to represent the mathematical situation in different ways, and to determine similarities/differences between these representations. So, whilst this may not necessarily involve proof, the emphasis is on understanding mathematical processes, and seeing topics and concepts not in isolation but as part of the wider subject.

Both procedural fluency and conceptual understanding are clearly desirable, and arguably completely useless without each other. But what is not so obvious is which we as teachers should try to help our students develop first.

The NCTM is very clear on this: ‘Procedural fluency builds on a foundation of conceptual understanding, strategic reasoning, and problem solving … Therefore, the development of students’ conceptual understanding of procedures should precede and coincide with instruction on procedures’.

I am no longer convinced. I think such conclusions come from the perception of explicit instruction as a passive form of lecturing about meaningless rules that students are expected to memorise and not question, combined with a rather naive appreciation of inflexible knowledge and the importance of motivation.

Willingham (2002) distinguishes between flexible, inflexible and rote knowledge. Knowledge is flexible when it can be accessed out of the context in which it was learned and applied in new contexts. Inflexible knowledge is meaningful, but narrow – it is narrow in that it is tied to the concept's surface structure, and the deep structure of the concept is not easily accessed. Rote knowledge, on the other hand, is devoid of meaning. Willingham's point is that most of what we as teachers consider to be rote knowledge may in fact be inflexible knowledge, and the development of inflexible knowledge is a necessary step along the path to expertise. To illustrate this point, Willingham gives the following example: 'your knowledge of calculating the area of rectangles may have once been relatively inflexible; you knew a limited number of situations in which the formula was applicable, and your understanding of why the formula worked was not all that clear. But with increasing experience, you were able to apply this knowledge more flexibly and you better understood what lay behind it'.

If we insist that students have a conceptual understanding of every concept and process we teach them, we would never get anywhere. Are we really saying, for example, that unless students can explain exactly why multiplying the length and width of a rectangle gives us the area, they are not allowed to use that method? Appreciation of why a method or process works often only comes after using it many times. John von Neumann said: 'In mathematics you don't understand things. You just get used to them'. I would modify this to the (slightly less catchy, but possibly more accurate): in mathematics you don't understand things straight away. You just get used to them, and then you understand them.

Likewise, a more thorough understanding of the process to find the area of a rectangle, including its limitations, only becomes apparent through exposure to carefully chosen examples and non-examples by the teacher (see Chapter 7), such as including all four dimensions or presenting the width and length in different units. In this way the development of procedural fluency aids conceptual understanding; but presented the other way round, the students could end up developing neither.

This point is summed up nicely by Rittle-Johnson (2016), a developmental psychologist who has extensively researched the relationship between

procedures and concepts in math learning, who said:

> Actually, I think it's a silly argument because the evidence is pretty clear that children really need to do both things. Understanding is super-important, but understanding relies on knowing enough that you can understand it. If you have to spend all your time figuring out what two plus three is, then you can't notice relationships between number pairs.

Then there is the issue of motivation. We have seen in Chapter 2 that a key determinant of motivation is students' belief that they are or can be successful. When attempting to develop conceptual understanding ahead of procedural fluency, I used to think to myself: *let's try to show them why something works, and if they get it, great; but if they don't, then no damage done*. So, I would muddle through a proof of Pythagoras's theorem, or an explanation of exactly why the process for dividing fractions works, all to a sea of rather confused faces. Ten minutes later, when I came to actually explain how to do the process – which was much simpler than explaining why it worked – the damage most certainly had been done. Students had decided that the topic was difficult and confusing, and many had resigned themselves to never understanding it. Indeed, Skemp (1976) – a proponent of teaching what he calls relational understanding ahead of procedures – himself concedes that doing it the other way around has its benefits: 'the rewards are more immediate, and more apparent. It is nice to get a page of right answers, and we must not underrate the importance of the feeling of success which pupils get from this'.

So, are we to conclude that we can just teach procedures with no reference to conceptual understanding at all?

Definitely not. Not only is the development of conceptual understanding in our students one of the ultimate goals of teaching, it is also necessary in order to learn new concepts and processes. But conceptual understanding of the new concept or process *itself* does not necessarily need to come first. In a review of empirical evidence for mathematics learning, Rittle-Johnson *et al* (2015) conclude that 'procedural knowledge supports conceptual knowledge, as well as vice versa, and thus that the relations between the two

types of knowledge are bidirectional'. Conceptual understanding can (and often *should*) be developed alongside or later after procedural fluency has been achieved.

Let me try to explain what I mean by looking at a few topics.

Pythagoras's Theorem
Before going anywhere near the actual theorem, students need conceptual understanding of the following:

- Right-angled triangles
- Hypotenuse
- What it means to square a number and find the square root of a number
- The components of an algebraic formula

Students also need procedural fluency of the following:

- Identifying the hypotenuse on a right-angled triangle
- Substituting numbers into a formula
- Squaring and square rooting numbers, both mentally and on a calculator
- Rounding numbers to a given degree of accuracy

I would teach these procedures explicitly, using the techniques described in the remainder of this book. I would then assess their understanding and deal with any misconceptions using the principles described in Chapter 11.

But do students need to know *why* Pythagoras's theorem can be used to calculate the length of a missing side of a right-angled triangle? I don't think so. Many of the proofs (the webpage www.cut-the-knot.org/pythagoras/ lists 117 of them) are far more complicated than the formula itself, relying on concepts such as similarity and ratio, and could always be introduced much later to benefit from interleaving (Chapter 12). For the reasons given in

Section 3.4, I would be reluctant to embark upon a guided discovery activity involving cutting out shapes or measuring lengths. Instead, I would present a demonstration of the theorem, using something like the interactive Geogebra file on the next page (Figure 3.9).

I am using technology to help students have a visual representation of the theorem. I am making use of Chapter 4's Modality Effect and Silent Teacher, and I can change things immediately based on their questions or the examples we work through. But I am not trying to convince my students this is why Pythagoras's theorem works.

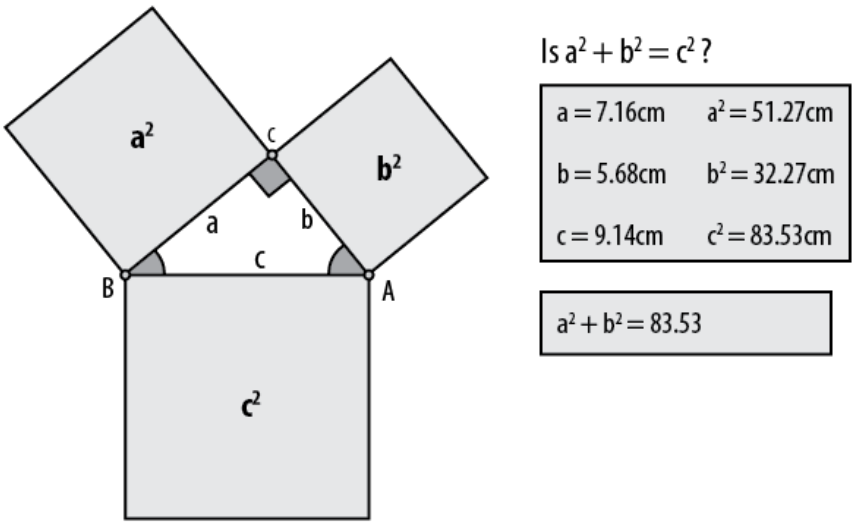


Figure 3.9 – Source: Craig Barton, created using Geogebra

Following this, I would proceed through a series of carefully chosen example-problem pairs, making use of non-examples and Supercharged Worked Examples (Chapters 6 and 7). This would be followed by plenty of Intelligent Practice of the basics (Chapter 7), culminating in the three-stage process to enable students to solve complex problems that is described in Chapter 9.

At the end of this process, do my students know why Pythagoras's theorem works? No. Do they understand it? Well, that depends what you mean by *understand*. They understand that the theorem describes the relationship between the length of sides in right-angled triangles, and they are able to apply this knowledge to solve a wide variety of problems across a number of

contexts. They know when to use the theorem and when not to use it. They have had success in the topic, and hence feel motivated. Just because they cannot prove it (yet, anyway), it is by no means a meaningless formula to them. They have procedural fluency in Pythagoras's theorem, and that is good enough for me, for now.

Multiplying Fractions
Conceptual understanding is needed of the following:

- Multiplication
- Fractions – including their different representations

Students also need procedural fluency of the following:

- Mentally multiplying two numbers together

Assuming that is sorted, would I now go on to develop conceptual understanding of the algorithm to multiply two fractions together? No I wouldn't. Just like with Pythagoras, I may offer an alternative representation of the algorithm, such as that below.

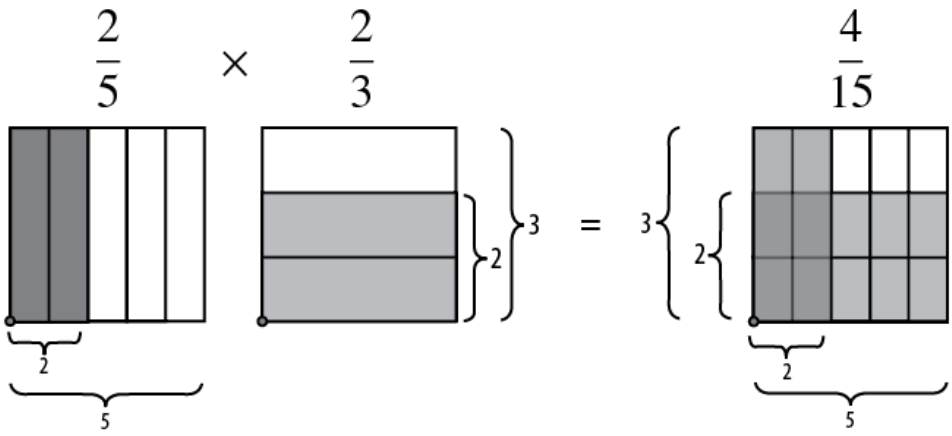


Figure 3.10 – Source: Area Model for Fractions, created by Celia Jiminez on Geogebra,

But I would not want my students to get too attached to this representation. After all, it certainly does not explain why the arithmetic process for multiplying fractions works. As Kris Boulton points out in his 2017 blog post 'Maths: conceptual understanding first, or procedural fluency?', it is just

an alternative means of conceptualisation. In addition, it is arguably a more complicated process than simply multiplying the numerators and denominators together, as it requires an understanding of area. It can also be painfully slow, as those diagrams take a while to draw. Furthermore, it is rather limited. Can it cope with the multiplication of mixed number fractions? Possibly, but it is a bit of a pain. How about when algebra is involved?

No, very quickly I would move students onto the standard algorithm for multiplying fractions. Through a careful choice of examples – involving mixed numbers, negatives, algebra *etc* – I can develop students' procedural fluency, all whilst building up their conceptual understanding.

Solving Equations

When learning to solve equations, conceptual understanding *must* come before procedural fluency. I am absolutely certain of that. Why? Well, look again at the definition of procedural fluency. To be able to apply the procedure for solving equations 'accurately, efficiently, and flexibly', students really need to know what they are doing to a far greater extent than when applying Pythagoras's theorem or multiplying fractions. In Section 1.4 I argued that a method of balancing both sides of the equation was a durable and mathematically valid way to teach solving linear equations. Without a solid grounding into why this method works, there is a danger that the examples we present to students will develop neither their conceptual understanding nor procedural fluency.

Let's imagine we present our students with the following example before we attempt to explain why the balance method works:

$2x - 7 = 9$

An attentive student may speculate that the order of operations determines the first step of the solution, combined with some application of inverse operations.

Now we go through this example:

$9 - 2x = 3$

This time the student notices that the order of operations does not justify that first step, and so concludes that something else must be going on.

Then we hit them with:

$5x - 3 = 2x + 12$

Now the student is really confused. In the first example we dealt with a $2x$ by dividing by 2, now we are dealing with a $2x$ by subtracting $2x$.

The method being used is mathematically valid and durable. However, without a careful explanation of why it works, students are unlikely to develop the procedural fluency that can in fact aid the later development of conceptual understanding. We can use a balancing scale and the power of analogies explained in Section 3.7. We can use numbers, and experiment using legal and illegal moves and discuss why some work and others don't. Once students have a solid conceptual understanding of the principles of balancing an equation, we can then proceed with the careful presentation of examples (Chapter 6), and an equally careful selection of practice questions (Chapter 7).

What I do now

It is obviously very important that students know why they are doing something, and hence it sometimes makes perfect sense to develop conceptual understanding before procedural fluency. Diving straight into a procedure with no context, justification or – and this is by far the most dangerous – assessment of understanding and fluency of prerequisite knowledge, is unlikely to be successful. Mathematics becomes a mere assortment of unconnected, confusing, meaningless rules.

However, when done badly, trying to develop conceptual understanding first can confuse students, be demotivating, and ultimately inhibit the development of both conceptual understanding and procedural fluency. Likewise, as discussed by Rittle-Johnson *et al* (2015), the development of procedural fluency (via the careful choice and presentation of examples described in Chapters 6 and 7) can aid the development of conceptual understanding.

So, how do you decide the best order?

Well, when assessing whether or not I should teach the How before the Why, I need the following criteria to be met:

1. The students lack the knowledge to understand the Why at the stage it is being taught.
2. Not knowing the Why does not inhibit their ability to do the How.
3. The How is a mathematically sound method, not a trick that has no mathematical validity.
4. The How needs to be a durable method that can be built upon.

It is by no means a perfect system, but it has meant the days of me feeling obligated to embark upon some painfully complicated explanation, before teaching something relatively straightforward are behind me.

OceanofPDF.com

3.10. Ending on a high

What I used to think

If a lesson was going well, and I had a sense that my students were understanding the core content of what I was teaching them, I would tend to end the lesson with a bit of a stinker. This could be a past exam question, a question in context, or a question with a particularly nasty twist. This would be the most difficult question I had asked all lesson, and its purpose would be to push my students to the limit – to provide evidence of the progress they had made. I would usually build it up, with something along the lines of: 'Okay, this is as tough as it gets. Can you do it?'.

More often than not, the result was that some students could and some could not, and I didn't really have time to go through the answer in great depth. But I was okay with that, as this was an extension question – I didn't expect all my students to get it right. The most important thing for me was what had happened in the previous 45 minutes.

Thinking back to that lesson later that evening over a cup of Mellow Birds, I was very happy how it went. The question I never considered was if my students were also feeling the same way.

Sources of inspiration

- Kahneman, D., Fredrickson, B. L., Schreiber, C. A. and Redelmeier, D. A. (1993) 'When more pain is preferred to less: adding a better end', *Psychological Science* 4 (6) pp. 401-405.

My takeaway

Kahneman *et al* (1993) discuss how people's memory of an experience is often dominated by the feelings of pain and discomfort during the final

moments, as opposed to what happened during the rest of the experience. They describe an experiment involving people placing their hands into ice cold water for different durations. The authors found that when given a choice between two equally painful experiences, with the second lasting longer but with a less painful ending (the temperature in the water is increased a little), participants preferred the latter, despite having to endure more pain overall. It was the memory of those final painful moments that lingered in their minds.

Now, clearly there are huge issues transferring these findings to the context of a maths lesson – indeed, the icy water could be more or less painful than a twisty exam question on 3D Pythagoras, depending on your view of mathematics – but it certainly has given me food for thought. What was my students' impression of that lesson, especially the students who could not do the final, challenging question? What will they remember? Will it be the 45 minutes of success they enjoyed at the start, or the 5 minutes of 'failure' at the end? According to this research, it may well be the latter.

Many of my students would have judged the lesson as unsuccessful because of that final, tricky problem. This negative emotion could potentially be swirling around in their heads for some time, maybe until their next maths lesson. In Chapter 2 we have seen how important students' perception of their ability to succeed in maths is for their motivation, and feelings like this can only do harm.

What I do now

The difficult questions now come in the middle of my lesson when I judge students to be ready for them. This gives me adequate time to work through them carefully, identifying and addressing specific issues that students had.

I always end my lesson with a question that is of mid-range difficulty, or maybe even easier. More often than not, it will be a diagnostic multiple-choice question (Chapter 11). The majority of the learning has happened in the first 45 minutes. My objectives in those last 5 minutes are for me to identify any key misconceptions that will inform my future planning and – the importance of this should never be understated – for my students to feel good about themselves ready for their next maths lesson.

Hopefully that wasn't too painful an ending to this chapter.

OceanofPDF.com

3.11. If I only remember 3 things...

- Explicit instruction, when done well, is the most effective form of instruction for everyone apart from experts in a given domain. Partial guidance during instruction could lead to the development of incomplete or erroneous knowledge.
- Inducing cognitive conflict can be an effective way to resolve lingering misconceptions, but students must have both the will and the ability to make the change, both of which are intimately related to domain-specific knowledge.
- It is often appropriate to develop conceptual understanding alongside or after the development of procedural fluency.

OceanofPDF.com

4. FOCUSING THINKING

Having already confessed my ignorance of how students think and what motivates them, I am acutely aware that I am not exactly coming out of this book smelling of roses. Well, here comes another one: for the vast majority of my teaching career I did not properly consider how I could help my students think more effectively. Indeed, my efforts were largely limited to saying things like, 'If you'd just shut up and listen, you might learn', whilst ploughing through a text-heavy PowerPoint and narrating over the top.

Cognitive Load Theory and the Cognitive Theory of Multimedia Learning have changed all that. I find they provide practical, easy-to-implement strategies to support the principles of Explicit Instruction introduced in the previous chapter, and my lessons have never been the same since.

OceanofPDF.com

4.1. Cognitive Load Theory and the Cognitive Theory of Multimedia Learning

What I used to think

In a tweet in early 2017, Dylan Wiliam described Cognitive Load Theory as 'the single most important thing for teachers to know'. I always take note of everything Dylan says, but this in particular caught my attention as I was only vaguely aware of Cognitive Load Theory. I then discovered Mayer's Cognitive Theory of Multimedia Learning – a related theory primarily concerned with the modality in which material is presented.

The more I read – and having interviewed the likes of Greg Ashman and Kris Boulton for my podcast – the more I was overcome with feelings of awe and frustration. Awe came from the simplicity of the theories – how relatively straightforward changes to the presentation of information and instructional design in general could have profound effects upon students' learning. Frustration came from the fact that I had not given sufficient attention to these theories before.

Sources of inspiration

- Barton, C. (2017a) 'Greg Ashman', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Kris Boulton – Part 1', *Mr Barton Maths Podcast*.
- Mayer, R. E. (2008) 'Applying the science of learning: evidence-based principles for the design of multimedia instruction', *American Psychologist* 63 (8) pp. 760-769.
- Martin, A. J. (2016) *Using Load Reduction Instruction (LRI) to boost motivation and engagement*. Leicester: British Psychological Society.

- Paas, F., Renkl, A. and Sweller, J. (2003) 'Cognitive load theory and instructional design: recent developments', *Educational Psychologist* 38 (1) pp. 1-4.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Cognitive Load Theory
The typical classroom is awash with information that competes for our students' attention. My voice, other students' voices, background noise, displays on the walls, windows, lights, any combination of text, numbers and symbols on the board and on worksheets, the layout of problems, gestures, movement, and of course anxiety, can all fill up our students' working memories. And, as we discussed in Section 1.1, working memory has a finite capacity.

Cognitive Load Theory (*eg* see Sweller *et al*, 1998) suggests that for instruction to be effective, care must be taken to design instruction in such a way as to not overload working memory's capacity for processing information. Martin (2016) explains that if working memory is overloaded, there is a greater risk that the content being taught will not be understood by the learner, will be misinterpreted or confused, will not be effectively encoded in long-term memory, and that learning will be slowed down. Cognitive overload is the enemy of learning. During such cognitive overload, there may not be sufficient spare cognitive capacity in working memory to process information and assimilate it into existing knowledge stored in long-term memory. For the authors of Cognitive Load Theory, learning is a change in long-term memory, so if nothing has changed in long-term memory, then nothing has been learned.

Cognitive Load Theory is not about making thinking easier. It is simply about teaching students in a way so that they can take the information in without overwhelming their working memories.

There are three types of cognitive load that we need to be aware of:

1. Intrinsic load
Intrinsic load relates to the inherent difficulty of the subject matter being learnt, and is determined by the complexity of the material and the prior knowledge of the learner. Specifically, the level of intrinsic load depends on the amount of element interactivity in the tasks that must be learned. Something like learning the names of polygons has a relatively low element interactivity, as knowing the name of a pentagon is not necessary to know the name of an octagon, hence they can be learned separately. However, something like solving a trigonometry problem has a relatively high element interactivity as you need to combine many dependent skills together to arrive at the answer. The degree of element interactivity also depends on the expertise of the learner, because what are numerous elements for a low-expertise learner may be only one or a few elements for a high-expertise learner. I see solving a pair of simultaneous equations as a single process, whereas for many of my students – especially the first time they meet the concept – it consists of many substeps. The most effective way to reduce this intrinsic load is to help students develop sufficient knowledge that can be stored and organised in their long-term memory as schemas. This allows students to work with a sizeable 'chunk' of information as if it was one item, which frees up capacity in working memory. Extensive practice also allows learners to automate key knowledge, so facts such as $6 \times 8 = 48$ and the procedure for adding fractions can be retrieved without imposing any strain on working memory.

2. Extraneous load
Extraneous load is load that is not necessary for learning – indeed, it is unhelpful. It typically results from badly designed instruction and the way material is presented to the learner. We have already met one type of extraneous load in Chapter 1 in the guise of maths anxiety, but there are many other forms. The majority of this chapter will look at how we can reduce this extraneous load.

3. Germane load
Germane cognitive load refers to the load imposed on working memory by the process of learning – that is, the process of transferring information into

the long-term memory through schema construction. In that sense, it can be thought of as 'good cognitive load' as it contributes to learning. We will discuss germane load in the final part of this chapter.

In incredibly simple terms, a key goal of instruction is to help students decrease intrinsic load through knowledge, decrease extraneous load through instructional design, and increase germane load to make use of the resulting spare cognitive capacity in a way that contributes to learning. Indeed, Paas *et al* (2003) explain that these three types of load are additive, so increasing germane cognitive load will only be effective if the total cognitive load remains within the limits of working memory.

Cognitive Theory of Multimedia Learning

Related to Cognitive Load Theory is the Cognitive Theory of Multimedia Learning (*eg* see Mayer, 2008). The theory also stresses the importance of reducing extraneous load, and is based upon three key assumptions concerning working memory:

1. Dual channel assumption – people have separate channels for processing visual and verbal material.
2. Limited capacity assumption – people can process only a limited amount of material in a channel at any one time.
3. Active processing assumption – meaningful learning occurs when learners select relevant material, organise it into a coherent structure, and integrate it with relevant prior knowledge.

This theory is both good and bad news for teachers. On the downside, our students can experience cognitive overload if just one of their channels gets filled up, for it is not possible to 'borrow' spare capacity from the other. But the good news is that we can make the most of students' limited working memories by carefully presenting information in different modalities.

When we combine these two theories together, we essentially have a blueprint for how materials should be presented to students in all phases of learning, but in particular during the early knowledge acquisition phase,

when it is of paramount importance that students' limited working memory resources are dedicated to processing exactly the right things.

What I do now

I am far more aware of both the importance and limitations of working memory. Working memory is the gateway to learning, and both Cognitive Load Theory and the Cognitive Theory of Multimedia Learning are about filling students' working memories up with the right things.

The rest of this chapter will look at how we do exactly that.

OceanofPDF.com

4.2. When silly mistakes may not be that silly

What I used to think

I used to think that students who made silly mistakes were lazy and careless. A missing unit of measurement here, an incomplete final step there, and before you knew it crucial exam marks were disappearing down the drain. Silly mistakes were the bane of my life. My technique for stopping my students making such mistakes? Repeatedly tell them not to be so careless, of course.

Strangely, this did not seem to work.

Sources of inspiration

- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Consider the following question:

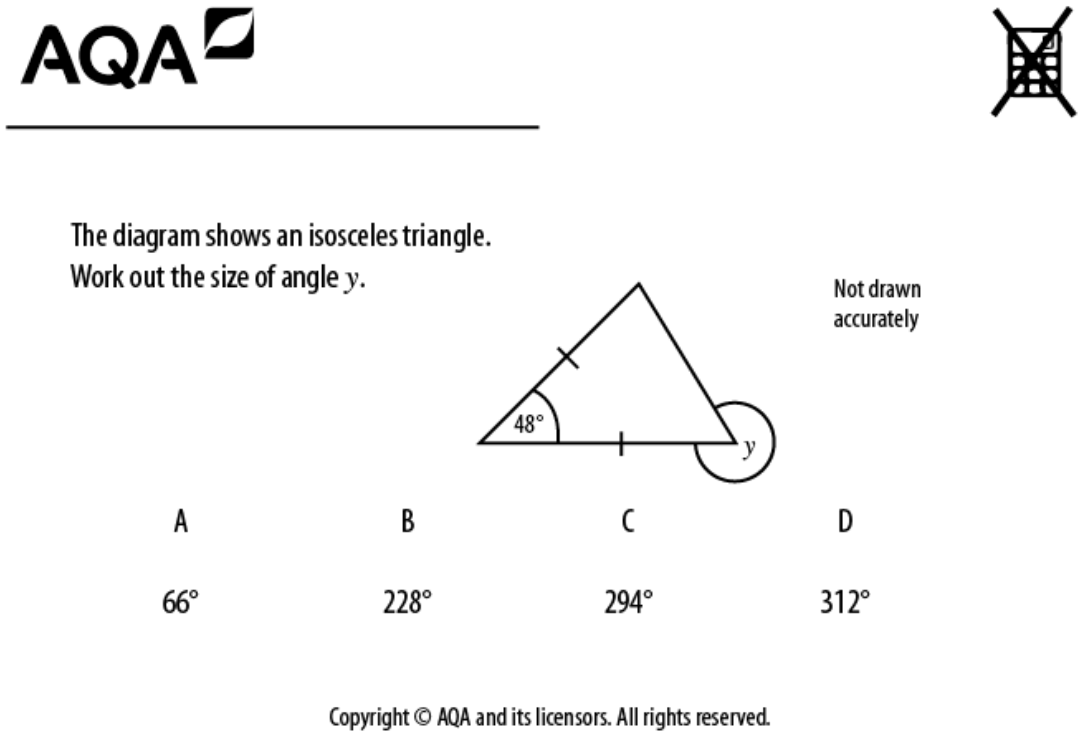


Figure 4.1 – Source: AQA for Diagnostic Questions

At the time of writing, this question has been answered well over 10,000 times on my Diagnostic Questions website.

Now, it may come as little surprise to you that the most popular incorrect answer is D (chosen by 22% of respondents), with students who are used to the conventional orientation of an isosceles triangle assuming that the bottom-right angle is also equal to 48°. However, would it surprise you that 18% of students answering this question went for A?

How on earth could anyone believe that the answer to this question was 66°? What a silly mistake!

Or is it?

Having considered Cognitive Load Theory and the workings of memory, I am not so sure any more. We can shed some light onto students' thought processes by reading some of the explanations they gave to accompany their choice of A on my Diagnostic Questions website:

180 take away 48 is 132, and 132 divided by 2 is 66, meaning both the sides are 66°.

I think that the answer is 66 because with isosceles triangles the two bottom angles are the same. So I turned the screen to find that the side two angles look closely identical, so I then took 48 from 180 because I knew that angles in a triangle add to 180, so then I got the answer 132, then I divided that by two, to get 66.

Both really nice pieces of thinking, but notice that the final step – namely subtracting their answer from 360 – is missing each time. Now this could be a silly mistake, but it could also be the result of cognitive overload.

Consider what a student might need to process in their working memories to solve this problem:

- What is the question asking?
- What do the angles in a triangle add up to?
- What type of triangle is it?
- How does that help me?
- Which are the base angles?
- What is 180 – 48?
- What is 132 ÷ 2?

Now, if the student's working memory is full to the brim by the time they reach the end of this chain of thinking, then 'What is the question asking?' may well have been bumped out, or there may simply be no spare cognitive capacity to process 360 – 66. In other words – and this is my technical term – they are cognitively knackered. In such a scenario, it is not surprising that the student settles upon an answer of 66°.

We can see here how domain-specific expertise, and the associated way of thinking that accompanies it, can help this process. For an expert, this question has a relatively low intrinsic load. They are likely to immediately

recognise this as an isosceles triangle problem – indeed, they may have solved an unconventional one like this in the past – and will easily be able to retrieve the answers to the calculations with little burden placed upon working memory. Hence, they have capacity left to complete the problem, check it makes sense, and also assimilate it into their long-term memories and hence learn from the process.

Here are a selection of other 'silly mistakes' that may indeed be the result of cognitive overload. In each case, notice how it is the final stage of the process where the mistake occurs – in other words at the very point when working memory is likely to be feeling the greatest strain.

- Doing a simultaneous equations problem and 'forgetting' to find y
- Estimating the mean from a grouped frequency table and 'forgetting' to divide by total frequency
- Completing an area problem and 'forgetting' to put the units in
- Doing a tree diagram question and 'forgetting' to add the fractions together at the end
- Solving a quadratic equation using the formula and 'forgetting' to round the answers to one decimal place
- Adding two fractions together and 'forgetting' to simplify the answer

What I do now

When a student makes what I would have previously dismissed as a 'silly mistake', I now ask myself if carelessness really is the cause, or could it be cognitive overload? If it is the latter, then telling the student to be more careful next time is likely to be a fruitless exercise. Instead, I need to equip them with the knowledge they need to successfully solve such problems, and ensure as much of it as possible is automated.

But this is a long-term solution, making use of the principles of explicit instruction together with the practical strategies described in the chapters that follow. However, there are more immediate steps we can take to reduce

extraneous load, and hence increase the chances of our students solving problems like this successfully.

OceanofPDF.com

4.3. The Modality Effect

What I used to think

I did not used to think a great deal about the difference between presenting an image or animation alongside text, versus alongside an oral description. Indeed, if forced to choose between the two, I would always opt for the image alongside text, as I perceived text more likely to stick in students' minds than narration.

Sources of inspiration

- Baddeley, A. (2012) 'Working memory: theories, models, and controversies', *Annual Review of Psychology* 63, pp. 1-29.
- Brzezinski, T. (no date) *Polygons: Exterior Angles – REVAMPED*. Available at: https://www.geogebra.org/m/mKzJCf5p
- Mayer, R. E. (2008) 'Applying the science of learning: evidence-based principles for the design of multimedia instruction', *American Psychologist* 63 (8) pp. 760-769.
- Mayer, R. E. and Anderson, R. B. (1992) 'The instructive animation: helping students build connections between words and pictures in multimedia learning', *Journal of Educational Psychology* 84 (4) pp. 444-452.
- Mousavi, S. Y., Low, R. and Sweller, J. (1995) 'Reducing cognitive load by mixing auditory and visual presentation modes', *Journal of Educational Psychology* 87 (2) pp. 319-334.
- Yorgey, B. (2017) 'Post without words 18', *The math less travelled blog*. Available at: https://mathlesstraveled.com/2017/05/31/post-without-

words-18/

My takeaway

The presentation of information in different modalities – referred to as The Modality Effect – is a feature of both Cognitive Load Theory (Mousavi *et al*, 1995) and the Cognitive Theory of Multimedia Learning (Mayer and Anderson, 1992).

To appreciate the importance of the Modality Effect we need to develop the basic model of working memory introduced in Section 1.1. Baddley and Hitch (*eg* Baddley, 2012) propose a multicomponent model of working memory, where working memory is not viewed as a single workhouse, but rather one divided up into separate components, each with their own specific role. The two most significant components for our discussion here are:

- The Phonological Loop – deals with speech and sometimes other kinds of auditory information
- The Visuospatial Sketchpad – holds visual information and the spatial relationships between objects

The key point here is that working memory gets overloaded if too much information flows into one of these components, but we can use the different components to aid processing. Hence, the capacity of working memory may be determined by the modality (auditory or visual) of presentation, and we can make more efficient use of the limited capacity of working memory by presenting information in a mixed (auditory and visual mode) rather than in a single mode.

Consider two slides to convey the message that exterior angles in a polygon add up to 360°:

Slide 1:

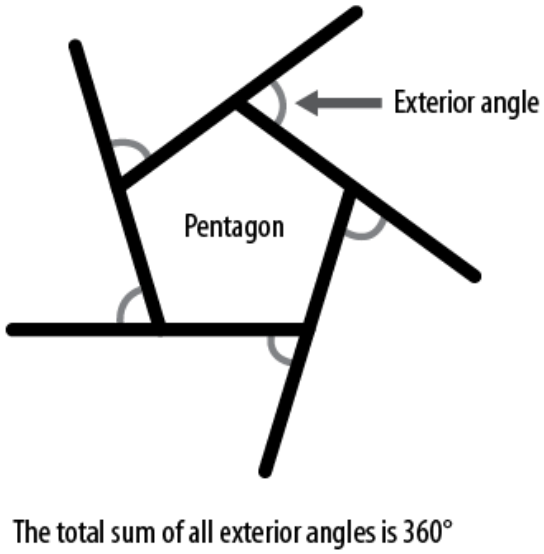


Figure 4.2 – Source: Craig Barton

Slide 2:

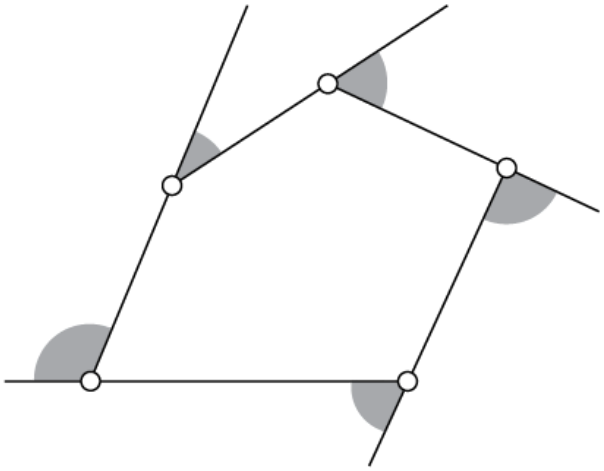


Figure 4.3 – Source: Tim Brzezinski created using Geogebra, available at https://www.geogebra.org/m/mKzJCf5p

Slide 1 looks friendly enough, especially to experts such as ourselves. It is an image surrounded by a small amount of text. But we must consider how working memory processes this. The image is held inside the Visuospatial Sketchpad. The text, however, is also initially held by Visuospatial Sketchpad before effectively being read aloud inside students' minds and thus held inside the Phonological Loop. It is this initial processing of the text as an image that could lead to cognitive overload for our students as the

Visuospatial Sketchpad reaches capacity, especially as students try to also process the meaning of the slide.

Now, Slide 2 is in fact a dynamic animation built in Geogebra by Tim Brzezinski. It can be presented on the board without text, together with an oral narration from me along the lines of 'Consider each of the exterior angles. We can see that together they form a full turn…'. In terms of working memory, we have freed up the Visuospatial Sketchpad to focus solely on the image, while making use of the Phonological Loop to deal with the narration.

Directly relevant to this is the Temporal Contiguity Principle from the Cognitive Theory of Multimedia Learning (Mayer, 2008). This explains that 'people learn more deeply when corresponding graphics and narration are presented simultaneously rather than successively'. This is interesting, as my intuition would be to present the image/animation, give students a chance to process it, and then begin the narration. This intuition is faulty due to my lack of understanding of how working memory functions. The image and the narration are processed in separate channels, and hence there is no advantage to be gained from separating them in time.

However, there are of course times where we may want to present an image and challenge the students to think what it represents. In such circumstances it clearly makes sense to separate the narration from the image. One of my favourites examples of this is the 'Post without Words' series from Brent Yorgey, which I would insist my students look at and ponder in silence before we discuss it:

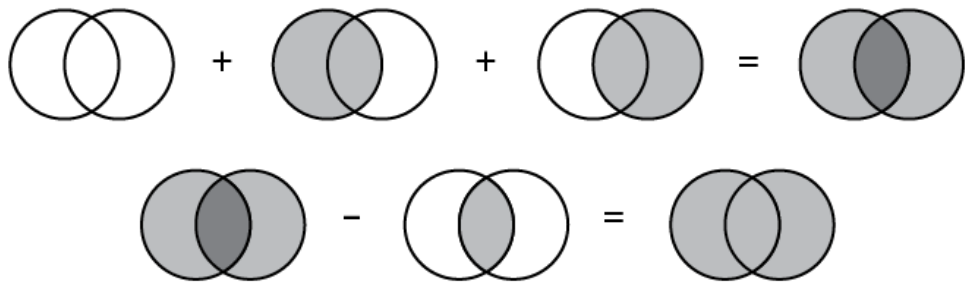


Figure 4.4 – Source: Brent Yorgey, available at https://mathlesstraveled.com/2017/05/31/post-without-words-18/

The presentation of images with or without narration and text may not seem like much. However, we need to remember that we are dealing with novice learners in the early knowledge acquisition phase of learning. A phrase I like is '**the cusp of understanding**'. Imagine you have a student who is close to making the key breakthrough needed to understand a concept – they are so close, but things could still go either way. Students on such a cusp of understanding need as much of their working memory resources as possible dedicated to the matter in hand. Hence, such small changes to the presentation of information could make all the difference, allowing for successful processing of information to instigate that all-important change to long-term memory. Indeed, when considering all the tweaks suggested in this chapter, it may be useful to think in terms of students on the cusp of understanding.

What I do now

Where possible – and when appropriate – I present images or animations without on-screen text, and narrate over the top instead. Obviously there are times when I do need to use text, and for those occasions I follow the guidelines suggested by the Split-Attention and Redundancy Effects discussed later in this chapter.

OceanofPDF.com

4.4. Learning styles

What I used to think

Seeing as we are discussing presenting information in different modalities, this seems like a perfect time to take a quick look at learning styles.

I used to think each student learned differently depending on their preferred learning style, and it was my job to cater to those individual preferences. I went on numerous (pretty expensive) courses to help me do just that. I needed to identify – usually by means of a questionnaire given to each student at the start of the year – who my Visual, Auditory and Kinesthetic learners were, and adapt my teaching accordingly. Details of exactly how I intended to do this had to be recorded on my lesson plans. If that meant writing a rap or devising a dance to help students understand how to add two fractions together, then so be it.

Yes, learning styles were all the rage throughout my first few years of teaching, and they are still around today. You only need to search for 'VAK' or 'Learning Styles' on *TES* to find a wondrous assortment of posters, stickers, questionnaires, activities and badges.

I never considered that it might, in fact, be a load of bollocks.

Sources of inspiration

- *Deans for Impact* (2016) Available at: www.deansforimpact.org
- Kirschner, P. A. and Van Merriënboer, J. J. G. (2013) 'Do learners really know best? Urban legends in education', *Educational Psychologist* 48 (3) pp. 169-183.

- Pashler, H., McDaniel, M., Rohrer, D. and Bjork, R. (2008) 'Learning styles: Concepts and evidence', *Psychological Science in the Public Interest* 9 (3) pp. 105-119.
- Willingham, D. T. (2005) 'Do visual, auditory, and kinesthetic learners need visual, auditory, and kinesthetic instruction?', *American Educator* 29 (2) pp. 31-35.

My takeaway

Okay, take a deep breath before we dive into this one…

It turns out that there is little to no evidence that students have preferred learning styles, or that we should attempt to cater to them. I have presented four sources above, should you wish to dig deeper into the matter. The findings are nicely summarised by Pashler *et al* (2008), who conclude:

> Our review of the learning-styles literature led us to define a particular type of evidence that we see as a minimum precondition for validating the use of a learning-style assessment in an instructional setting. As described earlier, we have been unable to find any evidence that clearly meets this standard…
>
> …The contrast between the enormous popularity of the learning-styles approach within education and the lack of credible evidence for its utility is, in our opinion, striking and disturbing. If classification of students' learning styles has practical utility, it remains to be demonstrated.

Willingham (2005) offers some useful advice. He explains that it is likely that some students have a relatively better visual memory or auditory memory, but that does not mean we should always teach to it. The key is that teachers should focus on the content's best modality, not the student's. Hence, we should teach geometry topics in a visual way, regardless of the preferred learning styles of our students, because even a so-called auditory learner will understand it better that way. Better still, we should look for ways to present topics in more than one modality, making use of the split-channel nature of working memory discussed in the previous section, but ensuring that the

different representations are complementary to each other. That way, all students, regardless of their perceived learning style, are likely to benefit.

By only presenting material to certain students in a single modality (for example, in an auditory way), just because they achieved a certain score on a dubious questionnaire, we risk kick-starting a self-fulfilling prophecy, whereby their apparent weakness in coping with visual and kinesthetic presentations becomes a reality through a lack of practice. Students need to be able to cope with material presented in all modalities, and it is our job as teachers to help – not hinder – that development.

What I do now

All mention of learning styles has been removed from my lesson plans and is considered akin to swear words in my classroom. Instead, I concentrate on finding the best modality for each concept I am teaching. Where possible, I use different representations to convey the same concept, but only if appropriate.

For example, when teaching sequences I always turn my attention to visualpatterns.org, allowing me to represent sequences visually as well as in their algebraic form:

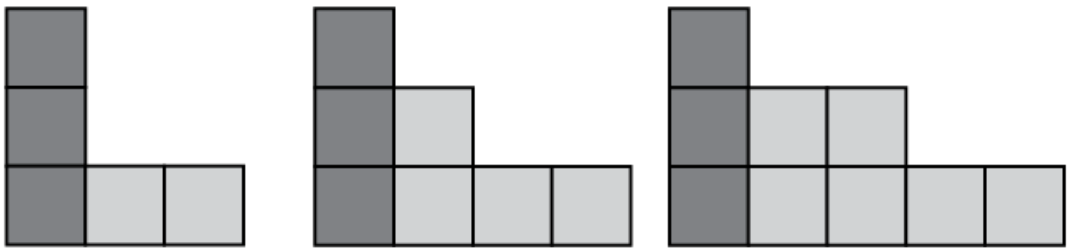


Figure 4.5 – Source: David Wees, post #25 available at http://www.visualpatterns.org/21-40.html

OceanofPDF.com

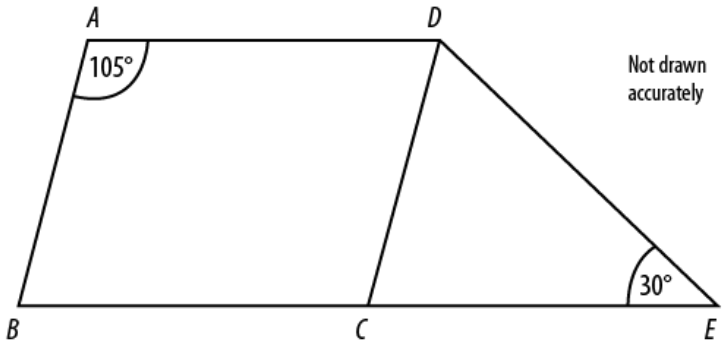
4.5. The Goal-free Effect

What I used to think

I used to think that giving students questions with a specific answer to find was perfectly fine. In fact, it was likely to be less confusing for them than a more open-ended problem. Indeed, that is how questions are presented in the exam, so what could be a better way to prepare students than to expose them to these kinds of questions as early as possible?

The problem was, many of my students did not know where to begin when faced with a question such as this:

A parallelogram *ABCD* and a triangle *DCE* are joined as shown.
BCE is a straight line.



Show that *DCE* is an isosceles triangle.
You must show your working.

[4 marks]

Figure 4.6 – Source: AQA November 2016 GCSE Maths Foundation Paper 1

Sources of inspiration

- AQA (2016) *Foundation Tier – Paper 1. 4365/1F.*

- AQA (2016) *Foundation Tier – Paper 2. 4365/2F.*
- Lovell, O. (2017) 'A Conversation with John Sweller', *ollielovell.com*. Available at: www.ollielovell.com/pedagogy/johnsweller/
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Sweller *et al* (1998) explain that goal-specific problems – like the one above – are those in which the top-level goals (*ie* final answer) can only be achieved by successfully completing the sub-goals (*ie* the steps leading up to it). To understand the impact of goal-specific problems on the working memories of students, we need to return to the distinction between how experts and novices think that was introduced in Section 1.2.

Consider for a moment how you would solve the problem above. Being immediately aware of the properties of isosceles triangles, you may reason that ultimately you will need to find the sizes of the missing angles in triangle DCE. In order to do this, you may first write down that angle DCB is 105° using your knowledge of angles in a parallelogram, and continue from there, working out the sizes of missing angles until you have found all you need.

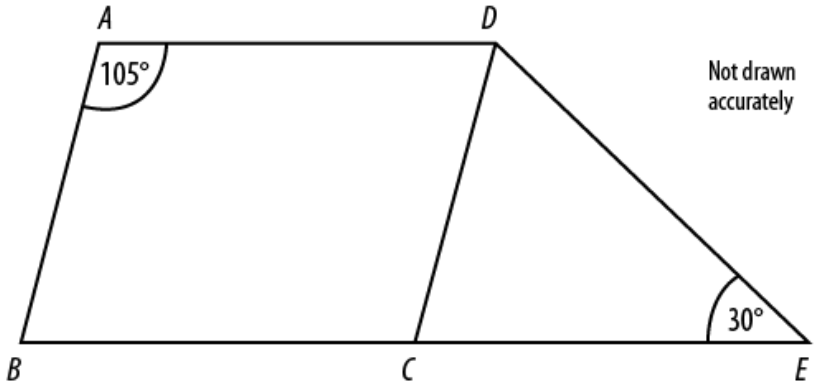
Is this how your students would solve the problem – especially the first time they have encountered angle facts? I know mine wouldn't, precisely because they are novice learners.

According to Sweller *et al* (1998), when faced with goal-specific problems, novice learners tend to embark upon a 'means-end analysis'. This involves the student attempting to juggle all the possible sub-steps that minimise the difference between the current state (the problem) and the end state (the goal, or final answer). Such an approach is cognitively demanding – especially given how far away the end goal can seem and how many sub-steps may be involved – and there are two likely outcomes. The first is that the student experiences cognitive overload and does not solve the problem.

The second is that even if the student solves the problem, they may not learn anything in the process. We will discuss this fascinating and worrying claim more in Chapter 9, but the essence of the argument is that the unstructured way novices approach problems is not compatible with the development of schemas. Sure, they may solve the problem at hand, but they are unlikely to learn anything that can be generalised and transferred to a new situation. For the proponents of Cognitive Load Theory, learning is a change in long-term memory, so if nothing changes in long-term memory, nothing has been learned. A relatively simple solution is to instead present students with 'goal-free' problems. These focus attention upon working forward from the information present one step at a time in a structured manner, rather than trying to hold multiple possible steps in working memory at once. A student can complete one step, reflect and process it, before moving on to the next. Each step feels in reach, rather than being some faraway, daunting final goal. This is less cognitively demanding and is more compatible with how an expert would attempt a problem. Crucially, it is more compatible with schema acquisition, and hence more likely to lead to learning.

So, we can adapt the question above to become:

A parallelogram *ABCD* and a triangle *DCE* are joined as shown.
BCE is a straight line.



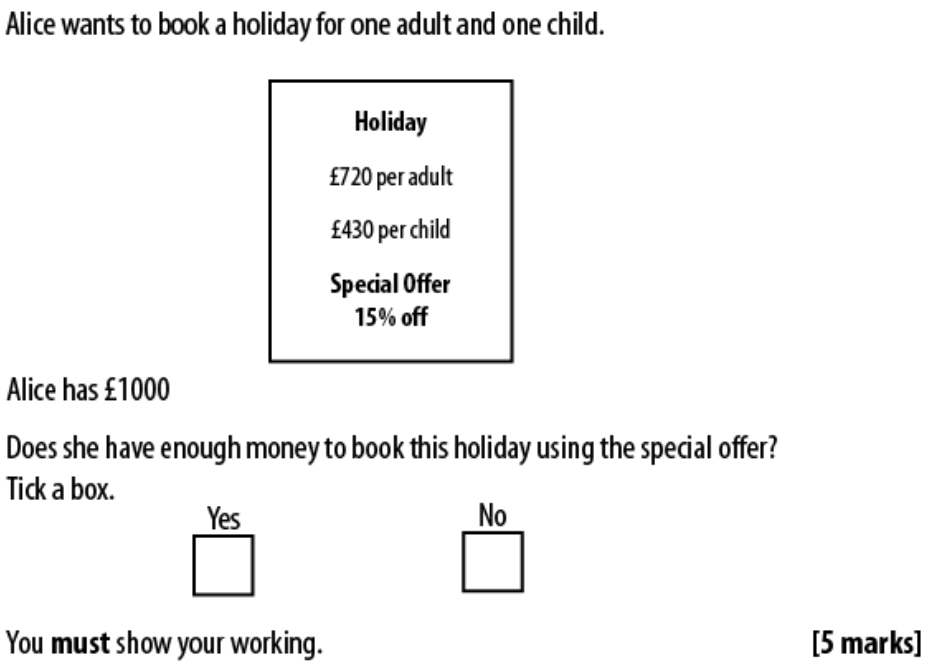
Find as many angles as you can, giving the reason for each one.

Figure 4.7 – Source: Craig Barton, adapting a question from AQA November 2016 GCSE Maths Foundation Paper 1

Students can proceed one step at a time. Once all angles have been found, then the question of what type of triangle DCE is can be addressed.

This does not need to be just geometry-based problems. Many questions that are multistep – as a rule, I tend to say 3 marks or more in an exam – can be turned into a goal-free problem. A key thing to bear in mind is that you don't want to remove too many constraints so that students are practising skills unrelated to the original question. So, if the original question involves angles or percentages, the goal-free version should encourage students to practise the same skills. Indeed, there are worse things to do in a department meeting (or a Friday night) than attempting to turn goal-specific problems into goal-free ones.

Take something like this:


Alice wants to book a holiday for one adult and one child.

Holiday

£720 per adult

£430 per child

Special Offer
15% off

Alice has £1000

Does she have enough money to book this holiday using the special offer?
Tick a box.

Yes ☐ No ☐

You **must** show your working. **[5 marks]**

Figure 4.8 – Source: AQA June 2016 GCSE Maths Higher Paper 2

With a quick tweak, we have:

Alice wants to book a holiday for one adult and one child.



What can you work out?

Figure 4.9 – Source: Craig Barton adapting a question from AQA June 2016 GCSE Maths Higher Paper 2

Again, once students have calculated all they can, we can reveal the original question, and show students that they have either already worked it out, or are incredibly close. This subtle, simple change means students make more progress through problems, learning all the time, instead of being overwhelmed by a single goal that seems far away and unattainable. I have found it makes students more willing to have a go at problems as they know they can make a start. That first taste of success is close at hand.

Now of course, the exam is going to be full of goal-specific problems. But that does not mean that the practice needs to be. Remember, we are talking about early knowledge acquisition phase here. We are trying to reduce the burden on students' fragile working memories to allow them to make changes to their long-term memories and hence learn something. The point about the format of the practice (goal-free problems) not matching the format of the exam (goal-specific problems) is a key component of Deliberate Practice, which is the focus of Chapter 8.

What I do now

In early knowledge acquisition phase, I make more use of goal-free problems. They are what I turn to the first time students apply a skill I have taught them. Moreover, I have found taking a past exam paper and converting it into a goal-free version is an ideal way to kick-start the revision process.

Interestingly, in a 2017 interview with John Sweller, Oliver Lovell suggested that we could encourage students themselves to turn any problem they are struggling with into a goal-free problem. So, when faced with a complex multi-step maths problem, students could essentially cover up the goal-specific question and just start working things out. Then, every few minutes, look back at the original problem and see if they are any closer to getting to where you need to be. If they can see how to get there, great. If they can't, they can just keep playing around and find some more stuff out. When asked what he thought about that idea, Sweller replied: 'That's smart. We haven't done that, but it ought to be done. Good idea. You should try it out. Yeah, I can't see anything but positives to that'.

OceanofPDF.com

4.6. The Split-Attention Effect

What I used to think

I'll once more admit that I did not think a great deal about the presentation of information involving text and diagrams. All my attention was on things like the quality of my explanations and the questions I gave my students to do. Little did I know that I was but a few simple tweaks away from making life a lot easier for me and my students.

Sources of inspiration

- Barton, C. (2017) 'Dani Quinn – Part 1', *Mr Barton Maths Podcast*.
- Mayer, R. E. (2008) 'Applying the science of learning: evidence-based principles for the design of multimedia instruction', *American Psychologist* 63 (8) pp. 760-769.
- Mayer, R. E. and Anderson, R. B. (1992) 'The instructive animation: helping students build connections between words and pictures in multimedia learning', *Journal of Educational Psychology* 84 (4) pp. 444-452.
- Mousavi, S. Y., Low, R. and Sweller, J. (1995) 'Reducing cognitive load by mixing auditory and visual presentation modes', *Journal of Educational Psychology* 87 (2) pp. 319-334.
- Sherwood, J. (2017) 'Building effective learning strategies into a mathematics curriculum', *Learning Scientists Blog*. Available at: http://www.learningscientists.org/blog/2017/6/13-1
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational*

Psychology Review 10 (3) pp. 251-296.

My takeaway

The Split-Attention Effect occurs when learners are required to process two or more sources of information simultaneously in order to understand the material. If students must split their attention between the two forms of presentation, this will increase the cognitive load imposed upon working memory (Mousavi *et al*, 1995). This is also related to the Spatial Contiguity Principle from the Cognitive Theory of Multimedia Learning (Mayer, 2008), which states that people learn more deeply when corresponding printed words and graphics are placed near rather than far from each other on the page or screen.

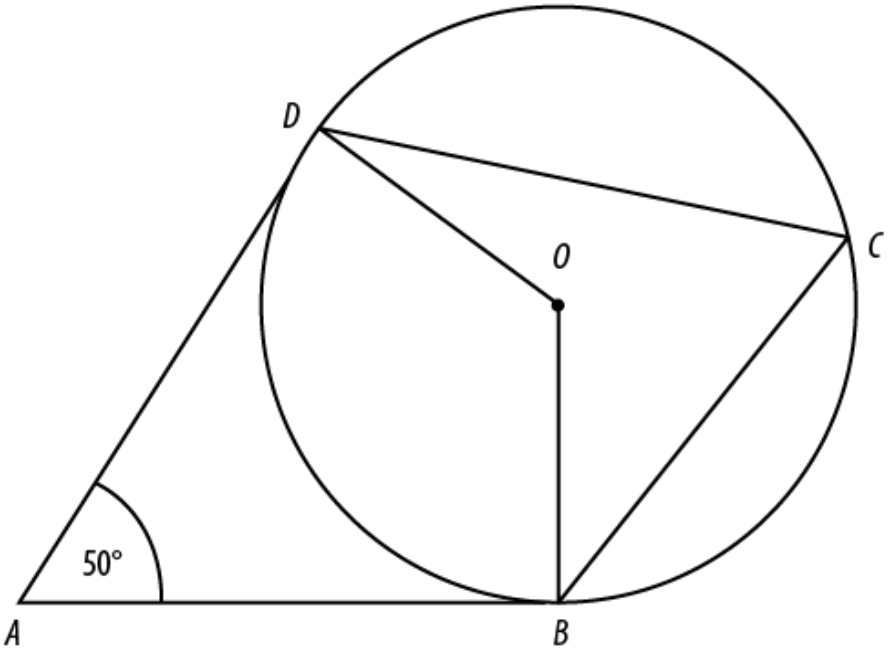
The Split-Attention Effect occurs all over the place during instruction, but fortunately it is relatively easy to fix.

Text and Diagrams

The Modality Effect above extolled the virtues of narration over text, but there are times when an image or diagram cannot be understood without text. When presenting students with examples where a diagram is involved (for example, most geometry topics), the text should be carefully integrated within the diagram, and not separate from it. This will prevent students from having to deal with these two forms separately and reduce the extraneous cognitive load.

A nice example of this is the following from Jemma Sherwood's 2017 blog post for the *Learning Scientists*.

Instead of giving the question to students like this:



B, *C* and *D* are points on the circumference of a circle, centre *O*.
AB and AD are tangents in the circle.

Angle *DAB* = 50°

Work out the size of angle *BCD*.
Give a reason for each stage in your working.

Figure 4.10 – Source: Jemma Sherwoord for Learning Scientists Blog, available at http://www.learningscientists.org/blog/2017/6/13-1

Give them the question on the next page.

All the key information is included, but it is better integrated, thus less of a burden on working memory, and hence cognitive capacity will be freed up to think more deeply about the problem in hand.

The same principles apply to worked solutions of problems. Having a diagram at the top of the page and the solution below is more cognitively demanding than it needs to be, as students must switch attention between the two presentations of information. Integrating key elements of the solution within the diagram makes it far easier to follow.

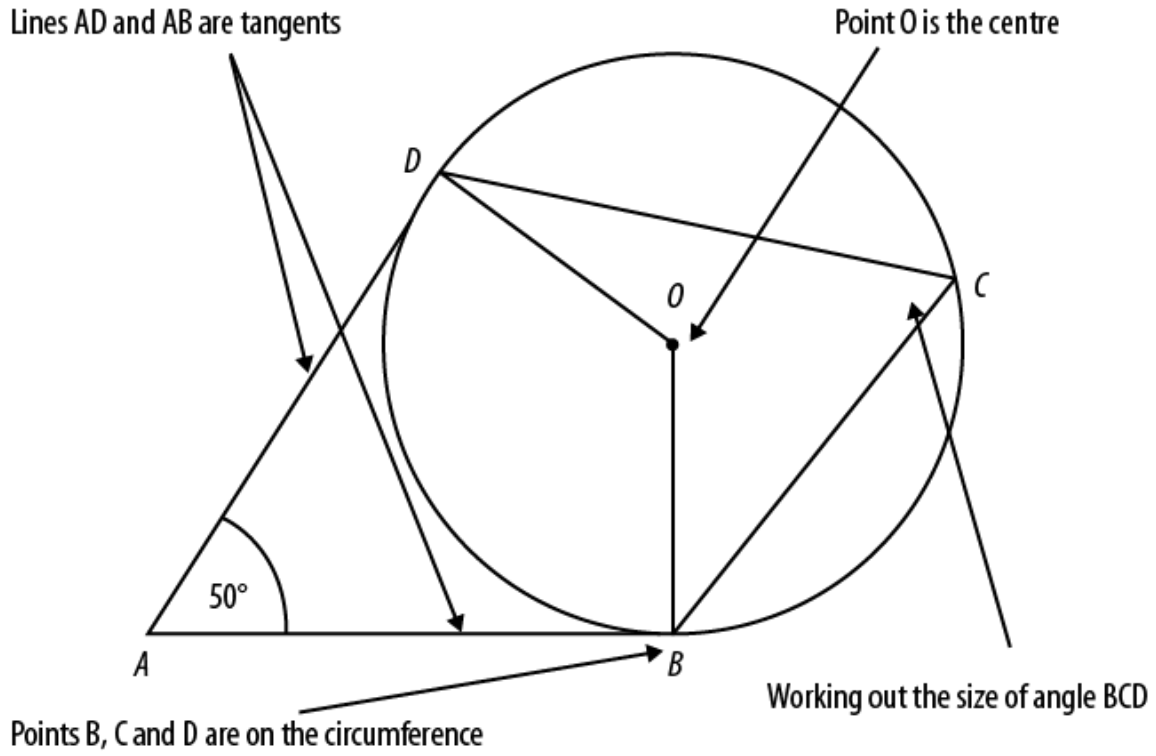


Figure 4.11 – Source: Jemma Sherwood for Learning Scientists Blog, available at http://www.learningscientists.org/blog/2017/6/13-1

Location of Questions

Keeping all aspects of a question visible at the same time (*ie* without having to constantly turn the page over) will also help reduce this unnecessary load. Some textbooks, worksheets and exam papers are notoriously bad for this, but that does not mean we need to subject our students to it, especially in the early knowledge acquisition phase of learning.

Useful talk

Often students will be busy copying down a worked example, or beginning to work on a problem, and I will have something important to say. This might be something along the lines of 'remember to write your units', or 'check your answer by using substitution at the end'. Referring back to the separate channels of working memory introduced earlier in this chapter, my words will need to be stored and processed in my students' Phonological Loops. The problem is, so too will the words they are reading on either the whiteboard or the problem in front of them as they effectively say them aloud in their heads. Once again, we risk cognitive overload.

I am a bugger for doing this – I just cannot keep quiet in lessons. I think it is due to some long-standing, misguided notion that silent classrooms are a bad thing, whereas noise is clearly a sign of learning. It took me 12 years, hundreds of pages of research, and a three-hour conversation with Dani Quinn on my podcast, but finally I now know that quite the opposite is true. When students are trying to concentrate, silence is golden. So, I have made myself up a rule that I try my very best to stick to. If I am putting some text up on the screen (and this includes algebraic equations, as they need to be read aloud in students' heads just like words), then I will give my students adequate time to read it themselves. The room will be silent. Then, if I have something important to say, I will stop the class, clear the screen on my interactive whiteboard (either by sliding the cover across the projector, or hitting 'B' during a PowerPoint presentation) or ask my students to turn the sheet or book they are working from face down. That way I reduce the chance of their attention being split across multiple sources of presentation. In general, I try to cut down the times I speak when text is visible.

And if the thing I have to say is not important at all…well, then we have the Redundancy Effect, which is coming up next!

Keeping still

As well as being an incessant talker in the classroom, I am also a mover. I wear a FitBit on my wrist, and if I have not clocked up 15,000 steps across a six-period day, then something has gone wrong. And moving is brilliant as it allows me to get to all corners of the classroom, neglecting no one and keeping an eye on all.

But such movement comes at a price.

Students' attention will naturally be drawn my way. The mere sight of me moving, combined with internal questions of 'why is he moving there?' and 'where is he moving to next?', as insignificant as they sound, still take up precious space in students' fragile working memories. The same is also true of any gestures I make.

So, once again, I have another rule. When I have something to say to the whole class, I keep incredibly still, with my hands behind my back. This

actually has a double advantage. As well as focusing the students' attention on me and my words, it also helps me think better. We have all experienced walking somewhere, having an idea, and then either slowing down or completely stopping to think about it. By instinct, we know we need all the capacity our working memories can muster to ponder the problem at hand. So, I now do exactly that in the classroom. By standing still, I can focus all my attention on the words and explanations I am saying. And then, when my students are working, I can – as quietly as possible – move around the room to monitor behaviour and offer one-to-one help.

Keeping still also means that when I do move and make gestures, it is more effective. For example, when explaining the nature of stationary points to my Year 12s, I like to get my arms swooping up and down the graph of a cubic equation a to draw students' attention to the slope of different sections of the curve. This now has more of an impact in enticing students' attention than when I was on the move the whole lesson.

What I do now

I carefully consider what my students' attention will be drawn towards, and if I can help improve the situation. Carefully integrating complementary information, keeping talk and text separate, and not moving around quite as much all sound like insignificant things, but they can have a big impact on the fragile working memories of our students who are on the cusp of understanding.

OceanofPDF.com

4.7. The Redundancy Effect

What I used to think

I used to believe that redundant information had – at worst – a neutral effect. In fact, that effect was more likely to be positive. After all, surely saying the same thing a few times is the best way to make something stick in, and adding a bit of flavour to a problem (or indeed a classroom) never did anybody any harm…

Sources of inspiration

- Mayer, R. E. (2008) 'Applying the science of learning: evidence-based principles for the design of multimedia instruction', *American Psychologist* 63 (8) pp. 760-769.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

We must be careful to distinguish the Redundancy Effect from the Split-Attention Effect. Split-attention occurs when learners are faced with multiple sources of information that must be integrated before they can be understood. The individual sources of information cannot be used by learners if considered in isolation, hence the need for integration. The Redundancy Effect occurs when learners are presented with additional information that is not directly relevant to learning, or with the same

information in multiple forms but in a single modality. The reason the Redundancy Effect is detrimental to learning is simply because redundant information is very difficult to ignore, and hence must be processed in students' limited working memories. Indeed, even if students manage to successfully filter out this redundant information, filtering is still effortful, and thus may divert attention away from thoughts that contribute to learning.

And, it turns out, my lessons were abundant with the redundant.

Redundant information in questions

Once you are on the lookout for redundant information, you see it everywhere. Exam papers, textbook questions, lessons from *TES* – you name it, it is there. And at the risk of repeating myself, just because students will be exposed to redundant information in exams does not mean they should be exposed to it during the early knowledge acquisition phase. Students' working memories are fragile, and we should do all we can to ensure that what is being processed in there is absolutely essential.

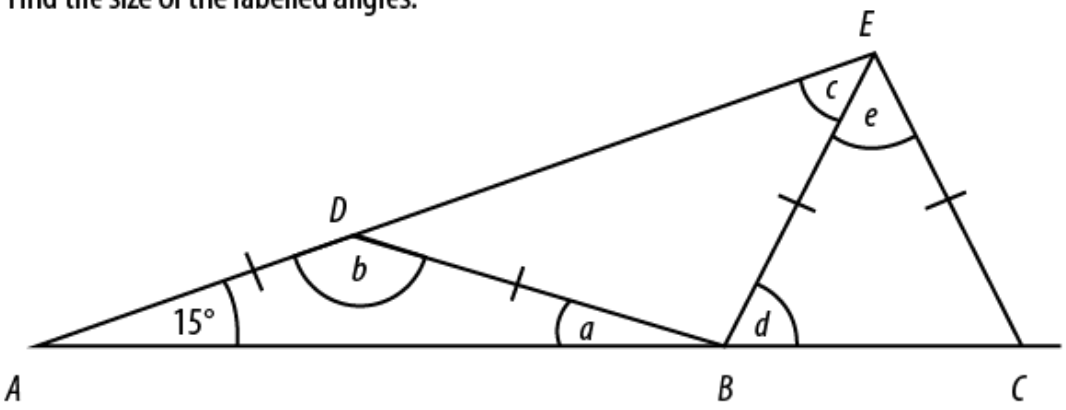
Consider this fairly standard angles question:

In the diagram below, ABC and ADE are straight lines.

AD = BD = BE = CE

Angle BAD is 15°

Find the size of the labelled angles.



Answers: a= _____ b= _____ c= _____ d= _____ e= _____

Figure 4.12 – Source: Craig Barton

Harmless enough, right? But if we put ourselves into the shoes (or the minds) of a novice learner who needs every spare bit of their cognitive capacity focused on the essential elements of the question, then this could be cognitively demanding in an unhelpful way. Remember, even redundant information needs to be processed. So, are there things we can get rid of?

In the diagram below, ABC and ADE are straight lines

Of course they are straight lines. We can see that! And I know that strictly we need to state this in order to be able to answer the question, but remember that we are in early knowledge acquisition phase, and our aim to to give students the very best chance of learning – in this case – about angles. What is the best that can happen with this opening phrase? Students ignore it. And the worst? Students ask about it, and then we end up either embarking upon a long (no doubt interesting) explanation of why it is necessary to define the lines in this way, or simply telling students it doesn't matter. Either way, I would argue that it is best to leave it out.

AB = BD = BE = CE

So long as students understand the dash notation to represent equal sides – and we can assess this baseline knowledge at the start of the lesson using the principles detailed in Chapter 11 – then this statement is redundant. Remove it.

Angle BAD is 15°

This is clearly labelled on the diagram. Now, of course, being able to label angles given their three-letter definition is a vital skill, but it is not necessary in order to answer this question. Considering it can only take up valuable space in working memory, it has to go.

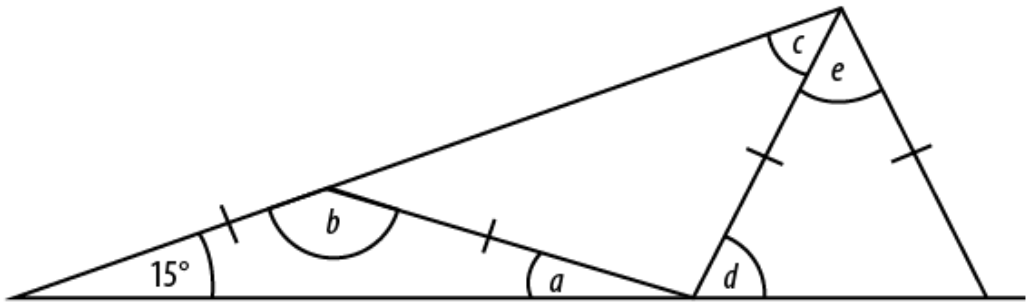
Find the size of the labelled angles

Even this can be removed given the space for answers underneath the question.

So, in the end we are left with the question below – a version of the question that retains all the key elements without the redundancy.

Reading out slides

Have you ever been in a presentation where the presenter read aloud the words on the slide? If you have been to any training given by me in the last ten years, then the answer is yes. It is infuriating for the very reason that it is cognitively draining in a completely unhelpful way.



Answers: a= _____ b= _____ c= _____ d= _____ e= _____

Figure 4.13 – Source: Craig Barton

As we have seen, text on slides is first stored in the Visuospatial Sketchpad, before moving to the Phonological Loop as it is essentially read aloud in the mind. The problem is, any form of talking over the top also needs to be stored and processed in the Phonological Loop. Hearing the slides being read out adds nothing at all to the information that can be gleaned from just reading them. Indeed, it can only interfere with the processing of the written information by clogging up vital space in working memory as you try to process it or filter it out. If you were one of my students, the best thing you could do would be to either close your eyes or cover your ears, neither of which is ideal.

As a general rule, if two pieces of identical information need to be processed by the same part of working memory, then one of them has to go. If reading and literacy are an issue, then narration without the text could be the way forward. If not, then I tend to favour just text, giving the students chance to read and process the information in silence at their own speed.

If what I am saying is different from what is written on the slides then we have the Split-Attention Effect, and I should follow the guidelines outlined in the previous section. But if what I am saying is literally repeating what is written, then no good at all can come from it, and I need to learn to shut up.

Useless talk

I have always found silence in a classroom a little disturbing. Hence, I tend to fill the void with various utterances. At times these will be borderline relevant – in which case I now know to follow the rule I have given myself in the Split-attention section – but more often than not they will be quite the opposite: 'Jenny, how is your sister getting on at uni?', 'Cal, any chance of you getting a haircut soon?'.

Now, there is little doubt that this is redundant information. It is not necessary in order for students to understand or solve what they are working on, and yet it must be processed by their fragile working memories. However, isn't this the kind of thing that helps us to get to know our students better, helps them feel comfortable, helps the development of

relationships that lead to them enjoying our lessons, trusting us, feeling comfortable making mistakes, and being willing to take risks? When I interviewed him on my podcast, Dylan Wiliam explained that forming strong relationships with students is one of the most important things a teacher can do.

But here is the thing – there is a time and a place. And during the early knowledge acquisition phase, or when students' working memories are pushed to the limits working on a tricky problem, is definitely not it. And I know I am more guilty than anyone of this sin because the first day I made a conscious effort to cut down on my redundant talk, one of my Year 11s asked me if I was feeling okay.

I now choose moments for my mathematically irrelevant talk very carefully indeed.

Fluff

Back in Section 2.3, we looked at some of the dangers of shoehorning real-life contexts into the maths classroom. One of those dangers is directly related to the Redundancy Effect. If a problem is wrapped up in a real-life context, then the specifics of that real-life context must be processed in working memory. Jo might be going to the shop, or Bruno might be paving his driveway, and our students have to take all that in, or effortfully filter it out.

Now, that is all well and good, but in early knowledge acquisition phase I feel we need to ask ourselves, what are we actually hoping to achieve? Do we want our novice learners to attempt to navigate their way through the largely irrelevant surface features of a problem and identify the deep structure? Or do we want them to focus entirely on working with the deep structure of a problem in order to build up fluency and mastery of the skills needed to solve it? Sure, the former ability is vital, but as we shall see in Chapter 9, it must be preceded by the latter.

So, in the early knowledge acquisition phase of learning, I like to remove all context, all fluff, all attempts at a real-life scenario, and focus purely on the mathematics. The same is true of the stars, banners and quirky clip-art that

used to litter my PowerPoint slides. This is related to the Coherence Principle (*eg* Mayer, 2008) which states that 'people learn more deeply when extraneous material is excluded rather than included'. Leaving out this fluff gives our students, and their fragile working memories, the best chance of acquiring the inflexible knowledge that for Willingham (2002) is a necessary step along the road to expertise. After all, students remember what they attend to.

Classroom Displays

I am trying to start a movement, and so far there is only me in it. Part of the problem is I don't really have a catchy acronym – the best that I can come up with is, rather ironically, BAD. The other issue is that it involves taking something from teachers that they hold very dear.

Yes, I am proposing we Ban All Displays.

Now, I love nothing more than classroom walls overflowing with wonderful examples of the bright and colourful work my students have done. And my students love it too. Even the most mathematically reluctant take a certain amount of pride from having their finest work on display for all to see.

However, all that bright and colourful work is pretty hard to ignore, especially when it is either in the direct or peripheral vision of students looking towards the front of the classroom, and even successfully filtering it out is effortful. So if, during a particularly tricky worked example of how to draw a sample space diagram, Matthew's attention gets drawn towards Olivia's lovely piece of work on visual representations of sequences, then Matthew may find he suddenly lacks the cognitive capacity in his working memory to successfully process all he needs to about sample space diagrams.

The message is simple and unpopular: fill your boots on corridors, but classroom displays should be placed at the back of the classroom, or ideally removed completely. We need to reduce all forms of information that are not pertinent to the matter in hand.

When I present this final implication of the Redundancy Effect to teachers, there are three common responses:

1. Things like number lines, fraction walls and mathematical definitions are really useful.
But are they useful for all topics? Are they useful in a lesson on angle facts, estimating the mean and indices? If not, then there is a chance students will be thinking about material not directly related to the concept at hand. If they are relevant, then we can simply put the display back up, or project it on the board.

2. Good displays help students remember key information.
I am not entirely sure this is true. Sure, it is great to have a display containing the first 100 prime numbers, or a poster with key area formulae on it. But I think these can also make students dependent on having such information readily to hand. In Chapter 12 we will look at the crucial importance of retrieval for the learning process. If students are using these displays as a crutch – they know they are there, so what is the point in trying to remember the information? – then they are not forcing retrieval from long-term memory. It is like students revising by going over notes and examples – the reading of them feels familiar, so students think they have committed them to memory. It is only when they find themselves in a situation where that support is not available (like an exam) that students realise they have not committed the information to memory at all. We need to give our students opportunity to go through the effortful process of attempting to retrieve information from their long-term memories, because it is this process of retrieval that leads to learning and retention. Hence, I prefer to have the information from the displays handy – print-outs or Knowledge Organisers are great for this – so I can give them out as and when they are needed. But I want students to have the opportunity to induce retrieval of the information first.

3. My Senior Leadership Team will not like it.
They do not like a lot of things, but that does not make it right.

I much prefer 'Whiteboard Walls'. They can be as simple as whiteboards stuck around the classroom, or as sophisticated as special paint that means you can write directly on the walls with pens and then rub off. Students have an opportunity to showcase their work, I can model work to small groups, the work is always directly relevant to the lesson, it can be easily rubbed off

when it is not needed, it is always changing and hence always fresh, and perhaps best of all it allows students to take advantage of Chapter 6's Worked Example Effect.

What I do now

In short, I look to identify all sources of redundant information and remove them from my instruction and classroom materials. Sure, students will need to deal with certain types of redundant information at some stage, but that does not and should not be in the early knowledge acquisition phase of learning.

OceanofPDF.com

4.8. Silent teacher

What I used to think

I first became aware of the principle of The Silent Teacher in about 2012. Basically, it involves a teacher modelling how to do something – for example, solving a linear equation – but in absolute silence. The students would watch, also in silence, and at the end of the demonstration there would be an opportunity for the students to ask questions such as, 'Why did you divide by 2 at that stage?'.

For a while, Silent Teacher was all the rage. Everyone was doing it: maths teachers, English teachers, even the school chaplain. And then, just like Collective Memories, PLTS, and numerous other 'next big things', it seemed to disappear into the educational abyss.

Has the time come to bring it back?

Sources of inspiration

- Mousavi, S. Y., Low, R. and Sweller, J. (1995) 'Reducing cognitive load by mixing auditory and visual presentation modes', *Journal of Educational Psychology* 87 (2) pp. 319-334.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Considering our work on the Split-attention and Redundancy Effects, it turns out that the Silent Teacher approach has significant merit in terms of Cognitive Load Theory and the Cognitive Theory of Multimedia Learning.

Let's take the example of solving a linear equation, such as $3x - 4 = 17$

Now, a traditional way to model this may be to write down the steps and narrate over the top. Something along the lines of:

'Okay, so we have $3x - 4 = 17$. Now, I need to get the x all on its own, so the first thing I am going to do is I am going to add 4 to both sides of the equation. This leaves me with $3x$ on the left-hand side, and 21 on the right.'

Let's analyse what is likely to be happening in the students' working memories whilst this is happening. Students are looking at a board with $3x - 4 = 17$ written on it. This is processed exactly as text would be, first in the Visuospatial Sketchpad, and then essentially repeated aloud in the students' minds as they say '*three-ex-minus-four-equals-seventeen*' where is it dealt with by the Phonological Loop. Now, when considering my narration and writing down of the steps, we see we in fact have instances of both redundancy and split-attention. Me reading out the equation is redundant, and students alternating their attention between processing the written words and symbols, together with my narration, is likely to cause split-attention. All of this is more cognitively demanding than it needs to be.

Instead, I could present the worked solution in absolute silence, thus reducing the load imposed on students' Phonological Loops. Moreover, it allows me to divert all of my attention to what I am writing instead of concentrating on writing and talking concurrently.

Having finished my silent demonstration, I then have the option to either encourage my students to ask questions about my worked example, or to add my narration over the top. As we will see in Chapter 6, I favour the latter.

When I pitched the idea of Silent Teacher to Peps Mccrea during our podcast interview, he likened it to being at the cinema. Before the film starts, people are chatting away, playing with their phones, spreading their limited attention across numerous mediums. But once the lights go down, silence descends, and attention is firmly fixed on the screen in front. Presenting examples in silence can have the same effect. Indeed, to take Silent Teacher

to the next level, we could close the blinds, turn off the lights, and let the show begin!

An added advantage of the Silent Teacher approach is with regard to behaviour. When presenting worked examples where student discussion is permitted, it can often be difficult to discern if a discussion going on between two students is related to the content of the example or not. 'Sir, I was just helping Josh with the question', is a response that leaves us poor teachers in a bit of a quandary. With Silent Teacher there is no grey area. Students know they have to be in absolute silence. I am certainly no behaviour expert, but the creation of that clear boundary seems to make everyone's lives a lot easier.

What I do now

I now make extensive use of Silent Teacher, particularly when first introducing a concept. It fits in really nicely with the process for presenting worked examples that I will discuss in Chapter 6, and the students seem to enjoy it. It appears this forgotten 'next big thing' is due a resurgence.

OceanofPDF.com

4.9. Germane Load

What I used to think

So far the focus of Cognitive Load Theory has been on making thinking as easy for students in the sense that all their limited working memory capacity should be focused entirely on the thing we want them to attend to in order to prevent cognitive overload occurring. In other words, as teachers we should try to do two things:

1. Reduce intrinsic load by helping students acquire the background knowledge necessary so that they may chunk and automate the sub-components of a complex task.
2. Eliminate unhelpful extraneous load via the careful presentation of information and Chapter 6's worked examples.

However, if we successfully reduce these two types of loads, what exactly do our students' working memories do with that spare capacity?

Sources of inspiration

- Mccrea, P. (2017) *Memorable teaching: leveraging memory to build deep and durable learning in the classroom*. CreateSpace Independent Publishing Platform.
- Paas, F., Renkl, A. and Sweller, J. (2003) 'Cognitive load theory and instructional design: recent developments', *Educational Psychologist* 38 (1) pp. 1-4.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

- Van Merriënboer, J. J. G., Kester, L. and Paas, F. (2006) 'Teaching complex rather than simple tasks: balancing intrinsic and germane load to enhance transfer of learning', *Applied Cognitive Psychology* 20 (3) pp. 343-352.

My takeaway

Mccrea (2017) sums up this issue beautifully:

> Our working memory is a high maintenance mechanism. Give it too little to play with and it begins to look for more interesting fodder. Give it too much to juggle and it'll drop all the balls.

If we have successfully freed up space in students' working memories, we had better fill it up with something useful, otherwise information not related to the concept at hand may fill the gap and all our efforts will have been in vain.

Van Merriënboer *et al* (2006) address this problem by introducing the concept of 'germane load'. Germane load can be viewed as 'good cognitive load', in that it directly contributes to learning by aiding the construction of cognitive structures and processes – the schemas from Chapter 1 – that improve knowledge.

The authors found that whilst reducing extraneous load is effective in producing high retention of the material, these techniques alone do not allow students to transfer their knowledge to new situations. They argue that there is a need to do things like vary the conditions of practice and only give limited guidance and feedback in order to induce germane cognitive load and improve transfer. In other words, in order to improve learning – in particular the transfer of knowledge to new contexts – we need to make learning more difficult…but difficult in the right way! This concept is closely related to Bjork's fascinating idea of 'desirable difficulties' that will be discussed in Chapter 12.

There is the danger that including the concept of germane load makes Cognitive Load Theory impossible to falsify. For example, assuming that the overall load is kept constant, a decrease in performance could be attributed

to a rise in extraneous load that impairs germane cognitive processes. Conversely, if the performance increases, it could be attributed to a germane load enhancement made possible by a drop in extraneous load. The theory can't lose!

Building in the concept of germane load might well be making Cognitive Load Theory unnecessarily complicated, but it feels important to include some measure of positive cognitive load. My view is this: at all times, but particularly during early knowledge acquisition, we need to ensure our students' thinking is as focused on the task at hand as possible, and we can achieve this using all the principles we have discussed in the sections above. But we need to ensure that thinking is not too easy. If students are cruising through lessons on autopilot, then their learning is unlikely to be deep, and learning without the ability to transfer it to new situations is not really learning at all. If we have taken measures to free up capacity in working memory, but that capacity is not subsequently used to think hard about the task or concept in hand, then we have essentially wasted our time.

It is a fine balance. Make thinking too easy, and the necessary changes to long-term memory may not occur for learning to take place. Likewise, make thinking too hard and we risk cognitive overload, where again no learning takes place. Flipping heck, being a teacher is a tough job.

It is also worth reiterating the view of Paas *et al* (2003) that intrinsic, extraneous and germane load are thought to be additive. Hence, the approach of decreasing extraneous cognitive load while increasing germane cognitive load will only be effective if the total cognitive load remains within the limits of working memory.

What I do now

The first thing to say is that I always, always take steps to reduce extraneous load. Such load fills up students' fragile working memories in a way that is not conducive to learning, and hence is completely unnecessary at all stages of the learning process. But I am also conscious that students need to be thinking hard to be learning. This is the essence of Cognitive Load Theory – getting students to think hard about the right things in order to facilitate the change in long-term memory necessary for learning to occur.

Sweller *et al* (1998) explain:

> The combination of decreasing extraneous cognitive load and at the same time increasing germane cognitive load involves redirecting attention: Learners' attention must be withdrawn from processes that are not relevant to learning and directed towards processes that are relevant to learning and, in particular, toward the construction and mindful abstraction of schemas.

Practical ways we can achieve the right balance of germane load will be addressed throughout the remainder of this book.

OceanofPDF.com

4.10. If I only remember 3 things…

1. We want students to be thinking hard about the right things, so we should take steps to reduce extraneous load.
2. Students can be thinking hard but not learning if they experience cognitive overload.
3. Sounds and images are processed and held in separate challenges in working memory, which we can use to our advantage via the careful presentation of information.

OceanofPDF.com

5. SELF-EXPLANATIONS

This brief chapter serves two purposes. It provides the rationale for the Supercharged Worked Examples that I will introduce in Chapter 6. But it also stands alone in its own right, for until I began reading about both the power and subtleties of student self-explanations, I had no idea just what an important part of the learning process they can be.

OceanofPDF.com

5.1. The Self-Explanation Effect

What I used to think

Student explanations have always featured heavily in my lessons. If you had asked me a couple of years ago who the main beneficiaries of these explanations are, I would have said me, their teacher. Asking students to expand upon an answer was a way of me assessing the depth of their understanding – a tool of formative assessment. And if you had asked me the most effective way to elicit these explanations, I would have said you need to ask students to articulate their thoughts, obviously.

It turns out that I was making not one, but two mistakes.

Sources of inspiration

- Barton, C. (2017) 'Ed Southall – Part 2', *Mr Barton Maths Podcast*.
- Chi, M. T. H. (2000) 'Self-explaining: The dual processes of generating inference and repairing mental models' in Glaser, R. (ed.) *Advances in instructional psychology: educational design and cognitive science, Vol. 5*. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 161-238.
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J. and Willingham, D. T. (2013) 'Improving students' learning with effective learning techniques: promising directions from cognitive and educational psychology', *Psychological Science in the Public Interest* 14 (1) pp. 4-58.
- Renkl, A. and Atkinson, R. K. (2003) 'Structuring the transition from example study to problem solving in cognitive skill acquisition: a cognitive load perspective', *Educational Psychologist* 38 (1) pp. 15-22.

- Rittle-Johnson, B. (2006) 'Promoting transfer: effects of self-explanation and direct instruction', *Child Development* 77 (1) pp. 1-15.
- Siegler, R. S. (2002) 'Microgenetic studies of self-explanation' in Granott, N. and Parziale, J. (eds) *Microdevelopment: transition processes in development and learning.* Cambridge: Cambridge University Press, pp. 31-58.

My takeaway

Before we delve too deep into the benefits of student self-explanations, we must first establish exactly what we mean by the term. According to Chi (2000), self-explaining refers to:

> The activity of generating explanations to oneself … It is somewhat analogous to elaborating, except that the goal is to make sense of what one is reading or learning, and not merely to memorise the materials. In this sense, self-explaining is a knowledge-building activity that is generated by and directed to oneself.

This is the first thing I had been missing. Self-explaining is not explaining concepts to others – indeed, later on in this chapter we will see why explaining to others is potentially a dangerous thing to do during early knowledge acquisition phase – nor is it necessarily saying anything out loud. It is the simple act of pausing and reflecting on a single step in a solution, a concept or an explanation. It is asking yourself, 'what does this mean?', 'why am I writing this?' and 'how does this step follow on from the last?'.

Secondly, the impact on learning of such self-explanations can be profound. Indeed, the Self-Explanation Effect – whereby learners who attempt to establish a rationale for the solution steps by pausing to explain the examples to themselves appear to learn more than those who did not – has been observed and replicated across multiple domains and ages of students.

In their review of effective learning techniques, Dunlosky *et al* (2013) rate self-explanation to have moderate utility. They conclude that self-explanation effects 'have been shown across an impressive range of learning

outcomes, including various measures of memory, comprehension, and transfer'.

In a study involving students studying maths from ages 8 to 11, Rittle-Johnson (2006) found that self-explanation had a particularly profound effect on problem-solving. She found that students who self-explained were more likely to invent new problem-solving approaches, apply and transfer correct procedures to a broader range of problems, and retain correct procedures over a two-week delay.

Similarly, Siegler (2002) suggests one contributing factor to the high levels of maths performance of Japanese students compared to their US and English counterparts is the emphasis on generating explanations for why mathematical algorithms work. In Japanese classrooms, both teachers and students spend considerable time trying to explain why solution procedures that differ superficially generate the same answer, and why seemingly plausible approaches yield incorrect answers. Indeed, on my podcast you can hear Ed Southall discussing in detail a lesson on sequences he observed during his trip to Japan that is filled with discussion and self-explanations. Encouraging children to explain why the procedures work appears to promote deeper understanding of them than simply describing the procedures, providing examples of how they work, and encouraging students to practise them.

Hence, self-explaining can actually *cause* learning, and is not just an indication of understanding.

But what actually explains the Self-Explanation Effect?

Well, in the language of Cognitive Load Theory (Chapter 4), the act of self-explaining is likely to increase cognitive load – as it is more cognitively demanding than simply copying down or carrying out a procedure – but in a way that contributes to schema acquisition. Hence, self-explanations may be described as *germane load* (see Section 4.9). For Siegler (2002), self-explanations lead to students searching more deeply for explanations when such efforts are made, increasing the accessibility of effective strategies

relative to ineffective ones, and increasing the degree of engagement with the task.

Chi (2000) identifies two main causes of the Self-Explanation Effect, both of which have practical applications for teachers of mathematics:

1. Generating inferences
This involves the learner inferring information that is missing from an example's solution. This will form the basis of the Supercharged Worked Examples that we will look at in the next chapter.

2. Repairing the learner's own mental model
Here it is assumed that the learner engages in the self-explanation process if he or she perceives a divergence between his or her own mental representation and the model conveyed by the example's solution. This then helps students correct any erroneous reasoning. This is one of the reasons I always give my students access to answers to classwork, homework and low-stakes quizzes, and will be discussed further in Section 8.4.

What I do now

I now pay much more notice to self-explanations as a learning tool, not just as an indication of student understanding. I use the findings discussed in this chapter and the next to make self-explanations as beneficial as possible.

OceanofPDF.com

5.2. Making the most of self-explanations

What I used to think

Seeing as I had never really considered the power of self-explanations as a learning tool, it is of little surprise that I did not consider ways to make self-explanations as effective as possible. As it happens, I was missing out, big time.

Sources of inspiration

- Chi, M. T. H. (2000) 'Self-explaining: The dual processes of generating inference and repairing mental models' in Glaser, R. (ed.) *Advances in instructional psychology: educational design and cognitive science, Vol. 5*. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 161-238.
- Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P. and Glaser, R. (1989) 'Self-explanations: how students study and use examples in learning to solve problems', *Cognitive Science* 13 (2) pp. 145-182.
- McEldoon, K. L., Durkin, K. L. and Rittle-Johnson, B. (2013) 'Is self-explanation worth the time? A comparison to additional practice', *British Journal of Educational Psychology* 83 (4) pp. 615-632.
- Renkl, A. (1997) 'Learning from worked-out examples: a study on individual differences', *Cognitive Science* 21 (1) pp. 1-29.
- Siegler, R. S. (2002) 'Microgenetic studies of self-explanation' in Granott, N. and Parziale, J. (eds) *Microdevelopment: transition processes in development and learning.* Cambridge: Cambridge University Press, pp. 31-58.

My takeaway

Just like we observed in Chapter 4 when looking at the presentation of information, it turns out that with a few relatively straightforward tweaks we can make the Self-Explanation Effect even more powerful.

1. Explaining someone else's answer
Because I used to believe that the main benefit of self-explanations was for me to get a sense of my students' understanding, I thought that it was better for students to explain their own answers. After all, who better than the students themselves to explain their own thinking, and thus give me an insight into their thought processes? Therefore, I would let students explain their own answers far more than I challenged them to explain others.

However, Siegler (2002) found that children who were randomly assigned to explain the experimenter's reasoning learned more than children who explained their own reasoning. Studies conducted in other laboratories have shown that the people whose reasoning is being explained need not be present for the positive effects to emerge. Encouraging children and adults to explain the reasoning that they encounter in textbooks has similar benefits.

Hence, now I devise answers and examples that I have created and challenge my students to explain them – firstly to themselves during a moment of silent contemplation, and only then as part of a wider discussion. This has the added advantage of allowing me to carefully construct and present examples exactly how I need in order to draw out their key features and highlight common misconceptions. Sometimes I challenge students to explain each other's answers and methods, but I tend to anonymise the work to avoid any unnecessary anxiety and distractions.

2. Explaining if something is wrong or right
I was a little wary about using incorrect examples in case an incorrect procedure was reinforced.

However, one of Siegler's (2002) key findings is that 'explaining why correct answers are correct and why incorrect answers are incorrect yields greater learning than only explaining why correct answers are correct'.

So now I encourage my students to explain why things are both right and wrong, using examples I have carefully prepared. However, I tend to leave this discussion until students have practised the correct method so they have a sense of what is right and what is wrong, and I always make it crystal clear what the correct answer is at the end of this process lest it lead to confusion.

Further discussion and considerations about the use of incorrect examples and practical techniques for their use can be found in Chapters 6, 7 and 11.

3. Explaining to someone else

I used to think it was a good idea to get students to explain their reasoning to someone else. After all, surely both parties benefited – the student on the receiving end of the explanation essentially had a personal tutor helping them understanding a concept, and the person giving the explanation was compelled to explain, discuss and argue their way of thinking, thus strengthening their own understanding. Hence, I would challenge students to 'convince someone else you are right' on a regular basis.

Chi (2000) makes the point that the focus when self-explaining is simply to understand or make sense of something, whereas the purpose of talking or explaining to others is to convey information to them. Talking or explaining to others adds the requirement to the explainer of monitoring the listener's comprehension, essentially assessing their levels of understanding and adjusting and adapting where needed. It seems reasonable to assume that cognitive capacity is consumed through talking, noticing and responding, especially in a novice learner who has just encountered a concept, which may prevent that student from engaging in critical self-explaining behaviours and thus benefiting from the process. In such a scenario, the explanation given is likely to be muddled and unclear, so neither party may benefit.

These days I am much more selective in my use of 'convince me'. I like the process of 'convince yourself, convince a friend, convince a sceptic', but only if I am sure students have actually convinced themselves. A well-designed diagnostic multiple-choice question (Chapter 11) might be one way to quickly assess this, followed by asking students not just to indicate who got it right, but also who feels confident explaining to others.

The same cognitive strain is likely to be experienced when asking students to explain their answer to me. Hence, in order to increase the chance of students having engaged in a process of personal, silent self-explanation, I force myself to pause after asking a question, calling for a brief period of silence, before eliciting any responses.

4. Self-explanations v more practice

We have seen the benefits of student self-explanation throughout this chapter, but there is no doubt that it takes longer to explain and answer a question than just to answer it. When you combine this observation with the clear benefits of regular practice, it raises the obvious question: is self-explanation worth the time, or should we just get our students to practise more instead?

McEldoon *et al* (2013) attempt to provide an answer. The authors compared the effectiveness of self-explanation prompts to the effectiveness of solving additional practice problems in primary school maths students. Students were placed into three groups: the *self-explain group* solved six problems and were prompted to self-explain after each; the *control* group also solved six problems but were not prompted to self-explain; and the *additional practice* group solved 12 problems and were also not prompted to self-explain.

The authors found that relative to the control condition, the self-explain condition supported greater conceptual understanding, particularly of equation structures, and greater procedural fluency, particularly for procedural transfer. However, relative to additional practice, the self-explain condition had modest benefits. The findings suggest that self-explanation prompts have some small unique learning benefits, but that greater attention needs to be paid to how much self-explanation offers advantages over alternative uses of time.

The obvious conclusion is that a balance is needed, but it is worth bearing in mind that self-explanation does not need to be overly time-consuming. It is the act of pausing and reflecting on an answer or part of a procedure – often for no more than a few seconds. It does not necessarily need articulating or discussing – and in fact we have seen that doing so may indeed impose a significant load upon working memory. Hence, if we can encourage our

students to pause briefly to reflect, then they may enjoy the many benefits of self-explaining while still having plenty of time for additional practice. In Section 7.8 we will see how Intelligent Practice is designed precisely for this, and in Section 8.4 why it is vital students have answers to classwork for the same reason.

5. Spontaneous Self-Explainers
Renkl (1997) finds that the majority of learners do not spontaneously engage in successful self-explanation strategies. He explains:

> The finding that more than half of the subjects had to be assigned to the group of unsuccessful learners, reaffirms research findings that learners, left to their own devices, typically fail to show effective learning behaviors when no external support (*eg*, teacher guidance or scaffolding) is present.

A key factor may be prior knowledge, with a study by Chi *et al* (1989) concluding that high-ability students tend to spontaneously self-explain more often than low-ability students.

This is crucial, as there is a danger in assuming that any spare working memory capacity will be used up for things that contribute towards learning – *ie* germane load. If this is not the case, and students do not naturally engage in self-explanations, then their learning will not be as effective as it could be.

This has significant implications for the classroom, and suggests that we should prompt students to provide self-explanations during instruction. This important finding will play a key role in the Supercharged Worked Examples that follow in the next chapter, and is also the rationale for the Counter of Hope…

6. The Counter of Hope
If I set my students off on a challenging activity, I like to give them each a counter – the Counter of Hope. This counter can be exchanged at any time during the activity for a hint from me. If students are working in a group, then it tends to be one Counter of Hope per group. I make it crystal clear to students that once they have used up this counter, then there is no help from

me at all. This simple act has quite a profound effect on students – it causes them to think really carefully about the help they seek. In the past, when many of my students have encountered a problem, their immediate response was to either give up or to ask me for help. The fact that they know help is available via the Counter of Hope alleviates the issue of giving up, but because they now know they can only access such help once during the activity, they are far more likely to pause and reflect when they are stuck. In other words, they are more likely to self-explain. The effect is even stronger when working in groups. Here, I have seen one student put up their hand to ask for help, and another member of the group virtually climb across the desk to yank her teammate's hand back down, uttering 'ask me first, ask me first, we don't want to waste our counter'. Anything that compels students to persevere, think hard and work together positively is good with me, and I am perpetually surprised how much influence a small circular piece of plastic has over Year 7s and Year 13s alike.

What I do now

In order to take full advantage of the Self-Explanation Effect, I:

- ensure I provide students with opportunities to self-explain carefully chosen worked solutions.
- make sure that some of these solutions are incorrect, clearly indicating when this is the case.
- only challenge students to explain to someone else later on in the learning process.
- pause before asking students to explain to me.
- ensure that self-explanations do not slow down the pace of the lesson and do not come at a price of additional practice.
- prompt students to self-explain using Chapter 6's Supercharged Worked Examples.
- use the Counter of Hope to prompt students to self-explain and be more selective over the requests for help.

OceanofPDF.com

5.3. If I only remember 3 things…

1. The Self-Explanation Effect suggests that self-explanation can be a powerful learning tool.
2. There are several easy-to-implement steps we can take to make the most out of these self-explanations.
3. Students are not natural self-explainers and hence may need to be prompted.

OceanofPDF.com

6. MAKING THE MOST OF WORKED EXAMPLES

Eagle-eyed Cognitive Load Theory connoisseurs may have noticed that in Chapter 4 I left out a major finding from the theory – The Worked Example Effect. But I did so with good reason. The increased emphasis of examples, both in my choice of them and how I present them to students, has been one of the most profound changes in my teaching. As such, it deserves not just one chapter dedicated to it, but two.

In this chapter we will look at why worked examples are so important, and how to present them to students for maximum effectiveness. In Chapter 7, I will turn my attention towards the choice of worked examples and student practice questions.

Just before we begin, I want to emphasise that my lessons do not start with worked examples. They always start with a low-stakes quiz (Chapter 12) followed by a diagnostic assessment of prerequisite knowledge (see Chapter 11). Then, I will look to introduce the concept. This may be via the kind of interactive demonstration described in Section 3.4, or by providing a purpose as discussed in Section 2.5, or utilising key elements of a story structure (Section 3.6). If maths is seen as an endless cycle of worked examples and practice, presented cold and in isolation from the rest of the subject, it is unlikely to be a pleasant experience for anyone.

OceanofPDF.com

6.1. The Worked Example Effect

What I used to think

I used to believe that worked examples had limited use, and it was important to get students practising on their own as quickly as possible. My reasoning was – at least in my own head – very sound. As a trainee teacher I was told of two cardinal sins: talking too much and not differentiating. A lesson crammed full of worked examples would fall victim to both of these. It would be me, the teacher, dominating the lesson, while my students passively moved through the lesson all at the same pace.

No, too much emphasis on worked examples was definitely a bad thing. Far better to whizz through a few and get students onto their own independent work as quickly as possible. That way, the high-flyers could race ahead, and I would be freed up to wander around the class and help the students who were struggling. Surely that was good – if not outstanding – teaching?

The problem was, many of the high-flyers were perhaps not flying quite as high as either they or I perceived them to be, becoming stuck as soon as things got a little tricky. Likewise, I ended up repeating the same things to the students who were struggling – things, in fact, that the whole class would probably have benefited from hearing.

Sources of inspiration

- Atkinson, R. K., Derry, S. J., Renkl, A. and Wortham, D. (2000) 'Learning from examples: Instructional principles from the worked examples research', *Review of Educational Research* 70 (2) pp. 181-214.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational*

Psychology Review 10 (3) pp. 251-296.

My takeaway

The 'Worked Example Effect' is the name given to the widely replicated finding that novice learners who try to learn by solving problems perform worse on subsequent test problems, including transfer problems different from the ones seen previously, than comparable learners who learn by studying equivalent worked examples (Sweller *et al*, 1998; Atkinson *et al*, 2000).

When I first came across this, I didn't believe it. Surely you learn more from trying to solve problems than by just studying the solution? But the key is in the *trying* to solve problems. As we have seen in Section 4.5 on goal-free problems (and will revisit in Chapter 9), when novice learners try to solve problems they tend to do so using a cognitively demanding strategy that places a heavy burden upon their fragile working memory. They get so bogged down in the finer details of the problem that they are unable to appreciate more global issues, such as the problem's deeper structure or the successful strategy they employed, thus inhibiting the ability of the learner to successfully process and then transfer the information into their long-term memory.

In contrast, studying a worked example reduces the burden on working memory, because the solution only has to be comprehended, not discovered. Moreover, worked examples direct attention – or, in other words, spare cognitive capacity in working memory – toward storing the essential relations between problem-solving moves. Hence, students learn to recognise which steps are required for particular problems, which is the basis for developing knowledge that can be transferred to other situations.

So, by moving students away from worked examples and on to problems to solve independently too soon, I was in danger of inhibiting their learning.

What I do now

Three things:

1. Most importantly, I pay a lot more attention to worked examples, both in terms of their delivery (this chapter) and my choice of those examples (Chapter 7).

2. I am a lot more reluctant to move students away from worked examples and on to independent work as quickly as I used to do – we will address the implications of this for differentiation and holding students back at the end of this chapter.

3. As well as having the answers to the classwork available to the students (Section 8.4), I also like to have a few full worked examples of the solutions to hand. Section 4.7's Whiteboard Walls are ideal for this. That way, if students are struggling they can read over the worked example at their own pace, observing not just the solution, but the structure and path to the solution, which should aid in the formation of those all-important schemas needed for learning to take place.

OceanofPDF.com

6.2. Example-Problem Pairs

What I used to think

For many years my so-called standard lessons followed this format:

1. Starter
2. Worked examples
3. Practice of basics
4. Application problems
5. Plenary

Of course, if I was being observed, then I would whip out something much more fancy, but that is beside the point.

I would batch the worked examples together at the start of the lesson and then set my students off on independent work. My rationale was two-fold: get the worked examples out of the way as soon as possible so the students can get cracking at their own pace; and batch them so students could experience a coherent sequence of examples through a given concept from start to finish. Likewise, I would involve the students heavily during the presentation of worked examples, asking them what to do next, provoking discussion and debate.

However, I tended to find that by the time I had gone through four or five worked examples – something that could easily take up to 20 minutes or more because of student involvement – and then set the students to work, many had no clue whatsoever how to do the first question. The remainder of the lesson would be spent going round from group to group, having the same conversations over and over again.

Not ideal.

Sources of inspiration

- Atkinson, R. K., Derry, S. J., Renkl, A. and Wortham, D. (2000) 'Learning from examples: Instructional principles from the worked examples research', *Review of Educational Research* 70 (2) pp. 181-214.
- Barton, C. (2017a) 'Doug Lemov', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Greg Ashman', *Mr Barton Maths Podcast*.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Rosenshine, B. (2012) 'Principles of instruction: research-based strategies that all teachers should know', *American Educator* 36 (1) pp. 12-39.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.
- Trafton, J. G., and Reiser, B. J. (1993) 'The contributions of studying examples and solving problems to skill acquisition', *Proceedings of the 15th Annual Conference of the Cognitive Science Society*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 1017-1022.

My takeaway

Trafton and Reiser (1993) present two major findings, which, whilst they may seem obvious, are both of crucial importance. They conducted experiments designed to answer two key questions they had about the most effective way to learn from worked examples.

1. Does separating source examples from target problems hamper learning?
Participants who solved a target problem immediately after the source example (Alternating Example) were compared to those who studied a block of source examples followed by a block of solving target problems (Blocked Example). Subjects who solved problems interleaved with examples took less

time on the target problems than subjects who studied a block of source examples and a block of target problems. Crucially, participants in the Alternating Example condition also submitted more accurate solutions than subjects receiving blocked examples.

2. Is solving problems better than studying examples if the examples are not accessible during subsequent problem-solving?
There is a danger that we can get too caught up in the power of examples, and think that students do not need to solve any problems at all. In fact – and unsurprisingly – the researchers found that participants who attempted to solve problems as well as studying worked examples performed better than those who merely studied worked examples. As we shall see in Chapter 12 when considering the power of tests, it is this retrieval process induced when answering a question that is so important to learning, and if students are never compelled to access their memories of those examples, then they will never benefit from what we will come to call the Testing Effect. Hence, we can conclude that subsequent independent practice appears to be required to derive the full benefit from studying examples.

Atkinson *et al* (2000) provide a practical solution to these to questions. The authors recommend interleaving worked examples with related questions for students to solve alone. It matches Rosenshine's (2012) second principle of instruction: 'Present new material in small steps with student practice after each step'. It is an approach discussed in detail by Greg Ashman when I interviewed him for my podcast, and which I will offer my own take on below.

What I do now

When doing a worked example in class, I follow the **Example-Problem Pair** approach.

I split my board in two, having the worked example on the left, and a mathematically similar example for my students to try themselves immediately afterwards on the right.

My board initially looks something like this:

Worked Example	Your Turn
$\frac{3}{5} + \frac{1}{4} =$	$\frac{2}{3} + \frac{1}{5} =$

Figure 6.1 – Source: Craig Barton

I am then very particular about how I deliver this example-problem pair, and I fear some of you are not going to like what I am about to say.

In the past, I have involved students lots during the course of a worked example, asking them questions like: 'What should I do next?', 'What do you think the next line I write is going to be?', 'How do I know this is correct?', and so on. But now I don't. Why? Because these are great questions to induce retrieval of something students already know, but terrible questions when attempting to teach something for the first time. To see what I mean, consider the three possible scenarios that can result from any of these questions:

1. The student you ask gets it exactly right.
Brilliant, but so what? What does that tell you about the other 29? The perils of asking one-to-one questions will be discussed further in Chapter 11.

2. The student gets it wrong.
I used to love this, because it signalled the start of a whole-class discussion. I would put on my best poker face, write the given answer down on the board, and then see if anyone had anything to add. Students would suggest other answers, which I would also write down. Every answer was treated equally. I would then invite students to argue their case, and put it to a class vote.

The first thing to say is that this took ages. But I could cope with that if I felt it lead to better learning. However, I think it just led to confusion. Students struggled to articulate what they meant. Questions, statements and accusations bounced around the room. And during the discussion, something interesting happened. I could see some students – those who were pretty sure they knew how to do it before – either getting frustrated or growing a little puzzled, and the students who were confused to begin with were having their fears and negative perceptions of maths further stoked. Indeed, during one of these discussions, a girl who found maths particularly challenging once screamed at me: 'Just tell me the right answer, sir, please!'. I would like to take this opportunity to say: Georgia, I am sorry.

At the end of the process I would inevitably end up explaining the answer myself, but for many students the damage was already done.

3. The student sort of gets it right.
Now this one is interesting. Taking the example of the fractions above, a student may say something like, 'You have to make the bottoms the same. I just times them together'. Now, this response is not incorrect by any means, but there are two issues with it.

First, I would prefer the student said 'denominators' as opposed to 'bottoms'. This returns us to the fascinating topic of technical language in mathematics that we touched upon in Section 3.4. If I am going to insist on the use of such technical language, I need to be sure that every single child in the class can link 'denominator' to the 'the number on the bottom' as quickly, directly, and with as little thought (*ie* none) as they can link the word 'bottom' to… well, 'bottom'. Otherwise, imagine if a student in the class had to pause – for even just a second or two – in their mind to recall whether that meant bottom or top. Best-case scenario, they are probably going to fail to follow what comes next due to cognitive overload. Worst case, they misremember it as meaning 'top', and things go really wrong. This has serious implications for how and when we teach technical language in general. Obviously we want our students to be comfortable with terms such as numerator and denominator, but also we cannot allow learning those to impede learning to add fractions, especially when it is something that can be taught completely separately to that skill. Therefore, in some situations, I may actually prefer

the use of 'bottoms', as I can be sure every student knows what that means, so there will be no barrier to me teaching them how to add fractions. I can then deal with terminology issues at a later date once students are developing a good understanding and experiencing a growing confidence in adding fractions. Either way, that is my call to make, and I would prefer to be in control of the technical language used, rather than to leave it to chance based on what comes out of a student's mouth.

The second issue is that the method described by the student is fine for this particular example, but is problematic when the denominators are large numbers which have a more manageable lowest common multiple, such as 12 and 18. As such, I don't really want other students to hear the relatively inflexible procedure suggested by the response. For those students who are not comfortable adding fractions – and if this is the first lesson on the topic, that is likely to be most of the class – it is of paramount importance that they hear a correct procedure, preferably one that has been carefully planned out long before the lesson.

But I cannot let this incomplete explanation go. After all, it is out in the open now. So what can I do?

Lemov (2015) calls a common teacher technique in dealing with such responses 'rounding up'. Rounding up involves a teacher responding to a partially or nearly correct answer by affirming and repeating it, but then also adding critical detail (perhaps the most insightful or challenging detail) to make the answer fully correct. Lemov explains the consequences of this beautifully:

> The teacher has set a low standard for correctness, and explicitly told the class that a partial answer was fully right when it wasn't. She has crowded out students' own thinking by doing cognitive work that students could – and should – do themselves. The student who answered thinks, 'Good, I did it,' when in fact she didn't. The teacher's crediting her has eliminated the opportunity for the student to recognize the gap between what she said and what would have constituted a top-quality answer. It flatters to deceive. The result, over

> time, is that students believe they're right, prepared, and up to standard, when in fact they're not.

I have done this so many times in the past! 'Yes, Molly, I see what you are saying, you are saying…', followed by something that definitely wasn't what Molly was saying. Lemov's solution is to hold out for that perfect right answer. But is it fair to expect perfection from our students during the early knowledge acquisition phase of learning? And what damage is being done in our desperate quest for them to say exactly what we want – what we need – them to say? If there is one particular response I have in mind, why on earth should I play a game of 'guess what is in my head' with students, with each swing-and-a-miss potentially causing more confusion? Why not just cut out the middleman and deliver the explanation myself?

The key is, I don't want an incorrect or half-right explanation during this phase of knowledge development, and I don't want a debate or a discussion. I have both delivered and observed too many lessons that fall apart in this way. A discussion beforehand to address key misconceptions that could inhibit the understanding of the new concept is brilliant, and will be the focus of Chapter 11. Likewise, discussions with students after the presentation of a worked example and them having tried themselves are also brilliant, because I have given students the very best possible chance to understand the concept and fill in any gaps on their own. But not in the middle.

No, I have come to the conclusion that minimal student involvement during the presentation of a worked example is crucial. So, my process looks like this:

1. Use diagnostic multiple-choice questions at the start of the lesson to assess baseline knowledge, identify and resolve misconceptions. The full process for doing this will be covered in Chapter 11, but it needs making clear that there is plenty of opportunity for discussion at this stage. If the new concept I am intending to teach relies on specific baseline knowledge, and that baseline knowledge is not secure, then there is no point me going on, and I will intervene accordingly. This stage of the process is very interactive.

2. But, during the introduction of the new concept and the subsequent worked example, students are silent. I want them to be focused, and more importantly I want to ensure that what students hear and see is as clear, unambiguous and correct as possible.

3. I model the solution in silence first (Section 4.8's Silent Teacher), pausing briefly after each step, paying full respect to both the Split-Attention and Redundancy Effects and the burden they place upon students' fragile working memories. Students do not write anything down whilst I am doing this – they just watch.

4. Once I have finished my silent solution, I pause. I then narrate and/or annotate over the top.

5. When I have finished this stage, I ask students to copy down my solution into their books. This is important because the Worked Example Effect results from studying worked examples and reflecting on them. I want students to have access to this correct worked example both when they are doing the practice questions later that lesson, and so they can refer to it in the future. I also feel it is important that students experience writing down correct mathematics – setting it out exactly as I have. It is a habit I need them to develop, and one which comes more quickly the more correct maths they write.

6. At the end of this process, I do not ask if there are any questions, I simply ask students to try the paired problem. Again, I do not want confusion spreading around the class, and I do not want to encourage students to give up without trying. If I have presented my worked example clearly enough, then most students should be able to have a good stab at the paired problem. I do not want to slow down the pace of the lesson for a few students who are struggling – students who might in fact be fine once they actually try the problem.

7. Students try the paired problem in absolute silence. I am a fan of mini whiteboards for this part. I find students are more willing to have a go at something if they know they can rub it out without leaving a permanent record.

8. Whilst students have been working on the problem, I have been walking around the class. If I see an example of a student's work that is set out really well, I will use *show-call*. When I interviewed Doug Lemov for my podcast he described show-call as his favourite of all the techniques in *Teach like a Champion*, and we discuss it in depth during the interview. In this instance, show-call involves me borrowing the student's mini whiteboard and using a visualiser to project the work onto the board. Alternatively, taking a photo of the student's work, saving it in Google Drive and opening it up on my classroom computer works really well, and has the added advantage that I have a permanent record of interesting answers to use with future students or even my colleagues. If I do not have a visualiser or the appropriate technology, I can always get the student to come to the front and write their work up, but this is more time consuming. I find it incredibly powerful for students – especially those who may be struggling – to see an example of one of their peers' success, far more so than if they only see my work. It conveys the message that the standards I expect are in their reach.

9. I then say: 'If anyone is stuck or confused, don't worry, I will be round to help you soon'. In the past, I would have invited students who were struggling to explain why in front of the whole class. But this can be a daunting experience for students, and also it risks slowing down the pace of the process.

10. Then I set students off to practise. The sequence of practice questions has been chosen carefully to enable students to develop both procedural fluency and conceptual understanding following the principles of Intelligent Practice that will be discussed in Section 7.8. During this stage my time is freed up, so I can go and help out any students who are struggling, either individually or in a small group. If it becomes clear there is a whole-class issue at this stage, then I can return to the front of the class to sort it out. Once again, show-call comes into its own here. I like to borrow a student's work and use it as a focal point of class discussion. I thank the student for providing such an interesting and important thing to discuss, explaining that if we hadn't picked up on this now, it could have been a disaster. I always ask the

class what is right about the answer first, before we dig into the error. I thus normalise the error, and turn it into a positive experience for the student. Once again, if I have used the photo/Google Drive option, I have a record of interesting errors and misconceptions that I can use in future.

11. I usually end this section of the lesson with a carefully chosen diagnostic multiple-choice question designed to identify any lingering misconceptions. Alternatively I may simply ask students to indicate any questions they got wrong, and intervene accordingly – either with the whole class there and then, or with small groups or individuals at an appropriate point in that lesson or in the future.
12. I then return to the board for another example-problem pair, and the process begins again.

I cannot find the words to describe how different this approach is to how I used to present worked examples. It was a scene of discussion and debate, where I would encourage students to share their thoughts, and where every suggestion and answer was considered worthy of consideration and further investigation, no matter how potentially damaging to the novice learner it could be.

When I describe my new approach to teachers, many are horrified. Here are five things they say:

1. Telling students exactly how to do something without asking them first feels like cheating.
I know what they mean. Telling students how to do something used to be the thing I fell back on once all my efforts to prise the correct answer from my students had failed. But that is ridiculous. I am their teacher, and I am teaching them something that I know a lot more about than them. Now I cut out the middleman and opt straight for the method and explanation I have carefully planned before the lesson – something I believe will give my students the greatest opportunity of understanding a concept – and save the discussion for later, once students have been given the very best chance to understand.

2. That must take ages.
Not at all. Depending on the complexity of the problem, I would say from the start of my silent run-through, up to the students completing their paired example, maybe three minutes at most. So much time is saved by cutting out the discussions and debates that used to infect my worked examples. Also, I find behaviour a lot less of an issue, because students know they must remain entirely silent – there are no grey areas.

3. Your students are not engaged or active during this process.
The battle between engagement and disengagement, activity and passivity, has plagued my teaching for years. Engagement and activity were what I strived for precisely because they were things I could observe. Noise – in the form of discussion and questions – was good, whereas silence was a sign of disengaged, passive robots who simply could not be learning. This was a view reinforced by hundreds of lesson observations that I both received and gave.

Now I know that was simply wrong. Questions and discussions are great (and indeed fundamental to assessing understanding, as we shall see in Chapter 11), but they are cognitively demanding, both to ask and interpret. Given what we know about the limits of working memory, they may simply be too much in the initial knowledge acquisition stage, risking cognitive overload and no learning taking place. At this early stage of knowledge acquisition, I want to ensure that as much of each student's limited working memory capacity as possible is directed towards the matter in hand. I want a calm, focused, controlled environment. Often that leads to prolonged periods of silence from my students. But this is not passivity or disengagement. This is the first stage of learning.

Moreover, such silent focus has an additional benefit that I had not considered until I was delivering training on this process recently. A teacher from a school with a high proportion of English as an Additional Language (EAL) students explained to me that often her EAL students really struggle to follow class discussions, finding multiple voices and rapid dialogue overwhelming. Two weeks later, she could not believe the difference that the Silent Teacher approach – followed by carefully chosen narration and annotation – made to all her students' ability to grasp concepts, but to her

EAL students in particular. Maths finally began to make sense for them. The phrase she used was 'they began to shine'. It is students like this that may suffer in periods of apparent activity and engagement.

4. Where on earth is the differentiation?
Ah yes, the eternal issue of differentiation. What about the student who knows how to add fractions – aren't I holding them back by making them sit through my worked example? Indeed I am, and Cognitive Load Theory itself explains that once students reach a certain level of expertise, then worked examples become less beneficial to learning than solving problems independently (see Section 6.7's Expertise Reversal Effect). But here's the thing – the process of going through an example-problem pair takes around three minutes depending on the complexity of the question. If I have to hold someone back for three minutes, then so be it. Those three minutes will not be wasted – students will get to see exactly how I want the question set out, and have an opportunity to fill in any lingering gaps in their knowledge. They will also have an opportunity to show what they can do during the paired problem, and will certainly be challenged during the careful selection of practice questions that will follow (see Section 7.8), and any subsequent extension material. And anyway, the example-problem pair approach is a damn sight quicker than the way I used to present examples. Those discussions could easily last 10 or 15 minutes, and whilst confusion spread around the room, I could certainly see those students who understood becoming frustrated and even a little bit confused themselves.

Likewise, the student who needs extra support will get it soon enough when my time is freed up during the Intelligent Practice phase, but not before they have had an opportunity to try. More often than not, temporarily removing the ever-present safety net of being able to ask for help leads to students surprising themselves about exactly what they are capable of.

5. I am doing that already.
I hear this a lot. But it is only when I dig a little deeper that I discover differences. On the surface, it can seem that the way I am advocating is something as simple as, 'You watch me do an example, and then you try one yourself'. But, for me, the subtleties are the key. It is presenting the example first in silence, followed by narration; it is the way my board is set out and

the equivalent complexity of the follow-up question; it is students not writing anything down, followed by the use of books and mini whiteboards; it is the use of show-call; it is the absence of questions or discussion until the opportune moment; it is the careful choice of the practice questions that follow. Each of these things is done for a definite and carefully considered reason, and the absence of any of these factors is likely to make the process far less successful. This isn't chalk and talk, nor is it lecturing. It is my take on the most effective means to help students acquire structured, organised and automated knowledge that will help achieve whatever they want in maths.

I'll be honest – I am sold on this approach. Hence, in the initial knowledge acquisition phase, many of my lessons contain sequences of example-problem pairs and carefully designed practice, all completely teacher-led.

OceanofPDF.com

6.3. Labels

What I used to think

I was always a fan of annotating various parts of a worked example. I believed that applying labels to key steps was a good way to help students break down more complex questions into easier, more manageable chunks.

Imagine I was teaching mean from a frequency table, and a particular question involved calculating the mean number of pets in a Year 8 class:

Number of Pets	Frequency	Pets × Frequency
0	5	0
1	8	8
2	12	24
3	7	21
TOTALS	32	53

Figure 6.2 – Source: Craig Barton

Now, when presenting my solution, I have three options:

Option 1: No label
Simply write the final calculation as 53 ÷ 32

Option 2: Superficial label
Have something like total number of pets ÷ total number of students, followed by 53 ÷ 32

Option 3: Abstract label
Have something like total of data times frequency ÷ total frequency, or even $\frac{\Sigma fx}{\Sigma f}$, followed by the calculation.

I would tend to favour Option 2. My logic being that it made the numbers in the question make more sense to the students as they could relate them to the question itself. Surely that was the best thing I could do for their understanding?

Sources of inspiration

- Catrambone, R. (1998) 'The subgoal learning model: creating better examples so that students can solve novel problems', *Journal of Experimental Psychology: General* 127 (4) pp. 355-376.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

Before we consider the nature of the labels themselves, it is worth reflecting on why such labels may help students' understanding.

Catrambone (1998) explains something us teachers know all too well – learners have difficulty solving problems that require solutions different from those demonstrated in examples. In other words they suffer from a failure to transfer. However, Catrambone found that if the worked examples learners study are organised by subgoals (defined as a meaningful conceptual piece of an overall solution procedure), then learners are more successful. Therefore, for multi-step procedures, teachers can identify and label the subgoals required for solving a problem, and encourage students to do the same. This practice makes students more likely to recognise the underlying structure of the problem and be able to apply the steps necessary to solve the problem to other problems that are not identical to the example they are studying.

But what do we need to consider when using such labels?

Well, the first point is related to the Split-Attention and Redundancy Effects discussed in Chapter 4. The labels should be carefully integrated into the solution. Likewise, if we are carrying out a worked example on the board,

whilst at the same time narrating over the top and writing down a label, then it is likely to place a significant strain upon students' fragile working memories, especially in the early knowledge acquisition phase, and could result in cognitive overload. Hence, I employ the Silent Teacher approach (Section 4.8) when first presenting an example, and carefully add any labels and narrations in afterwards.

But, there is a second point related to the nature of the labels.

In his experiment involving statistics worked examples, Catrambone (1998) found that participants in both label conditions outperformed participants in conditions without labels. However, students exposed to superficial labels (labels tied to the surface structure of the particular problem, such as Option 2 above) were less successful than those who experienced abstract labels (Option 3 above) in transferring their knowledge to other contexts. Students can find the mean number of pets all day long, but how about the mean number of siblings, amount of pocket money, or occurrences of volcanic eruptions?

However, before we are tempted to jump to the conclusion that 'abstract labels are best for everyone', there is a problem as explained by Catrambone:

> Unfortunately, a label that is related to surface features of a problem will be more likely to lead a learner to form a solution procedure that is tied to those features. An abstract label is less likely to lead a learner to make this mistake, although the learner must have relevant background knowledge in order to take advantage of an abstract label. These results suggest that cues such as labels can play a strong role in the formation of solution procedures. Because of this, care must be taken to construct cues in a way to aid the formation of structured solution procedures. For learners with weaker backgrounds these cues might need to be tied at least partially to example features despite the danger that this may lead the learner to form representations that have erroneous surface ties. However, for learners with stronger backgrounds, the cues can be constructed more abstractly, thus helping them to form appropriate subgoals.

Hence, we once again see the importance of domain-specific knowledge, not only in allowing learners to circumvent the distracting features of the surface structure in the first place, but in fully benefiting from the power of labels.

What I do now

It's a tricky one.

I am certainly more conscious of the need to include labels and annotations throughout my worked examples, presenting them in a way so as not to overburden students' fragile working memories. However, careful consideration needs to be given as to the nature of those labels.

Superficial labels are likely to lead to a lack of transfer, and yet always presenting students with abstract labels runs the risk of treating novice learners as if they were experts; and as we have seen in Chapter 1 and will continually see throughout this book, experts and novices think differently. Willingham (2002) argues that the development of inflexible knowledge is a necessary step on the journey from novice to expert, so tying the label's detail to the specific surface features of the problem (Option 2) may make sense during the early knowledge acquisition phase. But I am now careful to move on to more abstract labels (Option 3) whenever possible in order to facilitate better transfer.

OceanofPDF.com

6.4. Supercharged Worked Examples

What I used to think

As soon as I had discovered the power and simplicity of example-problem pairs, I didn't think worked examples could get any better. Indeed, I believed that my worked example life had peaked. But in the dark of night I couldn't help but dream: if only there was a way to combine the effectiveness of example-problem pairs with the power of the Self-Explanation Effect…

Sources of inspiration

- Booth, J. L., Cooper, L. A., Donovan, M. S., Huyghe, A., Koedinger, K. R. and Paré-Blagoev, E. J. (2015) 'Design-based research within the constraints of practice: AlgebraByExample', *Journal of Education for Students Placed at Risk* 20 (1-2) pp. 79-100.
- Chi, M. T. H., Bassok, M., Lewis, M., Reimann, P. and Glaser, R. (1989) 'Self-explanations: how students study and use examples in learning to solve problems', *Cognitive Science* 13 (2) pp. 145-182.
- Renkl, A. (1997) 'Learning from worked-out examples: a study on individual differences', *Cognitive Science* 21 (1) pp. 1-29.

My takeaway

We have already discussed the power of both the Worked Example effect and its effective deployment under the structure of example-problem pairs. In the previous chapter we met the Self-Explanation Effect, whereby students who are prompted to self-explain learn more than those who do not self-explain. When individuals self-explain, they generate inferences that are missing from an example's solution and, where appropriate, repair their own mental model when there is divergence between his or her own mental

representation and the example's solution. The problem, as identified by Renkl (1997) and Chi *et al* (1989), is that the majority of students do not spontaneously self-explain, with high-ability students tending to spontaneously self-explain more often than low-ability students.

Fortunately, Booth *et al* (2015) suggest a way we can harness both effects in one fell swoop, by using what I am calling Supercharged Worked Examples.

The study in question came about from a challenge to identify an approach to narrowing the minority student achievement gap in the US high-school Algebra 1 course without isolating minority students for intervention. The authors attempted to do this by designing and testing 42 Algebra assignments with interleaved worked examples that targeted common misconceptions and errors. The worked examples contained three parts:

1. The worked example
2. A section to reflect on what certain aspects of the worked example meant, and how and why they were carried out that way
3. A related problem to complete

Hence, it is the middle section – the opportunity for self-explanation – that distinguishes this approach from the example-problem pair approach. Notice also how the students are prompted to self-explain. We are not relying on students to spontaneously self-explain, which we know is not guaranteed to happen, and even more unlikely for the lower-achieving students who need it most.

The results were impressive. The three-pronged approach led to a significant boost in performance across all measures – both conceptual and procedural – with the greatest impact on students at the lower end of the performance distribution. Interestingly, students using this approach outscored the control group, even though control students had double the practice solving problems on the assignments.

What I do now

I now make regular use of Supercharged Worked Examples, especially for questions that involve several substeps. The process is very similar to that of example-problem pairs, with a couple of key additions.

So, when carrying out a Supercharged Worked Example on rearranging formula, initially my board will look like this:

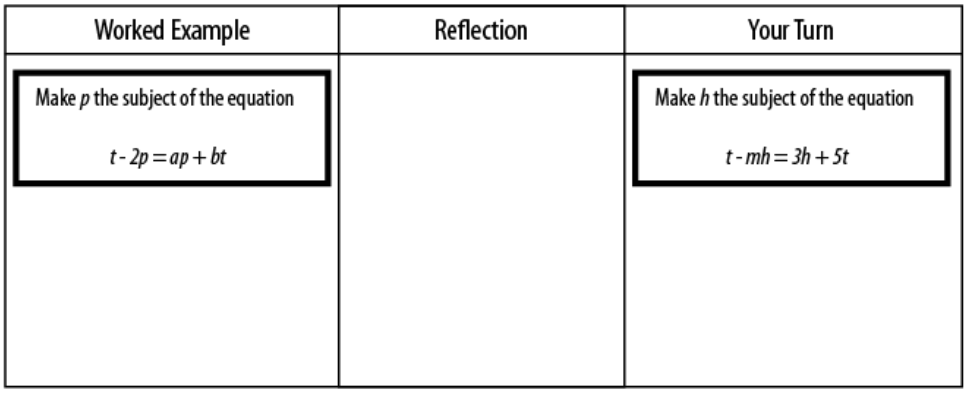


Figure 6.3 – Source: Craig Barton

I will complete the first line of working in silence. When I have finished, I will narrate over the top along the lines of: 'I want to get all the terms involving p on one side of the equation. So, a good way to do this is to add 2p to both sides of the equation, because this collects together the two terms that involve p on the right-hand side of the equation'.

My board will now look like this:

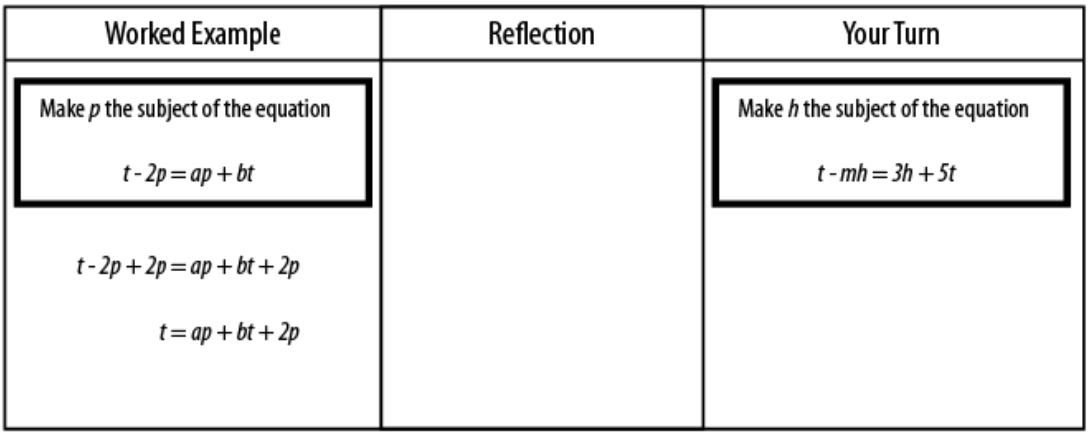


Figure 6.4 – Source: Craig Barton

Because there are several stages to this problem, I will ask students to copy this line of working down into their books now, as opposed to waiting until the end.

Then comes the self-explanation. I will ask students to pause and consider: why did I add 2p to both sides? Here, I don't need students to write anything down. Recalling our work from Chapter 5, self-explaining is all about pausing, reflecting and attempting to make sense of what one is hearing or seeing. Attempting to write down such an explanation is more difficult, and the additional strain placed upon working memory may tip students into a state of cognitive overload. It may also slow down the pace of the explanation. Hence, a verbal prompt from me and a few seconds of silent contemplation may be all that is required.

I may then ask a student their explanation, or I may simply explain it myself. As we saw earlier on in this chapter, I take a far more prominent role in the modelling of examples than I did previously, as I want to ensure that what students hear and see has been carefully planned and is under my control. Likewise, verbalising explanations can be as cognitively demanding as writing them down, and I am always conscious of avoiding the confused discussion and debate that used to plague my worked examples.

I then move onto the next line of working, again presenting it in silence, before adding narration, and prompting students to self-explain. This process continues, and my board may end up looking something like this:

Worked Example	Reflection	Your Turn
Make p the subject of the equation $t - 2p = ap + bt$		Make h the subject of the equation $t - mh = 3h + 5t$
$t - 2p + 2p = ap + bt + 2p$ $t = ap + bt + 2p$	*Why did I add 2p to both sides of the equation?*	
$t - bt = ap + bt + 2p - bt$ $t - bt = ap + 2p$	*Why did I subtract this time and not add?*	
$t - bt = p(a + 2)$	*How did I know to factorise?*	

Figure 6.5 – Source: Craig Barton

Notice that some of the prompts are more difficult than others to explain. 'How did I know to factorise?' is a tough one, and it would be unrealistic to expect students – especially if this is the first time they have encountered a question like this – to provide a perfect answer. But that is perfectly fine. The point of self-explanations is not to have a complete explanation, and certainly not to attempt to verbalise or write one down. The point here is to encourage students to pause, think and notice. They are unlikely to develop an understanding of when and how to factorise from this one example. But if they can get into the habit of pausing and reflecting at various stages in their working, their understanding will gradually become more complete.

Once I have gone through the example, it is time for students to try the related problem, followed by a carefully chosen sequence of questions as will be discussed in Chapter 7. It is at this stage where I can go around the classroom, dealing with any issues on a one-to-one basis. If I sense that in fact there is a wider issue, I can return to the front of the class and deal with it appropriately.

Supercharged Worked Examples take a little longer than example-problem pairs to complete, and students' ability to self-explain grows with experience in the topic or concept. Therefore, I tend to do one or two example-problem

pairs before hitting students with the Supercharged version. I tend to reserve them for the more complex, multi-step processes, where it is important that students pause to consider what they are doing and why they are doing it. As a way of tapping into the benefits of both the Worked Example and Self-Explanation effects, they are now a fundamentally important part of my teaching.

OceanofPDF.com

6.5. Mistakes in Worked Examples

What I used to think

I used to worry that it was dangerous to show students worked examples containing mistakes as it may lead to the development of misconceptions. Hence, I was more inclined to teach the correct procedures and deal with misconceptions as and when they became apparent.

Sources of inspiration

- Booth, J. L., Cooper, L. A., Donovan, M. S., Huyghe, A., Koedinger, K. R. and Paré-Blagoev, E. J. (2015) 'Design-based research within the constraints of practice: AlgebraByExample', *Journal of Education for Students Placed at Risk* 20 (1-2) pp. 79-100.
- Booth, J. L., Lange, K. E., Koedinger, K. R. and Newton, K. J. (2013) 'Using example problems to improve student learning in algebra: differentiating between correct and incorrect examples', *Learning and Instruction* 25, pp. 24-34.
- Große, C. S. and Renkl, A. (2004) 'Learning from worked examples: what happens if errors are included?' in Gerjets, P., Elen, J., Joiner, R. and Kirschner, P. (eds) *Instructional design for effective and enjoyable computer-supported learning*. Tübingen: Knowledge Media Research Center, pp. 356-364.
- Große, C. S. and Renkl, A. (2007) 'Finding and fixing errors in worked examples: can this foster learning outcomes?', *Learning and Instruction* 17 (6) pp. 612-634.
- Lutwyche, A. (2017) 'Clumsy Clive: averages and range', *TES* Resources. Available at: https://www.tes.com/teaching-resource/clumsy-clive-on-

averages-and-range-11570708

- Siegler, R. S. (2002) 'Microgenetic studies of self-explanation' in Granott, N. and Parziale, J. (eds) *Microdevelopment: transition processes in development and learning*. Cambridge: Cambridge University Press, pp. 31-58.

My takeaway

We have already discussed the potential dangers of dealing with misconceptions in Section 3.8 when we looked at cognitive conflict, and misconceptions will be at the heart of Chapter 11 on formative assessment. But there is more than enough controversy to go around, so let's dive into some more here.

Is it a good idea to present students with worked examples that contain mistakes?

We saw back in Section 5.2 the benefits of students explaining someone else's answer, together with Siegler's (2002) finding that 'explaining why correct answers are correct and why incorrect answers are incorrect yields greater learning than only explaining why correct answers are correct'.

Similarly, in their study of US high-school students learning how to solve algebraic equations, Booth *et al* (2013) sought to discover if there are differential effects on learning when students explain correct examples, incorrect examples, or a combination of the two. Incorrect worked examples were clearly labelled as such, with students asked to consider questions such as, 'Why is this a wrong step for Ben to take?'. Results indicated that students benefited significantly from exposure to incorrect examples, specifically with regard to the development of conceptual understanding.

According to Booth *et al*, a combination of correct and incorrect examples is beneficial because the incorrect examples force students to attend to the critical features of a problem (which helps them not only to detect and correct errors, but also to consider correct concepts), while the correct examples provide support for constructing correct concepts and procedures.

The authors explain that this finding is especially important to note because:

> When examples are used in classrooms and in textbooks, they are most frequently correctly solved examples. In fact, in our experience, teachers generally seem uncomfortable with the idea of presenting incorrect examples, as they are concerned their students would be confused by them and/or would adopt the demonstrated incorrect strategies for solving problems. Our results strongly suggest that this is not the case, and that students should work with incorrect examples as part of their classroom activities.

However, before we get carried away and start filling our boards with error-strewn worked examples, a word of caution.

According to Große and Renkl (2004), learning from worked examples where errors are included can indeed enhance learning and transfer, but only if students have good prior knowledge of the topic. They explain:

> Learning with incorrect examples poses challenging demands on the learners. They have to represent not only the correct solution in their working memory, but also the incorrect step with an explanation why it is wrong. Learners with low prior knowledge who cannot form larger chunks for information coding can easily be overtaxed.

This finding is supported by Große and Renkl's (2007) more recent paper which found that relatively novice learners cannot benefit from incorrect examples when they are expected to locate and identify the error in the example themselves.

This makes perfect sense. First, if students do not understand the topic, then how are they to spot the mistakes? Indeed, if students fail to distinguish between correct and incorrect procedures, then there is the possibility they may develop misconceptions based around the incorrect examples. Second, in the context of Cognitive Load Theory, if students' knowledge is not accessible and automated, then their working memories are likely to become overloaded whilst searching for the right and wrong answers simultaneously.

What I do now

It seems clear that students need a certain amount of domain-specific knowledge in order to benefit from exposure to mistakes in worked

examples. Hence, I am leaning towards correct worked examples as presented in this chapter during early knowledge acquisition, moving onto clearly labelled incorrect worked examples once students begin to get familiar with the concepts.

Good quality sources of 'spot the mistake' style are fortunately abundant in the world of maths. A search on TES for 'tick and trash' will return a lovely series of activities whereby students must decide between two solutions to the same problem, one of which is right and the other is wrong. Likewise, one of my favourite resource authors, Andy Lutwyche, has created a wonderful collection of resources called 'Clumsy Clive' (and the A level equivalent, 'Erica's Errors'). These consist of a series of worked examples written by Clumsy Clive which are scattered with mistakes. It is the students' jobs to find, explain, and correct these mistakes. 'Clumsy Clive: Averages and Range' looks like this:

Questions 4:

The following data was collected regarding the amount of post one street received on one day:

Number of letters	0	1	2	3	4
Number of houses	4	7	6	4	3

Calculate the mean number of letters received by the houses on that day.

Clive's answer:	*Your answer:*
Add up the total number of letters and divide. Total number of letters: $0\times4+1\times7+2\times6+3\times4+4\times3=43$ Answer: $\frac{43}{5}=8.6$	

What mistake has Clive made?

Figure 6.6 – Source: Andy Lutwyche, available at https://www.tes.com/teaching-resource/clumsy-clive-on-averages-and-range-11570708

Indeed, Clumsy Clive has become one of my favourite sources of end-of-topic assessments.

OceanofPDF.com

6.6. Fading

What I used to think

Okay, so having discovered the beauty of Supercharged Worked examples, this time I was 100% convinced that my worked example life had well and truly peaked.

Then I discovered fading…

Sources of inspiration

- Atkinson, R. K., Renkl, A. and Merrill, M. M. (2003) 'Transitioning from studying examples to solving problems: effects of self-explanation prompts and fading worked-out steps', *Journal of Educational Psychology* 95 (4) pp. 774-783.
- Renkl, A. and Atkinson, R. K. (2003) 'Structuring the transition from example study to problem solving in cognitive skill acquisition: a cognitive load perspective', *Educational Psychologist* 38 (1) pp. 15-22.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Renkl and Atkinson (2003) propose a challenge to the example-problem pair and Supercharged Example approach heralded in the sections above. They suggest that instead of providing students with a complete worked example, followed by a problem for them to solve, followed by a complete worked example, followed by problem to solve, *etc*, a *fading* procedure is more effective. This involves providing a complete worked example, and

then following this up with an almost complete worked example, but with one step removed that students need to complete themselves.

Specifically, Renkl and Atkinson suggest:

> First, a complete example is presented (model). Second, an example is given in which one single solution step is omitted (coached problem solving). Then, the number of blanks is increased step by step until just the problem formulation is left, that is, a to-be-solved problem (independent problem solving). In this way, a smooth transition from modeling (complete example) over coached problem solving (incomplete example) to independent problem solving is implemented.

This approach is closely related to the Completion Problem Effect, which forms a part of Cognitive Load Theory. Completion problems provide a bridge between worked examples and conventional problems. Students who are challenged to complete sections of worked examples tend to perform better than students who have the full solution presented to them, because generating part of the solution constitutes germane load (Section 4.9). Sweller *et al* (1998) explain that 'worked examples are completion problems with a full solution, and conventional problems are completion problems with no solution. A good progression may be to start with completion problems that provide almost complete solutions, and gradually work to completion problems for which all or most of the solution must be generated by the learners'.

In their study into such a fading procedure, Renkl and Atkinson (2003) found:

1. The fading procedure produced reliable effects on near-transfer items but not on far-transfer items. Near-transfer problems have the same deep structure but different surface structures. In other words, the strategy to solve them is exactly the same, but the context is different. Far-transfer problems have both a different context and a different deep structure.

2. It was more advantageous to fade out worked-out solution steps using a backward approach, by omitting the last solution steps first, instead of omitting the initial solution steps first (*ie* a forward approach).

Atkinson *et al* (2003) seek to discover how to improve performance on far-transfer problems. Just like we saw in the section on Supercharged Worked Examples, they argue we can improve the effectiveness of worked examples by prompting student self-explanations. The researchers conducted experiments that combined fading with the introduction of prompts designed to encourage learners to identify the underlying principle illustrated in each worked-out solution step. They set up four conditions: Example-Problem Pair (EP); Example-Problem Pair Plus (EP+), which contained the self-explanation prompts; Backwards Fading (BF); and Backwards Fading Plus (BF+). Crucially, the experiment was designed in a way so that participants had to complete the exact same number of solution steps in each condition – it was just in the EP approach they all came together, whereas in the BF model they were introduced more gradually.

The key findings were as follows:

1. The BF condition was associated with a higher solution rate of near-transfer problems than EP.
2. The prompting of self-explanations can substantially foster both near and far transfer. It is also notable that the advantage of such prompting could be achieved without significantly increasing learning time – as we have seen, it is the simple act of pausing and reflecting, not necessarily verbalising or committing to writing.
3. There was no evidence of an interaction between the use of fading and the use of self-explanation prompts on any of the measures.

For me, there are two key takeaways here. First, the backwards fading procedure must be taken seriously. Second, prompting students to self-explain is effective no matter how the worked examples are presented, and hence is something of a no-brainer when it comes to teaching and modelling.

What I do now

So, we have an alternative to the example-problem pair approach, which has been shown to outperform its rival.

This raises the obvious question: when do we use each one? For me, it depends on the complexity of the problems and the prior knowledge of the class. Relatively simple problems probably lend themselves better to a straightforward example-problem pair approach. This can be presented quickly, with little fuss. Likewise, for students with higher prior knowledge, or students who have seen the topic before (*eg* in a revision lesson), I would also favour the example-problem approach. However, for more complex problems – particularly the multi-mark epics that students have nightmares with in exams – or with students who are struggling to grasp basic concepts, then I turn my attention to Supercharged Worked Examples, and the backwards fading approach. That extra structured support, and those moments of silent contemplation, may be exactly what is needed.

OceanofPDF.com

6.7. The Expertise Reversal Effect

What I used to think

I don't know about you, but any time I discover something promising – be it Tarsia jigsaws, collective memories, revision clocks, or walking-talking mocks – I tend to get more than a little obsessed, quickly believing it is the answer to all my prayers. And so it was when I discovered the power of worked examples. Surely the effects were so powerful and so positive that every single student should be allowed to benefit from them?

Sources of inspiration

- Kalyuga, S., Ayres, P., Chandler, P. and Sweller, J. (2003) 'The expertise reversal effect', *Educational Psychologist* 38 (1) pp. 23-31.
- Kalyuga, S., Rikers, R. and Paas, F. (2012) 'Educational implications of expertise reversal effects in learning and performance of complex cognitive and sensorimotor skills', *Educational Psychology Review* 24 (2) pp. 313-337.
- Martin, A. J. (2016) *Using Load Reduction Instruction (LRI) to boost motivation and engagement*. Leicester: British Psychological Society.

My takeaway

A recurring theme throughout this book is that experts and novices think and learn differently, and it is our job as teachers to account for that. Whereas much of the emphasis so far has been to warn against the dangers of treating novices like they are experts – and indeed, we will see even more of this in Chapter 9 on problem-solving – we must also address the danger of treating experts like they are novices when it comes to the use of worked examples.

We have seen in this chapter how during early knowledge acquisition, learning from worked examples is very advantageous, whereas learning by solving problems is not. However, as Kalyuga *et al* (2003) describe, instructional techniques that are highly effective with inexperienced learners (novices) can lose their effectiveness and even have negative consequences when used with more experienced learners (experts). Hence we have the Expertise Reversal Effect.

The argument is that worked examples contain information that is easily determined by more experienced learners themselves and, therefore, can be considered redundant. We have seen in Section 4.7 that even redundant information needs to be processed, and devoting working memory to redundant information effectively takes away a portion of the learners' limited cognitive capacity that could be devoted to the more useful germane load. Indeed, having gone to the trouble of helping students acquire, organise and automate knowledge through the careful presentation of worked examples and Chapter 7's Intelligent Practice, we want to ensure that they use their spare cognitive capacity for something that is going to contribute to learning.

Moreover, this redundant information may even interfere with the schemas constructed by experienced learners, preventing them from seeing the deeper connections in problems that are essential for transfer. Kalyuga *et al* (2012) explain that the Expertise Reversal Effect is caused by a potentially conflicting overlap between available knowledge structures in experts' long-term memory and externally provided information that addresses the same situations or guides solving the same problems. In such cases, providing minimal external information might be more effective for these learners who can then take the full advantage of their available knowledge base.

For example, what if my students have solved a problem differently from how I have presented it in the worked example? Not only am I slowing them down by insisting they go through my carefully constructed worked example step-by-step, but I am also potentially confusing them. Their method may in fact be more efficient than mine, or they may be really happy with their approach but fail to understand mine. At this stage of

development, working through complex problems independently is likely to be more beneficial for long-term learning than studying worked examples.

Indeed, if we return to Martin's (2016) five principles at key points in the learning process, we see that less-guided approaches such an inquiry-based learning are an essential part of learning once students have benefited from carefully planned, teacher-led guidance:

1. Reducing the difficulty of a task during initial learning
2. Instructional support and scaffolding through the task
3. Ample structured practice
4. Appropriate provision of instructional feedback
5. Independent practice, supported autonomy, and guided discovery learning.

Of course, there are potential issues with this.

For a start, we may wish to expose students to different methods, and indeed our method in particular. This is certainly the case when students have to write down certain steps in order to get full marks on a mark scheme, or where a shortcut that a student may take does not have the solid mathematical foundations upon which we wish to build in the future. In these cases, it may still be good practice to let the student carry on, and then have a whole-class discussion of different methods and approaches towards the end of the lesson. I have already stated my (controversial and no-doubt highly unpopular) view that I do not favour much whole-class discussion when initially going through an example, but I am all in favour of it when the time is right, once students have had the chance to gain familiarity of the concept in the way I have presented it.

But perhaps the biggest problem is the challenge of recognising when students have made the transition from novice to expert and hence can start to be exposed to more complex problems independently. Too soon and we have novice learners struggling with problems and perhaps not learning

anything (Section 9.4); too late and we have the Expertise Reversal Effect. Being a teacher is flipping tough!

What I do now

The Expertise Reversal Effect suggests that a gradual release from teacher-directed to independent work is needed once students reach a certain level of domain-specific expertise. Exactly what this gradual release might look like is the subject of Section 9.3, where I describe how I try to help my students become the independent problem-solvers I want them to be.

Judging when students are ready for less teacher-led instruction is very tricky. The domain-specific nature of expertise, together with the difficulty in distinguishing between learning and performance (Section 12.2), are two contributing factors. But the techniques of formative assessment that we will cover in Chapter 11, together with regular low-stakes quizzes (Section 12.8), can improve the accuracy of these judgements, both for me as their teacher and for the students themselves.

Hence, I often make additional worked examples available for students once we have covered the basics of a topic and students are working through my carefully chosen practice questions (Section 7.8). I print these out and put them in piles at the ends of rows. That way students can judge for themselves whether they are at the stage where using worked examples will help or hinder them.

But given the difficulty in making accurate judgements of understanding, together with the significant benefits of worked examples outlined in this chapter and the speed in which these worked examples are completed, I always err on the side of caution. I will never skip through my carefully planned worked examples. I want every child to see and benefit from them. At the end of the day, I feel it is far better to mistakenly treat an expert as a novice than the other way around.

OceanofPDF.com

6.8. If I only remember 3 things…

1. Examples should not be bundled together at the start of a lesson, but should be presented as example-problem pairs, followed by carefully chosen questions for students to complete, before returning to another example-problem pair.
2. Prompting students to self-explain during the completion of a worked example (Supercharged Worked Examples) makes them an even more effective tool of instruction and learning.
3. Domain-specific experts do not benefit as much from worked examples as novice learners, but we need to be very careful when making the decision that our students have acquired sufficient expertise.

OceanofPDF.com

7. CHOICE OF EXAMPLES AND PRACTICE QUESTIONS

My approach to planning across the various phases of my teaching career can be summarised as follows:

2004 to 2007: The Resource Years
Hours would be spent trawling through *TES* in search of that magic activity, be it a mystery, a jigsaw, a card sort, or a PowerPoint with loads of fancy animations. And when I found it, by God I was going to use it, whether it was suitable for my class or not.

2008 to 2015: The Question Years
For the next seven years I became convinced that good questions were the key to planning and teaching, and my lesson plans were full of them. This began with a belief in the virtue of probing, open-ended questions – 'but how do you know that?', 'can you convince me?' *etc*. I then moved on to the use of diagnostic multiple questions, both for formative assessment purposes, and then later as a learning tool in their own right. This will be the subject of Chapter 11.

2016 onwards: The Example Years
Recently, however, my eyes have been opened up to the power and importance of examples. In Chapter 6 we looked at an effective way to present examples to students. Here I want to focus on the *choice* of examples, because having spoken to the likes of Daisy Christodoulou and Kris Boulton for my podcast, I am more convinced than ever that the choice of the examples and subsequent practice questions we give our students is the

single most important factor in determining their initial understanding of a skill or concept.

This chapter will be my attempt to explain why.

OceanofPDF.com

7.1. Examples v Definitions

What I used to think

I used to think that definitions and the subsequent explanation of that definition were the most important part of helping students understand a concept. Hence, I would start with the definition, and then follow it up with examples and practice.

The problem was, my students never really seemed to understand or use the definitions all that much.

Sources of inspiration

- Charles, R. I. (1980) 'Exemplification and characterization moves in the classroom teaching of geometry concepts', *Journal for Research in Mathematics Education* 11 (1) 10-21.
- Goldenberg, P. and Mason, J. (2008) 'Shedding light on and with example spaces', *Educational Studies in Mathematics* 69 (2) pp. 183-194.
- Wiemann, C. (2007) 'The "curse of knowledge", or why intuition about teaching often fails', *APS News* 16 (10) (no pagination).
- Wilson, P. S. (1986) 'Feature frequency and the use of negative instances in a geometric task', *Journal for Research in Mathematics Education* 17 (2) pp. 130-139.

My takeaway

Mathematics is full of definitions. Here are some of my favourites:

A polygon is a closed 2D shape with straight sides.

A geometric sequence is a sequence of numbers where each term after the first is found by multiplying the previous one by a fixed, non-zero number called the common ratio.

An equation is a statement that the values of two mathematical expressions are equal.

A simple random sample is a subset of a statistical population in which each member of the subset has an equal probability of being chosen.

There is something almost paradoxical about definitions. The more technical and complicated you make them, the less prone they are to ambiguity and misrepresentation, but unfortunately the less likely they are to be understood, most notably by novice learners like our students. The problem is exacerbated when definitions are used at the start of a topic, as this is precisely the time when students have the least chance of fully understanding the definition and its related concepts.

Look at the definition of an equation, for example. In order to fully appreciate its meaning, students need to be secure in their knowledge of the concepts of 'statement', 'value', 'expressions' and even 'equal'. Is it any wonder many students tell us they hate algebra?

But this can be hard for teachers to appreciate, burdened as we are by the *curse of knowledge*. Wiemann (2007) explains that the curse of knowledge occurs when you know something well, and it thus becomes extremely difficult to think about it from the perspective of someone who does not know it. So, whilst those definitions (and the subsequent explanations we may offer to further clarify them) make perfect sense to us – blessed as we are by hard-earned, deep, connected domain-specific knowledge – they may not to our novice learners. Indeed, students may not understand the definitions and explanations we give due to issues with knowledge, language, concentration, misconceptions, or a whole manner of other things. We need to keep in mind that students do not learn from what we say and do; they learn from their own individual interpretations of what we say and do. If these interpretations are obfuscated by misunderstood definitions and explanations, then learning is not likely to be as effective as it could be.

Goldenberg and Mason (2008) instead favour the use of examples before the presentation of any definition. They argue that students are more likely to attach meaning to the examples teachers go through, forming their own conceptions and conclusions, which can make up for the shortfalls in vocabulary and conceptual understanding of students that can render definitions alone pretty useless.

But we can go one better, by also exposing students to non-examples.

Charles (1980) argues that for relatively simple concepts, a sequence of examples from which to generalise may be sufficient, but for more 'difficult' concepts non-examples are necessary to delineate the boundaries of the concept. Wilson (1986) points out that learners can be distracted by irrelevant aspects of examples, so the presence of non-examples provides more information about what is, and what is not, included in a definition. So non-examples are important for allowing students to see the boundary of what works and what doesn't work. They make concepts less fuzzy, and make-up for shortfalls in definitions and explanations.

Goldenberg and Mason (2008) emphasise that the definition is still important for clarity and confirmation. This avoids you having to present a potentially infinite number of examples and non-examples. However, this definition needs to come after exposure to carefully chosen examples and non-examples, at a time when students have a much better idea of what is going on.

What I do now

Instead of starting with the definition and explanation of a concept, I start with examples. Once students have seen a reasonable number of carefully chosen examples and non-examples, they form their own interpretation of the concept, and hence are in a much better place to understand and appreciate the subsequent definition. A selection of what are and what are not polygons, equations, geometric sequences and random samples goes a lot further in enabling students to grasp the concepts during early knowledge acquisition phase than the most technically correct definition ever could.

So, instead of presenting students with a definition of a polygon – such as 'a polygon is a closed 2D shape with straight sides' – students could instead be presented with a series of examples of both polygons and non-polygons. For example:

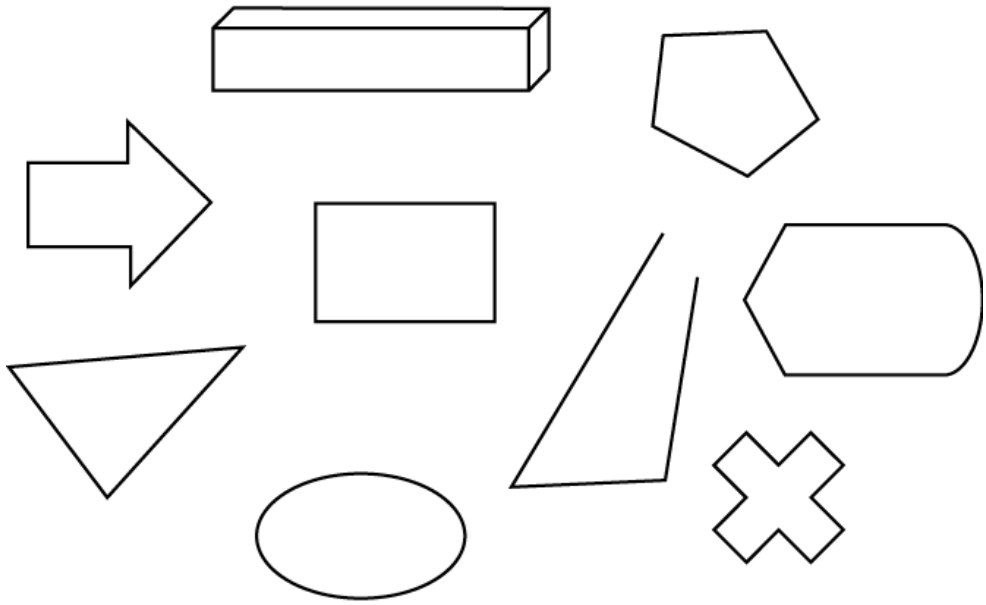


Figure 7.1 – Source: Craig Barton

I could silently place ticks next to polygons and crosses next to non-polygons as a means of enabling students to generate their own understanding of a polygon before introducing the formal definition. Or we could have a whole-class vote, with me subsequently placing a score next to each shape to indicate how many students think it is a polygon – this can lead to some brilliant discussions. Alternatively, I can print the selection of shapes out and challenge each student to 'circle all the polygons' as a means to assess understanding following my explanation.

The key point is that students' attention is brought to bear not just on the similarities between the shapes that are polygons, but crucially upon their difference from the shapes that are not. It is an appreciation of this difference that allows for a deeper understanding of what a polygon is.

Likewise, the best way to enable students to fully appreciate what we mean by an equation might be to present them with the following:

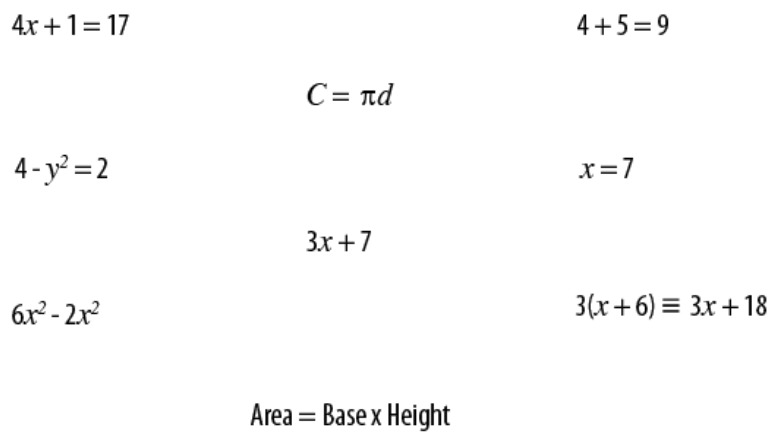


Figure 7.2 – Source: Craig Barton

There is a danger that the presentation of all these examples at the same time will be too cognitively demanding. An alternative is to present one thing at a time, indicate whether or not it fits into a certain group, make a single, deliberate change, and then reflect on the effect this change has had. This feeds into the concept of Section 7.8's minimally different examples.

So, in order to help students understand what we mean by an equation, I could start with:

x

I pause to give students time to reflect and predict, before silently indicating that this is not an equation by placing a cross next to it. I then make a copy of x below, then change this copy to $4x$, pause again and then also indicate that this is not an equation. I like to make a copy and change it (as opposed to just changing the original) so students have a visible record to help the development of their understanding.

After a while, my board may look like this:

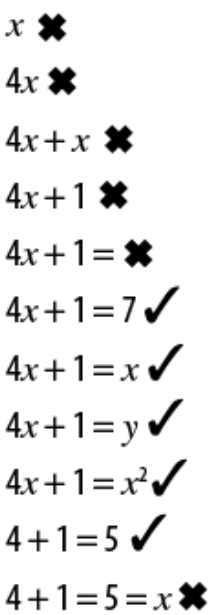


Figure 7.3 – Source: Craig Barton

OceanofPDF.com

7.2. Examples v Rules

What I used to think

Okay, so we have established that definitions and explanations may be difficult to understand, but surely nothing beats a clear, unambiguous rule for carrying out a procedure? After all, aren't such rules the foundations upon which maths is built?

Hence, I would happily present rules to my students, attempt to show where these rules came from, get them to neatly copy the rule down in their books, and give students plenty of practice of them.

What could possibly go wrong?

Sources of inspiration

- Anderson, J. R. (1996) 'ACT: a simple theory of complex cognition', *American Psychologist* 51 (4) pp. 355-365.
- Morgan, J. (2017) 'Topics in Depth', *Resourceaholic maths blog*. Available at: http://www.resourceaholic.com/2017/06/topics.html.
- Pearce, J. (2016) 'Reflections on #MathsConf5 – The dangers of students over- or under-generalising', *Must-be Maths blog*. Available at: http://www.mustbemaths.com/2015/09/reflections-on-mathsconf5-dangers-of.html
- Skemp, R. R. (1976) 'Relational understanding and instrumental understanding', *Mathematics Teaching* 77 (1) pp. 20-26.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.

My takeaway

Let's consider five common rules in mathematics:

- *To find the mean you add up all the numbers and divide by however many numbers there are.*
- *Angles on a straight line add up to 180°.*
- *Two minuses make a plus.*
- *To expand a bracket, multiply everything on the outside by everything on the inside.*
- *When you multiply two indices with the same base, you add the powers.*

These rules each have the advantage that they are short, written in relatively easy to understand language, and not too tricky to remember. The problem is, time and time again students misapply them due to problems of over- and under-generalisation.

Let's briefly look at each rule in turn.

To find the mean you add up all the numbers and divide by however many numbers there are

This rule is ideal if you do indeed have a list of numbers. But what happens when students encounter mean from a frequency table? Or estimating the mean from grouped data? Or working backwards from the mean to find a given value? All of a sudden our rule appears rather deficient and something else is needed. In the words of James Pearce, we have 'over-generalised'.

A common response is to argue that when students first encounter the mean it is usually only in the context of a list of numbers, and so that definition works fine. However, we return to Section 1.5 where we looked at building methods that last. What happens in Years 8 and 9? If students learn one rule for the mean for a list of numbers, then another for a frequency table, another for grouped frequency, another for backwards means, and so on, then all of a sudden we have several concepts, each of which is potentially

viewed as separate from the others, all prone to being misremembered or misapplied.

Angles on a straight line add up to 180°

This rule looks safe enough. But consider the following question below from my Diagnostic Questions website:

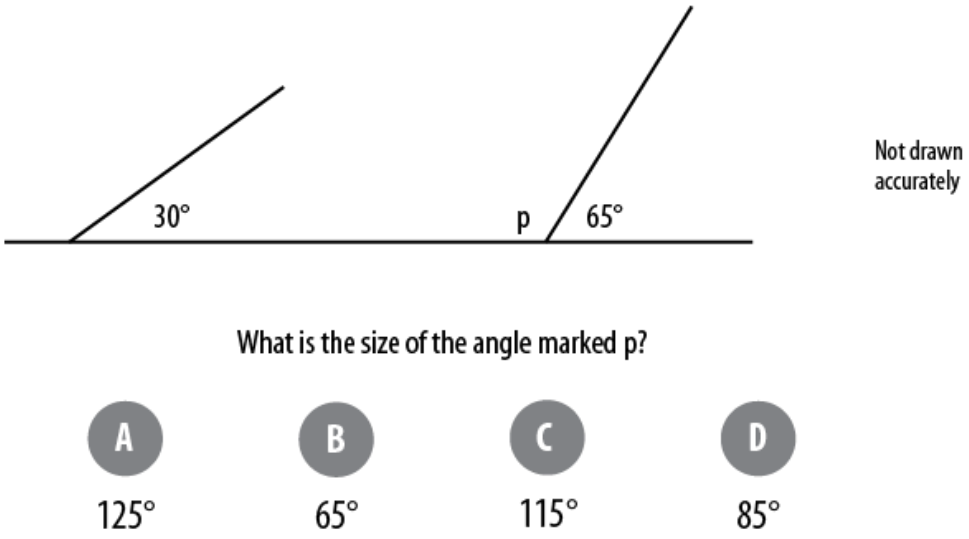


Figure 7.4 – Source: Craig Barton, created for Diagnostic Questions

If I told you that 35% of students at the time of writing opted for D) 85°, we can once again see the dangers of over-generalisation as *all* angles on the straight line are considered.

Two minuses make a plus

This is arguably the most dangerous rule in mathematics, prone to over-generalisation and misapplication by students of all ages. A student with this mantra echoing in their heads will inevitably stumble when faced with -7 – 5.

To expand a bracket, multiply everything on the outside by everything on the inside

This rule sounds perfectly safe, until, that is, you consider something like: $2 + 3(4x - 1)$. I certainly do not want everything on the outside being

multiplied by everything on the inside. Once more, the perils of over-generalisation are plain to see.

When you multiply two indices with the same base, you add the powers

This rule can lead to a nasty case of under-generalisation. If students only ever encounter this rule in the context of positive bases and powers that are integers, such as $8^3 \times 8^6$, then the rule implicitly becomes *when you multiply two indices with the same positive base, you add the integers in the powers together*. Hence, there is a danger that when faced with any of the following, students may come unstuck, fearing they need another rule to help them answer the questions:

$(-2)^3 \times (-2)^6$

$7^{1.4} \times 7^{0.3}$

$3^{\frac{1}{3}} \times 3^{\frac{2}{5}}$

$(x + y)^3 \times (x + y)^{-6}$

$1.8^{2n} \times 1.8^{3n}$

The problem is that it is virtually impossible to write each of these rules in a way that covers all the possible twists and turns that students could encounter, avoiding the traps with over- and under-generalisation. Try it with the rule for finding the mean to see how tricky it is. '*To find the mean you add up all the numbers, but if these numbers are in a frequency table then you need to multiply the frequency by the number to get the partial sum, oh and if the frequency table is grouped then you first need to get an unbiased estimate for the numbers by finding the midpoint…*'. Doesn't exactly roll off the tongue.

Indeed, in the context of arguing for Relational Understanding over Instrumental Understanding, Skemp (1976) points out that 'if the teacher asks a question that does not quite fit the rule, of course they will get it wrong'. But for me, that is not an argument against instrumental understanding – I discuss the argument further in Section 3.9 – it is an

argument for ensuring we give students examples such that they may develop a deeper understanding of the concept before the presentation of any rule.

From the models of thinking we looked at in Chapter 1, in particular those proposed by Anderson and Sweller, we can infer that it is better to focus on the ideas behind a concept rather than an algorithm for solving a particular subset of cases. Algorithms that only work in a few cases are less compatible with the development of the connected body of knowledge that form the long-term memories of experts. Likewise, continually having to decide which algorithm is suitable for a given case is likely to impose a significant burden upon working memory and threaten cognitive overload. Hence, we may surmise that it is good practice for students to encounter different types of cases as soon as possible in learning about a concept, to avoid incorrect over- and under-generalisations being embedded, which can be very difficult to resolve.

It is worth repeating that students do not learn from what we say and do; they learn from their own individual interpretations of what we say and do. Hence, as we saw with definitions and explanations above, presenting the examples and non-examples first should lead to a deeper understanding of a concept than a rule which may be applied either too narrowly or too broadly.

What I do now

Rules may be misunderstood or misapplied by students, leading to nasty cases of over- and under-generalisation. Well-chosen and well-presented examples and non-examples can illustrate a rule better than words ever can. They can help students arrive at a place where they already have a sound understanding of the rule before it is crystallised into words. Indeed, with enough good examples, followed by the Intelligent Practice we will look at later in this chapter, students may never need to rely on remembering an abstract rule at all.

Consider the rule: *angles on a straight line add to 180°*. To help address and hopefully rectify some of the misconceptions associated with this rule, students could again be presented with a selection of scenarios and asked to

identify when this rule can be applied. Jo Morgan put together the following collection as part of her excellent 'Topic in Depth' series, where the question posed is, 'Which pairs of angles will sum to 180°?':

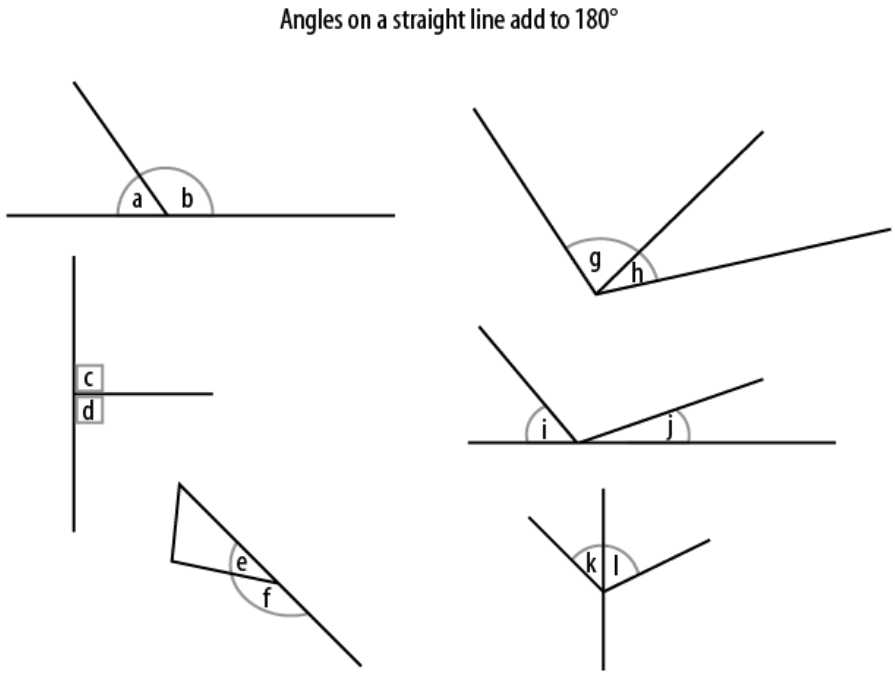


Figure 7.5 – Source: Jo Morgan, available at http://www.resourceaholic.com/2017/06/topics.htm

Finally, a quick word on negative numbers, for I feel this is a topic I have routinely failed to teach well for about ten years. I started with the *two minuses make a plus* rule, and then – as discussed in Section 3.7 – 'progressed' onto a dodgy analogy involving soup. Neither was successful.

Now I introduce negative numbers via sequences of carefully chosen examples, utilising humans' natural tendency to spot patterns discussed in Section 3.4. So, a sequence of examples to develop the ability to subtract negative numbers may look like this:

Example	Your Turn
3 - 3 = 0	2 - 3 =
3 - 2 = 1	2 - 2 =
3 - 1 = 2	2 - 1 =
3 - 0 = 3	2 - 0 =
3 - (-1) =	2 - (-1) =
3 - (-2) =	2 - (-2) =
3 - (-3) =	2 - (-3) =

Figure 7.6 – Source: Craig Barton

Next I would present a sequence that begins with a negative number, and I would do similar sequences for each of the operations.

I present these examples in silence, pausing briefly before writing each answer down, making little fuss, and then give my students similar patterns to complete themselves on mini whiteboards. Crucially, I then ensure students receive regular practice of isolated questions (*ie* not presented in a pattern) at spaced intervals throughout the coming weeks, months and year. These will appear in the homeworks and low-stakes quizzes that will be discussed in Chapter 12. This way, students develop fluency for dealing with negative numbers, which is far superior to a rule prone to misapplication, or a dubious analogy.

All of this is not to say that rules – just like definitions – do not have a place. Rules can clarify, summarise and be committed to memory. Just like definitions, rules also remove the need for an infinite number of examples and non-examples to understand a concept. However, rules that are not first preceded by well-chosen and well-presented examples and non-examples can often lead to trouble.

OceanofPDF.com

7.3. Boundary examples

What I used to think

I used to think that there was little point focusing on 'extreme' cases. Hence, I would concentrate on covering fairly standard examples. Take simplifying fractions. I may choose to present the following examples to my students, going through each one, following the principles of example-problem pairs or Supercharged worked examples as discussed in the previous chapter:

$\frac{9}{18} =$ $\frac{6}{42} =$

$\frac{70}{80} =$ $\frac{36}{45} =$

$\frac{6}{9} =$

Surely my students are now perfectly placed to deal with any fraction in need of simplification that should come their way?

Sources of inspiration

- Mason, J. and Watson, A. (2001) 'Getting students to create boundary examples', *MSOR Connections* 1 (1) pp. 9-11.
- Skemp, R. R. (1976) 'Relational understanding and instrumental understanding', *Mathematics Teaching* 77 (1) pp. 20-26.

My takeaway

Given the importance I am now attaching to examples over definitions, explanations and rules, it is vital that the examples I choose cover the full example space of the concept in question. Any omissions in the examples I choose – or the practice I subsequently give the students to do – will inevitably lead to gaps in their knowledge. Hence, the extreme examples, which I would previously leave out, hereby become increasingly important.

Mason and Watson (2001) coined the term 'boundary examples' to describe these extreme examples. They explain: 'We use the word 'boundary' because we see students' experiences of examples in terms of spaces: families of related objects which collectively satisfy a particular situation, or answer a particular mathematics question, or deserve the same label. Such spaces appear to cluster around dominant central images'. The point is that by not explicitly addressing examples 'on the boundary' there is the danger that key features will be missed at the expense of others.

Boundary examples are – for want of a better phrase – weird questions, or normal-looking questions with weird results. They are not necessarily more difficult examples. This is related to the concept of under-generalisation discussed in the previous section – we do not want our students forming narrow views of the rules or procedures they encounter. It is also related to the concept of self-explanation that was covered in Chapter 5 – the mere act of pausing to consider what they have done, why they have done it, and if it makes sense, has been shown to be beneficial for student understanding.

Let's return to the simplification of fractions. There is nothing inherently wrong with the examples that I have chosen to go through with my students. Sure, the progression of questions could be improved, but we will consider this more when we get to Intelligent Practice at the end of this chapter. But it is what is missing that is important.

A relatively obvious step would to be to include an example involving algebra, such as: $\frac{10a}{12a}$

But what about examples such as this:

$$\frac{-6}{15}$$

$$2\frac{2}{6}$$

$$\frac{10}{5}$$

These either look different, or in the case of the third example, have a weird result. Students are thus encouraged to pause, self-explain what is going on, and thus broaden their understanding of the concept of simplifying fractions.

And it is not just fractions. Polygons with convex angles, straight lines in the form $x = ?$, linear equations like $4 - 2x = 10$, the mean of algebraic terms, and quadratic expressions that do not factorise, are all possibilities that spring to mind. If these types of examples are left out, it may be possible for students to appear as though they have understood a topic, whereas in fact they only have a surface level of understanding.

Mason and Watson go further, extolling the benefits of challenging students to create their own boundary examples, explaining: 'If you cannot construct boundary examples for a theorem or a technique, then you do not fully appreciate or understand it'. This is likely to be challenging for novice learners, especially in the initial stages of skill acquisition, and we need to consider the fragile nature of their working memories. However, challenging students to construct a particular example, then a peculiar example (*eg* one which no one else in the class is likely to think of), and then (if appropriate) a general, or at least maximally general example, seems a really useful practice to develop, all whilst considering the burden on students' working memories.

What I do now

The inclusion of boundary examples is crucial in order for students to have a more holistic understanding or appreciation of a concept. Failure to include them can lead to the problems of under-generalisation discussed in the previous section. Hence, when planning the examples I will present to my students, I now always consider including the more extreme, strange cases that test the boundary of the concept under consideration.

There is a potential danger when it comes to boundary examples – that they will simply be too much for students to cope with. We have seen in Chapter 4, and throughout this book, the dangers of overloading students' working memories, and too-complex examples too soon risk doing exactly that. So, the temptation may be to leave their introduction until later in the learning process. But to do so may – perhaps counter-intuitively – make them into the weird, tricky examples we don't wish our students to perceive them as.

For example, consider straight-line graphs. If we start with the basics (lines in the form $y = mx + c$), and then only introduce boundary examples after students are comfortable, then these examples are likely to be viewed as out of the ordinary, and hence more difficult. However, if our sequence of examples looks like this:

$y = 4x + 7$

$y = 4x - 7$

$y = 7 - 4x$

Then the third one will appear no stranger than the first two, and a more holistic understanding of straight-line graphs should develop.

The dangers of excluding boundary examples in the development of a concept were really brought home to me when analysing the following question on my Diagnostic Questions website:

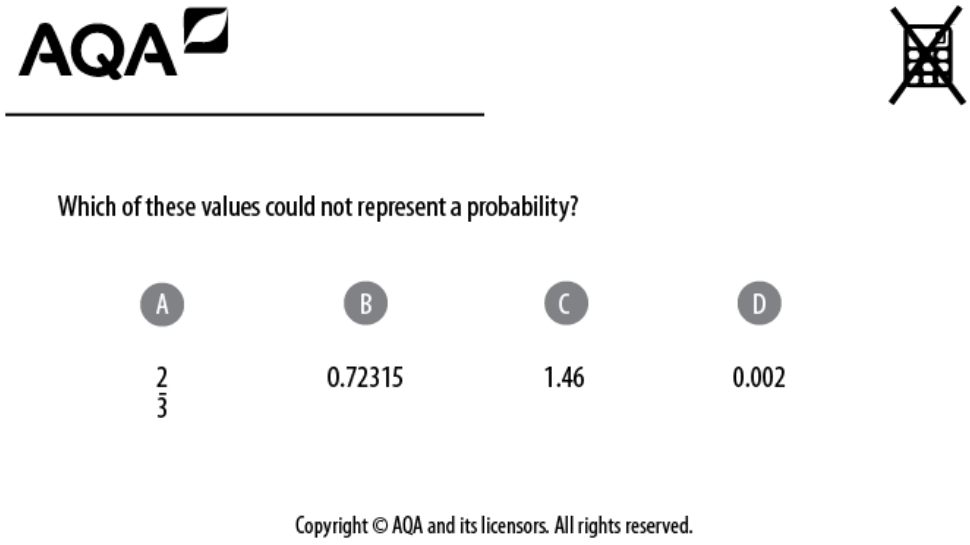


Figure 7.7 – Source: AQA for Diagnostic Questions

Surely the rule: *probabilities must be less than or equal to 1* is about as straightforward as it gets in maths? But why, then, did 47% of the 5000+ students who answered this question get it wrong?

A few students' explanations reveal all:

I think B because it's just a massive decimal and the rest look pretty legit. I also don't see how a number that big could be correct

I think B because you wouldn't see this in a probability question.

I think D because you can't have a 0.002 as an answer because it is too low.

If students are only used to meeting 'nice-looking' probabilities during examples and practice questions, then it is little surprise they come a cropper when they encounter strange-looking answers. Perhaps the solution is to present students with the following exercise, so they may form a more complete understanding of the concept of what is and what is not a probability:

Which of the following could represent probabilities?

1. 0.3

2. -0.3

3. 1.3

4. 0.03

5. 0.43045783

6. 1.43045783

7. -0.43045783

8. 0.4

9. 0

10. 1

11. 2

12. -1

13. $\frac{2}{3}$

14. $\frac{3}{2}$

15. $\frac{43}{51}$

16. $1\frac{2}{3}$

Figure 7.8 – Source: Craig Barton

A final point to make is that these boundary examples do not necessarily need to be covered as worked examples. As I will demonstrate at the end of this chapter when we look at intelligent practice, they can easily be incorporated into a well-designed sequence of questions that students answer independently. This gives students a chance to consider the strange question or outcome at their own pace, self-explain, and seek help if needed.

OceanofPDF.com

7.4. Same Surface, Different Deep Problems

What I used to think

Imagine I have just taught a lesson on the basics of Pythagoras's theorem to my Year 9s. Following a formative assessment strategy (Chapter 11), I am confident that the majority of my students can work out the hypotenuse and non-hypotenuse of a right-angled triangle. What would I do next? It's obvious! I would get the students practising applying Pythagoras's theorem in a variety of different contexts (leaning ladders, drifting boats, all the classics). This would all be topped off by their 'Pythagoras Homework' – a two-sided affair consisting of skill-based Pythagoras questions and then Pythagoras application questions. And each time my students successfully used Pythagoras to solve one of these contextual problems, I would feel content in the knowledge that we had nailed another topic. Not only could my students successfully do Pythagoras, they could apply it as well.

The strange thing was, a few weeks later when a Pythagoras application question came up in an exam, my students seemed to miraculously lose the ability to apply Pythagoras's theorem.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Birnbaum, M. S., Kornell, N., Bjork, E. L. and Bjork, R. A. (2013) 'Why interleaving enhances inductive learning: the roles of discrimination and retrieval', *Memory & Cognition* 41 (3) pp. 392-402.
- Goldenberg, P. and Mason, J. (2008) 'Shedding light on and with example spaces', *Educational Studies in Mathematics* 69 (2) pp. 183-194.

- Rohrer, D., Dedrick, R. F. and Burgess, K. (2014) 'The benefit of interleaved mathematics practice is not limited to superficially similar kinds of problems', *Psychonomic Bulletin & Review* 21 (5) pp. 1323-1330.

My takeaway

Rohrer *et al* (2014) argue that the correct solution to most mathematical problems involves two steps: identify the strategy needed to solve the problem, and then successfully carry out that strategy. By following up my Pythagoras lesson with a series of application problems, all of which require the use of Pythagoras, I am denying my students the opportunity to practise that first step – identifying the strategy. Hence, all I was really testing when giving my students contextual problems was whether or not they could use Pythagoras's theorem, not whether they could recognise when Pythagoras's theorem was (and crucially was not) needed.

The problems I was giving my students all had the same deep structure (they all required the use of Pythagoras to solve them), but had different surface structures – ladders, boats, diagonals of football fields, and so on. I believed that by presenting students with different surface structures and focusing on the similarities between them (*eg* where is the clue that a right-angled triangle is needed?), I was helping my students to develop the ability to apply Pythagoras's theorem in a wide variety of contexts. However, my students already knew what the deep structure of the problem was – it was the same as all the others! As we saw in Section 1.2, and will meet again in Chapter 9, one of the distinguishing features between experts and novice learners is the former's ability to recognise the deep structure of a problem. By denying students the opportunity to practise and develop that ability, it should have been of little surprise that my students were unable to transfer their knowledge a few weeks later.

There are two ways of dealing with this. In the long term we can make use of both the spacing and interleaving effects, ensuring that following the sequence of lessons on Pythagoras, students never go more than a few weeks without encountering a Pythagoras question – both straightforward and

contextual – in either a starter, low-stakes quiz or homework. Engineering such desirable difficulties is the subject of Chapter 12.

But in the short term we can make use of non-examples, whilst also harnessing the spacing and interleaving effects on topics studied in the past. As explained by Birnbaum *et al* (2013), instead of just focusing on similarities between members of a group, the presence of non-examples forces students' attention on the differences, thus allowing for a greater overall understanding.

What I do now

Now, doing this is a little bit tricky, but not impossible. It requires us to devise a series of problems that are the polar opposite of traditional application questions. We need a series of questions, all of which have similar surface structure but different deep structures. That way, students are forced not just to focus on the similarities between questions, but also the differences, to enable a deeper understanding of contexts.

So, having learned the basics of Pythagoras, we may present students with the collection of problems below (Figure 7.9).

The surface structure is identical in each case – an isosceles triangle – but the deep structures are very different. And, of course, once students have studied topics such as right-angled trigonometry, proof and further trigonometry, these also can be presented with similar-looking surface structures and added to the mix.

Or imagine you have just taught your students how to calculate frequency density from a grouped frequency table in order to draw a histogram. Sure, you could give them loads of frequency tables to practise from – but the danger here is that they develop the belief: 'Every time I see a grouped frequency table, I just need to go through the motions to find frequency density'. It is this cruising on autopilot, falling back on familiar routines regardless of the subtleties of the question, that can really hinder students' development in mathematics. But fortunately, some carefully designed questions can help, with the added bonus of revising previously studied concepts (Figure 7.10).

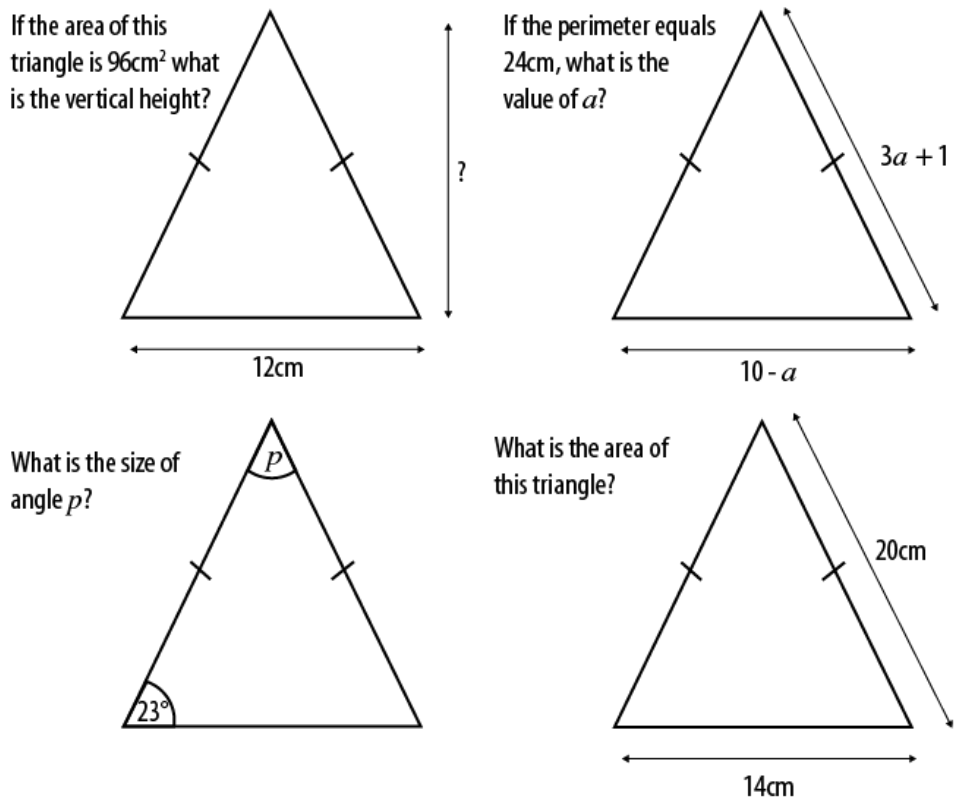


Figure 7.9 – Source: Craig Barton

Work out the information needed to allow you to draw a cumulative frequency diagram

Weight (Kg)	Frequency
40 ≤ w < 50	10
50 ≤ w < 60	18
60 ≤ w < 70	22
70 ≤ w < 80	6
80 ≤ w < 90	4

Work out the information needed to allow you to estimate the mean weight

Weight (Kg)	Frequency
40 ≤ w < 50	9
50 ≤ w < 60	15
60 ≤ w < 70	20
70 ≤ w < 80	11
80 ≤ w < 90	5

Work out the information needed to allow you to determine which group contains the median weight

Weight (Kg)	Frequency
40 ≤ w < 50	18
50 ≤ w < 60	14
60 ≤ w < 70	8
70 ≤ w < 80	4
80 ≤ w < 90	6

Work out the information needed to allow you to draw a histogram

Weight (Kg)	Frequency
40 ≤ w < 50	12
50 ≤ w < 60	20
60 ≤ w < 70	18
70 ≤ w < 80	8
80 ≤ w < 90	6

Figure 7.10 – Source: Craig Barton

I have decided to name these Same-Surface-Different-Deep problems, or SSDD for short – catchy, eh? I think this could be my next obsession. Exposing students to these batches of questions at the end of the initial

knowledge acquisition stage should help them better develop the all-important ability to transfer their knowledge to different situations, which is a vital component of problem-solving. We will delve into this more in Chapter 9, where I will share two more SSDD selections.

OceanofPDF.com

7.5. Ambiguous answers

What I used to think

I would often ask students to give me examples of the concept we were discussing. 'Draw me an example of a parallelogram' might be a challenge I would lay down to my Year 7 class before embarking upon a lesson on the properties of quadrilaterals.

What could possibly go wrong?

Sources of inspiration

- Goldenberg, P. and Mason, J. (2008) 'Shedding light on and with example spaces', *Educational Studies in Mathematics* 69 (2) pp. 183-194.
- Mason, J. and Watson, A. (2001) 'Getting students to create boundary examples', *MSOR Connections* 1 (1) pp. 9-11.

My takeaway

Asking students to provide an example of a concept is a powerful way to assess their understanding. However, Goldenberg and Mason (2008) point out a potential issue.

Take the challenge I issued my Year 7s. If a student responded to this by drawing a square, what would you think?

Interpretation 1: the student does not understand what a parallelogram is at all

Interpretation 2: the student has a really deep understanding of parallelograms and has chosen to demonstrate this understanding by drawing a special case.

Without further probing, it is impossible to know. Hence, if I am asking this question to instigate a discussion, then it is fantastic. If I am asking as a means of quickly assessing understanding, then I may have a problem on my hands.

What I do now

If I need a quick assessment of understanding, I try to reduce the potential for ambiguous answers that are thus difficult for me to interpret. My favourite way to do this is with a 'circle the…' activity, including a variety of examples and non-examples. So here I would have 'circle the parallelogram', and ensure the square was part of the selection of shapes. As discussed in Section 7.1, this activity can be done as a silent demonstration, whole-class votes, or printed out and completed individually.

And these are not just limited to geometry topic. Hours of fun can be had with the likes of 'circle the surd', 'circle the expression', and so on. I also like to use well-written diagnostic multiple-choice questions to help reduce the ambiguity of responses – something we will delve into in Chapter 11.

None of this is to say that asking students to produce their own examples is a bad thing. Indeed, in Section 7.3 we saw Mason and Watson (2001) recommend challenging students to come up with examples right on the boundary of concepts. My point is simply that such answers may prove difficult to interpret quickly, and as such we need to be careful if using them as part of a formative assessment strategy.

OceanofPDF.com

7.6. Ambiguous questions

What I used to think

I confess that I did not used to think too carefully about the examples I gave my students. Hence, I have blissfully used all of the following examples in the past:

Evaluate 2^2

Which is bigger: 0.7 or 0.85?

What is the area of a rectangle with base 6cm and height 3cm?

Simplify $\frac{16}{64}$

Write as an improper fraction: $2\frac{3}{10}$

Sources of inspiration

- Bills, L., Dreyfus, T., Mason, J., Tsamir, P., Watson, A. and Zaslavsky, O. (2006) 'Exemplification in mathematics education', *Proceedings of the 30th conference of the international group for the psychology of mathematics education, Vol 1*. Prague: Charles University, pp. 126-154.
- Rowland, T., Thwaites, A. and Huckstep, P. (2003) 'Novices' choice of examples in the teaching of elementary mathematics' in Rogerson, A. (ed.) *Proceedings of the international conference on the decidable and the undecidable in mathematics education*. Brno: Unknown publisher, pp. 242-245.

My takeaway

In the previous section we looked at the potential difficulties of ambiguities in student responses. However, unless we are careful, there is a real danger that we as teachers can be the propagators of a sea of ambiguity ourselves.

Bills *et al* (2006) explain that a teacher's poor choice of examples can have a detrimental effect on learning by making it more likely students will jump to the wrong conclusions. Likewise, a study by Rowland *et al* (2003) documents this for novice teachers in a primary setting, where the unintentionally 'special' nature of an example can mislead learners.

Consider the innocent-looking examples I shared above and think what misconceptions they may help to foster. Or how about this one from @MathsMrCox on Twitter: 'The median of 4 and 5 is 4.5, so the median of 3 and 6 must be 3.6'.

If students can get examples correct in multiple different ways, but only one of these ways is correct and transferable to other questions, then we could well be asking for trouble, as the dangers of over-generalisation once again rear their ugly heads.

What I do now

Subject knowledge is so much more than just being able to do maths. It is about predicting where and why students are likely to go wrong, and doing all that we can to help them. The careful selection of examples is a key component to this.

As we will see when we get to Chapter 11, I am more than a little obsessed with misconceptions. A key criterion I use when writing diagnostic multiple-choice questions is that students should not be able to get the question correct whilst still holding a misconception. When planning the (non-multiple-choice) examples I give my students, I apply the same criterion.

It is not easy to do, and inevitably I still slip up. Indeed, when presenting the following examples of solving linear equations at a conference, Senior Ofsted Inspector Jane Jones pointed out that I could inadvertently be

propagating the misconception that the coefficient of x always matches the solution!

$2x - 3 = 1$

$3x - 2 = 7$

Luckily, Ofsted do not grade conference presentations.

So, I am now acutely aware that an ambiguous example – or, more accurately, an example with multiple ways of arriving at the correct answer, only one of which is correct – can lead to over-generalisations and misconceptions that can be a nightmare to dislodge. Experience, running examples past colleagues, and learning from past mistakes seem to be the best ways to cope.

OceanofPDF.com

7.7. Extension questions

What I used to think

I used to think that the most effective thing to do was to go through basic examples of a concept in class, nice and slowly, nice and clearly, and then use more challenging questions as extensions for the students to try on their own, either in class or as part of homework. After all, when did my students need my help the most? For the basics, of course! And if students didn't get onto those extension questions in class, or left them out when doing their homework? Well, no big deal: they were just bonus work.

Sources of inspiration

- Barton, C. (2017) 'John Corbett', *Mr Barton Maths Podcast*.
- Frost, J. (no date) *Dr Frost Maths*. Available at: http://www.drfrostmaths.com/resources/
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.
- OCR (2015) *GCSE Check in Tests*. Available at: http://www.ocr.org.uk/qualifications/gcse-mathematics-j560-from-2015/planning-and-teaching/
- Quinn, D. (2017) 'Never let me go', *Until I Know Better blog*. Available at: https://missquinnmaths.wordpress.com/2017/05/31/never-let-me-go/

- Sweller, J. (1988) 'Cognitive load during problem solving: effects on learning', *Cognitive Science* 12 (2) pp. 257-285.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

Extension questions, whether they are part of classwork or homework, will inevitably be the most challenging questions. The chances are, they will also be the most interesting. Take, for example, the extension question at the end of the Fractions GCSE check-in test from OCR:

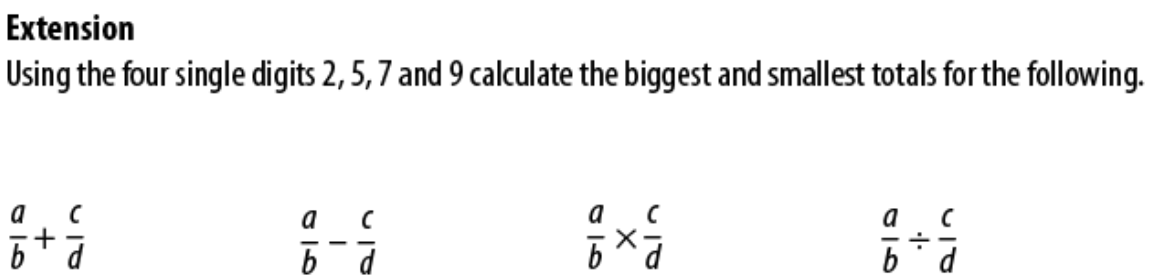


Figure 7.11 – Source: OCR, available at http://www.ocr.org.uk/qualifications/gcse-mathematics-j560-from-2015/planning-and-teaching/

What a great question!

Or how about this little beauty from Dr Frost's lesson on surds:

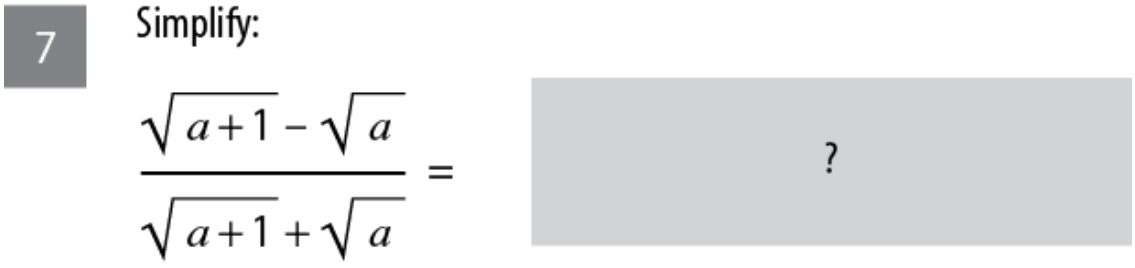


Figure 7.12 – Source: Jamie Frost, available at http://www.drfrostmaths.com/resources/resource.php?rid=242

A classic!

And yet, where does providing structured guidance and support for these types of questions rank in our priorities compared to the more straightforward questions? For me, it was pretty low down. If some students got them correct, great. If not, maybe I would quickly show the answers, but I certainly would not spend the kind of time and effort presenting the solution to problems like these in the carefully considered manner I would go through other examples, as explained in Chapter 6.

It was a 2017 blog post from Dani Quinn that first got me thinking about this. Dani argues that we model examples that are too easy, then leave pupils to 'discover' methods for the harder ones. This is tough enough for students in class, but including questions like this in homework, at a time when they are faced with their lowest point of support, could make homework incredibly frustrating. And on a purely practical level, when a child says they couldn't do this aspect of their homework because they did not understand it, it is pretty difficult to discern whether they are telling the truth or are making an excuse and have not even bothered attempting it.

Proponents of Cognitive Load Theory would certainly agree with this. Kirschner *et al* (2006) argue that unguided instruction is usually less effective than explicit instruction unless the learners are domain-specific experts. As discussed in Chapter 6, if left to their own devices, novice learners are likely to attempt to solve challenging problems by embarking upon a means-end analysis. This is a cognitively demanding strategy whose function is to find a solution to a specific problem, and not to alter long-term memory, which is necessary for learning to take place. In contrast, studying a worked example both reduces working memory load because search is reduced or eliminated and directs attention (*ie* directs working memory resources) to learning the essential relations between problem-solving moves. Students learn to recognise which moves are required for particular problems, which is the basis for the acquisition of problem-solving schemas. Hence, the kind of challenging, interesting extension questions that we expect students to do on their own are the very ones in which they both need and will benefit most from our structured support.

But where on earth do we find the time to dedicate to going through more challenging, interesting problems like the one above? For Dani, the answer is

to quickly model a basic example, get the students to practise a related problem (the example-problem pair approach described in Chapter 6), but then to move as fast as you can to the harder content. Dani is happy for the students to simply 'mimic' when going through the basics in the example-problem pair stage – in other words, as she explains, 'they do a very similar one on whiteboards so I can see if they can recall and follow the steps' – arguing that this is much more important than them being able to articulate the steps at this early stage.

There is another possibility. When I interviewed John Corbett for my podcast, he explained in depth how he has made flipped-learning work for his students. He describes a lesson on circumference of a circle, whereby all the basic procedural steps are completed by the students themselves prior to the lesson. John provides a link to a video (he has hundreds of top-quality ones on his website, corbettmaths.com) together with a worksheet. Students complete the work at their own pace at home, and email John their answers. The start of the lesson then comprises a quick assessment of understanding, and then there is a whole 50 minutes to dedicate to the kind of interesting, challenging problems that students really need support with. Now, there are obviously potential issues with non-compliance and technology, all of which we discuss in the interview, but there is little doubt that John has more time to support his students on complex questions than I do.

What I do now

Dani Quinn argues that we should plan harder examples in class, go through the basics at good pace and with little fuss, and use homework to practise and consolidate things students already know how to do. This has the double-advantage of freeing up class time to spend on more difficult examples, and ensures that non-compliance on homework is entirely the fault of the student – it is a lack of effort, not a lack of understanding.

I have certainly taken Dani's 'little fuss for the basics' point on board. In the past, I have been guilty of milking every last drop out of the most basic of examples – 'But *why* is $2x \times 5y = 10xy$, Jenny? I need convincing!'. Far better to go through these examples more quickly – mini whiteboards are ideal for this – making it crystal clear what the correct answer is, and if necessary

allowing students the opportunity to fill in any gaps in knowledge through the Intelligent Practice (Section 7.8) that follows the examples. As we have seen in Section 3.9, and in the work of Willingham (2002), the development of such inflexible knowledge is a necessary step on the path to expertise. A more complete understanding can come later.

However, I have one caveat. We have seen throughout this book – and will see again in Chapter 9 – just how important domain-specific knowledge is. It is the distinguishing feature between novice learners and experts, and it fundamentally changes the way students think and how they respond to different forms of instruction. Omitting or going too fast through basic examples runs the risk that students may miss something vital. I have witnessed too many students miss something in the early stages of learning a skill or concept, and then fake their way through the more complex stuff, either because they are too afraid to admit they are struggling, cannot be bothered, or do not know they are struggling. And, of course, if we are providing structured support for these more complicated questions, it is even easier for these students to slip under the net compared to if we were asking them to try them independently. We need to go quickly through the basics, but not too quickly!

All of this is not to say that I do not fully agree that we need to spend more time offering support on the types of questions where students will really benefit. Indeed, I specifically allocate time to go through these more interesting and challenging questions – and if I do not get time that lesson, I will make time the next.

To treat these questions as bonus material would be a waste, and asking students to struggle alone on them, either in class or at home, is potentially dangerous. I have said before, and I will say again before this book is through – too much struggle is bad, both for motivation and for learning. Moreover, challenging and interesting questions should neither be optional, nor reserved for the chosen few. If we have high expectations of all students, and we provide appropriate support and guidance, then I truly believe students will rise to meet them. Example-problem pairs or Supercharged examples provide a robust way to give students the very best possible chance to understand a difficult concept or a question. We just need to ensure the

basics are in place to enable students to reap the maximum benefit from this approach.

OceanofPDF.com

7.8. Minimally different examples and Intelligent Practice

What I used to think

As is becoming ever more apparent as we journey through this chapter, I never used to pay anywhere near enough attention to the examples I chose for my students. But if you think that was bad enough, I am only really getting started. Because my decisions about the sequencing of both worked examples and the subsequent practice I gave my students to do were – quite frankly – disgraceful. Strap yourselves in for this one, as there is a bumpy ride ahead.

Sources of inspiration

- Barton, C. (2016) 'Bruno Reddy', *Mr Barton Maths Podcast*.
- Barton, C. (2017) 'Kris Boulton', *Mr Barton Maths Podcast*.
- Butterfield, B. and Metcalfe, J. (2006) 'The correction of errors committed with high confidence', *Metacognition and Learning* 1 (1) pp. 69-84.
- Marton, F. and Pang, M. F. (2013) 'Meanings are acquired from experiencing differences against a background of sameness, rather than from experiencing sameness against a background of difference: putting a conjecture to the test by embedding it in a pedagogical tool', *Frontline Learning Research* 1 (1) pp. 24-41.
- Quinn, D. (2017) 'Never let me go', *Until I Know Better blog*. Available at: https://missquinnmaths.wordpress.com/2017/05/31/never-let-me-go/

- Renkl, A. (1997) 'Learning from worked-out examples: a study on individual differences', *Cognitive Science* 21 (1) pp. 1-29.
- Singh, A., Marcus, N. and Ayres, P. (2012) 'The transient information effect: Investigating the impact of segmentation on spoken and written text', *Applied Cognitive Psychology* 26 (6) pp. 848-853.
- Watson, A. and Mason, J. (2006a) 'Seeing an exercise as a single mathematical object: using variation to structure sense-making', *Mathematical Thinking and Learning* 8 (2) pp. 91-111.
- Watson, A. and Mason, J. (2006b) 'Variation and mathematical structure', *Mathematics Teaching incorporating Micromath* 194 (January) pp. 3-5.

My takeaway

Variation Theory

In order to understand just how poor my sequencing of examples and practice questions was, we first need to venture into the fascinating world of Variation Theory.

According to Marton and Pang (2013), the key conjecture of variation theory is that 'meanings are acquired from experiencing differences against a background of sameness, rather than from experiencing sameness against a background of difference'. I love that phrase so much, I am going to say it again, and in bold: ***meanings are acquired from experiencing differences against a background of sameness***. Hence, careful variation is seen as the key to learning.

We have seen from Willingham that students remember what they think about or attend to. In Section 3.3, I argued that a drawback of less-guided forms of instruction was that we have less control over what students direct their limited attention towards. With the freedom of approach that such forms of instruction bring, students are less likely to focus on the features that are the most important for developing conceptual understanding or procedural fluency. Incorporating the principles of Variation Theory into the examples and questions we give our students is the very opposite of this. By

holding everything else the same and varying one element, we can direct students' attention to that element so they can predict and observe the effect it has on the answer.Watson and Mason (2006a) extol the virtues of intelligently varied questions. They explain: 'if learners think that mathematical examples are fairly random, or come mysteriously from the teacher, then they will not have the opportunity to experience the expectation, confirmation and confidence-building which come from perceiving variations and then learning that their perceptions are relevant mathematically'.

Indeed, these expectations may prove even more valuable if they are not realised, due to the Hypercorrection Effect. The Hypercorrection Effect refers to the finding that errors committed with high confidence are more likely to be corrected than low-confidence errors (*eg* Butterfield and Metcalfe, 2006). In other words, when you find out that something you were convinced was correct is in fact incorrect, there is a bigger improvement in the change to your thinking than if you were not so sure. Hence, if students are expecting a particular answer due to the logical variation of the sequence of questions, and then discover that the answer is in fact something else, it may lead to a bigger improvement in their learning than if they simply arrived at that answer without forming that expectation. I will discuss this further in the examples that follow, and we will return to the power of the hypercorrection effect in Chapters 8, 11 and 12.

Watson and Mason (2006b) conclude: 'control of dimensions of variation and ranges of change is a powerful design strategy for producing exercises that encourage learners to engage with mathematical structure, to generalize and to conceptualize even when doing apparently mundane questions'.

For me, that is the true power of the Variation Effect. In isolation, each question is not at all special. A little dull, in fact. But when included in a carefully constructed sequence, these apparently mundane questions combine to produce something rather special. By working through carefully chosen sequences of questions, students have to carry out procedural operations, thus engaging in vital practice. But through connected calculations, they also have the opportunity to consider the deeper structure.

Such variation allows students to anticipate, notice and then generalise, instead of permanently playing catch-up.

If only I had known about this sooner...

Minimally Different Examples

Okay, so let's say I was teaching something like expanding double-brackets. My sequence of worked examples might be something like:

$(x + 2)(x + 3)$

$(x + 5)(x - 1)$

$(x - 4)(x - 3)$

$(x + 3)^2$

$(2x - 1)(x + 5)$

There was a logic here. Notice how the first example has two plus signs, the second has a minus, and so on? And then see how I introduce a coefficient of 2 on one of the x terms? Nice, eh?

Definitely not. There is simply too much variation. Between examples 1 and 2, signs and the numbers inside the bracket change. Hence, it is difficult for students to spot exactly why the answer to example 1 is different to example 2. What features are they supposed to direct their limited attention to? The examples become separate entities, completely unrelated to each other. No predictions can be made, thus coherent expectations are less likely to be formed, schemas are less likely to be developed, and the change in long-term memory that is necessary for learning may not materialise. This is a golden opportunity missed.

I was introduced to the concept of minimally different examples by Bruno Reddy when he gave a talk at the first ever MathsConf in Kettering in 2014, and I have become obsessed ever since. Bruno was a guest on my podcast where, amongst many other things, he discussed the use of minimally different examples in his lessons. The key is that just one aspect of a particular example is changed each time. I think of this like a good science

experiment – just one variable is changed, the effect on the outcome is observed, and that change can thus be attributed to that variable.

On the next page is the sequence of examples and practice problems, using Section 6.2's Example-Problem Pair approach, I now use when introducing expanding double brackets.

These examples are certainly not presented all at once. I write the first example on the board, solve it in silence using the model of Silent Teacher discussed in Section 4.8, and then narrate over the top of my workings, explaining how and why I have done each step. Then students try the related problem, usually on mini whiteboards. I then follow the advice of fellow podcast guest, and Bruno's former colleague, Kris Boulton: I make an exact copy of the first example underneath it, and make a big deal out of rubbing out the term that I changed to create the second example, so students are made explicitly aware of the change and can thus notice the subsequent effect it has on the answer. The reason I make a copy of Example 1 before making the change (as opposed to rubbing out a feature of Example 1 itself) is that I want students to observe the original example and the answer so they can compare and contrast the minimally different follow-up. I then write their second problem next to this, and so on.

Students remember what they think about or attend to, so by holding everything else constant and changing one thing, I am directing their attention towards the feature of the example that matters. That is the key to good variation.

Example	Your Turn
$(x+2)(x+3)$	$(x+4)(x+5)$
$(x+3)(x+2)$	$(x+5)(x+4)$
$(x-3)(x+2)$	$(x-5)(x+4)$
$(x+3)(x-2)$	$(x+5)(x-4)$
$(x+2)(x-3)$	$(x+4)(x-5)$
$(x-3)(x-2)$	$(x-5)(x-4)$

Figure 7.13 – Source: Craig Barton

You may feel that there is in fact far too little variation here, and we risk slowing students down. But that misses the point. In reference to this same concept, Watson and Mason (2006a) argue:

> At some stage students need to know how to multiply pairs of brackets containing two terms. Indeed they have already done so when learning two-digit multiplication. At some stage they also need to incorporate knowledge about signs and knowledge about multiplying more complicated algebraic terms into this process – but first let's build up fluency and familiarity with the basic process. If this doesn't happen, learners are likely to get stuck on exercises not because they do not know how to 'remove brackets' but because they muddle signs and coefficients.

But going through this sequence of examples may not take all that long at all. Following Dani Quinn's advice from when we looked at extension questions in Section 7.7, all of this can be done with little fuss and at a good pace on mini whiteboards. These are quick-fire questions that do not require much discussion. Indeed, any benefits from discussion should be trumped by the connections students are able to make via the carefully chosen sequence of examples. If the sequence of examples is planned well enough, students will be able to predict the answer before working it out. As Watson

and Mason (2006a) explain: 'seeing their expectations confirmed can inspire confidence for learners'. Likewise, seeing their expectations not confirmed can be a powerful learning moment via aforementioned The Hypercorrection Effect.

Technology can really help this process, both in terms of pace and clarity. Dynamic geometry, in particular, offers a wonderful way to utilise the principles of variance. Using Desmos to show the effect of changing one aspect in the sketching of straight-line graphs can be very effective, and is much quicker than sketching by hand, thus allowing me to demonstrate more examples in a shorter period of time. Here is a sequence of lines I may plot:

$y = x$

$y = x + 1$

$y = x + 2$

$y = x + 3$

$x + 3 = y$

$y + x = 3$

$y = 3 + x$

$y = x - 3$

$y = x - 2$

$y = 2 - x$

$y - 2 = x$

My screen on Desmos looks as follows:

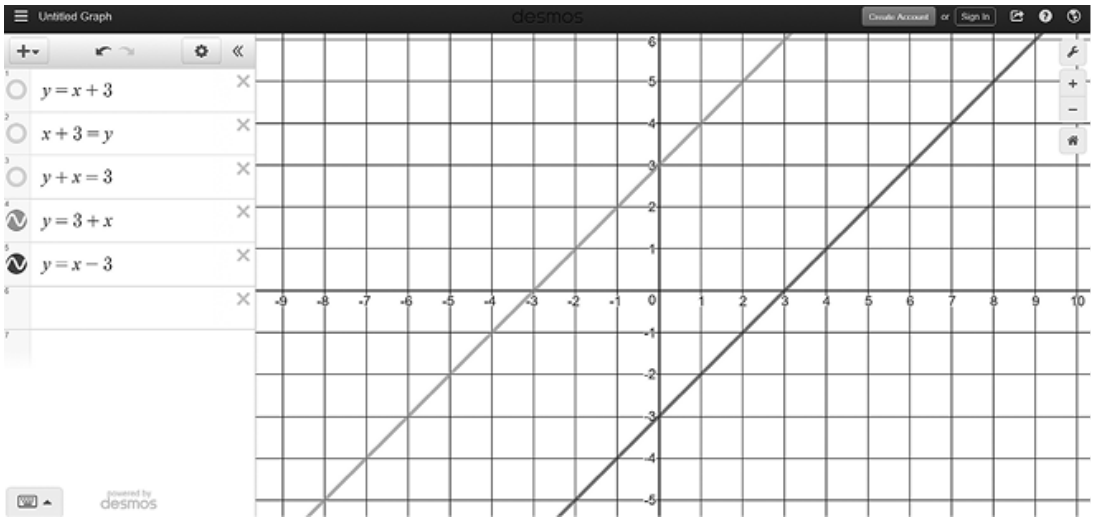


Figure 7.14 – Source: Craig Baton, created on Desmos

I recreate the effect of rubbing out by copying the equation of the line I have just drawn in the box on the left, highlighting the element I am going to change with my mouse, announcing the change so students have a chance to form an expectation of what will happen, and then typing in the new version. I also hide all but the previous line, thus having a maximum of two lines visible at all times to prevent the screen becoming too visually busy, thus making the effect of the variation crystal clear.

Of course, an alternative to copying out equations and manually changing one element can be achieved easily on Desmos via the animation of constants. So, I could type in the equation $y = mx + c$, and adjust the value of c by means of a slider. On the surface this seems like an excellent way to demonstrate variance, as everything else stays the same, so any movement must be down to changes in c. However, as a general rule I rarely do this – especially with younger students. I worry that the concept of a constant may be too difficult to comprehend for novice learners, and thus result in the overwhelming of working memories during early knowledge acquisition phase. I do not want students to view this as a demonstration where lines are whizzing magically about the screen. Likewise, the 'transient information effect' (*eg* Singh *et al*, 2012) tells us that when explanatory information disappears before it can be adequately processed it leads to inferior learning than more permanent sources of information.

No, I need my students to see this as a logical progression through a sequence of examples, where one thing is changing each time in a controlled manner, and we are pausing to observe and self-explain the effect.

Intelligent Practice

Just as I did not put too much thought into the worked examples I went through with my students, I was also pretty relaxed when it came to the choice of questions I would give my students to practise. Usually this would involve me grabbing any old selection of questions.

Hence, if I wanted some practice questions on finding the fraction of an amount, I might turn to the wonderful Mr Carter Maths site and project up something like the screen on the next page.

I would tell students to work their way down the Bronze questions, and then move onto Silver and Gold. And then at a click of a button, I could reveal the answers. Maybe teaching isn't that hard after all?

In Chapter 6, I outlined a teacher-led model for the presentation of worked examples that followed either an example-problem pair or a Supercharged approach. Both methods require independent student practice following the presentation of the worked example and related problem, and unless just as much care is taken over the choice of these questions as is over the examples we present, we risk undoing all our good work.

Well-designed practice questions can help students develop procedural fluency, whilst also allowing them to make the connections and cognitive leaps necessary to develop conceptual understanding. Furthermore, unlike teacher-led worked examples, students have an opportunity to work at their own pace throughout the practice questions. They have time to self-explain, go back over previous work, make and correct mistakes, and seek help from me or a peer, all whilst being unrestricted by the needs of the rest of the class. In short, the choice of questions we give our students to practise is of vital importance.

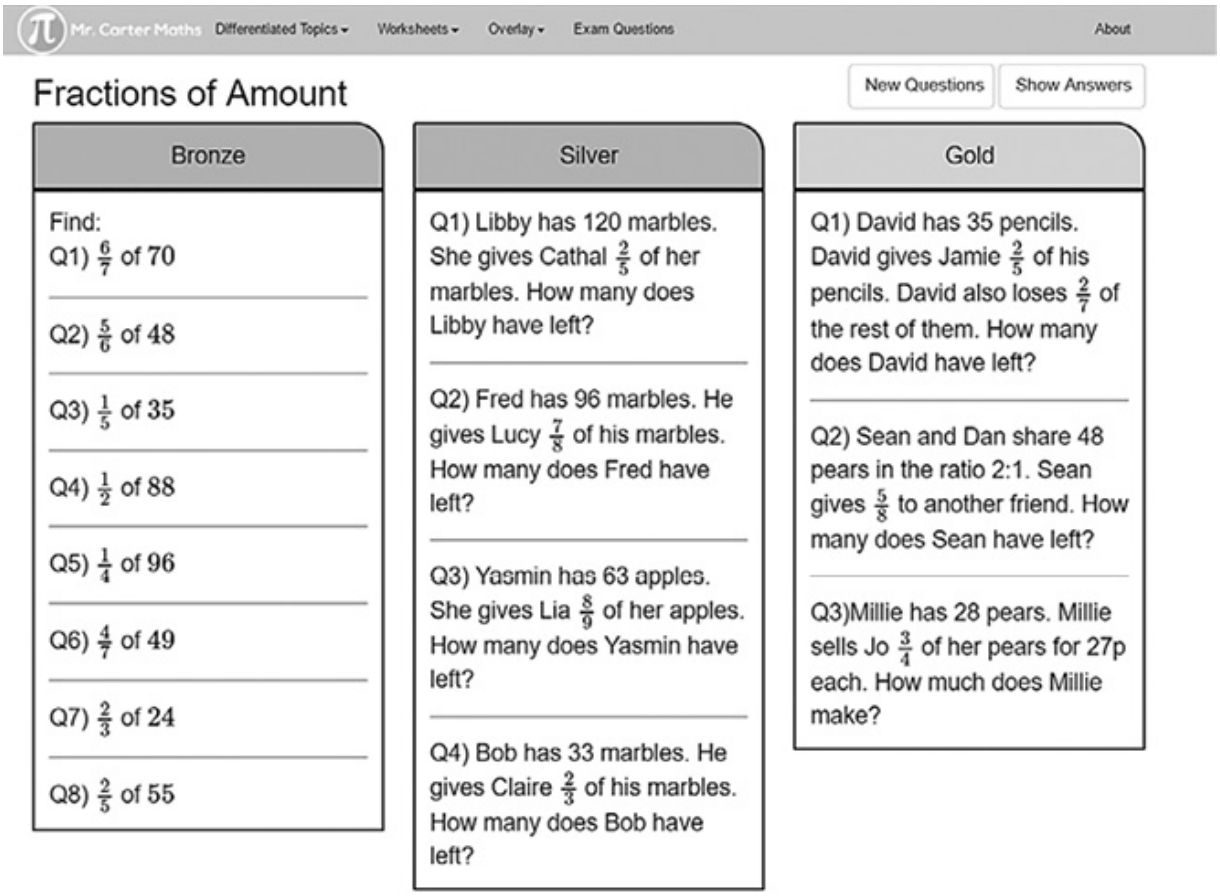


Figure 7.15 – Source: Darren Carter, available at http://www.mrcartermaths.com/

I favour the term *Intelligent Practice* to describe what I want to achieve in the choice of the questions I give my students. As well as the careful sequencing of questions, I also want to include the boundary examples discussed earlier in this chapter. These are questions that still flow naturally from the previous ones following the principles of variation theory discussed above, but which throw students off autopilot and cause them to think about their answers. Prompting students to pause and consider weird questions or results whilst they are working at their own pace is important because of the Self-Explanation Effect (Chapter 5), and to negate the dangers of under-generalisation identified earlier in this chapter.

It is early days in my journey into the fascinating world of intelligent practice, but here is my attempt to improve on the selection of practice questions for finding the fraction of an amount presented above. See if you

can discern my logic, before I give my commentary on the sequence of questions below:

1. $\frac{1}{5}$ of 60	**6.** $\frac{20}{50}$ of 300
2. $\frac{1}{5}$ of 30	**7.** $\frac{2}{5}$ of 3
3. $\frac{2}{5}$ of 30	**8.** $\frac{5}{2}$ of 3
4. $\frac{20}{5}$ of 30	**9.** $\frac{3}{2}$ of 5
5. $\frac{20}{50}$ of 30	**10.** $\frac{3}{5}$ of 2

Figure 7.16 – Source: Craig Barton

Generally, I want to direct students' limited attention to the things that are changing between questions, so they can predict and observe the effect it has on the answer. Students remember what they think about or attend to, and I have the opportunity to control that via this sequence of questions.

1. A simple unitary fraction of an amount to kick things off
2. The fraction remains the same but the amount halves. Students have the opportunity to notice that the answer halves as well.
3. This time the amount is constant but the fraction doubles. Students may predict the answer to this, and also have the opportunity to consider why the answers to Questions 1 and 3 are the same.
4. What happens if we multiply the numerator by 10, keeping everything else the same? Why is the answer the same as 4×30?
5. How about multiplying the denominator by 10? Does this surprise students?
6. Here all three elements from Question 3 have been multiplied by 10.

7. Now I am concerned that there is too much variation here, and maybe this question should follow directly from Question 3, but following our work on Boundary examples in Section 7.3, I want students to be aware that the answer does not always have to be a nice integer. I want them to encounter this now, in the early stages of knowledge acquisition, so it does not cause them to doubt their method at a later date. And I want to provide it here, amidst a logical sequence of examples, so students have an opportunity to reflect and self-explain where the answer has come from.

8. What happens if the fraction is flipped upside down? Students may predict that the answer is also flipped upside down, and be surprised when this is not the case. Here we again have the Hypercorrection Effect, and such a surprise may prove a valuable learning experience. The other thing I like about Question 8 is that the fraction is improper – in the past I have confined fractions of amount questions to nice proper fractions, for no good reason that I can think of.

9. What happens if the 3 and the 5 swap places?

10. Another possible surprise – why did the same thing not happen here as it did in Question 9?

I would hope that students are making these connections, and self-explaining as they go along. But we have seen from Renkl's (1997) work in Section 5.3 that the majority of students are not natural self-explainers. Hence, at the end of this sequence of questions I may have a statement along the lines of: 'look back at your answers and see if you spot any connections'. This is meant to serve as a prompt to allow students to take advantage of Chapter 5's Self-explanation Effect. If I consider some of these connections to be of particular importance, I may instigate a discussion about them following the presentation of the answers.

Similarly, I have created this sequence of questions on calculating the gradient between two points, which addresses a couple of other key features:

1. (0, 0) and (3, 6)	**6.** (0, 0) and (7, -2)
2. (3, 6) and (0, 0)	**7.** (0, 0) and (-2, 7)
3. (6, 3) and (0, 0)	**8.** (0, 0) and (-2, 0.5)
4. (-6, 3) and (0, 0)	**9.** (0, 0) and (-2, 2)
5. (-6, -3) and (0, 0)	**10.** (0, 0) and $(-\frac{2}{5}, \frac{2}{7})$

Figure 7.17 – Source: Craig Barton

In general, I have kept one point as the origin in each example. I feel varying this as well as the other features discussed below would be too much for students to attend to in one sequence of practice questions.

1. A relatively straightforward question with a positive integer solution to get the ball rolling.
2. This is intended to help students notice that the order that the pairs of coordinates are presented does not affect the size or value of the gradient. This is a feature that in the past I would have taken for granted, for the precise reason that I am an expert.
3. Changing the order of the pairs of coordinates does not matter, but changing the x and y values certainly does. This question also gives us a boundary example in the form of a non-integer solution.
4. What effect does one of the points having a negative value have if everything else stays the same?
5. How about if both values within a coordinate are negative?
6. This time we have a negative y-value, and also a nasty fractional gradient – students need to see that gradients are not always nice integers.
7. What if we have two negatives and a fractional gradient?

8. This is to show students that the values in the coordinates do not have to be integers.

9. Here we have a surprise, or a weird answer. The calculation is relatively straightforward, but what will students do with the result?

10. Finally, the division of fractions is interleaved. If all students do not get on to this question, then no problem. But for those who do, this represents a good challenge, as well as an opportunity to retrieve knowledge of a skill they may not have practised for a while.

Questions from teachers

When I present sequences of examples and practice questions like this, there are four common questions raised by teachers that I would like to address:

How do you differentiate?

Differentiation was once the bane of my life. I used to have about 18 different worksheets for each lesson, and my job quickly transformed from teacher into more like someone who hands out flyers on the street. Now I use my favourite form of differentiation – time. Differentiation is in the form of how long students take to answer these sequences of questions. The logic behind the sequencing dictates that the most important questions in building up understanding of a concept occur at the start. Extension material (often a related UKMT question) and interleaved concepts (see Chapter 12) will be included for those students who finish early. Just as important, time is made for students who do not manage to complete all the core questions. This could be at the end of class, five minutes over lunch, during an intervention session, or simply as part of homework. But this is rarely needed precisely because the questions have been designed in such a way that every student can succeed. The logical progression, the variation and invariance, gives students every possible chance of navigating their way successfully through the sequence of questions in a way that simply does not happen with sporadic, ill-thought-out exercises.

What if students do not spot the connections?

The first thing to say is that most students do, even if they do not realise it. Watson and Mason (2006a) explain that 'perceived variation and invariance

generates expectations which may be conscious or subconscious'. Second, the prompt to look back at their answers may help make the process of reflection and self-explanation more explicit. But, to be honest, I do not expect every child to spot every connection. The important ones I will discuss as a class, but some I will let pass. I expect every child to be able to do every single one of these questions, and eventually I expect every child to spot the connections that will enable them to answer the questions flexibly in the manner that I do. But this will take different amounts of time for different students, and I am okay with that. The key thing is that I provide all students with the opportunity to make those connections, and I can do that through intelligently chosen and sequenced questions.

Surely this takes ages?
It certainly does…if you drag out each question in the manner I used to do. I would squeeze every single last drop of life out of every question – 'convince me', 'but *why* is this the case?'. But these questions are presented with minimal fuss. More importantly, the sequencing lends itself well to students spotting connections and hence working faster than if the questions were less carefully chosen. Take the Fractions of Amount example. On average, I would expect a Year 7 class to take no more than five minutes to complete these ten questions. At the end I would simply project up the answers – practice makes permanent, after all – ask if there were any issues, discuss a couple of connections, and move on to another example-problem pair. If it takes longer, there is usually good reason for it, with an issue identified that needs resolving.

So, are you saying Mr Carter's questions are a load of rubbish?
Oh God, no. They are superb questions, and his site is simply a Godsend. The key point, however, is that I no longer use them during the early knowledge acquisition phase of learning. At this point students are very much novices and I want to carefully build up their conceptual understanding of the skill or concept we are covering, and carefully varied sequences of questions allow for this. Two weeks later, when I want an effective way to induce retrieval, making use of the spacing effect we will discuss in Chapter 12, Mr Carter's questions are exactly what I need. Indeed,

the fact that the questions are not related to each other is a strength, because each serves as its own independent test of retrieval.

Who needs discovery?

A final point to make is that all this is taking place during what is essentially a teacher-led form of explicit instruction. There are no open-ended, unstructured tasks, just really well-chosen and sequenced practice questions. And yet students are having the opportunity to make connections, form expectations, and think deeply. Watson and Mason (2006b) put it best:

> Careful choice of examples in which certain aspects vary and others remain unchanged can, in these ways, draw learners into engaging with mathematical structure and meaning throughout their mathematics lessons, not just in special lessons or tasks which are designed to encourage thinking. Furthermore, desire and need to cover the curriculum is not upset by this approach, because the mathematics being explored is curriculum content. But what is far more important is that learners' ability to see, expect and think is, through controlled variation, being used fruitfully to learn mathematical methods and concepts.

What I do now

When planning a lesson, the majority of my time is spent selecting both the choice and sequencing of examples and practice questions. Very rarely do I find sequences of questions that vary things in the way I want – often leaps are made which make the formation of expectations impossible – so I end up making more of my own. It is something that takes time at first, and I am by no means an expert at it. But the effect it can have on the volume of content students get through in lessons, and the connections they are able to make on the road to both procedural fluency and conceptual understanding, make it worth every single minute.

OceanofPDF.com

7.9. If I only remember 3 things…

1. Choosing examples is possibly the most important part of the planning process, making up for deficiencies and ambiguities in definitions, explanations and rules.
2. Non-examples can be an effective way of adding clarity and completeness to students' understanding of a concept.
3. Carefully designed progression and variation in examples and exercises can allow students to make the vital connections that develop both procedural fluency and conceptual understanding.

OceanofPDF.com

8. DELIBERATE PRACTICE

In 2016, I read the book *Peak: Secrets from the New Science of Expertise* by Anders Ericsson and Robert Pool. The book described how the concept of Deliberate Practice has been applied to fields as varied as chess and heart surgery in order to dramatically improve performance on the path to expertise. I loved the book, but the more I read, the more I wondered how much – if any – could be applied to the development of students' maths skills, specifically in the early knowledge acquisition phase of learning?

This chapter is an attempt to answer that question.

OceanofPDF.com

8.1. Breaking it down

What I used to think

I used to think that it was best to introduce complex processes as a whole. That way, students could get a better feel for the entire process from start to finish, experiencing the full logical development. Hence, I would take students step by step through how to complete a complex process, and then get them to practise.

A classic example of this is simultaneous equations. I would usually start with a relatively straightforward example, perhaps something like:

$4x + y = 17$

$2x + y = 9$

I would slowly take students through how to solve this, carefully articulating and annotating the thinking behind each decision I made.

Then, having covered that example, I would move on to something more complex, such as:

$12x - 2y = 8$

$5x + y = 18$

All building up to something like:

$5x + 8y = 11$

$7x + 3y = -1$

Aside from the obvious problem of batching all my examples together at the start of the lesson instead of following the example-problem pair approach

(Chapter 6) followed by a careful choice of practice questions (Chapter 7), surely there was nothing else wrong with what I was doing?

The problem came whenever a student would announce that they were no good at simultaneous equations. It was quite tricky to know where to begin to help them.

Sources of inspiration

- Barton, C. (2017) 'Kris Boulton – Part 1', *Mr Barton Maths Podcast*.
- Boulton, K. (2017) 'My best planning – Part 1', *To the real maths blog*. Available at: https://tothereal.wordpress.com/2017/08/12/my-best-planning-part-1/
- Ericsson, K. A., Krampe, R. T. and Tesch-Römer, C. (1993) 'The role of Deliberate Practice in the acquisition of expert performance', *Psychological Review* 100 (3) pp. 363-406.
- Ericsson, K. A. and Pool, R. (2016) *Peak: secrets from the new science of expertise*. Boston, MA: Houghton Mifflin Harcourt.
- Ericsson, K. A., Prietula, M. J. and Cokely, E. T. (2007) 'The making of an expert', *Harvard Business Review* 85 (7/8) pp. 114-121.
- Sweller, J., Van Merriënboer, J. J. G. and Paas, F. G. W. C. (1998) 'Cognitive architecture and instructional design', *Educational Psychology Review* 10 (3) pp. 251-296.
- Wiemann, C. (2007) 'The "curse of knowledge", or why intuition about teaching often fails', *APS News* 16 (10) (no pagination).

My takeaway

For Ericsson *et al* (1993), the principle of Deliberate Practice involves breaking down a complex process, isolating an individual skill and working on it, receiving regular and specific feedback so you can improve your performance. It is how professional musicians prepare for a recital – not by practising the piece from start to finish, but by concentrating on small sections of it at a time until they achieve mastery. It is how young footballers

are trained – not by playing 90-minute 11-a-side matches, but by focusing on specific drills, practising them over and over again until mastery is achieved.

While the worlds of football and music may seem miles away from our maths classrooms – and as such we need to be incredibly careful when drawing inferences – I could certainly see parallels. Specifically, I was drawn to the idea of breaking down a larger process, and using well-designed practice to help improve performance. Indeed, a term I particularly like from the Deliberate Practice literature is 'high-quality reps' to refer to this well-designed practice.

Hence, I began to consider just how many individual sub-processes or skills a relatively complex process such as 'solving simultaneous equations' could be broken down into. After a good deal of thought, and ignoring application and contextual problems (the difficulty with these will be covered in the next chapter), I came up with the following list of nine:

1. Decide if the equations are in the correct form.
2. Decide if we need to manipulate one or both equations.
3. Decide if we need to add or subtract.
4. Successfully add or subtract algebraic expressions, possibly involving negative numbers.
5. Solve a linear equation, possibly involving negative numbers.
6. Substitute the solution into an algebraic expression.
7. Solve another linear equation.
8. Substitute two solutions into an algebraic expression to check the answer.
9. Interpret the solution.

I was feeling quite happy with myself…and then I interviewed Kris Boulton for my podcast. During his fascinating two-hour description of how he

plans a sequence of lessons on simultaneous equations, he suggested *13* was a better number. Indeed, before going anywhere near deciding if the equations are in the correct form to be solved, Kris spends a significant amount of time helping his students develop the skill to know if in fact an equation or a pair of equations *can* be solved. Kris has since blogged about this planning process in a four-part series that I highly recommend.

I found it difficult to break the process of solving simultaneous equations down into sub-processes or individual skills precisely because that is not how I see them. I see solving a simultaneous equation as a single process. Why? Because I am an expert in the domain of solving simultaneous equations. I have mastered each of the individual sub-processes involved, and hence they are automated and easily retrievable from my long-term memory. I have also solved hundreds, if not thousands of simultaneous equations questions before. I can solve them without too much thought or strain on my working memory, and I have spare cognitive capacity to deal with any twists and turns. My students are novice learners, and as such they do not think the same way. Because I suffer from the curse of knowledge (*eg* Wiemann, 2007), for many years I failed to appreciate this crucial fact.

Unfortunately, the more subject knowledge we obtain, the greater our curse of knowledge. But that is precisely why subject knowledge alone is not enough to make a good teacher. Knowledge of where students go wrong in various processes is just as important. In Chapter 11 we will look at ways teachers can identify student misconceptions, but something I've found extremely useful is to talk to other people. When attempting to break down a process into smaller steps or sub-processes, I write down a list, ask colleagues to do the same, and then compare lists. It makes a great exercise in departmental meetings – slightly more productive than arguing about targets or the latest school marking policy. Better still, I challenge students to come up with a list. There is perhaps no better way to gain an insight into how novices think than to ask them.

Breaking a complex process like solving simultaneous equations down in this way has two key advantages. Our work throughout this book (Chapter 4 in particular), has stressed just how fragile the working memories of novice learners are. Trying to remember and carry our 9 or 13 steps in order to

solve a problem is likely to be cognitively demanding, and if cognitive overload occurs then no learning may take place. By breaking down a process into smaller steps and tackling each one at a time, we can ease the burden imposed upon our students' working memories. Eventually, when these skills have been mastered, the student can carry out more complex tasks as they can draw on the automated processes stored in long-term memory.

Secondly, it allows us to identify weaknesses. A student telling you they are no good at simultaneous equations does not give you any clue as to where in the process they get stuck. Likewise, filling out a Question Level Analysis (QLA) after an exam and finding out that some of your students scored 0 out of 5 on the simultaneous equations question and some got 2 out of 5 also tells you very little. Was it the substitution that they could not do, the solving of the equation, or do they in fact have a misconception with negative numbers? If the students who scored 0 had got the first part correct, would they in fact have solved the rest of the problem perfectly? By breaking down complex processes and tackling them using the five-stage process I will introduce in the next section, we are better able to identify and resolve areas of weakness before they jeopardise the entire process.

What I do now

When faced with introducing a complex process (in essence, one that involves multiple steps) to my students, I first break it down into as many separate steps as possible. I then embark on the five-stage process of Deliberate Practice that I will introduce in the next section.

OceanofPDF.com

8.2. The five stages of Deliberate Practice

What I used to think

As I have already confessed, I used to think the best way to introduce a complex process was to do it as a whole.

This time, let's take a process such as adding two fractions together. My examples would probably look something like this:

$$\frac{1}{5} + \frac{1}{4} = ?$$

$$\frac{2}{3} + \frac{1}{6} = ?$$

$$\frac{3}{7} + \frac{2}{5} = ?$$

I would emphasise the need to find a suitable common denominator, then look at the process of transforming the fractions so that the two denominators were the same, before finally explaining how we finish off the problem. My students would successfully answer my cue-laden prompts and questions, before embarking upon plenty of practice to ensure they too could carry out this process in one go, just like me.

Some students would inevitably struggle. But, I reasoned, that was probably because they had not been listening properly.

Sources of inspiration

- Barton, C. (2017) 'Kris Boulton – Part 1', *Mr Barton Maths Podcast*.

- Boulton, K. (2017) 'My best planning – Part 2', *To the real maths blog*. Available at: https://tothereal.wordpress.com/2017/08/19/my-best-planning-part-2/
- Ericsson, K. A., Krampe, R. T. and Tesch-Römer, C. (1993) 'The role of Deliberate Practice in the acquisition of expert performance', *Psychological Review* 100 (3) pp. 363-406.
- Ericsson, K. A. and Pool, R. (2016) *Peak: secrets from the new science of expertise*. Boston, MA: Houghton Mifflin Harcourt.
- Ericsson, K. A., Prietula, M. J. and Cokely, E. T. (2007) 'The making of an expert', *Harvard Business Review* 85 (7/8) pp. 114-121.
- Hin-Tai Ting (2017) 'Is this the best we can do? Part 7: the spacing effect', *Mathagogy Blog*. Available at: https://mathagogy.wordpress.com/2017/07/29/the-best-we-can-do-7-the-spacing-effect/

My takeaway

In the last section I discussed the need to break down a complex process into sub-processes for novice learners. Failure to do so means it is incredibly difficult to isolate the exact nature of difficulties students encounter, and tackling the process as a whole could lead to cognitive overload with no learning taking place. Now we will look at what we do next. Inspired by the work of Ericsson and others, I call this The Five Stages of Deliberate Practice:

1. Isolate the skill.
2. Develop the skill.
3. Assess the skill.
4. The final performance
5. Practise retrieval later.

Let's see this in operation when introducing how to add fractions.

1. Isolate the skill
As we saw in the previous section, we need to circumvent the curse of knowledge and think really carefully about the sub-processes involved. For example, the process of adding two fractions together could be broken down as follows:

1. Decide if the fractions are in a form ready to be added.
2. Decide on an appropriate common denominator.
3. Use knowledge of equivalent fractions to transform the fractions.
4. Complete the addition.
5. Simplify the answer.

2. Develop the skill
In order to develop these individual skills, we need to use the lessons learned from Chapters 6 and 7, specifically example-problem pairs, examples and non-examples, and Intelligent Practice.

So, for 'decide if the fractions are in a form ready to be added', I may start by writing up pairs of fractions on the board, one at a time, indicating in silence which ones are in a form to be added in the manner I described when introducing the concept of an equation in Section 7.1. Alternatively, I may look to develop conceptual understanding using diagrams of shapes split up into fractions to illustrate that it is the denominator that determines if the fractions are in the correct form.

I would then present students with the following set of questions:

Isolate Decisions: Can they be added?

$\frac{1}{3}+\frac{1}{3}$	Yes/No
$\frac{1}{3}+\frac{2}{3}$	Yes/No
$\frac{2}{3}+\frac{1}{5}$	Yes/No
$\frac{2}{3}+\frac{2}{5}$	Yes/No

$\frac{4}{5}+\frac{5}{4}$	Yes/No
$\frac{3}{8}+\frac{3}{4}$	Yes/No
$\frac{1}{5}+\frac{1}{15}$	Yes/No
$\frac{1}{50}+\frac{1}{100}$	Yes/No

Figure 8.1 – Source: Craig Barton

This is a simple yes/no activity. I would get students to complete it on mini whiteboards and it should take them no more than a minute. It has been designed to ensure that students know to focus upon the denominator and not the numerator, and avoid any traps concerning multiples. If the activity does take more than a minute, then it is probably because a misconception has been identified that needs addressing – and I would much rather address it here instead of trying to unpick it whilst wrapped up in the more complex process of adding fractions.

Similarly, when developing 'decide on an appropriate common denominator', I would complete an example-problem pair and then give my students something like:

Isolate Decisions: Which denominator?

$\frac{1}{5}+\frac{1}{5}$	Denominator:
$\frac{1}{5}+\frac{2}{5}$	Denominator:
$\frac{2}{5}+\frac{1}{5}$	Denominator:
$\frac{2}{5}+\frac{1}{3}$	Denominator:

$\frac{2}{5}+\frac{2}{3}$	Denominator:
$\frac{3}{8}+\frac{3}{4}$	Denominator:
$\frac{1}{8}+\frac{5}{12}$	Denominator:
$\frac{7}{16}+\frac{9}{24}$	Denominator:

Figure 8.2 – Source: Craig Barton

As discussed in Chapter 7, my aim is for a logical variation and progression through the questions, so students can focus on what has changed and what has stayed the same, and the subsequent impact upon their answers. Notice that the ability to successfully complete this activity requires students to understand and be able to find lowest common multiples. This would have been assessed at the start of the lesson as required baseline knowledge in the manner that will be discussed in Chapter 11, and reinforced during the example-problem pair.

3. Assess the skill

Assessment of the skill does not need to immediately follow the development stage. In fact, given the distinction between learning and performance that we will cover in Chapter 12, this may not be the optimal time to assess. Instead, I may do it at the end of the lesson or the start of the next.

My preferred forms of assessment are low-stakes quizzes (Chapter 12) or diagnostic multiple-choice questions (Chapter 11). An example of the latter

for the sub-process 'decide on an appropriate common denominator' might be:

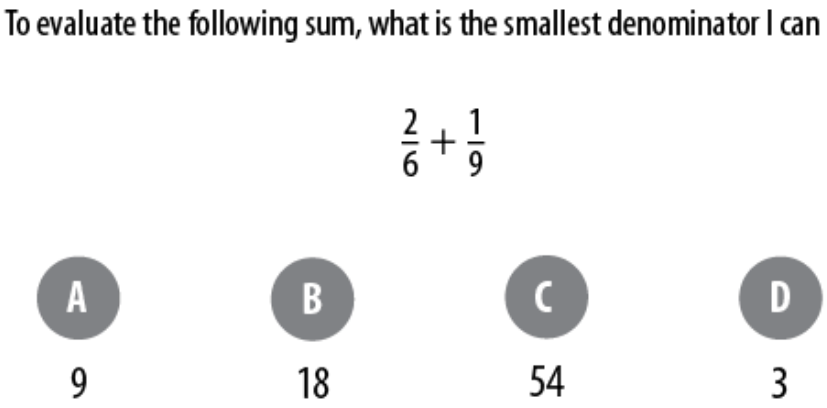


Figure 8.3 – Source: Craig Barton, created for Diagnostic Questions

We will cover the rationale and the different ways of using these types of questions in Chapter 11, but it is worth noting here that my aims are to identify any misconceptions and their specific nature, and to see if I have enough evidence to move on.

4. The final performance

Once all sub-processes have been isolated, developed and assessed, it is time for the final performance. I would use the Supercharged Worked Example approach (Chapter 6) to ensure that students had an opportunity to observe me model the entire process and self-explain at key points. Then it would be time to practise. This may simply consist of a question or series of questions that bring all the sub-processes together. In this case, a selection of standard questions on adding fractions is perfect. Exam questions are also fine to use, but it is worth noting that contextual questions can place an extra burden upon students' fragile working memories that I may wish to avoid at this stage of the learning process. Despite the intensive process students have been through to get to this stage, they are still novice learners, and I would be inclined to keep context out until they have mastered the entire process of adding fractions.

5. Practise retrieval later

The importance of inducing the retrieval of knowledge over different time intervals will be discussed in Chapter 12, but it is important here to note

that successful performance at the assessment stage may not be a reliable indication of learning. It is necessary to practise retrieval again at some points in the future when students have had an opportunity to start to forget. As we will see, this serves the dual function of making any assessment aimed at discerning understanding more reliable, as well as improving the storage strength of that knowledge in long-term memory.

This retrieval practice may come in the form of inclusion of a question on a low-stakes quiz, homework, or as part of a mixed-topic starter. It may consist of the final performance (in this case a pair of fractions to add), or any one of the sub-processes that have been identified.

A note about timing

When I interviewed Kris Boulton for my podcast, he described how he planned to cover the series of 13 sub-processes that he identified to master solving simultaneous equations over a sequence of lessons. For example, he may cover Sub-process 1 in the first lesson, then recap it by means of a test of retention in the second, before moving on to Sub-process 2. These two sub-processes would then be assessed in the same lesson where Sub-process 3 was developed, and so on, ensuring that all sub-processes are routinely revisited. Kris describes this sequence in more detail in a follow-up blog post, which can be found in the *Sources of inspiration*.

However, there is nothing to say that these individual skills and sub-processes need to be taught one after another. The very fact that they have been isolated means they can be tackled independently from each other and the process as a whole. Indeed, there may be a very good argument for separating them over time. We will discuss the power of the Spacing and Interleaving Effects far more in Chapter 12, but for now it is worth reflecting on a 2017 blog post by Hin-Tai Ting, a maths teacher at Michaela Community School. Hin-Tai describes his department's use of the *Connecting Math Concepts* curriculum, which is a direct instruction programme designed in consultation with Siegfried Engelmann. Here, a 60-minute lesson will consist of at least 8 different parts, with multiple concepts being covered. Indeed, the blog post showcases one particular part of a lesson that simply tests whether pupils remember the locations of the quotient, divisor and dividend – separate from any actual division – before

moving on to unrelated questions concerning written addition and place value. The process of written division will be subsequently developed over many lessons, alongside the revisiting of all the other skills students have been taught.

What I do now

I follow the five stages of Deliberate Practice when I am introducing a complex process. And as I have become more experienced doing this, I have come to the realisation that if I am defining complex processes as anything with multiple steps in it, then that is pretty much the whole of the mathematics I teach. Things that are single processes for me are not for my students. For example, before we go anywhere near any calculations, working out the area of a triangle needs to start by identifying what shapes are (and are not) triangles, and what lengths represent (and do not represent) the 'base' and 'perpendicular height'.

But by isolating, developing, assessing, performing and practising retrieval later, I can help students transition along the path to expertise so that they begin to see these as single processes like I do.

OceanofPDF.com

8.3. Practice v final performance

What I used to think

My rationale for not breaking complex processes down into sub-processes was that students would never be assessed that way on the final exam. No exam paper was going to ask students to decide if two fractions were in a form ready to be added together, or challenge them to find an appropriate common denominator. Instead, students would be presented with questions in their entirety, and expected to carry out the full process.

Hence, if that is what they needed to do on the exam, then surely the best preparation was to ensure that their practice was a similar as possible? This way of thinking would extend to: *the best way to prepare for exams is to do exam questions*, and *the best way to get good at solving problems is to solve lots of problems*.

In short, I believed that practice must resemble the final performance.

Sources of inspiration

- Ericsson, K. A., Krampe, R. T. and Tesch-Römer, C. (1993) 'The role of Deliberate Practice in the acquisition of expert performance', *Psychological Review* 100 (3) pp. 363-406.
- Ericsson, K. A. and Pool, R. (2016) *Peak: secrets from the new science of expertise*. Boston, MA: Houghton Mifflin Harcourt.
- Ericsson, K. A., Prietula, M. J. and Cokely, E. T. (2007) 'The making of an expert', *Harvard Business Review* 85 (7/8) pp. 114-121.
- McDermott, K. B., Agarwal, P. K., D'Antonio, L., Roediger, H. L. and McDaniel, M. A. (2014) 'Both multiple-choice and short-answer

quizzes enhance later exam performance in middle and high school classes', *Journal of Experimental Psychology: Applied* 20 (1) pp. 3-21.

My takeaway

Our first clue that practice does not need to look like the final performance can be found outside the domain of mathematics. As already discussed, when preparing for a recital a professional musician does not play the entire piece over and over again. If they did it would be difficult to identify the areas in which they need to improve. Instead they may practise scales, arpeggios, and small sections of the piece over and over again. Indeed, they may well do so in the manner described by both Kris Boulton and Hin-Tai above, interleaving the practice of past drills and sections alongside the study of new ones.

Likewise, the best way for a footballer to prepare for a 90-minute 11-a-side match is not to play lots of 90-minute 11-a-side matches. If they did, then it would be very difficult to discern how they could improve. Is their weakness short passing, dribbling, shooting, positional play, coping under pressure, acceleration speed, stamina, vision, heading, crossing? All of these skills contribute to the success of the overall performance, but are difficult to pinpoint and assess in the rough and tumble of an 11-a-side game. Similarly, the opportunities in which to practise these areas, and the ability to control the practice in order to make it as effective as possible, are lost during a big game.

Musicians and footballers prepare for the final performance using carefully designed, controlled and monitored drills, precisely because the purpose of practice is not to model the final performance, but to *improve* the final performance. And findings outside of the world of the classroom have repeatedly shown that the best way to do that is to make the practice itself very different from the final performance.

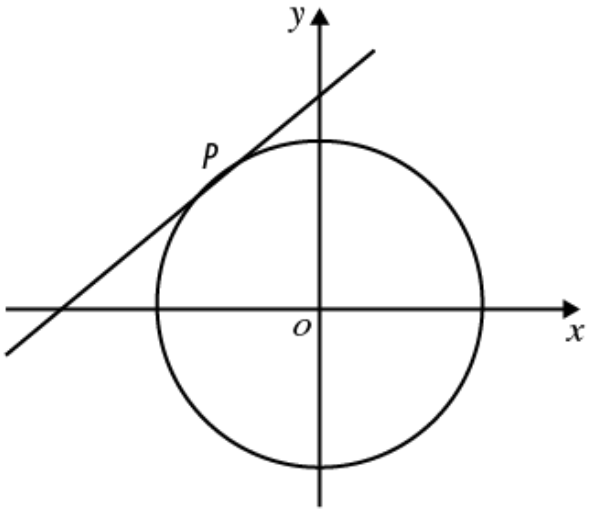
Of course, we need to be incredibly careful making inferences about how the techniques used in fields such as sport and music apply to the classroom. But fortunately, there is some evidence that the same principles apply.

McDermott *et al* (2014) found that the format of the practice quizzes did not need to match the format of the critical test for the benefits of testing and retrieval that we will discuss in Chapter 12 to emerge. In a study of history and science students, the authors found that both short-answer questions and multiple-choice quizzes were effective in preparing students for the final end-of unit exam where the questions were in a much longer form. What is also interesting is that even though the final exam contained no multiple-choice elements, students who were given multiple-choice questions during practice performed just as well – or even better in some cases – than those given short-answer questions.

The finding that short-answer and multiple-choice questions are effective in preparing students for exams that do not contain questions in that format is great news for us teachers. It means we can break down complex processes into sub-processes as described in the previous section. Having isolated the sub-processes, we can then plan explanations and worked examples, control practice through intelligent question choice, and crucially identify misconceptions far easier than if we had to tackle the processes as a whole. Using the techniques I will describe in Chapter 11, it also means we can quickly and efficiently assess understanding, whilst also identifying, understanding and hopefully resolving any misconceptions.

The benefits become even clearer when we consider multi-step and multi-topic exam questions. Consider the sub-processes that a student would need to go through in order to answer the following question from AQA's Summer 2017 GCSE Paper 1 Higher:

P (-1, 4) is a point on a circle, centre o



Not drawn accurately

Work out the equation of the tangent to the circle at P

Give your answer in the form $y = mx + c$ **[4 marks]**

Figure 8.4 – Source: AQA 2017 GCSE Maths Higher Paper 1

I got to 14 before I stopped counting – and that is ignoring the issue of deciding what the question is asking in the first place, which we will consider in the next chapter.

Repeatedly exposing students to exam questions like this is not the best way to help them get good at answering them if they have not yet mastered each of the sub-processes involved. Trying to answer these questions in one go is fine (and even beneficial) for an expert, but for a novice without the relevant domain-specific knowledge readily available and automated in long-term memory, it is likely to lead to cognitive overload and no learning taking place. Far better is to break questions like this apart, and employ the five-stage process described in the previous section to develop the skills students will need.

There is an additional advantage to the Deliberate Practice approach. Students will rarely have practised multi-mark questions in the exact same format that they appear in high-stakes exams. Variety in the context and the way topics are combined means that these types of questions on exams are highly unpredictable. By following a Deliberate Practice approach, students

gain access to a wide selection of individual skills that they have mastered and automated, and as such can develop experience in piecing them together in an appropriate manner to solve these problems, as opposed to hoping a question that looks exactly like one they have practised comes up.

Exam questions like this are the final performance. Practice does not – and should not – need to look the same.

What I do now

I practise complex skills using Deliberate Practice, and I am not at all worried if this practice looks nothing at all like the final process or exam we are working towards. More often than not, this practice will be in the form of short-answer, carefully varied questions as demonstrated earlier on in this chapter, or the diagnostic multiple-choice questions that we will cover in Chapter 11. I will also make use of goal-free problems and remove redundant information as discussed in Chapter 4. The full process and exam questions are the very last thing I look at once the sub-processes have been mastered

I also make sure I tell my students why we are doing it like this. Students can become frustrated that they are working on drills and multiple-choice questions that look nothing like what they will be studying on the exam. I have known parents get involved as well. I know students will reap the benefits in the long term when we put everything together, but that may not help in the here and now. Hence, I am keen to describe the rationale behind what we are doing, and putting it in the context of sport or music seems to help them understand my thinking. Of course, once students see how they perform at the end of the process and in subsequent tests of retention, the benefits speak for themselves.

OceanofPDF.com

8.4. Three reasons to always give students the answers

What I used to think

I used to be reluctant to give my students the answers to class exercises for fear they would simply copy them. Hence, I would let students work for a prolonged period of time – sometimes over 30 minutes – and then display the answers at the end of the lesson, or take the books in and mark the work myself.

Of course, if a student went wrong on Question 1, the chances are they would also have gone wrong on Question 2 onwards, and I now had a rather large problem on my hands.

Sources of inspiration

- Anderson, J. R. (1996) 'ACT: a simple theory of complex cognition', *American Psychologist* 51 (4) pp. 355-365.
- Barton, C. (2016) 'John Corbett', *Mr Barton Maths Podcast*.
- Butler, A. C. and Roediger, H. L. (2008) 'Feedback enhances the positive effects and reduces the negative effects of multiple-choice testing', *Memory & Cognition* 36 (3) pp. 604-616.
- Butterfield, B. and Metcalfe, J. (2006) 'The correction of errors committed with high confidence', *Metacognition and Learning* 1 (1) pp. 69-84.
- Chi, M. T. H. (2000) 'Self-explaining: the dual processes of generating inference and repairing mental models' in Glaser, R. (ed.) *Advances in*

instructional psychology: educational design and cognitive science, Vol. 5. Mahwah, NJ: Lawrence Erlbaum Associates, pp. 161-238.

- Coe, R. (2013) *Improving education: a triumph of hope over experience, CEM Inaugural Lecture*. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf
- Eich, T. S., Stern, Y. and Metcalfe, J. (2013) 'The hypercorrection effect in younger and older adults', *Aging, Neuropsychology, and Cognition* 20 (5) pp. 511-521.
- Ericsson, K. A., Krampe, R. T. and Tesch-Römer, C. (1993) 'The role of Deliberate Practice in the acquisition of expert performance,' *Psychological Review* 100 (3) pp. 363-406.
- Ericsson, K. A. and Pool, R. (2016) *Peak: secrets from the new science of expertise*. Boston, MA: Houghton Mifflin Harcourt.
- Ericsson, K. A., Prietula, M. J. and Cokely, E. T. (2007) 'The making of an expert', *Harvard Business Review* 85 (7/8) pp. 114-121.
- McDermott, K. B., Agarwal, P. K., D'Antonio, L., Roediger, H. L. and McDaniel, M. A. (2014) 'Both multiple-choice and short-answer quizzes enhance later exam performance in middle and high school classes', *Journal of Experimental Psychology: Applied* 20 (1) pp. 3-21.
- Metcalfe, J. and Finn, B. (2011) 'People's hypercorrection of high-confidence errors: did they know it all along?', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 37 (2) p. 437.
- Soderstrom, N. C., and Bjork, R. A. (2015) 'Learning versus performance: an integrative review', *Perspectives on Psychological Science* 10 (2) pp. 176-199.
- Wiliam, D. (2016) 'The secret of effective feedback', *Educational Leadership* 73 (7) pp. 10-15.

My takeaway

There are three main reasons why I now think giving students access to the answers to classwork at regular intervals throughout the lesson is a good idea:

1. Practice does not make perfect, practice makes permanent
This is never more important than in the domain of Deliberate Practice. Recall that for Ericsson *et al* (1993) the purpose of Deliberate Practice is to break down a complex process, isolate an individual skill, and work on it, receiving regular and specific feedback so you can improve your performance. Without that all-important feedback, learners in any discipline could be practising the wrong thing, making key mistakes, and crucially having those mistakes reinforced the more practice they undertake. The model of thinking proposed by Anderson (1996) in Section 1.1 explains that once something has been learned, it is incredibly difficult to rewire the path of production and hence unlearn it.

Players have coaches to give them immediate feedback on their performance in drills so they can adjust and improve, and not fall into bad habits. And even in the absence of a coach, a footballer does not have his or her eyes closed after hitting a cross-field pass, and a pianist does cover their ears when listening back to their piece. So why should a mathematician not be aware of the result of their attempt at answering a question? If students do not have access to some feedback on their answers, how do they know if they are practising correctly and hence committing the right knowledge to long-term memory?

2. The Self-Explanation Effect
We have seen in Chapter 5 the power of the Self-Explanation Effect, whereby learners who attempt to establish a rationale for the solution steps by pausing to explain the examples to themselves appear to learn more than those who did not. For Chi (2000), one of the main drivers of the effect is when the learner repairs their own mental model. Hence, when faced with an answer that does not match their own, students are forced to pause and self-explain. This self-explanation may be sufficient to repair their own mental model, or it may prompt them to seek help from a teacher or peer. Either way, it is likely to be beneficial for their learning, and can only happen if students have access to the answers.

3. The Hypercorrection Effect
As briefly mentioned in Chapter 7, the Hypercorrection Effect refers to the finding that errors committed with high confidence are more likely to be corrected than low-confidence errors (*eg* Butterfield and Metcalfe, 2006). In other words, when you discover that something you were sure was correct is in fact incorrect, the bigger the improvement in the change to your thinking. Eich *et al* (2013) found the effect to be stronger with young adults than older adults, Metcalfe and Finn (2011) identified it in children, and Butler and Roediger (2008) uncovered the effect in an experiment involving multiple-choice questions – something that will be of direct relevance in Chapter 11. Hence, we can tap into the benefits of the Hypercorrection Effect by giving students access to answers and prompting them to focus in particular on questions they were sure they got right.

So, I tell my students to give each question they have done a confidence score out of 10 before they check their answers. Then, if they have got any questions wrong, they should look at the confidence score they assigned to each question, and start thinking about those with a high score first. This moment of quiet contemplation makes use of the Self-Explanation Effect, whilst at the same time focusing on questions that have the biggest potential gains in learning. If there are questions that students still cannot figure out, then I ask them to place an asterisk by the side of them, and I set aside time in each lesson to focus on these asterisk questions. I do something very similar with the low-stakes quizzes we will discuss in Section 12.8, and it seems to be proving really effective. Students seem much more interested in the answer-checking process, and it has led to some fascinating discussions – and the exposure of significant misconceptions – when we consider the reasons behind these high-confidence errors.

But won't students just copy?
For Coe (2013), learning happens when people have to think hard. If students simply copy down the answers, then they are not likely to be thinking hard at all, and hence little learning may take place.

There are two ways we may counteract this:

1. Only give the answers

Just giving students access to the final answers, but insisting they show working out, may give us the best of both worlds. Wiliam (2016) suggests we should make feedback into detective work. Faced with a final answer that does not match their own, the student who gets something wrong is forced to look back at their own work and consider exactly where and why they have gone wrong. The best feedback – whether in written form or merely answers – is that which causes the students to think.

Indeed, in their study into the format of practice in the classroom, McDermott *et al* (2014) found that frequent classroom quizzing with feedback improves student learning and retention. Feedback was not given in the manner of paragraphs full of personalised evaluation of their work and targets to improve, but merely an indication if the answer they gave was right or wrong.

2. Create the right culture

In Section 2.4 we discussed the key role of the teacher in creating an environment where students feel safe, enjoy mathematics and do not fear making mistakes. A key follow-on from this – and one that will be discussed further in Chapter 12 – is that practice questions, homeworks and low-stakes quizzes should not be seen as tools of assessment, but tools of learning. The key to this is carefully managing the stakes so that the negative consequences for students getting things wrong are low, but the value placed upon learning is high.Is there a possibility that some students will still copy down the answers if they are so easily available? Of course there is. But do the costs of this outweigh the benefits of having students check their answers? Not at all.

The answers must be correct

I know this is a ridiculously obvious thing to say, but there are not many more damaging things to learning than an incorrect answer, especially when we consider the power of practice, self-explanations and hypercorrection. Imagine if a student gets a question correct, then looks up the answer and finds it to be different. Students may then embark upon a process of self-explaining and hypercorrection in order to try and justify this erroneous answer. They then may proceed to answer the next ten similar questions

using the same erroneous reasoning, hard-wiring the incorrect procedure. The resulting misconception may prove very difficult to resolve.

What I do now

In class
When I interviewed John Corbett for my podcast, he explained that he always provides his students with access to the answers to classwork. I now do the same. I tend to print out a few copies of the answers and leave them at the end of each row of desks. I tell students never to go more than ten minutes without having a quick check of the answers, and if they have made a mistake to:

1. See if they can figure out why.
2. Ask a friend.
3. Ask me.

Homework
As I will discuss in Chapter 12, the homeworks I give students are no longer topic-based. Hence, the danger of students practising lots of the same skill incorrectly is reduced. With homeworks, students do not have direct access to the answers, but I do employ a process of delaying feedback, which will be discussed in Section 12.10.

OceanofPDF.com

8.5. If I only remember 3 things…

1. Many processes that can be easily done in one go by expert learners actually consist of multiple sub-processes for novices.
2. The five-stage process of isolate the skill, develop the skill, assess the skill, the final performance, and practise retrieval later can serve as a model for Deliberate Practice.
3. Giving students access to the final answers during Deliberate Practice is important to ensure that they are practising the right things.

OceanofPDF.com

9. PROBLEM-SOLVING AND INDEPENDENCE

Indiana Jones – everyone's favourite fictional adventure-seeking archaeologist – spends much of his life on quests to find highly sought-after treasures, from the Holy Grail to the Ark of the Covenant. If they ever make a maths spin-off (it would have to be better than the fourth instalment), I may suggest *Indiana Jones and the Problem-Solving Skill* or *Indiana Jones and the Independent Learner.*

Much of my teaching life has been plagued by a seemingly never-ending quest to discover these most elusive of secrets. I don't think there has been a school report I have written, or a parents' evening I have held, that has not included the phrases 'he needs to develop his problem-solving skills', or 'she needs to become an independent learner'. The implicit assumption behind these phrases is that problem-solving is a skill that can be taught, in much the same way as adding fractions, and that becoming an independent learner is something that is within the full control of the student. I had never questioned those assumptions.

That is, until now.

OceanofPDF.com

9.1. What is a problem?

What I used to think

I'll be honest – for many years I have used the phrase 'problem' and 'problem-solving' without really questioning what indeed I meant by a *problem*.

It seems clear to me that the following is a problem, and a particularly nice one at that:

There is a school with 1000 students and 1000 lockers. On the first day of term, the headteacher asks the first student to go along and open every single locker, he asks the second to go to every second locker and close it, the third to go to every third locker and close it if it is open or open it if it is closed, the fourth to go to the fourth locker and so on. The process is completed with the thousandth student. How many lockers are open at the end?

But how about:

In a 20% off sale, a jacket costs £52.

What did it cost before the sale?

Is that a problem too?

Sources of inspiration

- AQA Problem-solving and Reasoning Guidance and Questions
- AQA (2005) *GCSE Mathematics Teaching Guidance Document, Version 2.0*.

- Schoenfeld, A. (2009) 'Learning to think mathematically: problem solving, metacognition, and sense-making in mathematics' in Grouws, D. (ed.) *Handbook for research on mathematics teaching and learning.* New York, NY: MacMillan, pp. 334-370.

My takeaway

Before we delve into the murky depths of the world of problem-solving, it is important to establish what we mean by a *problem*. This becomes especially important when we consider research relating to problem-solving, as the term can be applied to anything from the most straightforward question for the student to solve independently, right through to the kinds of unstructured puzzle I presented above.

Schoenfeld (2009) offers two contrasting definitions of the word 'problem':

Definition 1: 'In mathematics, anything required to be done, or requiring the doing of something.'

Definition 2: 'A question that is perplexing or difficult.'

I had always thought of problems as being more of the second variety, but it is worth bearing in mind that when much of the literature refers to 'problems', they mean the first type.

I think a useful way to think about problems is in terms of English GCSE Maths Assessment Objectives. These became of particular importance with the changing of the GCSE Maths specification (first examined in 2017), which had a greater emphasis on 'problem-solving'. The Assessment Objectives, together with explanatory notes and their relevant weighting across Higher and Foundation papers are:

AO1 – Use and apply standard techniques
Students should be able to:

- accurately recall facts, terminology and definitions.
- use and interpret notation correctly.

- accurately carry out routine procedures or set tasks requiring multi-step solutions.

Higher weighting: 40%

Foundation weighting: 50%

AO2 – Reason, interpret and communicate mathematically
Students should be able to:

- make deductions, inferences and draw conclusions from mathematical information.
- construct chains of reasoning to achieve a given result.
- interpret and communicate information accurately.
- present arguments and proofs.
- assess the validity of an argument and critically evaluate a given way of presenting information.

Higher weighting: 30%

Foundation weighting: 25%

AO3 – Solve problems within mathematics and in other contexts Students should be able to:

- translate problems in mathematical or non-mathematical contexts into a process or a series of mathematical processes.
- make and use connections between different parts of mathematics.
- interpret results in the context of the given problem.
- evaluate methods used and results obtained.
- evaluate solutions to identify how they may have been affected by assumptions made.

Higher weighting: 30%

Foundation weighting: 25%

It is AO2 and AO3 that I think of when I hear the terms ‘problem’ or ‘problem-solving’. It seems to me that AO1-style questions can be taught using the techniques of example-problem pairs (Chapter 6), intelligent and deliberate practice (Chapters 7 and 8) and careful use of desirable difficulties (Chapter 12). Whereas, to help students answer AO2 and AO3 questions, you perhaps need something more.

Here are some examples of AO2 and AO3 questions, taken from AQA’s very useful exemplar document. To prepare for what is coming during the rest of the chapter, it is worth solving these problems yourself and then considering:

- How did you know what each problem was about?
- How did you know how to solve them?
- At which state of the process might your students get stuck?
- How would you help them?

AO2

Foundation/Higher

Written as a product of prime factors $2014 = 2 \times 19 \times 53$

Work out all the factors of 2014

[3 marks]

Foundation/Higher

In a game you can choose to throw one die or two dice.

If you throw one die your score is the number you throw.

If you throw two dice your score is the sum of the numbers you throw.

It is your turn.

You need to score exactly 4 to win.

Should you choose to throw one die or two dice?

Explain your answer fully

[4 marks]

Higher only

Three people were born on the same date, but in different years.

The second person was born 5 years after the first.

The third person was born 7 years after the first.

Prove that the sum of their ages will always be a multiple of 3

[3 marks]

Higher only

$x^2 + 2ax + b \equiv (x - 5)^2 - a$

Work out the values of a and b

[3 marks]

AO3

Foundation/Higher

Seven pupils take a test marked out of 10

The mean mark is 8

The median mark is 8

The modal mark is 7

The range of the marks is 3

What are the marks in the test?

[3 marks]

Foundation/Higher

To make a shade of purple paint

Mix 2 litres of blue paint with 3 litres of red paint.

Fred accidentally mixes 3 litres of blue paint with 2 litres of red paint.

How much more red paint should he add to his mixture to get the correct shade of purple paint?

[3 marks]

Higher only

An amount is increased by 20%

40% of the new amount is 288

Work out the original amount.

[3 marks]

Higher only

Here is some information about the masses of Arnie, Beth, Chas, Dave and Ed
The mean mass of:

Arnie, Beth and Chas is 70 kg

Beth, Chas and Dave is 76 kg

Chas, Dave and Ed is 77 kg

Chas and Ed have the same mass.

Arnie is 75 kg

Work out the masses of the other four people.

[5 marks]

We can see there is an incredibly wide variety of questions in maths that could justifiably be classed as 'problems', both in terms of the breadth of mathematical content they cover and the variety of skills they call upon.

Whilst the some of the research may also refer to AO1 questions, the main focus of this chapter will be looking at AO2 and AO3 questions. Success in these questions is what I believe most teachers mean when they say they want their students to be good problem-solvers. They are the non-routine questions that are hard to predict and which students regularly struggle with.

What I do now

Given the fact that 60% and 50% of the questions that Higher and Foundation students respectively will be examined on at GCSE fit my definition of 'problem-solving' – with the proportion probably even higher at A level – it is obviously of great importance that we understand the best way to enable students to answer these questions successfully.

Doing exactly that is the focus of the rest of this chapter.

OceanofPDF.com

9.2. Why are some students bad at problem-solving…

What I used to think

Many of the students I have taught over my career have seemingly struggled with 'problem-solving'. Here are three different types of AO2 and AO3 problems, followed by some common scenarios I have experienced. My advice again is to try each problem yourself, and then see if my experiences ring any bells:

1. The Contextual Problem
Polly Parrot squawks every 9 seconds.

Mr Toad croaks every 21 seconds.

They both make a noise at the same time.

After how many seconds will they next make a noise at the same time?

One student (quite rightly, as it happened) pointed out that I had not taught him about parrots. Some students appear to latch onto the word 'seconds' and begin attempting to use some convoluted version of the speed, distance, time formula. Some multiply the numbers together, whereas others opt for division, subtraction or addition. Almost all of the students who struggle are drawn to the surface structure of the problem, and fail to appreciate the deeper structure – that this is really a problem about lowest common multiples.

2. The Subtle Contextual Problem
Ben and Zoe share some money in the ratio 2 : 5.

Zoe gets £210 more than Ben.

How much does Ben get?

A £60 B £84 C £105 D £140

This is a question that AQA created for my Diagnostic Questions website. At the time of writing, it has been answered over 20,000 times by students from all over the world, and 38% of respondents have opted for answer A. Why have they done this? Well, when students given an answer on Diagnostic Questions they are prompted to provide an explanation, and these two explanations certainly shed some light on the issue:

You add the ratios together and then divide it by the money. After, you times it by the ratio

2 + 5 = 7.

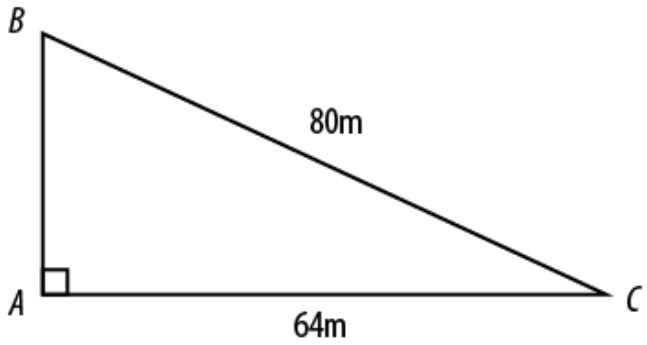
210 ÷ 7 = 30.

30 × 2 because that's what ben is. This = 60. So ben gets 60 pounds.

Students have again been drawn in by the surface structure. This time it is not some strange animal-related context, but the surface structure of *sharing in a ratio*. For students who got this question wrong, those three words meant one thing and one thing only: add up the parts, divide the total, and share accordingly. A well-rehearsed algorithm that they were more than willing to apply in this context, despite the fact that it led them to the wrong answer.

3. The Multi-Topic Problem

This triangular plot of land is for sale.



Not drawn accurately

The land is sold for £6400 per acre.
1 acre = 4047m^2
Work out the cost of the land.
Give your answer to 2 significant figures.

[6 marks]

Figure 9.1 – Source: AQA Topic tests – Higher tier – Real life, available at https://allaboutmaths.aqa.org.uk/attachments/6357.pdf

Once again students must determine what areas of maths this question is testing, but this time we have several surface features to contend with – triangles, lengths, money, area, weird units and rounding. These multi-topic problems are perhaps the most challenging of all. When students answer questions like this, anything could happen – and the three hours it took to mark a set of attempts to this question following the official mark scheme are testament to the fact that anything usually does. Moreover, these types of problems are inherently unpredictable – it is highly unlikely students will have encountered this particular combination of topics, presented in this particular way before.

Why is it that some students can answer these questions and some (often many more) cannot? Well, if you had asked me that question any time in the last 12 years I would have told you that it was because those students were simply not good at problem-solving. They lacked the strategic skill necessary to realise what a question was asking, decide upon a strategy, and then go ahead and solve it.

And whilst these three types of problem are obviously very different, my solution to dealing with students who lacked this problem-solving skill was always the same: show them how to solve problems! Specifically, having chosen a problem, I would follow a three-stage procedure:

1. Articulate clearly how I first decide what area (or areas) of maths the question is concerned with, and specifically what the question is asking. Often I would highlight key words, make nice little annotations in bubbles, and so on.

2. Then I would plan out a strategy. Where do I need to get to? What information do I need to work out to solve this problem? How can I get this information? Again, all the time verbalising the thoughts in my head, making it crystal clear to my students how my brain was working.

3. Finally, I would put my plan into action, ensuring I showed all stages of my working out, arriving at a final answer, and checking my answer made sense in the context of the question.

For me, only the last of these stages was concerned with mathematical knowledge. The first two were generic skills: spotting what the question was asking and devising a strategy. And as they were skills, I would teach them in exactly the same way as I would teach the skill of adding fractions – careful modelling in the manner described above, exposing students to lots of different examples. My logic was, if you watch an 'expert' solve enough problems, then you can become an expert problem-solver yourself.

And this approach always seemed to work. The struggling students would quietly copy down my methods, successfully answering my questions, and even utter the magic phrase, 'Sir, I get it now, this is dead easy'. We would all leave the lesson feeling very happy with ourselves.

The only problem was, whenever my students met similar questions again, it was as if they had not understood a word I said.

Sources of inspiration

- AQA (2015) *GCSE Mathematics: 90 maths problem-solving questions.*
- Ashman, G. (2017) 'What is the most useless problem-solving strategy?', *Filling the Pail blog*. Available at: https://gregashman.wordpress.com/2017/07/20/what-is-the-most-useless-problem-solving-strategy%EF%BB%BF/
- Barton, C. (2017) 'Ed Southall – Part 2', *Mr Barton Maths Podcast*.
- Catrambone, R. (1998) 'The subgoal learning model: creating better examples so that students can solve novel problems', *Journal of Experimental Psychology: General* 127 (4) pp. 355-376.
- Chi, M. T. H., Feltovich, P. J. and Robert Glaser (1981) 'Categorization and representation of physics problems by experts and novices', *Cognitive Science* 5 (2) pp. 121-152.
- Gick, M. L. and Holyoak, K. J. (1980) 'Analogical problem-solving', *Cognitive Psychology* 12 (3) pp. 306-355.
- Hill, N. M. and Schneider, W. (2006) 'Brain changes in the development of expertise: neuroanatomical and neurophysiological evidence about skill-based adaptations' in Ericsson, K. A., Charness, N., Feltovich, P. J. and Hoffman, R. R. (eds) *The Cambridge handbook of expertise and expert performance*. Cambridge: Cambridge University Press, pp. 653-682.
- Lester, F. K., Garofalo, J. and Kroll, D. (1989) *The role of metacognition in mathematical problem-solving: a study of two grade seven classes, final report*. Bloomington, IN: Indiana University Bloomington.
- Pak-Hong, C. G. (2012) *Problem-Solving Strategies: Research Findings from Mathematics Olympiads*. Available at: http://sms.math.nus.edu.sg/smsmedley/Vol-20-2/Problem-solving%20strategies%20-%20research%20findings%20from%20Mathematics%20Olympiad(PH%20Cheung).pdf

- Polya, G. (1945) *How to solve it*. Princeton, NJ: Princeton University Press.
- Schoenfeld, A. (1983) 'Episodes and executive decisions in mathematical problem solving' in Lesh, R. and Landau, M. (eds) *Acquisition of mathematics concepts and processes*. Cambridge, MA: Academic Press, pp. 345-395.
- Schoenfeld, A. (1985) *Mathematical problem-solving*. Cambridge, MA: Academic Press.
- Schoenfeld, A. (1987) 'What's all the fuss about metacognition?' in Schoenfeld, A. (ed.) *Cognitive science and mathematics education*. Hillsdale, NJ: Lawrence Erlbaum Associates, pp. 189-215.
- Schoenfeld, A. (2009) 'Learning to think mathematically: problem solving, metacognition, and sense-making in mathematics' in Grouws, D. (ed.) *Handbook for research on mathematics teaching and learning.* New York, NY: MacMillan, pp. 334-370.
- Tricot, A. and Sweller, J. (2014) 'Domain-specific knowledge and why teaching generic skills does not work', *Educational Psychology Review* 26 (2) pp. 265-283.
- Wiemann, C. (2007) 'The "curse of knowledge", or why intuition about teaching often fails', *APS News* 16 (10) (no pagination).
- Willingham, D. T. (2006) 'How knowledge helps: it speeds and strengthens reading comprehension, learning – and thinking', *American Educator* 30 (1) p. 30.

My takeaway

It seems sensible to assume that to be able to successfully solve problems, students need to be able to do the following things, which are directly related to the first three of Polya's (1945) classic four stages of problem-solving:

1. Understand what the problem is about.

2. Devise a strategy to solve the problem.

3. Carry out the strategy.

Likewise, it seems sensible to assume that when solving a problem, students must carry out these stages in the given order as I outlined above.

However, the assumption that these stages can be taught independently from each other, and that the first two are dependent upon some generic, easily defined set of skills, deserves closer scrutiny, especially given the difficulty many students seem to have getting through the stages unaided. Furthermore, once I began to question this assumption, it led to the rather counter-intuitive conclusion that maybe, just maybe, students do not get good at solving problems by solving problems.

What follows – both in this section and the next – is my attempt to justify this claim, and what we should do instead.

1. Understand what the problem is about

Have another look at the three problems from the start of this section.

How do you know what each of these problems is about? How do you know that the first one is about lowest common multiples, the second is sharing in a ratio but not in the 'conventional' sense, and the third is a combination of Pythagoras, area of a triangle and unit conversions? Is it because you have some magic skill that allows you to identify what any problem is about?

I don't think it is, and indeed to claim so is doing you a disservice.

The reason you can (probably relatively quickly) discern what each problem is about is because you are an expert. We saw back in Section 1.2 that one defining feature of an expert is their ability to recognise the deep structure of a problem. You can overcome the potentially distracting surface features, and get right to the heart of the issue. But this expertise has been hard-earned. It has been built up over many years. It is a result of you working hard to acquire deep, domain-specific knowledge of the subject matter involved, and having answered hundreds, if not thousands, of related problems. You make connections between previous problems that you

solved, not just because you've solved lots of them, but also because you have the knowledge to make those connections.

Many of our students cannot do this.

Glick and Holyoak (1980) demonstrate that even when participants are presented with two problems with the exact same deep structure, but two different surface structures, and these problems are given one after the other, the vast majority fail to spot the connection as they are drawn towards the distracting surface features. Likewise, Chi *et al* (1981) asked physics novices (undergraduates) and experts (PhD students) to sort a selection of physics problems into categories. They found that the novices sorted by the surface features of a problem – whether the problem described springs, an inclined plane, and so on. The experts, however, sorted the problems based on the physical law needed to solve it (*eg* conservation of energy), demonstrating awareness of the problems' deep structures. Referring to these two studies, Willingham (2006) explains that experts do not just know more than novices, they actually see problems differently. They are able to recognise a problem's deep structure.

But surely we can help our students recognise what problems are really about using two tried and tested techniques:

1. Highlight key words
When faced with a problem, and they do not know what it is about, I would always tell my students to highlight key words. But it is only now I realise this advice might not be the best, because it rests on the assumption that our students know which words to highlight. Consider the problem about parrots. What are the key words? Without knowledge of lowest common multiples and experience of solving similar problems, highlighting may only serve to draw students' attention even more to the distracting surface features of the problem. The reason we know which words to highlight is thanks to our hard-earned expertise.

2. Label and annotate solutions
Perhaps a more sophisticated form of highlighting is to label and annotate key sections of the solution to a problem. As discussed in Section 6.3,

Catrambone (1998) found that such labels and annotations are an effective way of organising the solution to problems into more manageable and meaningful subgoals, which make it more likely students can transfer these skills to a different context. Brilliant! But those of you paying attention may remember a fly in the ointment. The most effective labels to facilitate such transfer were those that were not tied to the surface feature of the problem, but in order to take advantage of such abstract labels, students needed strong background knowledge. As Catrambone concluded: 'for learners with weaker backgrounds these cues might need to be tied at least partially to example features despite the danger that this may lead the learner to form representations that have erroneous surface ties'.

So why do many of our students struggle to identify what a problem is about? Quite simply, because the ability to do so is a feature of hard-earned expertise. It is inherently dependent upon domain-specific knowledge, and the connections such knowledge allows experts to make between the many similar questions they have solved before. It is not a generic skill, and there is no shortcut to it.

2. Devise a strategy to solve the problem

Okay, so if understanding what a problem is about is not a generic skill that can be taught, but in fact something dependent upon the level of domain-specific knowledge, then surely the same is not true of devising a strategy to solve a problem? Well, let's take a look at three 'categories' of problem-solving strategies:

- Generic strategies
- Topic-specific strategies
- Metacognition

Generic Strategies

If we take a look at the eight problems posed at the start of this chapter – or indeed, if we pick up any exam paper – we will quickly reach the conclusion that there is no single 'problem-solving strategy' that can be applied to all problems. That, of course, is no big surprise. But if, say, there was a relatively

small number of problem-solving strategies we could teach our students, then that would make everybody's lives so much better.

AQA produced a wonderful booklet of 90 problems for the updated 2015 GCSE maths specification full of lovely AO2 and AO3 questions. In it, they attempt to tag each of the questions with one of five problem-solving strategies which could be used to solve them. These strategies are:

1. *To set out cases systematically*
2. *To work backwards from a value given in the problem*
3. *To find one or more examples that fit a condition for the answer*
4. *To look for and represent relationships between elements of the situation*
5. *To find features of the situation that can be acted on mathematically*

This all sounds promising. That is, until you read the paragraph below this list:

> However, it should be noted that these labels are for the convenience of reference in the resource, and should not be used in the classroom context. To use labels like this in class is to invite students to misunderstand these possibilities for action as distinct sequences of activity that can be learned and applied as procedures. Problem-solving is not a matter of identifying a problem as a 'type 3' and solving it by applying the 'find an example to fit' routine.

Indeed, you only need to look at three questions tagged in the same category (4. *To look for and represent relationships*) to see that AQA are spot on with their warning:

1. Bouncing Ball

A ball is dropped and bounces up to a height that is 75% of the height from which it was dropped.

It then bounces again to a height that is 75% of the previous height and so on.

How many bounces does it make before it bounces up to less than 25% of the original height from which it was dropped?

2. Mean

A set of a thousand numbers has a mean of zero.

All but two of the numbers are 1.

What is the mean of the other two numbers?

3. Youth Club Trip

40 members of Arwick Youth Club go on a trip to a leisure centre.

They go in minibuses that can seat up to 15 people each.

It costs £30 for each minibus and £150 for the group to have the use of the leisure centre.

How much will the trip cost per person?

Despite being classed as the same type of problem, the three examples here have three completely different deep structures. Being told to 'look for and represent relationships between elements of the situation' is unlikely to be of much use to the confused student faced with the bouncing ball problem. Far more useful is a sound knowledge of calculating repeated percentages of an amount – something that will be of little use calculating the mean of the two numbers or the cost of the trip to the leisure centre.

But maybe this is an issue with AQA's categories and labelling. So, let's turn our attention to Polya, whose classic 1945 work *How to solve it* is often held up as the seminal guide to mathematical problem-solving.

The second of Polya's four stages to solving a problem is 'Devise a Plan', and contains suggested strategies that include: 'Guess and check', 'Look for a pattern', 'Make an orderly list', 'Draw a picture', 'Eliminate possibilities', 'Solve a simpler problem', 'Use symmetry', 'Use a model', 'Consider special cases', 'Work backwards', 'Use direct reasoning', 'Use a formula', 'Solve an equation' and 'Be ingenious'. I regularly see these and similar strategies adorning the walls of classrooms, meant to stimulate the perplexed student and help them crack the tricky problem they are working on.

The first question to address is: *can such strategies be taught in the absence of domain-specific knowledge?*

I don't think they can. Take 'solve a simpler problem'. Unless a student can identify and understand the deep structure of the problem at hand, and have knowledge of that deep structure, how on earth can they solve a simpler one? Similarly, 'use direct reasoning' and 'consider special cases' are as useless to a novice learner as they are meaningless; and I don't think I could bring myself to suggest to a struggling middle-set Year 11 lad, becoming increasingly frustrated with a multi-mark GCSE question, that all he needs to do is be a little bit more ingenious.

Even advice such as 'draw a picture' – something which may indeed help students focus on the deep structure of a whole series of problems, including many of those encountered in trigonometry and probability – relies on a significant amount of domain-specific knowledge to construct an appropriate diagram, incorporating the key features of the problem. Without that knowledge the student may either be unable to produce the diagram, or may produce one that focuses on the problem's unhelpful surface features.

But there is a second question we need to ask: *are such generic strategies useful at all?*

On this, I am not so sure. To me, they feel more like labels we apply to strategies after their successful implementation than useful guidelines to follow when faced with a problem. When I solve a problem like the ones presented in this chapter, I do not consciously run through Polya's list of strategies. I use my domain-specific knowledge together with my vast bank

of experience of solving similar problems. If, on reflection, it turns out that I solved one percentages problem and one volume problem both by 'using direct reasoning', that is fine, and indeed my expertise may enable me to make that connection and learn from it. But telling a student faced with the same two problems – specifically a student who lacks my domain-specific expertise – to 'use direct reasoning' is likely to be no help at all.

Topic-specific strategies

Pak-Hong (2012) analysed a large selection of tricky Mathematical Olympiad questions and attempted to categorise them by the generic strategies suggested by Polya. The conclusion he reached was that 'the most effective problem-solving strategies are topic-specific'.This feels about right to me. I am more likely to form connections between the solution to two ratio problems than, say, a ratio and an averages problem that have both been labelled with the same generic strategy,

Indeed, in my podcast interview with Ed Southall, I asked him if he believed students could be taught to solve problems, and he said yes without the slightest hesitation. Ed talked of 'learned behaviours'. Specifically, he described how, via repeated exposure to geometry problems, he has learned to draw in a radius when faced with an angle problem involving a circle, and how he believes this is something we can teach our students to do. Similarly, I often tell my students when faced with an unfamiliar problem on Pythagoras or trigonometry in 2D or 3D to look for opportunities to create right-angled triangles. Likewise, when faced with a quadratic expression – be it on the denominator of an algebraic fraction, referencing a sketch of a quadratic curve, or engulfed in some strange context (Hannah's sweets, anyone?) – I advise that it is usually a good idea to factorise it.

These strategies are effective, but they are also topic-specific. Drawing in radii and factorising quadratics are unlikely to help when it comes to a problem involving sequences or ratio. More importantly, they also require knowledge – what are students going to do once you have drawn in the radius, or factorised the expression?

Some go further. In his summary of relevant research into problem-solving in mathematics, Schoenfeld (2009) includes the rather bleak quote from

Beagle: 'there are enough indications that problem-solving strategies are both problem- and student-specific often enough to suggest that finding one (or few) strategies which should be taught to all (or most) students are far too simplistic'.

I am not sure I would go quite that far. I believe there are topic-specific strategies we can teach our students to help solve the problems they may encounter, but that these are again rendered impotent without relevant domain-specific knowledge.

Can we teach these strategies alongside the acquisition of knowledge?
This is a key question to ask, because if the answer is yes, then students may indeed learn from solving problems, regardless of the level of domain-specific knowledge they possess. But I think the answer is no. Considering the fragility of novice learners' working memories and the debilitating effects of cognitive overload discussed in Chapter 4, trying to learn core skills whilst at the same time attempting to develop higher-level approaches to problems is likely to be too much for most students, and they may end up learning neither the skills nor the strategy. I will provide a concrete example of this in Section 9.4. I firmly believe that knowledge needs to come first.

Metacognition
There is evidence that we can help students develop generic metacognitive strategies (metacognition may be considered 'thinking about thinking') that will allow them to approach unfamiliar problems in a structured, systematic way, and hence devise effective strategies to solve them.

Schoenfeld (*eg* 1983, 1985, 1987) demonstrated that expert problem-solvers frequently engage in metacognitive acts in which they step back and reflect on the approaches they are using. They ask themselves planning and monitoring questions, such as: 'Is this going anywhere? Is there a helpful way I might represent this problem differently?'. They bring to mind alternative approaches and make selections based on prior experience. In contrast, novice problem-solvers are often observed to become fixated on an approach and pursue it relentlessly, however unprofitably.

So, can we teach these metacognitive strategies to our students to help them develop into expert problem-solvers? Well, a specific approach for mathematics is described by Schoenfeld (2009):

> Such [metacognitive] skills such can be learned as a result of explicit instruction that focuses on metacognitive aspects of mathematical thinking. That instruction takes the form of 'coaching,' with active interventions as students work on problems. Roughly a third of the time in the problem-solving classes is spent with the students working on problems in small groups. The class divides into groups of three or four students and works on problems that have been distributed, while the instructor circulates through the room as 'roving consultant.' As he moves through the room he reserves the right to ask the following three questions at any time:

- What (exactly) are you doing? (Can you describe it precisely?)
- Why are you doing it? (How does it fit into the solution?)
- How does it help you? (What will you do with the outcome when you obtain it?)

He begins asking these questions early in the term. When he does so the students are generally at a loss regarding how to answer them. With the recognition that, despite their uncomfortableness, he is going to continue asking those questions, the students begin to defend themselves against them by discussing the answers to them in advance. By the end of the term this behavior has become habitual.

Indeed, when I saw Alison Kiddle from NRICH at a recent conference and asked her whether you could teach problem-solving (she was trying to eat a piece of lasagna at the time), she referred to metacognition. She explained that it was not solving problems themselves that allows students to develop problem-solving skills, but specifically in the moments of reflection both during the process and afterwards. This reflection is unlikely to come easy to students, and indeed the approach described by Schoenfeld above could prove stressful and frustrating to the students in question, especially in the

early stages. Alison instead spoke about how praise can be used by the teacher to reinforce these behaviours. So, a teacher may say to a student something along the lines of 'I love the way you have systematically organised your work in a table', which allows the student to reflect on the specific element of their work that the teacher knows will be useful going forward.

A similar approach is put forward by Lester *et al* (1989). They suggest practical strategies that teachers can use to help students reflect upon their thinking and hence develop a more systematic approach to problem-solving. These include whole-class discussions, observations and relations to previous problems. In summarising their findings, the authors conclude:

- Metacognition instruction is most effective when it takes place in a domain-specific context.
- Metacognition instruction is likely to be most effective when it is provided in a systematically organised manner under the direction of the teacher.
- It is difficult for the teacher to maintain the roles of monitor, facilitator, and model in the face of classroom reality, especially when the students are having trouble with basic subject matter.

Helping students develop metacognition is clearly important – vital, in fact. The teacher will not always be there to help and guide students through the problem-solving process, so students must learn to pause, reflect and strategise. However, successful metacognition relies on one thing above all else: domain-specific knowledge. Without comprehensive knowledge of the content involved in the problem stored, organised and automated in long-term memory, the most effective metacognitive strategies in the world will be no help at all in helping the student arrive at the solution.

Summary

So why do many students fail to devise a strategy to solve problems? Firstly, because such a set of generic, teachable strategies does not exist – or at least not in a way that is likely to be of any use to a novice learner. The most

effective problem-solving strategies are likely to be topic-specific (or even question-specific) and their successful implementation relies first on the development of domain-specific knowledge, followed by careful exposure to related problems as will be discussed in Section 9.3. Metacognition is important, but this too is rendered impotent without the bedrock of knowledge. Domain-specific knowledge is the fuel that powers this second stage of the problem-solving process, and without it students are going nowhere.

3. Carry out the strategy

The requirement for success in this third and final stage of the problem-solving process is perhaps the easiest to determine. Having identified what the problem is about and devised a strategy to solve it, students clearly need domain-specific knowledge of all the topics involved. However, as we have seen, far from being merely being a factor in the final stage in the problem-solving process, such knowledge is the key to everything that comes before it. I cannot articulate the importance of knowledge to the problem-solving process any better than the following three quotes from authors who have dedicated most of their professional lives to studying how students think and solve problems:

Daniel Willingham (2006):

> What do all these studies boil down to? First, critical thinking (as well as scientific thinking and other domain-based thinking) is not a skill. There is not a set of critical thinking skills that can be acquired and deployed regardless of context. Second, there are metacognitive strategies that, once learned, make critical thinking more likely. Third, the ability to think critically (to actually do what the metacognitive strategies call for) depends on domain knowledge and practice. For teachers, the situation is not hopeless, but no one should underestimate the difficulty of teaching students to think critically.

André Tricot and John Sweller (2014):

> At any given time, we are unaware of the huge amount of domain specific knowledge held in long-term memory. The only knowledge

> that we have direct access to and are conscious of must be held in working memory. Knowledge held in working memory tends to be an insignificant fraction of our total knowledge base. With access to so little of our knowledge base at any given time, it is easy to assume that domain-specific knowledge is relatively unimportant to performance. It may be difficult to comprehend the unimaginable amounts of organised information that can be held in long-term memory precisely because such a large amount of information is unimaginable. If we are unaware of the large amounts of information held in long-term memory, we are likely to search for alternative explanations of knowledge-based performance. Those alternatives tend to consist of domain-general strategies.

Alan H. Schoenfeld (2009):

> In sum, the findings of work in domains such as chess and mathematics point strongly to the importance and influence of the knowledge base. First, it is argued that expertise in various domains depends of having access to some 50,000 chunks of knowledge in long-term memory. Since it takes some time (perhaps 10 seconds of rehearsal for the simplest items) for each chunk to become embedded in long-term memory, and longer for knowledge connections to be made, that is one reason expertise takes as long as it does to develop. Second, a lot of what appears to be strategy use is in fact reliance on well-developed knowledge chunks of the type 'in this well-recognized situation, do the following'.

What I do now

The three stages students must go through when solving a problem are:

1. Understand what the problem is about.
2. Devise a strategy to solve the problem.
3. Carry out the strategy.

Domain-specific knowledge is the key to all three – not just for successfully executing the final stage of the problem-solving process, but for getting to

this stage in the first place. The reason many of my students were not good problem-solvers was not because they lacked some magic problem-solving skill, it was because they lacked strong domain-specific knowledge, stored, organised and automated in their long-term memories. Without such knowledge, students simply cannot identify what the question is asking, select an appropriate strategy, or carry out that strategy, no matter how many times they watched me do exactly that.

Why did I not realise this before when it seems so obvious now? Well, precisely because I thought problem-solving was a skill, and when students were solving problems, the first thing that I observed going wrong was the first stage of the problem-solving process. Think back to Schoenfeld's (2009) observation that students spent 18 out of 20 minutes exploring, with no time left for analysing, planning, implementing and verifying. So it was my instinct to fix that. By focusing on trying to develop their ability to identify what problems were about, followed by the formulation of a strategy to solve problems – assuming both were skills that could be taught independently from knowledge – I was denying students of their best chance to become good problem-solvers. The crucial role that knowledge was playing in the process lay hidden at the bottom of the pile, and yet it was the one thing keeping everything else from tumbling to the ground.

As Schoenfeld (1992) pointed out, clearly successful problem-solving is not just about students' knowledge – it is about how, when and whether they decide to use it. We can help students do this by providing the structured systematic support described in the next section. But the point is that knowledge must be there in the first place. Trying to teach generic problem-solving strategies and skills without it is likely to be a fruitless, frustrating experience, and trying to develop this knowledge alongside is likely to be cognitively overwhelming. Knowledge must come first.

OceanofPDF.com

9.3. …and what can we do about it?

What I used to think

Quite simply, *I thought the best way to help students become better problem-solvers was to give them problems to solve.*

This would usually take place in the context of some 'problem-solving lesson', or when going through an exam paper. The types of problems and the content these sessions covered would vary greatly, much like the selection of eight problems presented at the start of this chapter.

I believed that by ploughing through problem, after problem, after problem, I could help my students spot what questions were asking, develop strategies to solve them, and carry out those strategies. I thought repeated exposure to problems was kind of a three-for-the-price-of-one mega deal.

In short, I thought I could fast-track my students on the path towards expertise.

Sources of inspiration

- Barton, C. (2017) 'Doug Lemov', *Mr Barton Maths Podcast*.
- Emeny, W. (no date) 'The Elements: Interweaving', *Great Maths Teaching Ideas*. Available at: http://www.greatmathsteachingideas.com/the-elements-interweaving/
- Kalyuga, S., Ayres, P., Chandler, P. and Sweller, J. (2003) 'The expertise reversal effect', *Educational Psychologist* 38 (1) pp. 23-31.
- Rohrer, D., Dedrick, R. F. and Burgess, K. (2014) 'The benefit of interleaved mathematics practice is not limited to superficially similar

kinds of problems', *Psychonomic Bulletin and Review* 21 (5) pp. 1323-1330.

- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

As we have seen throughout this chapter – and indeed, as a glance through most assessments will reveal – many of the problems students encounter are novel and largely unpredictable. The chances are that students will never have met nor practised problems with exactly the same combination of topics, or set in exactly the same context. This causes a bit of a problem for the explicit model of teaching that I have developed so far throughout this book. I am convinced that well-planned and presented examples and explanations, combined with intelligently sequenced practice questions, is by far the best way to help students acquire and automate the core knowledge of all concepts. But does that help students answer these unpredictable problems? Likewise, we have seen in Section 6.7 how the Expertise Reversal Effect (*eg* Kalyuga *et al*, 2003) suggests that a gradual release from teacher-directed to independent work is needed once students reach a level of domain-specific expertise. But what does this gradual release look like?

Back in Section 3.9 we met Willingham's (2002) distinction between flexible and inflexible knowledge. Knowledge is flexible when it can be accessed out of the context in which it was learned and applied in new contexts – in other words, it is the exact kind of knowledge needed for solving these unpredictable, novel problems. Inflexible knowledge is meaningful, but narrow in that it is tied to the concept's surface structure, and the deep structure of the concept is not easily accessed. Willingham argues that the development of inflexible knowledge is a necessary step along the path to expertise. We cannot miss out this step and simply teach our students to solve problems.

So, how might we facilitate the development of this flexible knowledge?

What I do now

Let's assume we want our students to be able to solve AO2 and AO3 questions that involve the area of simple shapes – questions such as these from Corbett Maths (Figures 9.2 and 9.3).

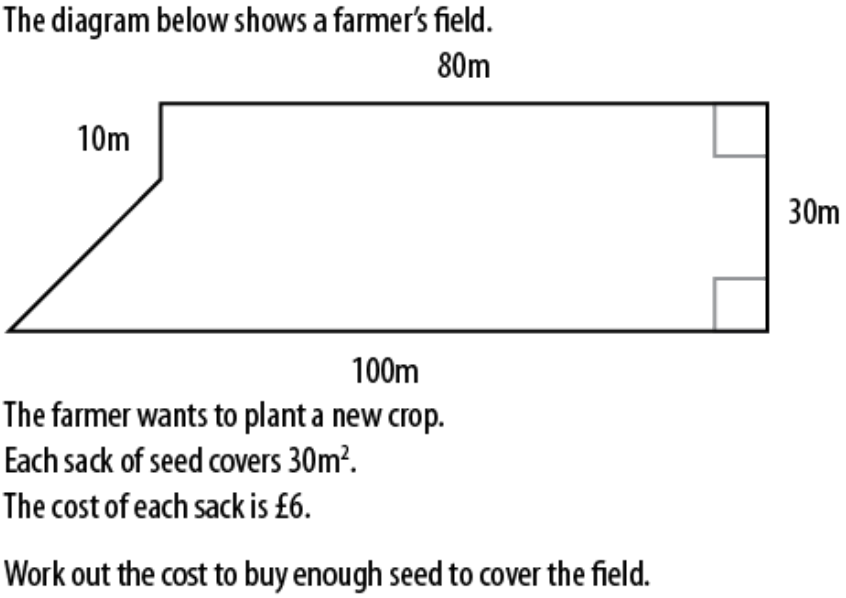


Figure 9.2 – Source: Corbett Maths, Area of Triangles, available at: https://corbettmaths.files.wordpress.com/2013/02/area-of-a-triangle-pdf.pdf

I now adopt the following three-stage process:

1. Teach and assess domain-specific knowledge

I have argued in the previous section that knowledge is the key to everything, and hence before students attempt problems like these, they need to be secure in the relevant domain-specific knowledge. This will involve knowing the formulae for the area of simple shapes, recognising when to use these formulae, and being able to use them. Ideally much of this will be automated in order to free up cognitive capacity in working memory in order to dedicate attention to the more global issues of the problems students will eventually encounter.

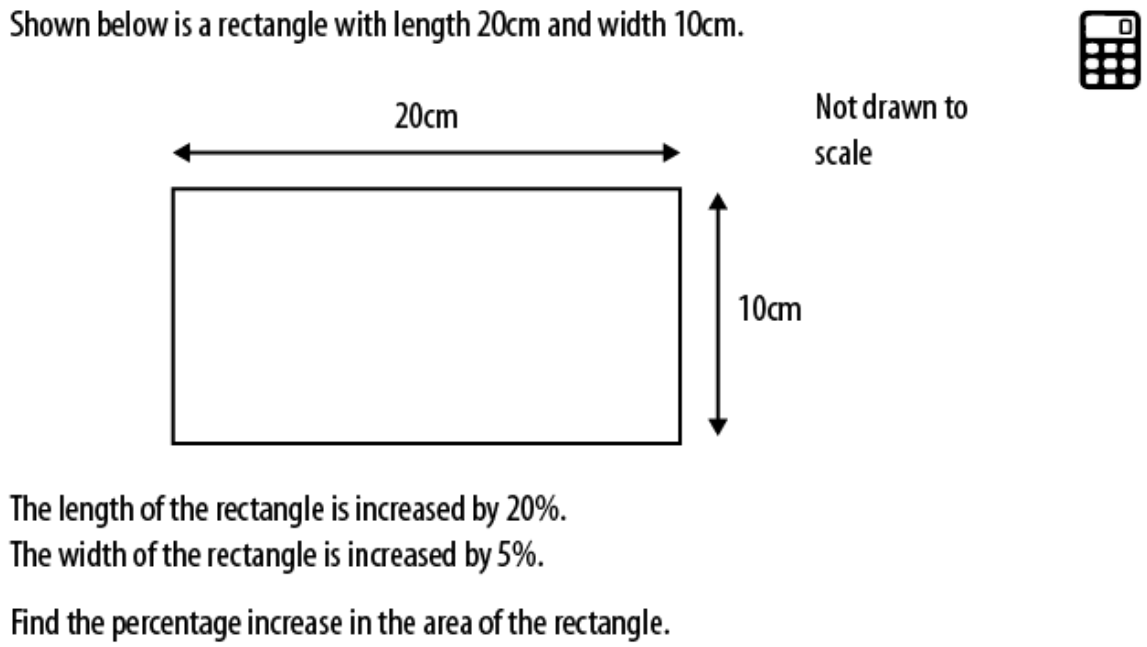


Figure 9.3 – Source: Corbett Maths, Area of Rectangles, available at: https://corbettmaths.files.wordpress.com/2013/02/area-of-rectangles-pdf.pdf

For me, the best way to help students acquire and automate this domain-specific knowledge is via the model of explicit instruction described so far in this book. Students are presented with a series of example-problem pairs and carefully chosen sequences of practice questions. Crucially, among the examples will be non-examples. So, if students are learning how to work out the area of a triangle, they will be presented with triangles where the two marked lengths are not the base and perpendicular height, or with shapes that are not triangles. We have seen how being able to identify the deep structure of a problem is key to the success of any problem-solving process, hence as early as possible students should be compelled to think hard during the development of these basic skills – to consider differences as well as similarities, and not just cruise through on auto-pilot.

I then use the principles of formative assessment (Chapter 11) and the power of tests (Chapter 12) to assess when students are ready to move on.

2. Batch related problems

Once this basic knowledge is secure, it is time to introduce students to some problems.

Students need experience of solving lots of problems. It is this bank of experience that experts rely on when encountering new problems – recognising key features that they can relate to problems solved in the past, and then accessing their deep pool of domain-specific knowledge to finish the job. But it is in the *selection* of these problems where I have previously gone wrong. Previously any old problem would do. But this was made on my misguided belief that problem-solving was a generic, transferable skill.

There are two ways I now like to batch related problems. In both cases, I will give my students an opportunity to try these problems first. Then I will either use the technique of 'show-call' (discussed in Section 6.2 and in my interview with Doug Lemov) as a means of showcasing successful and interesting attempts by students, or I will carefully model the solution myself, clearly articulating my thoughts. It is here that I will make explicit any topic-specific strategies we can use to solve the problems, at a stage in the learning process when students are best placed to benefit from them.

1. Interweave prior topics
We will discuss the Interleaving Effect in the context of memory in Chapter 12, but here it is enough to point out that 'interweaving' – Will Emeny's term for the process of building previously learned skills into questions on the current topic being studied – can be of great benefit to students. So, presenting students with problems on the area of simple shapes that include, say, decimal, fractional or algebraic dimensions, require a conversion of units, involve ratio, or require students to perform written multiplication or division, can be an effective way to deepen understanding of the concept in hand whilst also inducing retrieval of previously studied material.

2. Same Surface Different Deep (SSDD) Problems
Just like during the development of domain-specific knowledge, I want students to be able to recognise both the similarities and differences between problems in order to correctly identify the relevant deep structure. By presenting students with batches of problems in different contexts, but which all have the same underlying deep structure, we rob them of the opportunity to recognise these differences. Back in Section 7.4, based on the work of Rohrer *et al* (2014), I discussed a specific selection of non-examples that I use to do this – those with the same surface structure but different

deep structures (SSDD problems). We can also use SSDD problems to help our students become better problem-solvers. So, in the context of, say, the area of a rectangle, I might present students with a selection of problems that appear visually similar, but whose deep structures are very different. All questions are taken from the wonderful Corbett Maths.

We may have a loci problem involving a rectangular hall:

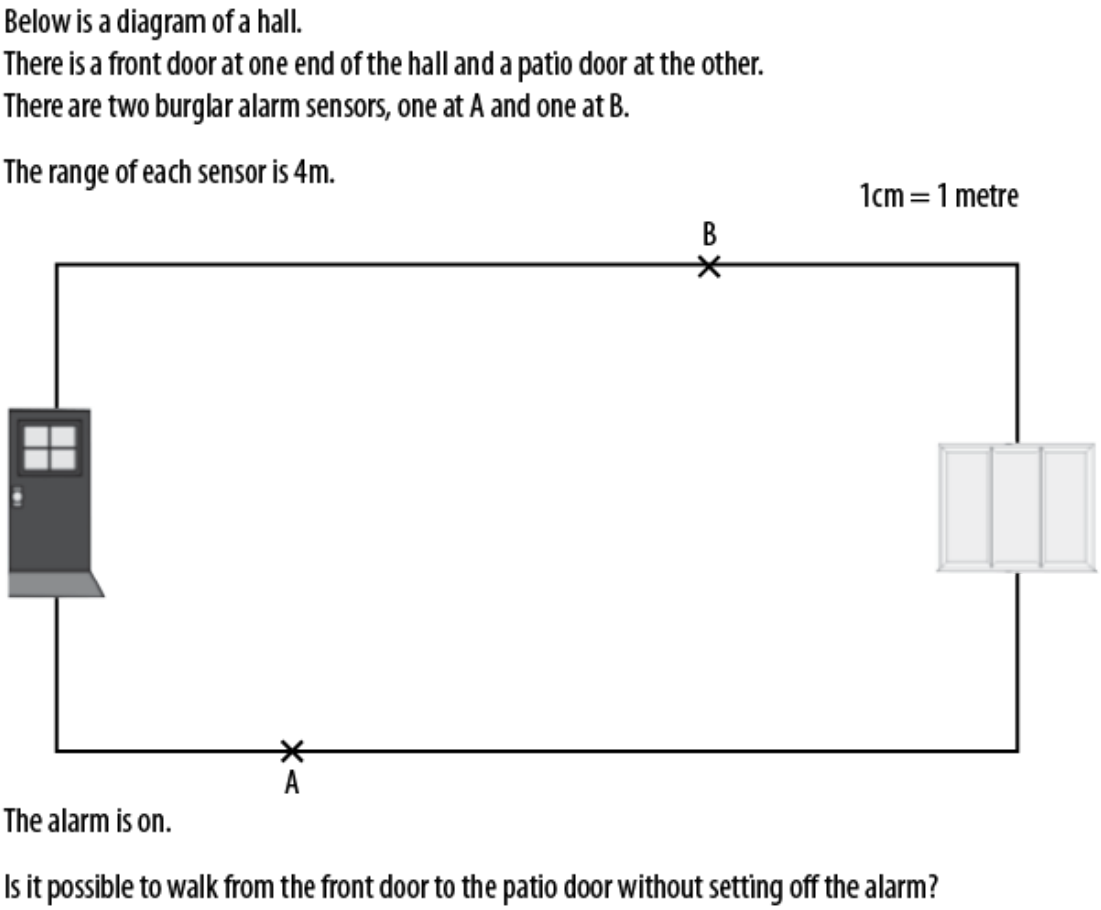


Figure 9.4 – Source: Corbett Maths, Loci, available at: https://corbettmaths.files.wordpress.com/2013/02/loci-pdf1.pdf

A question involving forming an algebraic expression for the perimeter of a rectangle:

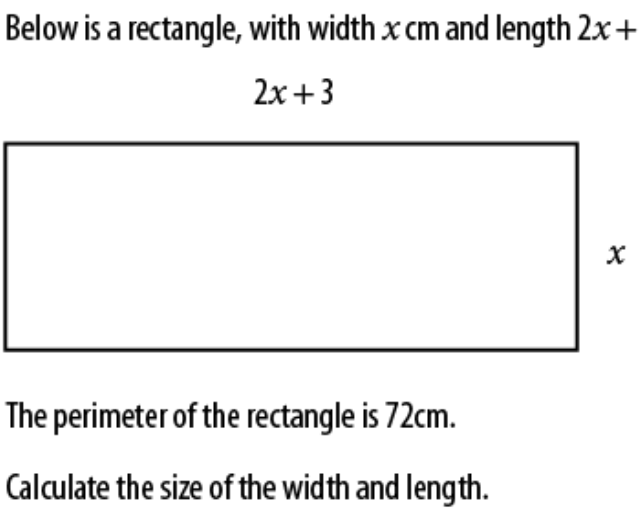


Figure 9.5 – Source: Corbett Maths, Forming and Solving Equations, available at: https://corbettmaths.files.wordpress.com/2013/02/forming-and-solving-equations-pdf1.pdf

The transformation of a rectangle on a coordinate grid:

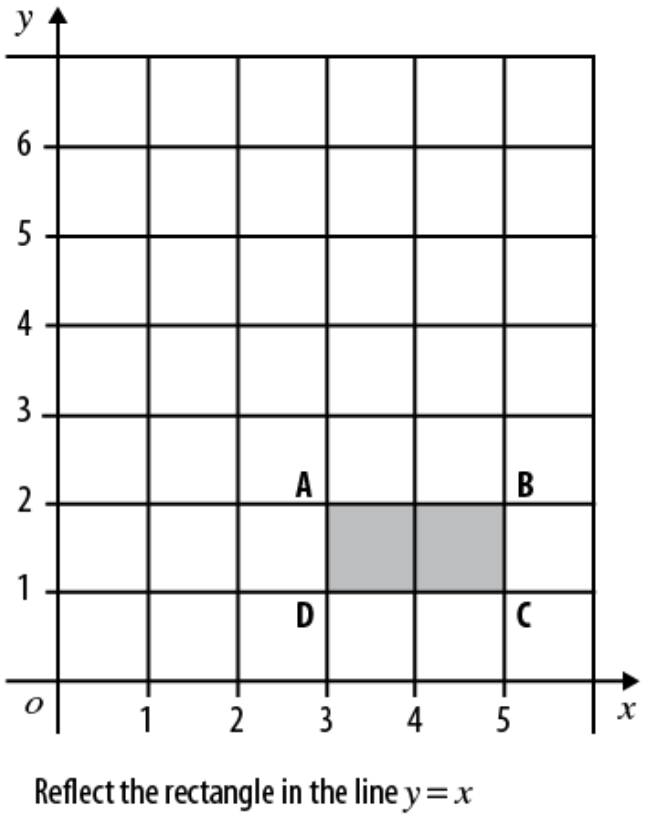


Figure 9.6 – Source: Corbett Maths, Reflections, available at: https://corbettmaths.files.wordpress.com/2013/02/reflections-pdf.pdf

All mixed in with problems concerning the area of a rectangle:

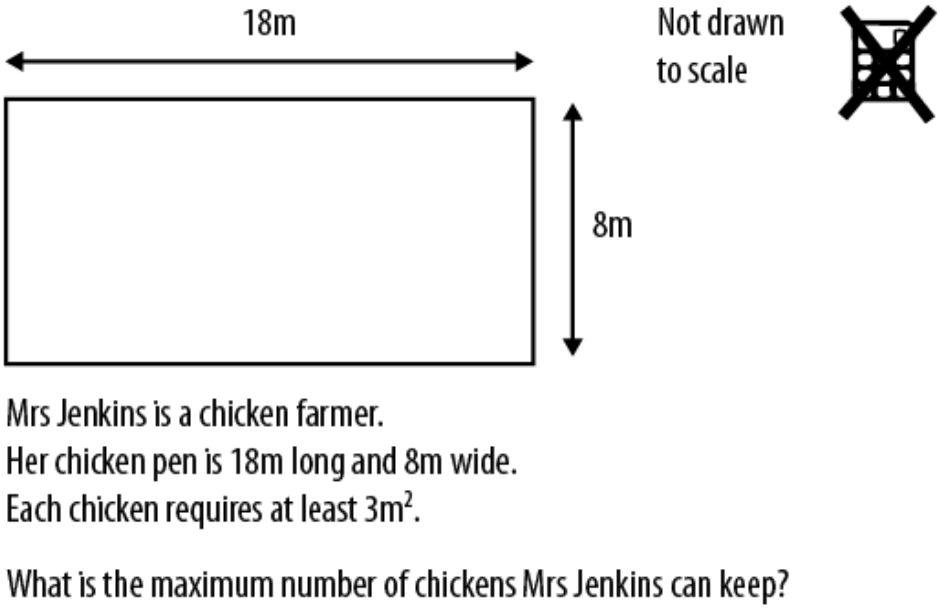


Figure 9.7 – Source: Area of Rectangles, available at: https://corbettmaths.files.wordpress.com/2013/02/area-of-rectangles-pdf.pdf

Students have the opportunity to revisit topics they have not seen in a while, thus benefiting from Chapter 12's Spacing Effect, whilst at the same time being forced to consider each problem's deep structure instead of automatically assuming that each time they see a rectangle they just multiply the dimensions together regardless of the question.

SSDD problems could well provide the key to helping students successfully navigate the kind of subtle contextual problems I discussed in Section 9.2. For example, having taught sharing into a ratio, I would now present my students with the following four problems:

Sarah and Carl share some money in the ratio 2 : 3

They have £350 in total

How much does Carl get?

Ben and Zoe share some money in the ratio 2 : 5

Zoe gets £210 more than Ben.

How much does Ben get?

Matthew and Rachel share some money in the ratio 3 : 5

Matthew gets £120

How much does Rachel get?

Mollie and Laura share some money in the ratio 5 : 4

Mollie gets £150

How much do they get all together?

Presenting these related problems together should help focus students' attention on the differences between them – in other words, the forming of discriminations. Hopefully this will help them avoid simply dividing the quantity by the sum of the two parts every time they encounter a ratio question.

As previously mentioned, the development of bundles of SSDD problems is one of my next big projects. If you want to help me out, then please get in touch.

The key to both of these techniques – interweaving and SSDD – is that students are not simply being bombarded with a random assortment of problems as might occur in a 'problem-solving lesson'. The problems have been carefully chosen. They are related in a way designed to help students make connections, think hard, and build up a bank of experience which they can draw from when presented with novel problems in the future. They should lead to the development of organised knowledge in long-term memory. In short, they are intended to slowly and systematically help students view problems how we do, and hence transition along the path from novice to expert.

3. Present problems in isolation

Only after we have successfully navigated through the first two stages will I present these kinds of problems in isolation across different time periods, again to take advantage of Chapter 12's Spacing Effect. This tends to happen

in homeworks, low-stakes quizzes or when students complete practice exam papers in the final push for revision. I ensure I keep track of which skills I have tested to ensure students never go too long without practising retrieval of a key skill.

This is by no means a revolutionary way to help students develop into problem-solvers. Indeed, it may seem a little dull. But it is a way that seems to be supported by experience, research, and what we know about how students think and learn. By previously assuming problem-solving is a generic skill – and is hence something that I could teach – I would jump straight to this third stage in the process. I would present students with a wide variety of problems, across a whole host of topics, hoping that if I explained things clearly enough when going through the problems, my students would become expert problem-solvers like me. What I know now is that this simply does not work. You cannot jump straight to the third stage, just as you cannot fast-track expertise.

OceanofPDF.com

9.4. Why struggle and failure aren't always good – Part 2

What I used to think

I used to love the sight of my students struggling through problems. Scratching heads, heavy sighs, and even the snap of a pencil thrown down in frustration were the soundtrack to learning. Even if my students didn't get the right answer, they had surely learned lots from the experience? But if they did get the right answer – even if it was not perhaps the way I would have taught it – then even better! They had struggled their way through, come up with a strategy, and got there in the end.

Isn't that what maths, and indeed life in general, is all about?

Sources of inspiration

- Clark, R., Kirschner, P. A. and Sweller, J. (2012) 'Putting students on the path to learning: The case for fully guided instruction', *American Educator* 36 (1) pp. 6-11.
- Christodoulou, D. (2017) *Making good progress?*. Oxford: Oxford University Press.
- Kapur, M. (2014) 'Productive failure in learning math', *Cognitive Science* 38 (5) pp. 1008-1022.
- Kapur, M. and Toh, L. (2015) 'Learning from productive failure' in Cho, Y. H., Caleon, I. S. and Kapur, M. (eds) *Authentic problem-solving and learning in the 21st century*. Singapore: Springer, pp. 213-227.
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of

constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.

- Lovell, O. (2017) 'A Conversation with John Sweller', *ollielovell.com*. Available at: http://www.ollielovell.com/pedagogy/johnsweller/
- Lovell, O. (2017) 'ERRR #009. Andrew Martin, Load Reduction Instruction, Motivation and Engagement', *ollielovell.com*. Available at: http://www.ollielovell.com/errrpodcast/errr-009-andrew-martin-load-reduction-instruction-motivation-engagement/
- Sweller, J. (1988) 'Cognitive load during problem-solving: effects on learning', *Cognitive Science* 12 (2) pp. 257-285.
- Sweller, J. (2017) '*TES* talks to… John Sweller', *TES.com*. Available at: https://www.tes.com/news/tes-magazine/tes-magazine/tes-talks-john-sweller
- Sweller, J., Mawer, R. F. and Howe, W. (1982) 'Consequences of history-cued and means-end strategies in problem-solving', *The American Journal of Psychology* 95 (3) pp. 455-483.

My takeaway

Back in Section 2.7, I outlined one reason why struggle and failure may not always be good – students are likely to need to believe that they can be successful in order to be willing to expend sufficient effort on a problem or activity, and too much struggle and failure prevents that. Here we have another reason.

Sweller *et al* (1982) make the rather worrying claim that '**learners can engage in problem-solving activities for extended periods and learn almost nothing**'. To understand why this may be the case, let's look at two examples.

Consider the Higher Tier GCSE question on the next page.

Before reading on, why not try the problem yourself, and compare your approach and thought processes to what follows?

Paul has 8 cards.
There is a number on each card.



Paul takes at random 3 of the cards.
He adds together the 3 numbers on the cards to get a total T.

Work out the probability T is an odd number.

Figure 9.8 – Source: Edexcel November 2014 GCSE Maths Higher Paper 1

Imagine we give this question to two groups of students. The first group are experts in the sense that they have a good understanding of the basics of combined probability, they have solved related problems before, they are fluent in their mental addition abilities, and can systematically list outcomes. How would they approach the problem?

Expert learners may quickly identify this as a question about combined events, and use their knowledge and experience of similar probability questions to decide upon a suitable tool to use to solve it. Despite no mention of a tree diagram, experts may recognise 'Paul takes 3 cards at random' as a cue to employ this particular tool. Now comes some struggle, however, with how to construct the tree diagram. Experts may initially draw four branches, one for each of 2, 3, 4 and 5. However, seeing just how big this tree diagram is going to be may cause them to pause. Careful consideration of the question may lead them to realise that two branches will be enough – 'odd' and 'even', and write the respective probabilities. Next we have the issue of the second set of branches, and their probabilities. Again, despite the absence of the phrase 'without replacement', experts will recognise that taking three cards at random implies no replacement, and hence correctly deduce probabilities for each of the branches. Experts will continue in this way, all the time struggling, but calling upon their past knowledge and experience of probability problems to help them through. Crucially, their fragile working memories are never overloaded with issues such as 'how do I draw a tree diagram?', 'what is odd + odd', or 'what probability goes here?', because such domain-specific knowledge is stored and automated in long-term memory and readily accessible. This frees up

cognitive capacity for thinking about the global picture, allowing experts to take a step back from the problem, rising above the tiny details, and plot a path towards the solution. Such students will not find this strategy easy, but through solving the problem in such a manner, they learn. Learning is a change in long-term memory, and now students have even more experience to add to their large, well-organised schemas. When faced with a related problem in the future, these experts will be better placed to deal with it. Likewise, taking on this problem with minimal guidance from the teacher is likely to be far more beneficial to their learning than exposure to worked examples and rigorous modelling – hence Section 6.7's Expertise-Reversal Effect.

How would a group of novice learners cope with the same problem?

Without sufficient domain-specific knowledge and experience, they are unlikely to recognise that a tree diagram is needed. Immediately they are at a disadvantage. Failure to recognise the deep structure of the problem causes novices to cling to the surface structure. 'He adds together 3 numbers' may signal a way in. Novices may choose three of the cards – possibly a 2, 3 and 4 – write them down and add them together. Without fluency in mental addition, even this process proves cognitively demanding. But students get through it, and next choose three more cards, this time 3, 4 and 5. They continue in this way, writing down three cards and adding up their totals. Crucially, there is no system to the way these cards are chosen, and inevitably some combinations are repeated and many more are missed. The task of choosing cards and calculating their totals may prove so cognitively demanding that novices do not have any spare cognitive capacity to recognise patterns. They do not realise that it is not the actual totals that matter, but whether those totals are odd or even. They just carry on regardless. Moreover, students are so consumed with the minutiae of the problem that no cognitive capacity remains to consider the global picture – why are they doing this? The result is that the novices may end up with an assortment of lists and totals, but not actually do anything with it – the fact that this is a probability question was pushed out of working memory long ago when the first set of cards was being processed. Just like the experts, novice students have struggled through this problem; but unlike the experts,

they may not have not learned anything from it. Nothing may have changed in long-term memory because no cognitive capacity was available in working memory to make sense of the problem and hence instigate a change. Crucially, when faced with a related problem in the future, these novices will be no better placed to deal with it. Their struggle has been in vain, their time has been wasted, and their confidence is likely to have been hit hard.

In order to fully grasp the potential fruitlessness of novices solving problems in this manner, we need to think of an aspect of our lives where we are novices. Fortunately, for me, there are many.

Not long ago, whilst travelling on the motorway, I got a burst tyre. My phone was out of battery, and I found myself stranded on the hard shoulder, armed with nothing more than a tattered car manual, no external help available and zero experience. After concluding that my technique of staring at the tyre for a few minutes hoping it would miraculous inflate was probably not going to work, I began leafing through the manual. After ten frustrating minutes I managed to locate my spare tyre, and after a further 20 minutes I actually got it out of the boot. I then discovered some weird tool called a 'jack' that my mum had got me for Christmas years ago and was still in its packaging. After a frustrating hour I managed to get the jack in place, remove the wheel and make the change. Along the way my working memory had been so swamped considering each minor detail ('what does this do?', 'where does this go?', 'has that last move really helped me?') that by the time I finished I had no real clue how I did it. Sure, there is a chance that if I found myself in that situation again I might be better equipped and hence do the job in, say, 45 minutes. However, there is little doubt that I would learn far more from an expert carefully demonstrating the steps involved, breaking down the complex skill following the principles of deliberate practice, articulating their thoughts, giving me opportunity to practise, assessing my understanding, providing feedback and so on. Being forced to solve a problem like this unaided was definitely not the best way to help me acquire the domain-specific skills needed to solve a problem like this. Imagine an expert in the same situation – someone who had knowledge and experience of changing lots of tyres. Sure, they may have never changed a tyre on my

particular make and model of car, and hence they may struggle in places. But they can call upon past experiences, and crucially maintain some spare cognitive capacity in working memory to keep an eye on the global problem. If and when they successfully change my tyre, they will have learned from the experience.

I am painting an extreme picture here to make my case, but the point is this – the struggle students endure when solving problems may not be worth the benefits to their learning.

According to Sweller (1988), for students to become good problem-solvers they need to form mental schemas from domain-specific knowledge which they can then apply to different situations. Unlike experts, novices lack the appropriate schemas to recognise and memorise problem configurations. So, just like me with the car tyre and the novice learners with the probability question, they set about solving problems by focusing on the detail and ignoring structure. In other words, they embark upon a means-end analysis.

Such a strategy has two drawbacks.

First, as can be seen in the examples above, trying to solve problems in this manner is cognitively demanding. The learner has to continually maintain and process the following aspects of the problem in his or her fragile, limited working memory:

- The current problem state – Where am I right now in the problem-solving process? How far have I come toward finding a solution?
- The goal state – Where do I have to go? What is the solution?
- The relations between the goal state and the problem state – Is this a good step toward solving the problem? Has what I've done helped me get nearer to where I need to go?
- The solution steps that could further reduce the differences between the two states – What should the next step be? Will that step bring me closer to the solution? Is there another solution strategy I can use that might be better?

- Any other subgoals along the way

This process risks overloading working memory, and hence is unlikely to lead to a successful solution.

The second drawback is potentially more serious. The problem-solving process novices employ is designed to solve a single problem, not to learn from the experience. The fact it is so cognitively demanding means precious working-memory resources are consumed by the minutiae of the problem and hence diverted away from transferring and storing information in long-term memory. The result is that the all-important mental schemas necessary for successful transfer are not developed. Local goals and relationships may swamp and obfuscate the bigger picture. Indeed – and this is a big one – even if the learner somehow manages to solve the individual problem they have been presented with, they may not be able to solve a related one as they have not learned anything systematic and transferable in the process. The learner may win the battle, but will probably lose the war.

In contrast, as we have seen in Chapter 6, studying a worked example reduces working memory load because students are only being asked to understand the solution, not solve the problem and understand the solution at the same time. This directs attention (*ie* working memory capacity) to learning the essential relations between problem-solving moves. Students learn to recognise which moves are required for particular problems, the basis for the development of schemas.

Moreover, solving problems independently without sufficient domain-specific knowledge is an ineffective way of teaching the fundamental skills and procedures required to solve the problem when compared to modelling, example-problem pairs and Intelligent Practice. Christodoulou (2017) puts it like this:

> Paradoxically, methods of teaching which ask pupils to do real and complex tasks will prevent pupils from developing the mental models they need to actually be able to solve those real and complex tasks. Pupils will be caught in a chicken-and-egg scenario: unable to solve complex problems because they do not have the necessary models, but

> unable to develop those models because they spend all their time trying to solve those complex problems.

The key point is that problem-solving is not a learning device – at least not for novice learners. Problem-solving must come at the end of the process, after the necessary domain-specific knowledge has been acquired. At that stage, when basic skills have been automated, there is capacity in working memory to direct towards appreciating the global features of a problem and hence learn from the experience. Indeed, this supports the conclusions from the Expertise Reversal Effect discussed in Section 6.7, whereby learning from problems eventually becomes more beneficial to students than learning from worked examples. But – and it is a big but – this is only true for experts in a particular domain. And how many of our students can we apply that tag to, especially towards the start of each academic year?

A twist: productive failure

There is an interesting body of research that offers an alternative view. Kapur and Toh (2015) describe a process of 'productive failure', whereby students who are left to first struggle at a problem, and then receive explicit instruction, show greater long-term conceptual understanding and transfer than students who receive the explicit instruction without first enduring the struggle. They explain: 'What is critical is not the failure to develop the canonical solution per se but the very process of generating and exploring multiple representations and solution methods, which can be productive for learning provided that direct instruction on the targeted concepts is subsequently provided'.

Kapur (2014) conducted a maths-specific study, involving 75 ninth-grade students studying standard deviation – a concept they had never encountered before. Students who engaged in problem-solving before being explicitly taught demonstrated significantly greater conceptual understanding and ability to transfer to novel problems than those who were taught first.

The logic seems somewhat similar to the idea of a Pretest that we will discuss in Chapter 12, whereby exposure to the concept before instruction leaves students better prepared to attend to the critical features during

instruction. Likewise there are parallels with Dan Meyer's headache-aspirin approach discussed in Section 2.5.

It is important to note that in separate 2017 interviews with Oliver Lovell, both John Sweller and Andrew Martin outline serious concerns they have with the experimental design used in these and other productive failure studies, citing the presence of multiple variables which make it impossible to determine cause and effect.

But I have another concern.

Productive failure is fascinating, and makes me think back to my tyre example. What would have helped me learn better – struggle followed by instruction, or instruction from the start? Quite possibly the former. But perhaps the bigger question is – would the struggle have caused me to have given up before the help came along? Indeed, key features of Meyer's headache-aspirin approach were that initial success preceded the struggle – students' methods were not wrong, they were just inefficient – and the struggle was brief.

Struggling and failing to solve problems may not be all bad news, as long as students battle through the problem, receive good explicit instruction afterwards and – this might be the most important part – so long as they are then responsive to that instruction. Whether this process is successful – and thus the failure is indeed productive – may well be down to the students themselves, and specifically their prior experiences with mathematics. The students in the Kapur (2014) study were all from a co-ed private school in India, with high levels of prior knowledge in associated concepts (mean, range, *etc*). They were all willing to have a go at the unfamiliar problem, and were responsive to the subsequent teaching. I saw parallels with a Year 7 class I once had and was lucky enough to teach for the next few years. They were not scared of anything – and that was a mindset that I carefully constructed over the long-term, built upon many moments of success. But I have seen too many middle-set Year 10s struggle when presented with something unfamiliar, give up, and then not want to know when I subsequently try to teach them. This may be more to do with my failings as a teacher, or the low levels of self-efficacy of these students based on their past

experiences. But for those students – and I fear there are many like them – I feel it is better to present them with carefully planned guidance first, so they are ready and willing to take on more complex problems as their learning develops.

What I do now

Is it correct to say that novice learners never gain anything from struggling on complex maths problems? I don't think so. Just as I am likely to change a tyre more quickly next time, so too are our students likely to recall something from the process of attempting to solve a problem that may transfer to other situations. So, their struggle is not entirely in vain.

But will they learn as much, in the same amount of time, as they would following the key principles of explicit instruction described in this book, comprising the careful selection and presentation of examples, followed by deliberate and intelligent practice, effective formative assessment and the considered implementation of Chapter 12's desirable difficulties? In my opinion, definitely not.

Problem-solving must come at the end of the three-stage process described in the previous section – specifically after domain-specific knowledge has been explicitly taught, practised and assessed.

It is worth reminding ourselves that even within the domain of mathematics, it is possible (and, indeed, likely) for students to be relative experts in one topic, but relative novices in another. All students have strengths and weaknesses. I was always good at algebra, but a bit dodgy with certain aspects of shape and space. Hence, just because students in our top-sets are good mathematicians on average, that does not mean that they benefit from being given problems to solve in areas where they are weakest. Sure, they may well have an advantage acquiring new knowledge based on the interconnected nature of mathematics and what they already have stored and organised in their long-term memories – and hence progress on the journey to expertise more quickly – but their domain-specific knowledge needs to be built up in just the same way as any other novice.

None of this is meant to devalue the use of problems. Indeed, one of the key advantages expert learners have over novices is their vast experience of different problems, which enables them to recognise common features in previously solved problems and apply them to novel situations. Likewise, they build up a repertoire of domain-specific problem-solving techniques. I see this all the time in my best A level students – by the time they sit the exam they have done every single past paper ever written, and as such can recognise similarities between problems with frightening speed. But the point is, before embarking upon this voyage through the world of problems, a certain amount of domain-specific expertise needs to be in place. Otherwise, it could be a load of struggle for not much reward.

OceanofPDF.com

9.5. Independent learners

What I used to think

Alongside wanting my students to become better problem-solvers, I also wanted them to become independent learners. I never really stopped to consider exactly what that phrase meant, but I guess if I was asked to describe the key characteristics of an independent learner, it would be something along the lines of list produced by the Alpha Omega Academy (2012):

1. Curiosity
2. Self-motivation
3. Self-examination
4. Accountability
5. Critical thinking
6. Comprehension with little or no instruction
7. Persistence

Yes, if my students had all those attributes, I would be laughing.

And how best to help my students develop the characteristics necessary to become independent learners? Well, give them very little support, of course. After all, if you create problem-solvers by getting students to solve problems, then surely you create independent learners by encouraging students to become independent?

Sources of inspiration

- Alpha Omega Academy (2012) *7 Characteristics of Independent Learners*. Available at: https://www.aoacademy.com/7-characteristics-of-independent-learners/
- Kalyuga, S., Ayres, P., Chandler, P. and Sweller, J. (2003) 'The expertise reversal effect', *Educational Psychologist* 38 (1) pp. 23-31.
- Kirschner, P. A., Sweller, J. and Clark, R. E. (2006) 'Why minimal guidance during instruction does not work: an analysis of the failure of constructivist, discovery, problem-based, experiential, and inquiry-based teaching', *Educational Psychologist* 41 (2) pp. 75-86.

My takeaway

A quick search for 'independent learning' on *TES* returns a whole host of lesson plans, posters and PowerPoints aimed at helping students become independent learners. They are full of advice along the lines of 'read the question again', 'think for a minute' and 'be brave'. This is fine, so long as the students have the knowledge to be independent. Otherwise, this advice is about as useful as the list of generic problem-solving strategies considered earlier on in this chapter is to the novice learner.

We have already encountered the Expertise-Reversal Effect in Section 6.7. Kalyuga *et al* (2003) argue that experts reach a point when they no longer learn from worked examples and benefit more from solving problems independently. The major difficulty is recognising when students have made the transition from novice to expert and hence can start to work more independently. Too soon, as argued by Kirschner *et al* (2006), and novices may not learn anything due to the cognitively demanding way they approach problems described in the previous section. By encouraging these novice learners to be more independent, and hence withdrawing guidance and structure, we are in essence removing the very thing that will enable them to become more independent.

Hence we are left with a conclusion similar to that reached earlier in this chapter with regard to problem-solving: just as the best way to help students develop into problem-solvers is not necessarily to give them problems to

solve, the best way to help students develop into independent learners is not to expose them to conditions where they need to be independent.

A recurring theme throughout this book is the difference between expert and novice learners, not just in terms of how much they know but also in how they think differently. Experts may well benefit from the conditions suited to independent practice, but only because they are experts. They have well-developed and organised bodies of knowledge stored in long-term memory that they can use in novel situations to solve problems without guidance from the teacher. Novices lack this. And here is the key point – we cannot force the transition from novice to expert by simply exposing novice learners to the same conditions in which experts thrive. There is no shortcut to expertise, and there is no shortcut to independence. Students must develop that large body of domain-specific knowledge first. And the rather counter-intuitive conclusion from this is that the best way for students to gain this knowledge is for us to create conditions that seem the antithesis to independent learning, namely guided instruction, example-problem pairs, and deliberate practice. Only with carefully planned and substantial support can we help our students become the independent learners we want them to be.

What I do now

Before I set students off to work independently, I ensure they have enough domain-specific knowledge to solve problems on their own. Counter-intuitively, I help them become more independent by using the techniques of teacher-led explicit instruction discussed throughout this book.

OceanofPDF.com

9.6. If I only remember 3 things…

1. Problem-solving is not a skill. In order to successfully solve problems, students need domain-specific knowledge.
2. The three-step procedure I use to help my students become problem-solvers is: develop inflexible knowledge via explicit instruction; carefully batch related problems together; present those problems in isolation over time.
3. Students can be struggling, but not learning.

OceanofPDF.com

10. PURPOSEFUL PRACTICE

Throughout this book I may have given the impression that I am anti-problem-solving. I certainly am not. I love a problem, and my students to need to build up the experience of solving them. It is just that in the past I have forced my students into problem-solving activities too early on in their development of a skill or concept. Findings from Cognitive Load Theory and my own experience of watching students struggle fruitlessly have led me to the conclusion that this is not a good idea. Novice learners may not learn from solving problems.

I suggested in the previous chapter that a model to help our students become the problem-solvers we all want them to be is to first develop inflexible knowledge via explicit instruction, then carefully batch related problems together, before finally presenting those problems in isolation over time.

But there may be another way. There is a branch of problems that fit under the category of what I will call Purposeful Practice. They enable students to develop that all-important inflexible knowledge and procedural fluency that is key in the transition from novice to expert, but also provide opportunities for students to make connections and think a lot deeper about concepts.

Where Deliberate Practice is my preferred model for introducing a concept, Purposeful Practice has become my go-to strategy for recapping a topic, which, as we will find out, is how we spend the majority of our time as a teacher.

OceanofPDF.com

10.1. The most difficult part of teaching

What I used to think

If you had asked me a few years ago to name the things that make teaching difficult, top of my list would have been the topic, the ability of the students and their behaviour. Some topics are inherently more difficult and potentially boring than others, some students struggle mathematically for a whole host of reasons, and without good behaviour all other techniques and strategies fade into insignificance.

Absent from the top of my list would have been the prior experience of my students in the concept I was about to teach them. Sure, I would assess baseline knowledge – although not as effectively as I will describe in Chapter 11 – and adapt my teaching accordingly. Sure, I would attempt to differentiate – although looking back I wished I had done this more in terms of the time students spent on tasks as opposed to the tasks themselves. But when planning a lesson, any consideration of students' feelings towards the topic or concept was pretty low down the list. I approached most topics as if it was the first time students had ever seen them.

I can illustrate this no better than the following admission: my Year 11 Foundation GCSE series of lessons on factors and multiples looked very similar to my Year 7 lessons on the same topic.

Sources of inspiration

- Department for Education (2014) *National curriculum in England: mathematics programmes of study*. Available at: https://www.gov.uk/government/publications/national-curriculum-in-england-mathematics-programmes-of-study/national-curriculum-in-england-mathematics-programmes-of-study

- Kalyuga, S., Ayres, P., Chandler, P. and Sweller, J. (2003) 'The expertise reversal effect', *Educational Psychologist* 38 (1) pp. 23-31.
- Nuthall, G. (2007) *The hidden lives of learners*. Wellington: NZCER Press.
- Soderstrom, N. C. and Bjork, R. A. (2015) 'Learning versus performance: an integrative review', *Perspectives on Psychological Science* 10 (2) pp. 176-199.
- Sweller, J. (1988) 'Cognitive load during problem-solving: effects on learning', *Cognitive Science* 12 (2) pp. 257-285.

My takeaway

Let's play a game. I am going to give you five concepts, and I want you to tell me at what school year (or age) students first meet these in the 2014 English Mathematics Programmes of Study that form the national maths curriculum. Answers are on the next page, so no cheating!

1. Working with basic angle facts, including angles on a straight line and around a point
2. Identifying factors and multiples
3. Rounding decimals to the nearest whole
4. Recognising equivalent fractions

Got your answers? Okay, here we go:

1. Working with basic angle facts, including angles on a straight line and around a point – Year 5 (age 9 to 10)
2. Identifying factors and multiples – Year 5 (age 9 to 10)
3. Rounding decimals to the nearest whole – Year 4 (age 8 to 9)
4. Recognising equivalent fractions – Year 2 (age 6 to 7)

How many of those topics are on your Key Stage 3 scheme of work? How many of them have you gone over with Year 10s and Year 11s over the years?

Now, depending on the specific scheme of work students are following, there is of course no guarantee that students will meet these concepts at this age. Some may meet them a year later, some a year earlier. Likewise, it is very unlikely that students will have mastered these concepts or tackled them in any great depth, whatever you define those terms to mean. But students have met them. They have experienced them. And hence, they have some feelings, thoughts and preconceptions about them. And this changes everything.

Nuthall (2007) makes the extraordinary claim that students already know 40-50% of whatever we are trying to teach them. The problem is that the 50-60% of things they do not know is likely to be different for every student in our class. Likewise, as we discussed when considering how students think in Chapter 1, if you attempt to teach students something about which they know absolutely nothing, they will have nothing to connect the concept to, which makes it less likely to make the journey from working memory to long-term memory which is necessary for learning to take place. The important distinction between learning and performance that we will encounter in the final chapter means that it is incredibly difficult – if not impossible – to discern exactly what our students really understand. And even if we could, what on earth would we do with that information? Plan 30 personalised lessons for each of our students, six times a day?

There are clear benefits to going over concepts again, either to fill in gaps in knowledge or – as we shall see – even to overlearn beyond the point of mastery. But that does not make the revisiting of concepts any easier to teach.

To illustrate this, imagine you are teaching your middle-ability Year 10 class angle facts. We know from what we have seen above that they first encountered this concept five years ago, and have met it every year since. Consider for a moment what students in your class might be thinking when you announce that we are looking at angles today. I reckon it could be any of the following:

- *I'm okay at it, but I probably need some more practice.*
- *I struggle with this one, but I will give it another go.*
- *I feel pretty confident with this topic.*
- *I hate this topic. I don't get it. I give up.*
- *This topic is dead easy, sir. I nail it every time (no you don't).*

Just like our students all know different things, they will all be thinking different things. Within a class of 30 students, it is possible – and indeed likely – that you will have at least one student pondering thoughts similar to each of the above. And that makes our lives as teachers a bit of a nightmare.

Let's look at how we might respond.

Starting with the first two:

- *I'm okay at it, but I probably need some more practice*
- *I struggle with this one, but I will give it another go*

Can we just give these students complex problems and exam questions to do to help them along the road to expertise? Not really. From the work we have done on Cognitive Load Theory and problem-solving (*eg* Sweller, 1988), we know novices do not learn best from solving problems. So can we just give them more routine practice? Possibly, but it depends on the students themselves. If they are happy doing questions that look the same as those they have done the last five years, then great. If not, then there may be an absence of purpose that could lead to an absence of effort. And how do we move them further along the road to expertise if we confine them to such routine practice?

How about:

- *I hate this topic. I don't get it. I give up.*
- *This topic is dead easy sir. I nail it every time. (no you don't)*

These students definitely need more practice, but they are unlikely to react well to being presented with a topic in the same way that they have seen it many times before. The first may give up because they have experienced previous failure, and the second may wonder why on earth you are offending them with something so ridiculously easy. So can we perhaps motivate them by approaching the topic in an exciting real-life context? Almost certainly not. As we have seen in Chapter 2, real-life maths contexts can often be more trouble than they are worth, and a more reliable driver of motivation is students' achievement and belief they can achieve.

And then there is:

- *I feel pretty confident with this topic*

For a start, how do they know? The key distinction between learning and performance that we will cover in the remainder of this book (see Soderstrom and Bjork, 2015, for a sneak preview) means that neither the students themselves nor we (their teacher) may be able to accurately assess understanding. And as we have seen, exposing novice learners to complex problems may be detrimental for their learning. But on the flipside, if in fact they are on the path to expertise, then the type of instruction that is most suited to novice learners, such as worked examples and Deliberate Practice, may not be appropriate for them due to the Expertise-Reversal Effect (*eg* Kalyuga, 2003), and these students would in fact benefit more from being exposed to more complex problems to solve independently.

And we are likely to have a good smattering of all these thoughts across our class!

And that is why I believe that teaching topics and concepts that students have encountered before is quite simply the most difficult part of teaching. I'd take a topic that students have never met – such as trigonometry or vectors, where I can happily plan how I am going to carefully introduce the key concepts and develop students' skills – any day over revisiting something students have seen many times. And yet, one look at any secondary scheme of work I have ever encountered suggests that a large proportion of our teaching time is spent revisiting topics.

Likewise, I always find it easier to plan the first lesson of a sequence of lessons than the second or third. Once students have been introduced to the key concepts (via example-problem pairs and Intelligent Practice), it is then that divergences start to appear. Some students have grasped the concept more quickly and are ready for something more challenging, whereas others are in need of further practice. Hence, the further we venture through any sequence of lessons, the more instances of the five cases of students outlined above we are likely to encounter, and hence the need for something different grows.

So, what on earth are we to do?

What I do now

Quite simply, when teaching a topic I know students have met at some stage before – either last lesson, or some other point in the past – I turn my attention towards Purposeful Practice.

OceanofPDF.com

10.2. What is Purposeful Practice?

What I used to think

I used to think there were two types of activities you could give students: practice and problems to solve. Practice, as I have argued throughout this book, was necessary in the early knowledge acquisition phase of learning. Problem-solving came later, when students had reached a certain level of domain-specific expertise, and would help push their understanding of concepts to the next level. I did not think there was anything in between these two extremes.

Fortunately, I was wrong.

Sources of inspiration

- Barton, C. (2017) 'Colin Foster', *Mr Barton Maths Podcast*.
- Driskell, J. E., Willis, R. P. and Copper, C. (1992) 'Effect of overlearning on retention', *Journal of Applied Psychology* 77 (5) pp. 615-622.
- Foster, C. (2013) 'Mathematical études: embedding opportunities for developing procedural fluency within rich mathematical contexts', *International Journal of Mathematical Education in Science and Technology* 44 (5) pp. 765-774.
- Foster, C. (2014) 'Mathematical fluency without drill and practice', *Mathematics Teaching* 240, pp. 5-7.
- Foster, C. (2017a) 'Mathematical études', *NRICH*. Available at: https://nrich.maths.org/13206
- Foster, C. (2017b) 'Developing mathematical fluency: comparing exercises and rich tasks', *Educational Studies in Mathematics*. Available

at: https://link.springer.com/content/pdf/10.1007%2Fs10649-017-9788-x.pdf

- Rohrer, D. and Taylor, K. (2006) 'The effects of overlearning and distributed practise on the retention of mathematics knowledge', *Applied Cognitive Psychology* 20 (9) pp. 1209-1224.
- Stewart, D. (no date) *Median Maths Blog*. Available at: http://donsteward.blogspot.co.uk/
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Inflexible knowledge: the first step to expertise', *American Educator* 26 (4) pp. 31-33.

My takeaway

Given the different experiences students have had with many of the topics they encounter in mathematics that I outlined in the previous section, there is a need for another type of activity. Teacher-guided instruction is ideal for novices, less-guided approaches are more beneficial for experts, but in most classes we have a mixture of the two, along with everything else in between. As such, we need something that gives the students sufficient practice to develop knowledge and procedural fluency, whilst at the same time providing opportunities to develop conceptual understanding and the problem-solving attributes discussed in Chapter 9 that experts possess.

The name I will give to the special type of activities or tasks that enable this is *Purposeful Practice*.

Let me show you what I mean with an example.

Say I am teaching my Year 9 class and I want to review their ability to add fractions. Maybe this is the first lesson in a Fractions Review unit, or maybe I taught them how to add fractions yesterday and now I want to give them more practice. I would start with an initial assessment of understanding, making use of diagnostic questions as discussed in Chapter 11. This may identify some major misconceptions that need dealing with there and then. Or it might reveal that some students simply need a reminder of the basic

procedure, which can be achieved efficiently by means of Chapter 6's example-problem pairs. But what do we do next?

Sure, I could give my students a series of fractions to add. But even if I carefully sequenced the questions using the principles of Section 7.8's Intelligent Practice, they are still questions, presented in the same format as they always have been. And for the reasons outlined in the section above, this is unlikely to be suitable for many of my students.

So instead, I could give them the following:

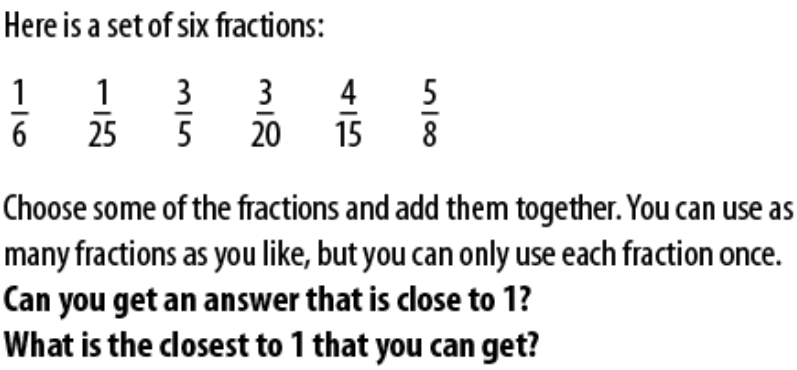


Figure 10.1 – Source: Colin Foster for NRCIH, available at https://nrich.maths.org/13205

This problem gives students plenty of practice of the key skill of adding fractions, but with a wider purpose. How close can the students get to 1? How have they approached this? How do they know if they are closer to 1 than their friend? Is there a quicker way to do this? And then if some students settle upon an answer, we can push them further – what if they could only use 3 fractions? How about 4? What if they can subtract fractions as well? Can they create a set of fractions that will allow them to get to 1 in multiple ways? Students end up adding more fractions than they ever would if they were presented on a worksheet and in a way that rarely causes a moan, and they have the opportunity to develop wider skills along the way. This is an example of Purposeful Practice.

Purposeful Practice – at least in the way I define the phrase – is based around five key principles:

1. Students need to experience early success

In Chapter 2, I argued that too much struggle and failure can be a bad thing for learning, and furthermore that achievement and students' perception that they can succeed are key drivers of motivation. Hence, any task or activity needs to give students the sense that they are making progress, or that they *can* make progress, pretty early on. This is the problem with many open-ended, unstructured tasks – the barrier to entry is just too high.

2. There must be plenty of opportunities to practise the key procedure
We have seen in the chapters on examples and Deliberate Practice just how important practice is. Whether students are responding to example-problem pairs or working through a carefully planned sequence of questions, it is the practice at their own pace and the opportunities to self-explain that helps knowledge transfer from working memory to long-term memory which is necessary for learning to take place. And we need them practising a specific procedure. One of the oft-cited strengths of unstructured, open-ended tasks is that students can tackle them lots of different ways. But I do not always want that. If I am revising a certain procedure or concept, then I want to make sure that students are practising the thing I intend.

3. The practice should feel different
Students have met many of these concepts and topics before. Presenting them again with another worksheet of boring questions, or even another Tarsia jigsaw to cut up, is likely to be demoralising. The student who did not understand it first time around is unlikely to believe that things will be any different this time. Likewise, the student who did understand it is likely to question why on earth they are being asked to do the same thing again.

4. Opportunities must exist for students to make connections, solve problems and think deeper
Success in maths is not just about answering routine questions. Students need to be able to transfer their knowledge to novel situations, spot connections, solve complex problems and think creatively. Indeed, once students reach a certain level of expertise, the sorts of activities that enable them to do this are the very ones that are needed to push their understanding on to the next level.

5. The focus is always on the practice

But there is no fast track to expertise. Students have to develop inflexible knowledge first. They have to have knowledge stored, available and automated in long-term memory in order to free up the cognitive capacity needed to make these connections, solve these problems, and crucially to learn from the experience. And it is worth pointing out again that students must be practising the thing I intend. Sure, they can make connections, solve problems and think deeper, but carrying out the intended procedure must always be at the forefront of the task.

Purposeful Practice is powerful because:

- *Everyone is working on the same task.*

There is no need to create a different set of worksheets for all the perceived different abilities of students in your class. There are no big, potentially dangerous decisions to be made as to when to move students on to the next task. Differentiation is done by time and the connections students make.

- *Students who are not expert are able to practise the key procedure.*

Because the practice feels different, and because there is a larger goal in mind, students are more prone to engage in the sort of practice that is crucial in order to develop the knowledge and procedural fluency they need in order to master the topic.

- *Students who are (or appear) expert have the chance to overlearn.*

Driskell *et al* (1992) along with Rohrer and Taylor (2006) suggest that over-learning – or learning beyond the point of mastery – can be beneficial. This is related to the distinction between learning and performance that we will consider in the next few chapters, but the key point is that learning can still improve after performance has seemingly peaked. Hence, even experts are likely to benefit from the sort of practice of key skills involved in Purposeful Practice activities.

- *These students can also develop deeper levels of understanding through self-explanation.*

The point of these tasks is that the practice has a wider purpose. There is a connection to spot, a relationship to explain, a shortcut to discover, a challenge to be solved, a surprise to rationalise. All of these things cause students to pause and consider what they are doing far more than if they were simply working their way through a series of unconnected questions on autopilot. This is linked to our work on Intelligent Practice in Chapter 7. By working through the calculations, pupils develop procedural fluency, but also have the opportunity to develop wider conceptual understanding. If they are secure in their own knowledge, they may also be able to share their thoughts and help others.

- *Both novices and experts benefit.*

This is the key to it all. Throughout this book we have focused on the difference between novices and experts. Novice learners benefit from teacher guidance and routine practice, whereas the Expertise-Reversal Effect suggests that experts may not. Conversely, expert learners benefit from independent problem-solving activities, whereas Cognitive Load Theory suggests that novices may not. Everyone benefits from Purposeful Practice, both experts and novices and the large contingent of students who are making the journey between the two.

Whilst developing my own ideas for Purposeful Practice, I came across the fantastic work of Colin Foster and his 'mathematical études', which we discussed in depth when I interviewed him for my podcast. The key to an étude – a term of musical origin – is that a self-imposed constraint brings out creativity. In an article for NRICH, Foster (2017a) explains:

> These are intended to be not just lovely rich tasks, but lovely rich tasks that force students' attention onto a particular, important mathematical procedure. So etudes don't allow students a choice of method. Allowing students to choose what method they use is often a great idea, but if you want them to develop their skill at a specific technique, then you have to force them to use that technique, otherwise they can end up unwittingly avoiding their areas of weakness by (perhaps cleverly) solving the problem some other way than you intended. So, mathematical etudes are designed with a specific procedure in mind.

> But, rather than simply asking students to perform that procedure repeatedly (as with tedious exercises), in an etude there is some wider problem-solving goal – something a bit more interesting is going on at the same time!
>
> The answers obtained to the procedure matter for some wider purpose, and so students are more likely to carry out the procedure carefully and check what they have done. Wrong answers are more likely to be challenged by a peer if they matter for some wider purpose. And having something else to focus on beyond the procedure should help to promote the desired automaticity of carrying it out.

Foster (2017b) presents evidence that mathematical études are comparable to traditional exercises in their effects on procedural fluency, going on to argue that this could make études a viable alternative to exercises, since they offer the possibility of richer, more creative problem-solving activity, with comparable effectiveness in developing procedural fluency. This is a key finding, as it means with these kinds of activities you get the best of both words. Indeed, Foster (2013) presents the diagram on the next page, which illustrates the vital role that the very best études and Purposeful Practice activities play.

Too often I would move from the top-left to the bottom-right before my students were ready, and completely neglect the possibility of something existing in the top-right.

What I do now

When introducing a topic or a concept for the first time, I follow a model of Deliberate Practice, encompassing Chapter 4's presentation of information, with Chapters 6 and 7's choice and delivery of examples and practice questions. This is true of completely brand new concepts, such as histograms and graphs of trigonometric functions, and also concepts that are essentially extensions of concepts students have previously studied, such as enlargements with negative scale factors and solving more complex linear equations. Here I want to provide plenty of structured support for my students, and I can do this because their prior experience and knowledge of the topic are similarly limited. I will also do the same if my formative

assessment strategy (see Chapter 11) reveals significant difficulties across the majority of the class with any topic students have previously met before. Whilst not ideal, here I feel it is probably best to start from square one.

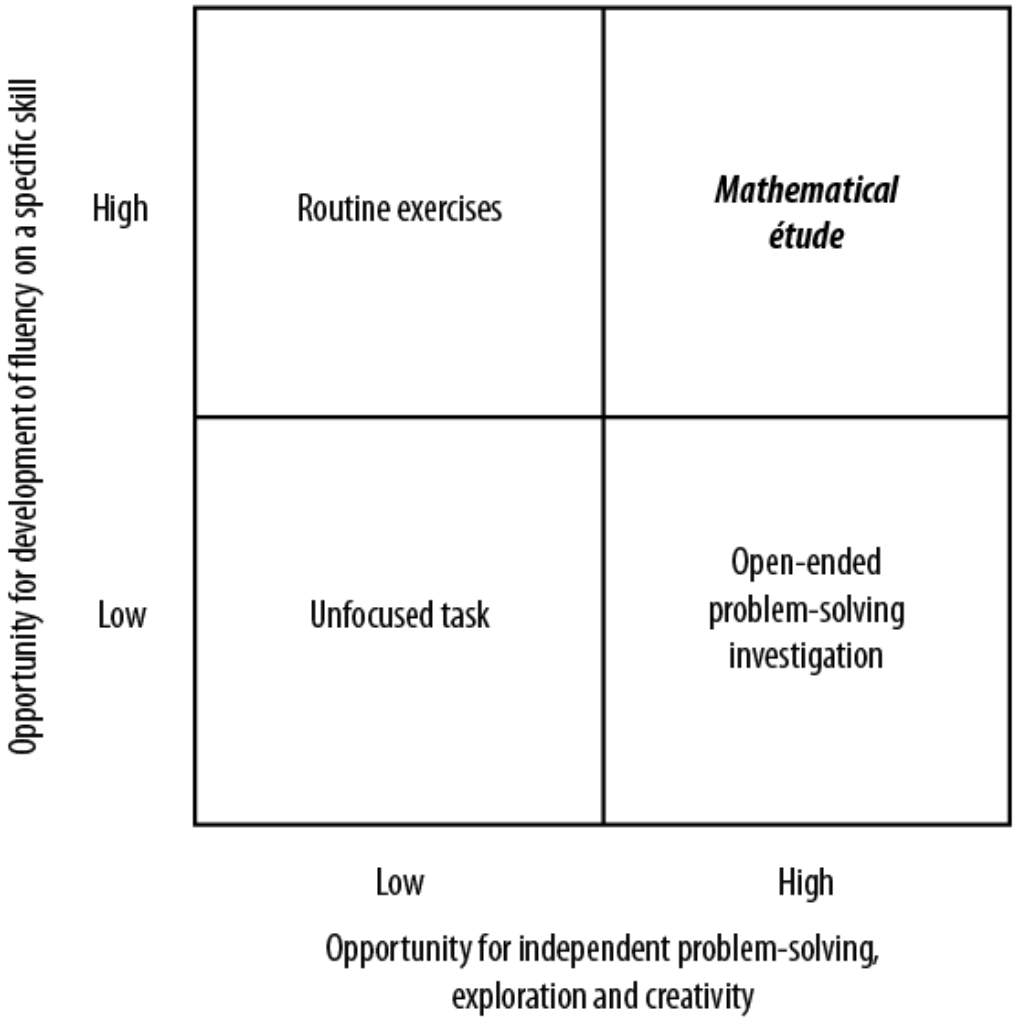


Figure 10.2 – Source: Colin Foster (2013)

However, when reviewing a topic or concept that I know students have encountered in its entirety, and that the majority of students clearly have some existing knowledge of, I need something else. The need is similar in the lessons immediately following the introduction of a concept, where some students need more straightforward practice and others are ready to push on. Here I will do a quick example-problem pair, both to show students how to set their work out and provide support for those struggling, followed by Purposeful Practice.

There are plenty of places to find Purposeful Practice activities, but no place (at the time of writing) that just contains them. For an activity to fit underneath the umbrella of Purposeful Practice, it must meet the five key principles I discussed above:

1. Students need to experience early success
2. There must be plenty of opportunities to practice the key procedure
3. The practice should feel different
4. Opportunities must exist for students to make connections, solve problems and think deeper
5. The focus is always on the practice

Don Steward's outrageously good Median Maths blog, NRICH, the Shell Centre, and Open-Middle problems are my top 4 places to visit in search of these magical tasks.

Keep the above criteria in mind when considering some of my favourite Purposeful Practice activities:

Factor Facts

Basic procedure practised: Working out the factors of a number

Factor facts

Establish, beyond doubt, that there are three pairs of consecutive numbers, less than 100, both of which have exactly 6 factors.

1	2	3	4	5	6	7	8	9	10
11	12	13	14	15	16	17	18	19	20
21	22	23	24	25	26	27	28	29	30
31	32	33	34	35	36	37	38	39	40
41	42	43	44	45	46	47	48	49	50
51	52	53	54	55	56	57	58	59	60
61	62	63	64	65	66	67	68	69	70
71	72	73	74	75	76	77	78	79	80
81	82	83	84	85	86	87	88	89	90
91	92	93	94	95	96	97	98	99	100

Find the three trios of consecutive numbers less than 100 all of which have exactly 4 factors.

Figure 10.3 – Source: Don Steward, available at http://donsteward.blogspot.ae/2012/12/factor-facts.html

Don Steward is the king of Purposeful Practice, with many of his activities fitting my stringent criteria. For example, take this beauty on finding factors – something that students have been practising from age 9, and yet still often require help with. Here that practice is induced by means of a couple of simple looking challenges. There is no way students can complete this challenge without practising the exact skill I need them to practise. If they are struggling to start, I can help them experience immediate success by suggesting that they begin by writing down the number of factors each number has, starting with 1. At different stages through this activity, students begin to spot connections, have insights, and think deeper. They realise that they can save time by crossing out prime numbers and square numbers. They may notice a relationship between the prime factors of a

number and the number of factors of that number. All the while they are practising finding factors – lots of them – and hence the main purpose of the activity is never lost.

Creating equations

Basic procedure practised: Solving linear equations

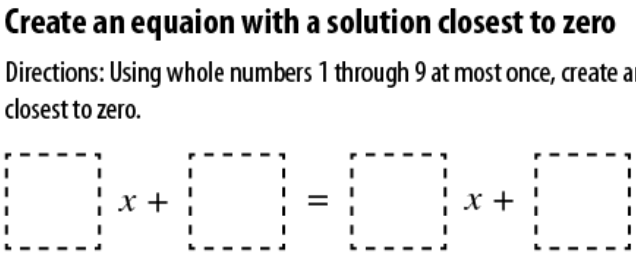


Figure 10.4 – Source: Robert Kaplinski for Open-Middle, available at http://www.openmiddle.com/solving-equations-with-variables-on-both-sides/

Here we have students practising the basics of solving linear equations, but again with a wider purpose. How close to zero can they get? How do they know how far off they are? And then we can push students further. How close to 1 can they get? What is the biggest value of x? How about the smallest? What if we could only use the numbers 1 to 5? What if the right-hand side of the equation was $\square - \square x$?

Area v Perimeter

Basic procedure practised: Working out the area and perimeter of rectilinear shapes

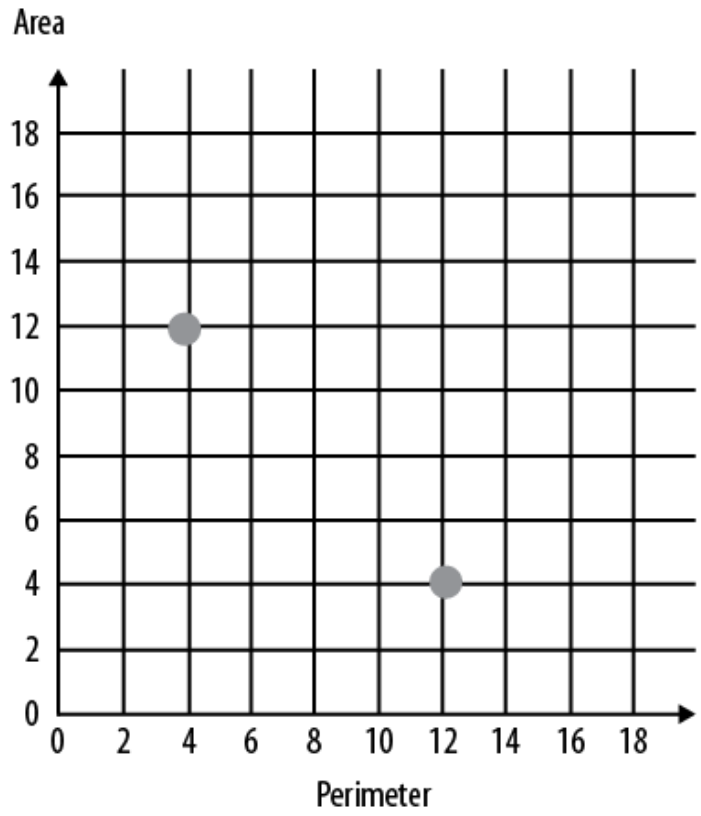


Figure 10.5 – Source: Malcolm Swan, available at http://blog.mrmeyer.com/2011/other-peoples-problems

When I first starting working on this book, I heard the sad news that Malcolm Swan had passed away. Malcolm was one of my true heroes, and through his incredible work at The Shell Centre and via the *Improving Learning in Mathematics* (also known as *The Standards Units*) materials, he has influenced generations of maths teachers and their lucky students. This is one of my favourite of Malcolm's activities. We will discuss when it is appropriate to combine related topics such as area and perimeter in Chapter 12, but for now let us revel in the practice students get at working out the area and perimeter of the various rectilinear shapes they create in their quest to fill up the grid with points. Which points on the grid represent squares, rectangles, *etc*? Can you draw a shape that may be represented by the point (4, 12) or (12, 4)? Which points are impossible? Thank you, Malcolm.

Venn Diagrams

Basic procedure practised: Working out mean, median mode and range from a list of numbers

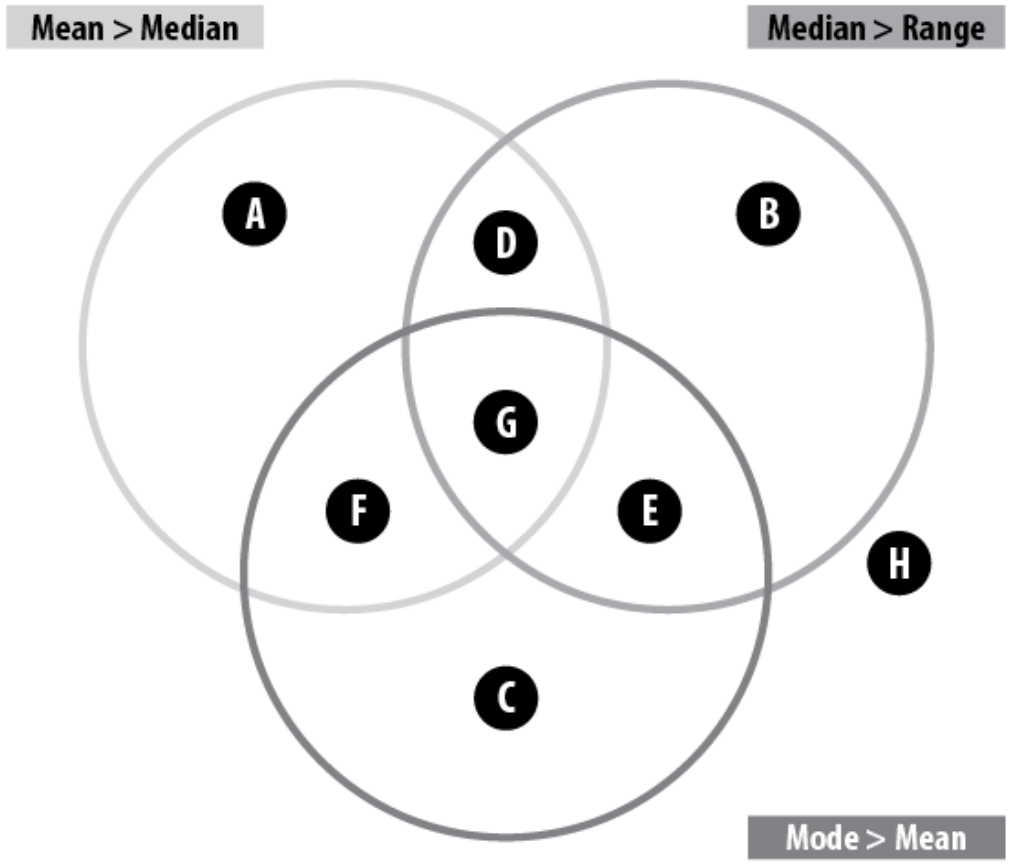


Figure 10.6 – Source: Craig Barton, available at http://mrbartonmaths.com/teachers/rich-tasks/venn-diagrams.html

Now, seeing as I created this activity, I don't really want to wax lyrical too much about its virtues. However, it is worth sharing one thing I absolutely love about activities structured around Venn Diagrams – they are easy to get started, and difficult to finish. If a student says they are stuck on the activity above, then tell them to think of any five numbers. The beauty is, these five numbers must belong inside one of the eight regions. And once they have figured out where they go, they are away! The activity then gets harder and harder as students strive to fill the remaining regions, all the time delivering valuable practice working both forwards and backwards to calculate mean, median, mode and range, as well as the effect changing a number has on each of the measures of average and range. Are some of the regions

impossible? If so, how do students know? And if students do finish, then no need to find them another activity. Simply challenge them to create their own Venn Diagram activity with a specified number of impossible regions.

OceanofPDF.com

10.3. If I only remember 3 things…

1. The hardest part of teaching is covering a topic or concept that students have had some prior experience of, and unfortunately that is the vast majority of secondary school maths teaching.
2. Purposeful Practice provides a way to ensure students are practising a key procedure, but with a wider purpose in mind.
3. It is this wider purpose that gives students the opportunity to solve problems and think creatively, and hence transition along the path from novice to expert learner.

OceanofPDF.com

11. FORMATIVE ASSESSMENT AND DIAGNOSTIC QUESTIONS

I am going to start with a rather big claim: asking and responding to diagnostic questions is the single most important thing I do every lesson.

This chapter will be my attempt to convince you why.

OceanofPDF.com

11.1. What is formative assessment and why is it important?

What I used to think

If you have been in teaching for any length of time, you will have witnessed the coming and going of many trends and fads. When I first started, it was all about Assessing Pupil Progress (APP). I spent around 93% of every lesson filling out grids about what students were supposedly learning, and 7% actually helping them learn things. Luckily, APP didn't last too long, but long enough for a frightening amount of time and money to be given to it. Then there were learning styles, and to a lesser extent growth mindset, both of which have been discussed already in the book and both of which, having being heralded as the next big thing, do not receive the prominence in schools that they once did.

For many years I lumped formative assessment (or *Assessment for Learning*, as it is otherwise known) in this same group. It was a phrase that was bandied around a lot – something all teachers were told we had to do – but without any real substance or conviction. It was also marketed as a generic teaching strategy – one that could be used across all subjects – and so it was accompanied by the usual whole-school training sessions, where us maths teachers were presented with examples from English, history and geography and persuaded that it definitely will work for equations, percentages and histograms.

So for half my career I steered clear of any mention of formative assessment. Then I came across the work of Dylan Wiliam. And it is a good job I did, because I am now convinced that teaching without formative assessment is like painting with your eyes closed.

Sources of inspiration

- Barton, C. (2017a) 'Daisy Christodoulou', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Kris Boulton', *Mr Barton Maths Podcast*.
- Black, P. and Wiliam, D. (1998) 'Inside the black box: Raising standards through classroom assessment', *Phi Delta Kappan* 80 (2) pp. 139-148.
- Black, P. and Wiliam, D. (2009) 'Developing the theory of formative assessment', *Educational Assessment, Evaluation and Accountability* 21 (1) pp. 5-31.
- Christodoulou, D. (2017) *Making good progress?*. Oxford: Oxford University Press.
- Cowie, B. and Bell, B. (1999) 'A model of formative assessment in science education', *Assessment in Education: Principles, Policy & Practice* 6 (1) pp. 101-116.
- Didau, D. (2014) 'Why AfL might be wrong, and what to do about it', *The Learning Spy Blog*. Available at: http://www.learningspy.co.uk/myths/afl-might-wrong/
- Kruger, J. and Dunning, D. (1999) 'Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments', *Journal of Personality and Social Psychology* 77 (6) pp. 1121-1134.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Rosenshine, B. (2012) 'Principles of instruction', *American Educator* 36 (1) pp. 12-19.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.
- Wylie, E. C. and Wiliam, D. (2006) *Diagnostic questions: is there value in just one?*. San Francisco, CA: Annual Meeting of the National

Council on Measurement in Education.

My takeaway

What is formative assessment?
In 2016, Dylan Wiliam sent the following tweet: *Example of a really big mistake: calling formative assessment 'formative assessment', rather than something like '**responsive teaching**'.*

Indeed, *responsive teaching* feels like a much better description to attach to the tools and strategies we will discuss in this chapter. The word 'assessment' conjures up visions of tests, marking and grades. For teachers it means more work, and for students more pressure. In Chapter 12 we will look at the importance of seeing tests as tools of learning, but the association with assessment has probably not helped the development and adoption of this most valuable of strategies.

Black and Wiliam (2009) explain that an assessment functions formatively 'to the extent that evidence about student achievement is elicited, interpreted, and used by teachers, learners, or their peers, to make decisions about the next steps in instruction that are likely to be better, or better founded, than the decisions they would have taken in the absence of the evidence that was elicited'.

Cowie and Bell (1999) define formative assessment as 'the process used by teachers and students to recognise and respond to student learning in order to enhance that learning, during the learning'.

Wiliam (2011) makes the point that any assessment can be formative, and that assessment functions formatively when it improves the instructional decisions that are made by teachers, learners, or their peers.

For me, formative assessment is all about responding in the moment. It is about gathering as much accurate information about students' understanding as possible in the most efficient way possible, and making decisions based on that. In short, it is about adapting our teaching to meet the needs of our students.

Why do we need it?
In his section on Checking for Understanding, Lemov (2015) argues that the most crucial task of teaching is to distinguish 'I taught it' from 'they learned it'. One of Rosenshine's (2012) 'Principles of Instruction' that we first met when discussing great teaching in Chapter 3 is 'check for student understanding'. Rosenshine explains: 'The more effective teachers frequently checked to see if all the students were learning the new material. These checks provided some of the processing needed to move new learning into long-term memory. These checks also let teachers know if students were developing misconceptions'.

So, how can we check for understanding?

Well, in a world without formative assessment, we are left with two options:

1. Rely on student self-report.
2. Rely on data from tests or homework.

Both of these have problems.

Lemov's (2015) very first principle is 'reject self-report'. He argues that asking students to judge if they understand something is flawed for a number of reasons. Take something like, 'Right, so that is how you work out the gradient of a line. Thumbs up if you understand'. First, it is the kind of question that demands a passive yes. My students will be able to tell from my body language and tone of voice that I am hoping for a yes so I can move on, and they are unlikely to want to be the one who slows me down by indicating their lack of understanding. Second, there is likely to be significant (explicit or implicit) peer pressure acting upon students to say they get it – 'If all my friends say they understand, then I better had as well, or I am going to look pretty thick'. Third, and possibly most importantly, there is this issue of students not actually knowing if they understand or not. The Dunning-Kruger effect (Kruger and Dunning, 1999) explains that people who are unskilled in a particular domain suffer a dual burden: they reach erroneous conclusions; and their incompetence robs them of the metacognitive ability to realise it. Hence, our novice learners may not

actually know they do not understand how to work out the gradient of a line – they may have a vague notion of what is going on and conclude that is enough – and so happily (and honestly) put their thumbs high into the air. We need more reliable evidence of understanding than can be gleaned from student self-report.

Sure, we may obtain that more reliable evidence about student understanding via a summative assessment somewhere down the line – possibly via a test or a homework – but at this stage it is likely to be far more difficult to undo the damage. Misconceptions may have been rehearsed, compounded and committed to long-term memory. Practice makes permanent after all. I have spent many a Sunday afternoon correcting the same mistakes over and over again in a class-set of books, writing the same detailed feedback to help the students understand. Wiliam (2011) explains that in a sense, these hours spent marking can be seen as the punishment given to teachers for failing to find out that they did not achieve the intended learning when the students were in front of them.

Teaching is only successful if students have understood and learned something. We can spend hours planning the very best explanations, examples, practice questions and problems for students to solve. But, as we all know, very rarely do things go exactly to plan. This may be because we are burdened with the curse of knowledge, or more likely it is due to the unpredictable nature of the 30 adolescents in front of us, with their 30 different minds and moods. Without evidence as to how students are responding to our teaching, the lesson could be going wildly out of control whilst we continue to revel in the bliss of ignorance. Successful formative assessment can help us identify problems and begin to fix things in the here and now much more effectively and efficiently.

Indeed, if we are to adopt the teacher-led strategies of explicit instruction described so far in this book, we are almost obliged to make full use of effective formative assessment strategies. If not, then explicit instruction can indeed become the lecturing to a group of passive, unengaged students that critics of the approach claim it to be.

All of this led Black and Wiliam (1998) to argue that 'formative assessment is an essential feature of classroom work and that development of it can raise standards. We know of no other way of raising standards for which such a strong *prima facie* case can be made on the basis of evidence of such large learning gains'.

But is formative assessment fundamentally flawed?
Before we go any further, we need to address a rather large elephant in the room. Some prominent people in the world of education believe formative assessment is fundamentally flawed, and with good reason. David Didau articulates this viewpoint very clearly:

> [Formative assessment] is predicated on the assumption that you can assess what pupils have learned in an individual lesson, and then adjust future teaching based on this assumption. But you can't. There's no meaningful way to assess what pupils have learned during the lesson in which they are supposed to be learning it. There's an impressive body of research that tells us that learning is distinct from performance. You cannot see learning; you can only see performance. If we measure performance then we may be able to infer what might have been learned, but such inferences are highly problematic for two reasons: 1) performance in the classroom is highly dependent on the cues and stimuli provided by the teacher and 2) performance is a very poor indicator of how well pupils might retain or be able to transfer knowledge or skills.

We will discuss the distinction between learning and performance and the significant implications for teaching in the following chapter, but it is enough to mention here that this is indeed a concern.

I have thought long and hard about this, I discussed it at length with both Daisy Christodoulou and Kris Boulton on my podcast, I changed my mind about 472 times, but I think I am finally happy with my viewpoint. It goes something like this.

Indeed, you may not be able to infer learning from the correct response to a single question (or even multiple questions) immediately following

instruction. In order to infer learning, you will need to ask the question – or variations of it – again at multiple points in the future. However, in that moment, a correct answer – especially if it is to a well-designed question and complemented by an explanation – is sufficient evidence to move on. For if we don't, how will we ever get anywhere? Furthermore, whilst you may not be able to infer learning from a correct answer, you can certainly infer a lack of learning and understanding from an incorrect answer. As Wiliam (2011) explains: 'The fact that learners know something now does not guarantee that they will know it in six weeks' time, but the fact that they don't know it now probably does mean that they will not know it in six weeks' time'.

Hence, I'm firmly of the belief that if used carefully, formative assessment is not flawed, but is possibly the most important and powerful tool a teacher can possess.

What I do now

Dylan Wiliam's book *Embedded Formative Assessment* lists around 50 practical, effective formative assessment strategies that both teachers and students can use, and there are more to be found in Daisy Christodoulou's *Making good Progress?*. Any attempt to replicate here would be but a pale imitation. So, instead, I am going to focus on one strategy of formative assessment that I use every day, every lesson with every class I teach. But before we get there, we need to take a brief diversion into the realms of classroom culture.

OceanofPDF.com

11.2. Classroom Culture

What I used to think

I never really gave that much attention to the culture of my classroom. As long as my students were working hard, not misbehaving too much and ideally laughing at my jokes, I was more than happy. Surely I did not need anything else to ensure that my teaching was effective?

Sources of inspiration

- Barton, C. (2016) 'Dylan Wiliam', *Mr Barton Maths Podcast*.
- Barton, C. (2017) 'Doug Lemov', *Mr Barton Maths Podcast*.
- Black, P. and Wiliam, D. (1998) 'Inside the black box: raising standards through classroom assessment', *Phi Delta Kappan* 80 (2) pp. 139-148.
- Coe, R. (2013) *Improving education: a triumph of hope over experience, CEM Inaugural Lecture*. Available at: http://www.cem.org/attachments/publications/ImprovingEducation2013.pdf
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Rosenshine, B. (2012) 'Principles of instruction: research-based strategies that all teachers should know', *American Educator* 36 (1) pp. 12-39.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

- Wiliam, D. (2016) 'The 9 things every teacher should know', *TES*. Available at: https://www.tes.com/us/news/breaking-views/9-things-every-teacher-should-know

My takeaway

It sounds ridiculously obvious to say, but for any assessment strategy – formative or summative – to work, students must actively and honestly participate. There are three factors that could prevent this, all of which we can change.

1. A fear of mistakes

If students are afraid of making mistakes, how can we learn from their misunderstandings?

We have probably all taught students who leave questions out in tests and homeworks for fear of being wrong, and we all know that such actions make it incredibly difficult to help them, as we have no indication of how much or in what areas their understanding is lacking. However, in my experience, far more common is a fear of making mistakes away from the written page. Many formative assessment strategies – and indeed the one I am going to focus on in this chapter – require students to be public about their answers, displaying their thoughts in front of me and their peers in the moment. If students fear making mistakes, and the consequences of those mistakes, then it is highly likely that they will fail to provide us with any useful information at all. After all, for the child who fears failure, not giving a response is far less daunting than having a go.

Lemov (2015) encapsulates this in the concept of creating a Culture of Error. He explains:

> This shift from defensiveness or denial to openness is critical. If your goal is to find and address the mistakes your students make, your task is far more difficult if your students seek to hide their errors from you. If, in contrast, they willingly share their struggles, mistakes, and errors, you can spend less time and energy hunting for them and more time fixing and learning from them.

So how do we create this culture of error, and hence overcome this potentially fatal problem?

I have discussed the crucial role the teacher plays in this process back in Chapter 2. Specifically, I discussed the need to ensure that students respect each other's answers, whether they are right or wrong, and that any deviation away from these high expectations is tackled head-on.

But such teacher actions can only take us so far. More than this, the questions we ask students need to be seen not as tools of assessment, but as tools of learning. This will be equally important when we look at the power of testing in Chapter 12. We can only hope to achieve this if there are no negative consequences for being wrong. We can do this by not grading or recording students' responses to the formative assessment questions we ask in class, for the presence of a grade or record puts a premium upon success, and they are not needed to inform our decisions in the moment.

There also needs to be positive consequences for honest participation. As I explained in Chapter 2, mistakes need to be embraced as learning opportunities. I know that sounds ridiculously clichéd, but it is true. When I interviewed Doug Lemov for my podcast, he shared some excellent advice for doing exactly this. If a student makes a mistake, we need to thank them. We need to explain that such a mistake is likely to be shared by lots of other students, and hence by making it, sharing it, and giving me a chance to fix it, they have genuinely improved the understanding of many of their peers.

The beauty of the strategy I am going to describe in this chapter is that built into to every multiple-choice diagnostic question are three mistakes. Even if students do not choose one of these options, I will often discuss them. This serves a dual purpose – it allows us to dig deeper into the skill or concept (more on this later) and it also makes mistakes and the discussion of them a regular and important part of the lesson.

2. Students opting out

The second factor that can render any assessment strategy – but in particular classroom-based formative assessment – limp and ineffective is the classic opt-out. Some students may choose not to give an answer, not for fear of

being wrong but, to put it bluntly, because they cannot be arsed to think. A shrug, an utterance of 'I don't know', or a wall of silence tells us absolutely nothing about a student's understanding of a given concept, and thus leaves us powerless to help.

Allowing such a response also conveys the message that non-participation is absolutely fine. Wiliam (2011) argues that engaging in classroom discussion really does make students smarter. So, when teachers allow students to choose whether to participate or not – for example, by allowing them to raise their hands to show they have an answer, or settling for a lack of response – we are actually making the achievement gap worse, because those who are participating are getting smarter, while those avoiding engagement are forgoing the opportunities to increase their ability.

Doug Lemov to the rescue once again. During our podcast conversation I asked Doug how to deal with a response of 'I don't know'. Doug explained that he would ask another student to help get them started and then return to the original student to see if they could carry on. If they still could not (or would not), then we can get another student to complete the answer, before returning to the original student, asking them to repeat the answer, and then giving them a follow-up question. The key point is that students need to see that a response of 'I don't know' is going to lead to just as much work, so they might as well actively and honestly participate from the start.

3. Finding comfort in one correct answer

Directly related to students themselves opting out is a common practice amongst teachers (myself very much included) that essentially does the student's job of opting out for them. See if this scenario rings any bells:

> Me: So, does anyone know what -5 – -2 is?
>
> *(3 hands go up, one of which is Josie. Josie always gets everything right)*
>
> Me: Josie, go for it
>
> Josie: -3, sir
>
> Me: And why is that, Josie?

> Josie: Because subtracting a minus is the same as adding a positive, and negative 5 plus 2 gives you negative 3
>
> Me: Loving your work as ever, Josie. Okay, let's move on.

Writing this makes me feel ashamed, as that is exactly how many of my early attempts to assess the understanding of my students proceeded. In Black *et al* (2004), a teacher is quoted as describing such a scenario as 'a small discussion group surrounded by many sleepy onlookers'. Likewise, when I interviewed Dylan Wiliam for my podcast and asked him to describe an approach in the classroom that he doesn't think is effective, he replied 'teachers making decisions about the learning needs of 30 students based on the responses of confident volunteers'. Rarely have truer words been spoken. I find solace in the fact that I am not alone. Dylan Wiliam (2016) himself describes a similar experience:

> When I was teaching full-time, the question that I put to myself most often was: 'Do I need to go over this point one more time or can I move on the next thing?' I made the decision the same way that most teachers do. I came up with a question there and then, and asked the class. Typically, about six students raised their hands, and I would select one of them to respond. If they gave a correct response, I would say 'good' and move on.

One of Coe's (2013) 'poor proxies for learning' is '(at least some) students have supplied correct answers', and it is easy to see why. I am seeking comfort in one correct answer. When Josie once again produces a perfect answer and a lovely explanation, I make two implicit assumptions: first, that this is down to my wonderful teaching; and second, that every other child in the class has understood the concept to a similar level. But, of course, I have no way of knowing that. By essentially opting out the rest of the class, the only information I am left with concerns Josie.

There are ways around this. We can use lollipop sticks, or random name generators to ensure each student has an equal chance of being selected. We can use Lemov's (2015) concept of 'cold-call', which, as we discussed during our podcast interview, Doug feels ascribes more meaning to each student's

answer as we have chosen them for a reason. We cannot tell Josie if she was right or wrong, but instead ask the rest of the class to indicate by raising their hand if they thought Josie was correct, or challenge another student to explain Josie's reasoning.

All these adaptations certainly improve my initial process, but they all suffer from the same fatal flaw. All students are not required to participate to the same degree, and so the only student's understanding I have anything resembling reliable evidence about is the student answering the question. Rosenshine's (2012) third principle of instruction is: 'Ask a large number of questions and check the responses of all students'. I am failing to do that.Lemov (2015) has a series of useful techniques under the umbrella of 'No Opt-Out' that can help counteract students opting out themselves as well as via the actions of the teacher. However, the strategy involving diagnostic questions that I am going to outline later in this chapter has the full participation of each and every student built in to its very core.

What I do now

I am always conscious of the need to ensure all my students actively and honestly participate when I ask questions. In Chapters 2 and 9 I argued that too much struggle and failure can be detrimental to both motivation and learning. However, that is not to say that mistakes should not be embraced. Mistakes – specifically how students cope with them and learn from them – are a key part of mathematical development. Likewise, without honest and active participation from students – *all* students – we are left to blindly guess at what they do and do not understand, and hence the help and support we offer may prove to be ineffective.

Fortunately, the use of diagnostic questions brings mistakes (and misconceptions) into the open, treating them as the learning opportunity that they are, whilst at the same time encouraging full class participation. Hence, I use them every single lesson, with every single class, every single day.

OceanofPDF.com

11.3. What is a Diagnostic Question?

What I used to think

I used to think two things that fundamentally dictated how I asked students questions and offered them support:

1. For any given question there were two groups of students: those that could do it and those that could not. Those that could do it were fine to get on with the next challenge, and those that could not needed help. Crucially, they needed the same help.

2. Closed questions are bad and open questions are good. Closed questions encourage a short response, whereas open-ended questions demand much greater depth of thought. Hence, I spent many years fighting the urge to ask students closed questions in class, and instead opted almost exclusively for things like 'Why do we need to ensure the denominators are the same when adding two fractions?' or 'How would you convince someone that $\frac{3}{7}$ is bigger than $\frac{4}{11}$?'.

Surely both of these beliefs were spot on?

Sources of inspiration

- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

- Wylie, E. C. and Wiliam, D. (2006) *Diagnostic questions: is there value in just one?*. San Francisco, CA: Annual Meeting of the National Council on Measurement in Education.

My takeaway

We will return to my first belief in a minute, but first let's deal with the issue of openness in questions.

The two questions above are great questions. They are important questions to ask students. But if our aim is to quickly and accurately assess whole-class understanding so we are able to make an informed decision on how to proceed with the lesson, they are not so good.

Their strength is their weakness. The fact that they encourage students to think, take time to articulate, and provoke discussion and disagreement makes them entirely unsuitable for effective formative assessment as laid out at the start of this chapter. How would we go about collecting and assessing the responses to 'Why do we need to ensure the denominators are the same when adding two fractions?' from 30 Year 8 students in the middle of a lesson as a means of deciding whether or not the class is ready to move on?

Open-ended questions like these are great for homework, tests, extension activities, and lots of other different situations. However, they are not great for a model of responsive teaching.

Nor is it the case that closed questions prevent thinking. Wiliam (2011) gives the example of asking if a triangle can have two right-angles. This is about as closed a question as you can get – the answer is either yes or no. But the thinking involved to get to one of those answers is potentially very deep indeed. Students may consider whether it is possible to have an angle measuring 0°, or if parallel lines will meet at infinity. But this closed question – whilst it is indeed a brilliant one – is equally unsuited for a model of responsive teaching. If a particular student answered 'no', would we be convinced that they understood the properties of triangles and angles fully? Or have they just guessed? Without further probing, it is impossible to tell, and hence we are back to the same issues with the more open fractions questions above.

So, if open-ended questions are unsuitable for this style of formative assessment, and not all closed questions are suitable, then what questions are left?

Step forward diagnostic multiple-choice questions – or just 'diagnostic questions' as I will refer to them for the rest of this chapter.

Diagnostic questions are designed to help identify, and crucially *understand* students' mistakes and misconceptions in an efficient and accurate manner. Back in Chapter 3, in the section on Cognitive Conflict, we discussed the distinction between the two, but it is worth repeating briefly here. Mistakes tend to be one-off events – the pupil understands the concept or the algorithm, but may make a computational error due to carelessness or, as we saw in Chapter 4, cognitive overload. Give the student the same question again and they are unlikely to make the same mistake; inform the student that they have made a mistake somewhere in their work, and they are likely to be able to find it. Misconceptions, on the other hand, are the result of erroneous beliefs or incomplete knowledge. The same misconception is likely to occur time and time again. Informing the student who has made an error due to a misconception is likely to be a waste of time as, by definition, they do not know they are wrong. Good diagnostic questions can help you identify and understand both mistakes and misconceptions.

The best way to explain a diagnostic question is to show you one.

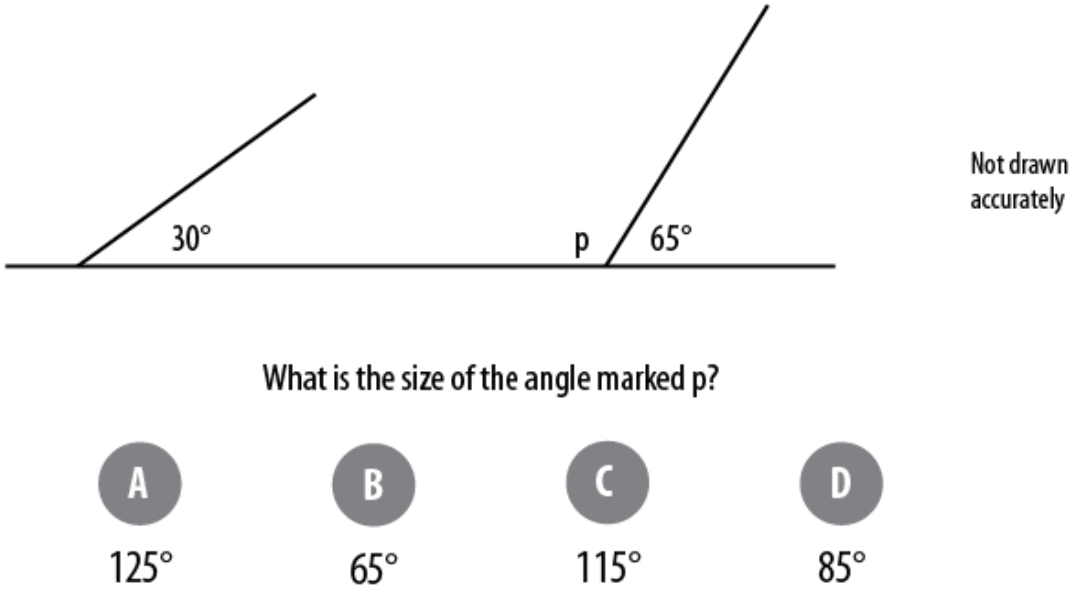


Figure 11.1 – Source: Craig Barton for Diagnostic Questions

Take a moment to look at the question, and in particular the four different answers. What would each of these answers tell you about the understanding of a student who gave them?

Answer A may suggest that the student understands that angles on a straight line must add up to 180°, and is able to identify the relevant angle, but has made a common arithmetic error when subtracting 65 from 180.

Answer B may be the result of the student muddling up their angle facts, mistakenly thinking this is an example of vertically opposite angles being equal.

Answer C is the correct answer.

Answer D may imply that the student is aware of the concept that angles on a straight line must add up to 180°, but has included all visible angles in their calculations.

Notice how each of these answers reveals a specific, and different mistake or misconception. Imagine you had a group of students who answered A, another group who answered B, and a final group who answered D (we will cover exactly how we go about collecting this information later in the chapter). Would all three groups require the same intervention from you, their teacher?

I don't think so. Which brings us to my second (erroneous) belief. It is not always the case that students either can or cannot answer a question correctly. Sure, there may be a group of students who get the question correct for the same or similar reasons. But there are likely to be students who get a question wrong for very different reasons, and it is the reason they get the question wrong that determines the specific type of intervention and support they require.

For example, students who answered B and D may both benefit from an interactive demonstration (for example, using Geogebra) to illustrate the relationship between angles on a straight line. Students who chose B could then be presented with an exercise where they are challenged to match up an assortment of diagrams with the angle fact they represent. Those who

selected D may benefit more from a selection of examples and non-examples of angles on a straight line, similar to those showcased in Section 7.2. But what about students who answered A? Their problem lies not with the relationship between the angles, but with their mental or written arithmetic. This may be a careless mistake, or it may be an indication of a more serious misconception with their technique for subtraction. Either way, it is not a problem that is likely to be solved by giving these students the same kind of intervention as everyone else.

However you choose to deal with these students, there is little doubt that there is an advantage to knowing not just which students are wrong, but *why* they are wrong. And I have never come across a more efficient and accurate way of ascertaining this than by asking a diagnostic question.

So, what makes a question a diagnostic question? For the way I define and use them, there needs to be one correct answer, three incorrect answers, and each incorrect answer must reveal a specific mistake or misconception, I can – and indeed do – ask students for the reasons for their answers (as I will explain soon), but I should not need to. If the question is designed well enough, then I should gain reliable evidence about my students' understanding without having to have further discussion.

The set of criteria that I impose in order to deem a question a good diagnostic question will be discussed next. For now it may be useful to see a few more of my favourite diagnostic questions, just to get into the swing of things. Each time, ask yourself what you would learn about your students from their choice of answers without them needing to utter another word.

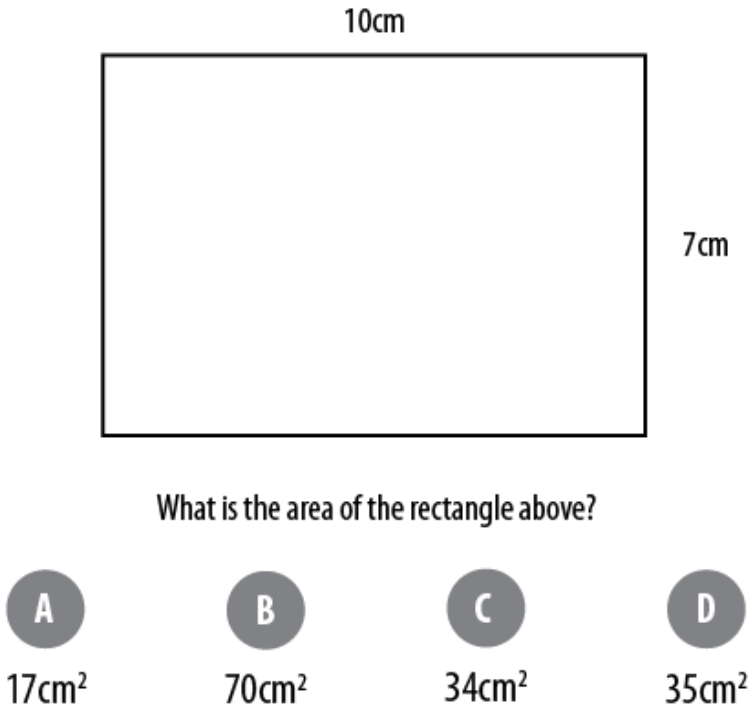


Figure 11.2 – Source: Craig Barton for Diagnostic Questions

$$\frac{x+2}{5} = x - 3$$

Which of the following is a correct next step in solving this equation?

A $x + 2 = 5x - 3$

B $x + 2 = 5x - 15$

C $5x + 10 = 5x - 15$

D $\frac{x}{5} = x - 5$

Figure 11.3 – Source: Craig Barton for Diagnostic Questions

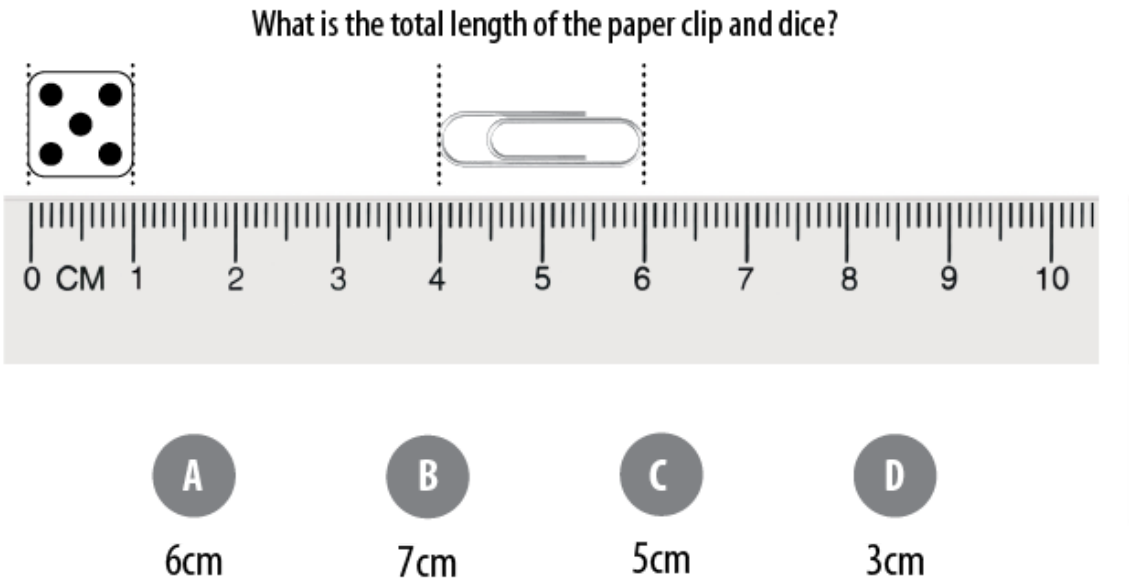
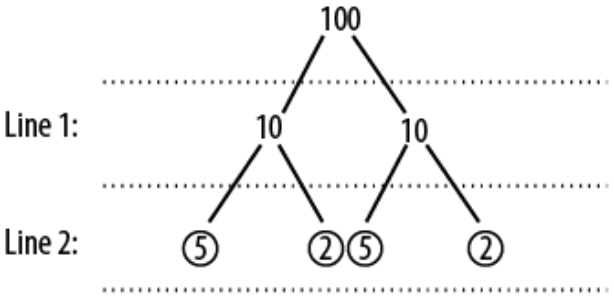


Figure 11.4 – Source: Craig Barton for Diagnostic Questions



Line 3: Answer: 2, 2, 5, 5

A student is asked to express 100 as a product of its prime factors.
On which line of working does the first error occur?

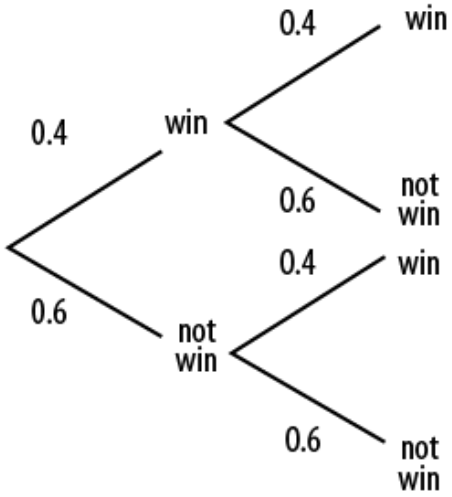
A Line 1

B Line 2

C Line 3

D No errors made

Figure 11.5 – Source: Craig Barton for Diagnostic Questions



What is the probability of winning two games in a row?

A 1.6 B 0.8 C 0.25 D 0.16

Figure 11.6 – Source: Craig Barton for Diagnostic Questions

£	Frequency	Cumulative Frequency
$0 < £ \leq 10$	6	
$10 < £ \leq 20$	10	★
$20 < £ \leq 30$	12	
$30 < £ \leq 40$	8	

What should replace the star in this table?

A 16 B 15 C 150 D 60

Figure 11.7 – Source: Craig Barton for Diagnostic Questions

Before we move on, it is worth noting that Wiliam (2011) gives examples of diagnostic questions with more choices and more than one correct answer. He explains how this can reduce the possibility of students getting the question correct through guesswork. This is certainly true, but I prefer to stick to four options, only one of which is correct. This is to make data collection as efficient as possible in the classroom, to make data analysis on my Diagnostic Questions website as easy as possible, and I find the consistency of approach helps when working with both teachers and students to create questions (the importance of which will be discussed later). To reduce the possibility of students getting the question correct by guessing, I often choose to ask a second question. The process for asking and receiving data from these questions in the classroom is so fast that there is plenty of time for a second question, which immediately reduces the odds of being correct through guesswork from $\frac{1}{4}$ to $\frac{1}{16}$.

What I do now

Every day I make use of a very specific type of closed question – a multiple-choice diagnostic question. I believe a good diagnostic question can give you a reliable insight into students' understanding of a concept far more accurately and efficiently than any other technique I have ever come across, as well as causing your students to think hard.

OceanofPDF.com

11.4. What makes a good question

What I used to think

When I first started creating diagnostic questions, I thought any old multiple-choice question would do. How wrong I was.

Sources of inspiration

- Kruger, J. and Dunning, D. (1999) 'Unskilled and unaware of it: how difficulties in recognizing one's own incompetence lead to inflated self-assessments', *Journal of Personality and Social Psychology* 77 (6) pp. 1121-1134.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.
- Wylie, E. C. and Wiliam, D. (2006) *Diagnostic questions: is there value in just one?*. San Francisco, CA: Annual Meeting of the National Council on Measurement in Education.

My takeaway

Not all diagnostic questions are born equal, and writing a good one is flipping hard. Indeed, the more I use diagnostic questions with my students and colleagues, the more I read about misconceptions in mathematics, and the more experience I get in writing them, the harder I am finding it! I take some solace from the fact that this could very well be the Dunning-Kruger effect (Kruger and Dunning, 1999) playing out, in that as I grow more knowledgeable I am also more aware of the difficulty of the challenge as well as my own considerable deficiencies.

At the time of writing I have written around 3000 diagnostic multiple-choice questions for mathematics. The vast majority of these I have used with my students either in the classroom context that I will discuss shortly, or as part of an online quiz on my Diagnostic Questions platform, and many have been tweaked, adjusted and binned over the years. Throughout that time, and inspired by the work of Wylie and Wiliam (2006) and Wiliam (2011), I have devised a series of golden rules for what makes a good diagnostic question:

Golden Rule 1: They should be clear and unambiguous
We will all have seen badly worded questions in exams and textbooks, but with diagnostic questions sometimes the ambiguity can be in the answers themselves. Consider the following question:

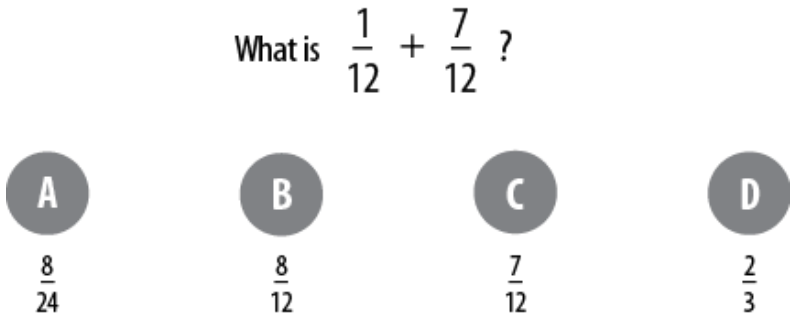


Figure 11.8 – Source: Craig Barton for Diagnostic Questions

At first glance, nothing may appear all that wrong. The wording of the question is clear, and the incorrect answers reveal specific misconceptions. But what is the correct answer? D is clearly correct, and is probably the author's intended correct answer. But how about B? Given that the question does not ask the student to simplify their answer, B is a perfectly legitimate correct answer. So, what do we infer if a student answers B? Is it that they cannot simply fractions, or they did not see D? Do they believe B is the only correct answer, or just one correct answer? The key point is that without asking them, we do not know for sure. And a key feature of a good diagnostic question is that we should be able to accurately infer a student's understanding from their answer alone without needing further student

explanation. In its current form, this question may be a good discussion question, but not a good diagnostic question.

Golden Rule 2: They should test a single skill/concept
Good questions test multiple skills and concepts. Indeed, a really effective way to interleave (see Chapter 12) is to combine multiple skills and concepts together within a single question. But good diagnostic questions should not do this. The purpose of a diagnostic question is to home in on the precise area that a student is struggling with and provide information about the precise nature of that struggle. If there are too many skills or concepts involved, then the accuracy of the diagnosis invariably suffers.

Consider the following question:

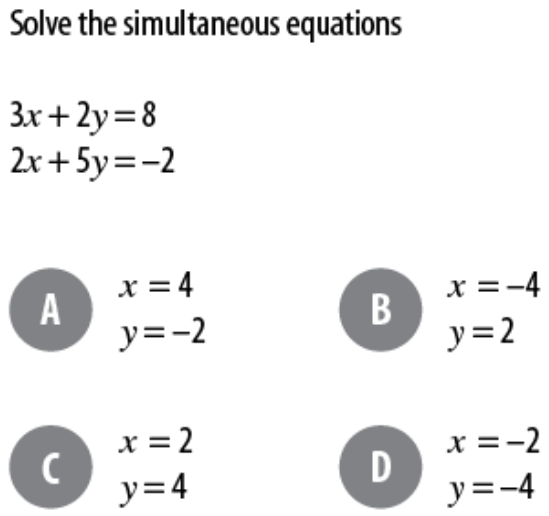


Figure 11.9 – Source: Craig Barton for Diagnostic Questions

What do we learn about a student's understanding of simultaneous equations from each of the wrong answers? As we saw in Chapter 8 on Deliberate Practice, there are many, many skills and concepts involved in answering simultaneous equations questions. Simply presenting students with a question at the start of the process and a series of answers for the end tells you very little about where in the process they are going wrong and why that is happening. Far better would be to break the question down into individual steps, following the principles of Deliberate Practice. So, the first two stages may look like this:

Solve the simultaneous equations

Step 1

$3x + 2y = 8$
$2x + 5y = -2$

Which of the following is a correct first step to eliminte y?

A $15x + 10y = 16$, $4x + 10y = -4$

B $15x + 10y = 40$, $4x + 10y = -4$

C $15x + 10y = 8$, $6x + 10y = -2$

D $6x + 4y = 8$, $6x + 15y = -2$

Figure 11.10 – Source: Craig Barton for Diagnostic Questions

Step 2

$15x + 10y = 40$
$4x + 10y = -4$

What is the correct next step of the method?

A $19x = 36$

B $11x = 36$

C $20y = 44$

D $11x = 44$

Figure 11.11 – Source: Craig Barton for Diagnostic Questions

Alternatively, it is possible to assess all of this in one cleverly designed question, such as the one on the next page.

Emma is attempting to solve this pair of simultaneous equations.

$3x + 2y = 9$ (i)
$4x - y = 1$ (ii)

Her working is shown in the four steps below, but her final answer is incorrect.

in which of the following steps **A**, **B**, **C** or **D** does her first error occur?

A. Multiply (ii) by 2: $8x - 2y = 2$ (iii)
B. Add (iii) and (i): $11x = 11$ (iv)
C. Divide both sides of (iv) by 11: $x = 1$
D. Substitute this value of x into (ii): $4 - y = 1$ gives $y = 5$

Figure 11.12 – Source: OCR for Diagnostic Questions

Golden Rule 3: Students should be able to answer them in less than ten seconds
This is directly related to Golden Rule 2. If students are spending more than ten seconds thinking about the answer to a question, the chances are that more than one skill or concept is involved, which makes it hard to determine the precise nature of any misconception they may hold.

Consider the following question from UKMT:

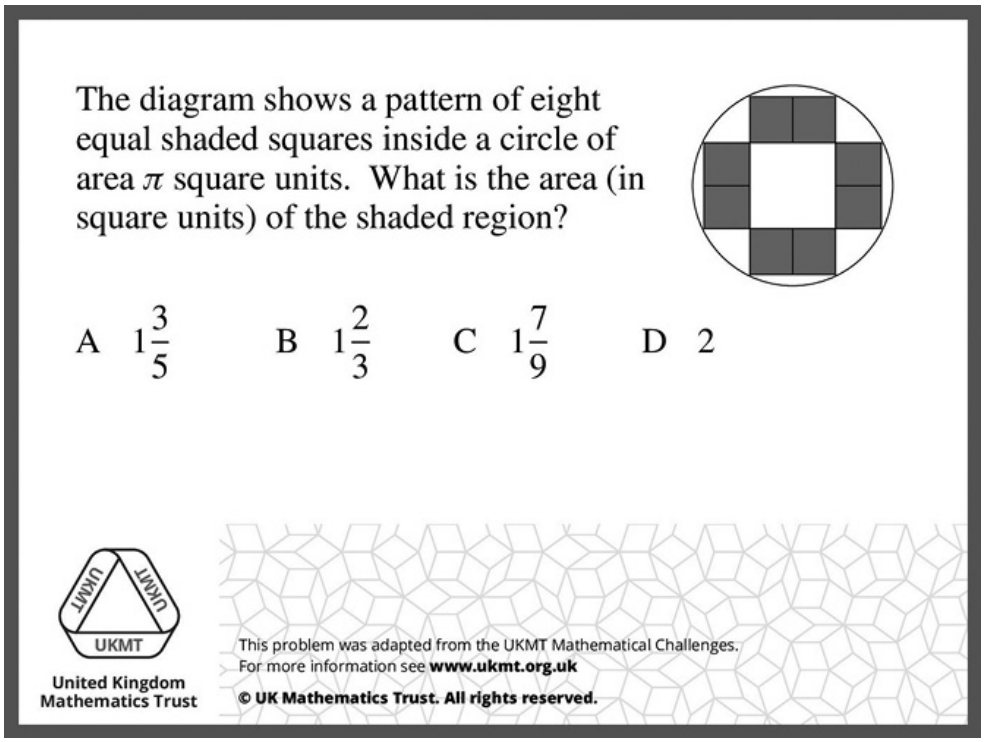


Figure 11.13 – Source: UKMT for Diagnostic Questions

Now, this is a brilliant question, but not a good diagnostic question. It is likely to take students a good few minutes to figure it out (or at least I hope it would – it took me ages!). During that thought process, lots of things are going on inside students' minds. The main purpose of a diagnostic question is to identify specific misconceptions, and that is difficult to do when there are so many steps and cognitive leaps to make in order to arrive at the final answer.

Golden Rule 4: You should learn something from each incorrect response without the student needing to explain

A key feature that distinguishes diagnostic multiple-choice questions from non-diagnostic multiple-choice questions is that the incorrect answers have been chosen very, very carefully in order to reveal specific misconceptions. In fact, they are often described as distractors – although I do not like this term as it implies they are trick questions, something I will be discussing later in this chapter. The key point is that if a student chooses one of these answers, it should tell you something.

Consider the following question:

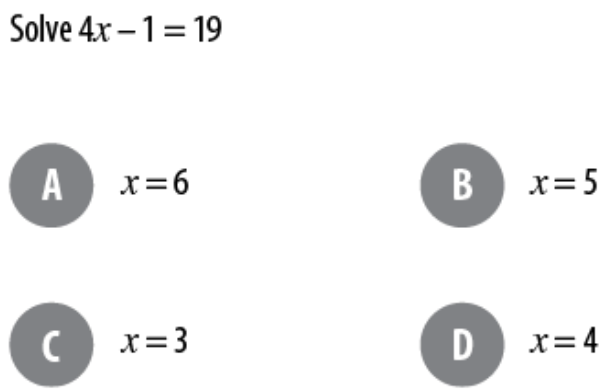


Figure 11.14 – Source: Craig Barton for Diagnostic Questions

B is the correct answer, but what do A, C and D tell you about the student's thinking? Not a lot, really. Far better would be to have something like 4.5 (student has subtracted 1), 80 (student has multiplied by 4) and 5.75 (student has divided by 4 first and then added 1).

Golden Rule 5: It is not possible to answer the question correctly whilst still holding a key misconception
This is the big one. For me, it is the hardest skill to get right when writing and choosing questions, but also the most important. We need to be sure that the information and evidence we are receiving from our students is as accurate as possible, and in some instances that is simply not the case.

Consider the following question:

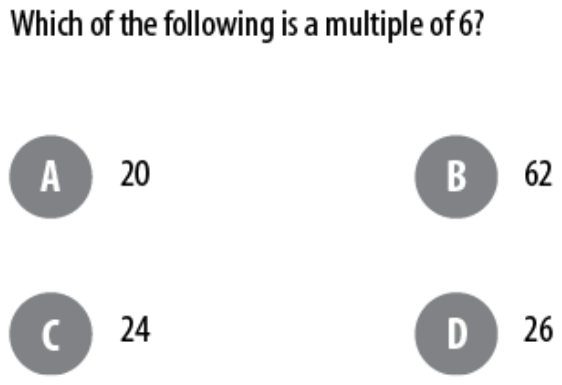


Figure 11.15 – Source: Craig Barton for Diagnostic Questions

On quick inspection, this question looks pretty good. C is the correct answer, B may indicate that students believe multiples start with the given number, and D may indicate they believe they end with that number. I am not entirely sure what A tells me – maybe an error with the 6 times tables – but apart from that I am pretty happy with this question.

Or am I? If I am going to use this question in class, presumably my purpose is something along the lines of assessing if students have a good understanding of multiples. And yet, something that is not assessed at all in this question is arguably the biggest misconception students have with the topic.

Imagine you are a Year 7 student coming into your maths lesson and you are told that today you are studying multiples. Oh no, you think, I always get multiples and factors muddled up – I can never remember which ones are the bigger numbers. And then you are presented with the question above, and a smile appears on your face. You can get this question correct without knowing the difference between factors or multiples as there are no factors present. And if I am your teacher, and several of your peers have the same problem, it could well be the case that you all get this question correct and I conclude that you understand factors and multiples, without ever testing to see if you can distinguish between the two concepts.

Interestingly, by presenting my students with this question, they may subsequently *infer* that multiples are 'the bigger numbers' due to the absence of any number smaller than 6, and hence may learn the difference between factors and multiples indirectly that way. However, this is something I would prefer to assess directly, especially if I am trying to discern in the moment if I have enough evidence to move on.

So, a better question might be something like this:

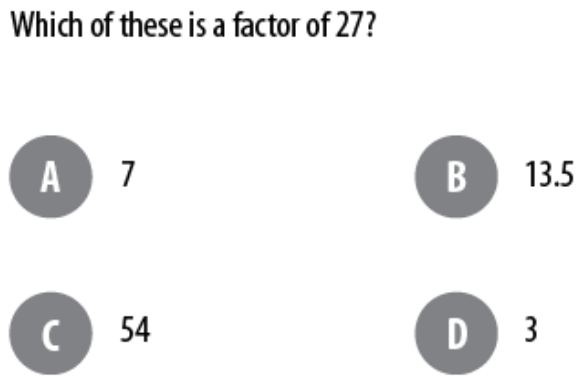


Figure 11.16 – Source: Craig Barton for Diagnostic Questions

I love this question – not just because it contains factors and multiples, but because of answer B. All of a sudden, dodgy definitions of factors such as *a number that goes into another number a whole number of times* are called into question.

As the possibility of getting a question correct whilst holding a key misconception is such an important one, I hope you will permit me one more example. Consider the following:

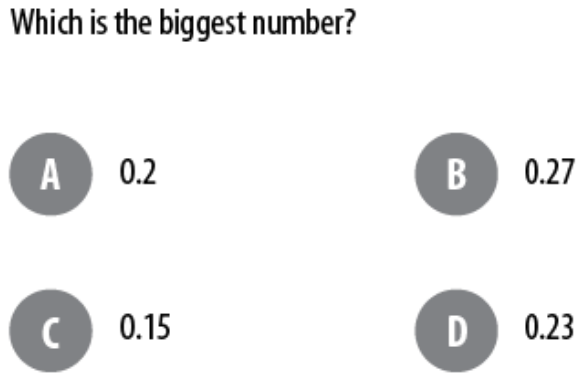


Figure 11.17 – Source: Craig Barton for Diagnostic Questions

Again, we have one right answer and three wrong answers, but this is a terrible diagnostic question. To see why, again pretend you are a Year 7, this time one who does not understand the difference between integers and decimals. You may look at this question and think well, I know that 27 is bigger than 2, 15 and 23, so 0.27 must be the biggest. You have a significant misconception when it comes to place value of decimals, and yet you have

managed to get this question correct. As your teacher, I am faced with false or incomplete information that I may interpret as evidence of your understanding.

A better question might be:

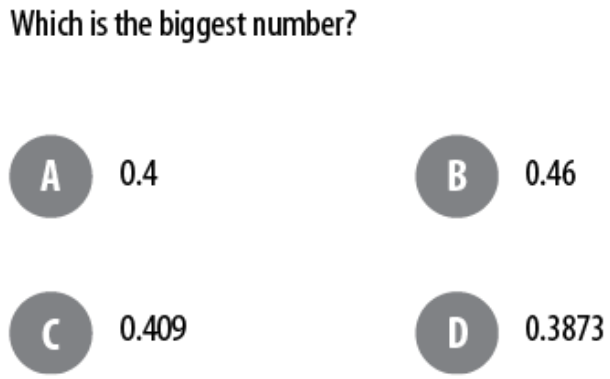


Figure 11.18 – Source: Craig Barton for Diagnostic Questions

Now, this is undoubtedly a harder question, and harder does not necessarily mean better, but students who have the same misconception with the place value of decimals would get this question incorrect, and hence would receive the help they need.

What I do now

Seeing as I make such extensive use of diagnostic questions, I want to ensure that the data I receive back from my students' answers is as accurate and valid as possible. Putting such time into the creation and selection of good questions is time well spent. As Wiliam (2011) puts it: 'Using multiple-choice questions provides a means for sorting all the students' responses ahead of time, so that precious classroom time is not spent trying to make sense of the students' answers'.

Hence, when writing or choosing diagnostic questions to use, I ensure they pass my Five Golden Rules:

1. They should be clear and unambiguous.
2. They should test a single skill or concept.
3. Students should be able to answer in less than ten seconds.

4. You should learn something from each incorrect response without the student needing to explain.
5. It is not possible to answer the question correctly whilst still holding a key misconception.

If they don't, then usually a little bit of tweaking can do the job. As I will discuss later in the chapter, such a process is an incredibly valuable thing for a teacher or group of teachers to do.

OceanofPDF.com

11.5. How to ask and respond

What I used to think

When I first came across the principle of diagnostic questions, I was immediately impressed. However, a nagging thought remained at the forefront of my mind: *this is going to be a bit of a hassle*. I pictured having to faff around with unreliable pieces of technology (voting devices, old laptops *etc*). Even the thought of getting the mini whiteboards out with my bottom-set Year 10s did not fill me with enthusiasm – armed with a pen and their adolescent imagination, A, B, C or D was the last thing some of those students were liable to write.

So, for the first year or so I did not make all that much use of asking diagnostic questions. I used them more as a one-off than a regular part of my lessons. Thankfully all that changed when I discovered that students actually have their own in-built voting devices…

Sources of inspiration

- Barton, C. (2016) 'Dylan Wiliam', *Mr Barton Maths Podcast*.
- Barton, C. (2017) 'Doug Lemov', *Mr Barton Maths Podcast*.
- Ciofalo, J. and Wylie, C. E. (2006) 'Using diagnostic classroom assessment: one question at a time', *Teachers College Record* January 10th.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- Tobin, K. (1987) 'The role of wait time in higher cognitive level learning', *Review of Educational Research* 57 (1) pp. 69-95.

- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

My takeaway

When I interviewed Dylan Wiliam for my podcast, he said many memorable things. But one that really stuck out in my mind was his observation that students rarely forget to bring their fingers to lessons. They may forget pens, pencils and homeworks, but fingers are always there. Likewise, fingers do not tend to be subject to the whims and fancies of unreliable technology, nor are they as easy to create inappropriate gestures with…well, apart from two notable exceptions. And so, when asking and responding to diagnostic questions in my classroom, I always get my students to use their fingers.

The process for asking and responding is one I have tweaked and developed over the last five years and – as I described when it comes to how I present worked examples – I am quite particular about it. It goes like this:

Phase 1 – The Initial Question

I project a question such as the following up on my whiteboard:

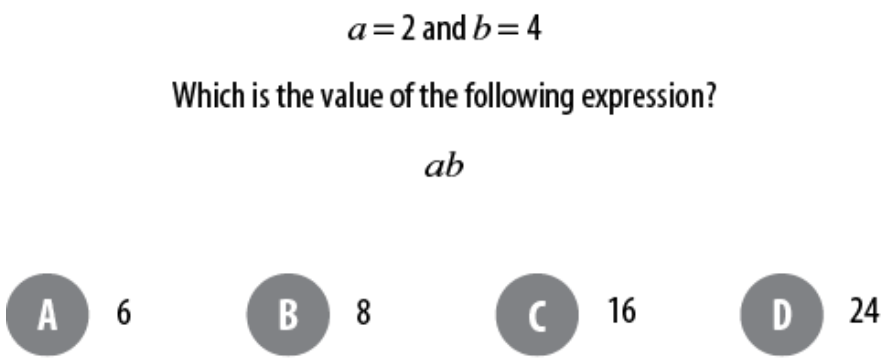


Figure 11.19 – Source: Craig Barton for Diagnostic Questions

The students are sitting in silence and I give them between five and ten seconds' thinking time. The exact time depends on many factors, ranging from the complexity of the question, the age and set of the students, and most importantly my gut feeling while observing them thinking. It is very rarely more than ten seconds – as discussed in the previous section, one of my criteria for a good diagnostic question is that students should be able to answer it in less than ten seconds – but it is always more than three seconds.

Wiliam (2011) explains how, when the question requires thought, increasing the time between the end of the student's answer and the teacher's evaluation from an average wait time of less than a second to three seconds produces measurable increases in learning. This supports our findings on the importance of silent self-explanation in Chapter 5. However, according to Tobin (1987), increases beyond three seconds have little effect and may cause lessons to lose pace. That is where my gut feeling and experience come into play. When I sense that students have had sufficient time to think and settled upon the answer, I continue.

I then say, '3, 2, 1, vote!', at which point my students indicate their choice by holding one finger in the air for A, two fingers for B, three for C, and four for D. Again, this is all done in silence.

Incidentally, my colleague Colm adds an extra option – if his students do not know the answer, then they are to hold up a fist. I sometimes use this, but I am also interested in what students think the answer is if they have to guess, so usually I stick to the 4 possibilities. Non-participation is not an option, and this all comes from building up the right classroom culture as described earlier in this chapter, ensuring students know that these questions are tools of learning, not tools of assessment. I have also noticed that students are more willing to offer an answer to a multiple-choice question than a non-multiple-choice question. Perhaps this is because working out what you think the answer is – even if you are unsure of it – and seeing it as one of the options gives you confidence to have a go. Indeed, the higher participation rate in multiple-choice questions was one of the reasons AQA opted to include them at various points throughout their Maths GCSE papers from 2017 onwards.

I should also say that I am not at all opposed to the use of mini whiteboards. I find they are an excellent way for students to illustrate their thinking, and the fact that students can rub out all record of their work in an instant means the less confident ones are more willing to participate. I use mini whiteboards for longer-form questions, and for the problem part of the example-problem pair approach described in Chapter 6, but not for answering diagnostic questions. Remember, diagnostic questions are short, snappy questions that test one skill or concept. There is rarely any need for

working out, and the time saved from leaving the whiteboards and temperamental pens in the cupboard can be significant.

Following the vote – and assuming that the question is a good one that I have studied in advance – a quick scan of the room immediately gives me some accurate and useful information about my students' understanding of interpreting algebraic expressions. I know how many have got the question right, how many have it wrong, and a good idea of exactly why they have gone wrong. I may choose to leave it at that, and simply use this information to dictate where I take the lesson next – do I need to go back over the basics of algebra armed with my new knowledge, or can I move on?

However, more often than not I like to dig deeper into students' answers. This is a relatively quick process, and it has a number of benefits. First and foremost, it provides me with extra information about how my students are thinking. This is particularly important if I suspect some students may be guessing, or if the question is not a particularly good one and an answer can be arrived at via two different ways of thinking. But there are also benefits for my students. In Chapter 5 we discussed the benefits of self-explaining. Now, whilst self-explanations do not need articulating – and indeed, doing so is potentially more cognitively demanding – being asked to explain their answer certainly compels students to reflect on their reasoning and organise their thoughts. Moreover, it provides a useful opportunity to hear how their peers have approached the question.

So, I ask for a student who has chosen A to give me their reason. I always start at A (and move through the alphabet), regardless of whether this is the correct answer or not, so students do not learn that I am always either starting with the right or wrong answer. Lemov (2015) might refer to this as 'managing the tell', and we discussed this very instance on my podcast.

The student may say something like, 'I added them together'.

I would simply say thank you, ensure the rest of the students in the class are listening and being respectful, and then ask for someone who has chosen B to give me their reason. This process would continue until I have collected a reason for each of the answers. If there is a particular answer that no one has

given, I wouldn't usually refer to it at this stage, but I will certainly use this in Phase 2.

An important point to note here is that this is in complete contrast to my approach for going through worked examples described in Chapter 6. There, my students are in absolute silence, and I am doing all the talking; whereas here I am ensuring the discussion takes place in an orderly manner and it is my students taking centre stage. Why the difference? Well, because the worked examples take place during the early knowledge acquisition phase of learning. There I want to take control to ensure that what students hear and see is as carefully planned as possible in order to give them a solid foundation upon which to build their understanding. When asking diagnostic questions I am reviewing their understanding of a skill or concept that they have previously learned. I want them to show and tell me what they do and do not know, and I want those who have not understood to benefit from the wisdom of those who have. The difference in approach is completely dependent on where we are in the phase of acquiring and mastering the knowledge in question.

After we have heard reasons for all the answers, it is time for a revote. Students can keep their original answers or switch. The majority of the time, most people now choose the correct answer. It is here, for the first time, that I announce what the correct answer is. Giving the answer straight away – for example, saying 'B is correct, now who can tell me why?' before any discussion – misses a golden opportunity. It stops my students thinking. Those who have got it right have no need to listen any more, and those that have got it wrong may feel disheartened. By doing what Lemov (2015) calls 'withholding the answer', I am giving students an incentive to keep listening and thinking throughout the whole process. If I have not been entirely happy with the clarity of the correct explanation given by my students, or I have something extra to add, I will do it at this stage.

Now, it may be tempting to move on here. Indeed, if most students got the question right first and second time, I may well do so, but not before having made a note of any students who are still struggling so I can go and help them out later on. However, if there are still plenty of students who have the wrong answer, I need to go into Phase 2. Likewise, if not many got the

question correct first time around, but then during the revote everyone got it right, I might be a little concerned that maybe they are just copying the perceived smartest student in the class, or they have picked up on some subconscious cues from me when students were discussing their answers. In that case, I may also opt for Phase 2.

Phase 2 – The Follow-up Question
The follow-up question should test the exact same skill as the first question. The idea here is not to progress the students onto something more difficult – that happens via the careful selection of worked examples and practice questions as discussed in Chapter 7 – but to assess understanding of a specific skill or concept.

So, I may ask something like the question over the page.

$a = 5$ and $b = 2$

Which is the value of the following expression?

ab

A 25 B 7 C 52 D 10

Figure 11.20 – Source: Craig Barton for Diagnostic Questions

Notice that the concept is the same, but both the numbers and the order of answers has changed.

Students are again given a short amount of thinking time in silence, and then asked to vote with their fingers. It is here things get interesting, as several scenarios could result. The following represents the way I would deal with each of these, developed from years of trying out various methods, both with my classes and working with other teachers' students. However, it is by no means meant as a definitive guide. Every class has varying needs, responds differently to certain kinds of approaches, and these can change from day to day, and lesson to lesson. But hopefully it will be of some use.

Scenario 1: *50/50 split*

When approximately half the class get the follow-up question right and half get it wrong, I like to pair up the students. Practically speaking I ask anyone who has got the question correct *and is confident that they can explain why* to raise their hand. I then ask the other students to move near to one of the hand-raisers, whereby some peer-to-peer coaching takes place. We have seen in Chapter 5 that explaining to others is likely to impose an additional cognitive load compared to self-explaining, so it is important that the hand-raisers feel secure in their knowledge. Following a few minutes of coaching, I will then project another follow-up question onto the board, and everyone votes independently.

Scenario 2: *Majority correct*

If most students have got the follow-up question correct and there are just a handful who are still struggling, I would be inclined to do some small-group intervention. I would ask all the students who got the question wrong to come to the front to work with me. Having seen their answers, I will be better informed as to where the difficulties lie, so hopefully I will be able to provide some focused support. Meanwhile, I will either ask the rest of the class to continue with the next piece of work, or, if I wish everyone to move on at the same time, I will set these students the challenge of analysing the question – in other words, writing down the reasons for each of the wrong answers – and then creating their own question. As we shall see throughout this chapter, it is not just answering diagnostic questions that brings great benefits to all concerned.

Scenario 3: *Majority incorrect*

This time we only have a handful of students who got the question right. Again, depending on the class and the situation in general, I may adopt a coaching model. Say there are six students who are confident in their answer. I would appoint them coaches of small groups of students and give them a few minutes to try to explain how they understand the question. Once again I would follow this up with a question that tested the same skill, and I would inform the coaches that I would be judging their success on

how the members of their team performed on this question. Alternatively, I may ask this group of six students to do some analysing and creating of questions, whilst I do some whole-class teaching with the rest of the class. I may also borrow one or two of this group to share their way of understanding and approaching the question in case it adds a viable alternative to mine.

Scenario 4: *All incorrect*

If everyone gets the question wrong, it is time to go back to basics. If the question involves some baseline knowledge that is crucial to the acquisition of the new skill that I intend to teach the students later in the lesson (see next section), then I will do the reteaching there and then, following similar approaches explained in the earlier chapters of this book. If the knowledge is not needed, then I may choose to make a note, and leave it until another time when I am better prepared. Either way, I am not back at square one because the makeup of my students' answers will give me valuable insight into the specific nature of their misconceptions, and thus enable my help to be more effective.

Scenario 5: *All correct*

If everyone has got the question correct, I may do quick check of explanations. I may also ask students why they think I have included the three wrong answers, and if there would be any other wrong answers that they may have chosen in order to get them thinking a bit harder. But crucially, I won't dwell too long on this – I will move on. Now, as we discussed at the start of this chapter, the distinction between learning and performance may suggest that I do not have sufficient evidence to infer what my students have learned and understood from their one correct answer in this single point in time. That is true. But I certainly do have enough evidence to move on in the here and now. If students get the question right, and it is a good question, then let's not hang about, let's crack on.

Two points need making. For starters, the first time you try any of the coaching approaches, it is likely to end in disaster. I know it did for me. But persist. Students need time to get used to how things work and your

expectations of their behaviour. When used consistently, responses to these scenarios become a regular and natural part of the lesson. But equally, for some classes, these may never work. If behaviour or confidence is an issue, then it may be best to engage in teacher-led questioning and explanations following the responses to the second question.

The second point is obvious, but very important. If you ask a diagnostic question, you must be prepared to respond to it. I have seen (and delivered!) too many lessons where a teacher asks a question at the start that they assume the students will get right, and then when the students in fact get it wrong, the teacher proceeds with the rest of the lesson regardless, usually because they have spent ages planning a lovely activity that they are determined to do no matter what. If knowledge of the skill or concept covered by that question is of fundamental importance in helping students understand what you are about to teach them (in other words, it is baseline/prerequisite knowledge), then it has to be dealt with, no matter what the rest of the lesson plan suggests. If it is not necessary – and the different types of questions you may ask at the start of a lesson will be discussed in the next section – then I advise telling the students not to worry about it, and coming back to this question at a time when you have had time to prepare explanations and supporting resources.

Likewise, all too common is the following approach to lesson planning, which often takes place when the lesson is being observed. You ask a tricky question at the start, the students cannot do it, you then teach them, and at the end of the lesson ask them that question again as a means of demonstrating the progress they have made. Simple! Whether you can infer learning or progress at all using this approach is a separate matter, but what is clear is that if in fact some or all of your students get that first question correct, then you have some decisions to make. As the tweet by Wiliam shared at the start of this chapter explained, formative assessment is responsive teaching. If you do not respond to your students' answers, then what is the point in asking the question?

Phase 3 – Ask the question again at a later date

The benefits of both testing and spacing will be discussed in the next two chapters, but it is important to note here that two questions, no matter how

well designed they are, will never be enough to ascertain the understanding of your students. Given how quick these questions are to ask and to receive data, it is relatively easy to build them in as a regular part of your teaching. So, having asked these questions today, I will make a note in my planner or calendar to ensure I ask a similar question in three weeks' time, and then in a few months' time. This could be in class, in a homework, or as part of a low-stakes quiz.

What I do now

I follow that three-phase approach to diagnostic questions every lesson, every day and with every class, asking a minimum of three questions per lesson. When and why I might ask these questions is coming soon to a section near you!

OceanofPDF.com

11.6. When to ask a diagnostic question

What I used to think

When I did use diagnostic questions, it would only be at the beginning of the lesson, mainly as a way of assessing some of the skills students would need for the rest of the lesson.

It turns out that this was like having one of the world's most powerful computers and only using it to play solitaire.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Ciofalo, J. and Wylie, C. E. (2006) 'Using diagnostic classroom assessment: one question at a time', *Teachers College Record* January 10th.
- Hodgen, J. and D. Wiliam. (2006) *Mathematics inside the black box: assessment for learning in the mathematical classroom*. Eastbourne: Gardners Books.
- Karpicke, J. D. and Blunt, J. R. (2011) 'Retrieval practice produces more learning than elaborative studying with concept mapping', *Science* 331 (6018) pp. 772-775.
- Kornell, N., Hays, M. J. and Bjork, R. A. (2009) 'Unsuccessful retrieval attempts enhance subsequent learning', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35 (4) pp. 989-998.
- Mccrea, P. (2017) *Memorable teaching: leveraging memory to build deep and durable learning in the classroom*. CreateSpace Independent

Publishing Platform.

- Richland, L. E., Kornell, N. and Kao, L. S. (2009) 'The pretesting effect: do unsuccessful retrieval attempts enhance learning?', *Journal of Experimental Psychology: Applied* 15 (3) pp. 243-257.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.
- Wiliam, D. (2016) 'The secret of effective feedback', *Educational Leadership* 73 (7) pp. 10-15.

My takeaway

There are four different times or scenarios that I now use diagnostic questions, with each serving its own purpose.

1. Start of the lesson

The start of the lesson is the most obvious time to use a diagnostic question, but for four different reasons.

a) To assess baseline knowledge

In a 2016 article on feedback, Wiliam wrote a line that really resonated with me: '**We need to start from where the learner is, not where we would like the learner to be**'.

One of the key features of the models of thinking that we discussed way back in Chapter 1 was that existing knowledge makes acquiring new knowledge a whole lot easier. This makes attempting to discover where students are – and in particular which misconceptions they hold – before you attempt to teach them new knowledge one of the most important parts of teaching. Hodgen and Wiliam (2006) advise that teachers should 'start from where the learner is, recognising that students have to reconstruct their ideas and that to merely add to those ideas an overlay of new ideas tends to lead to an understanding of mathematics as disconnected and inconsistent'.

But deciding which baseline knowledge to assess can be easier said than done.

A few years ago, I inherited a wonderful Year 11 class who suffered from a severe and debilitating lack of confidence. When I taught them vectors, I put in extra effort as it was one of the few topics they did not have any negative preconceptions about because they had not encountered it (and failed at it) before. However, all my hard work was undone when I gave my students an exam-style question like this:

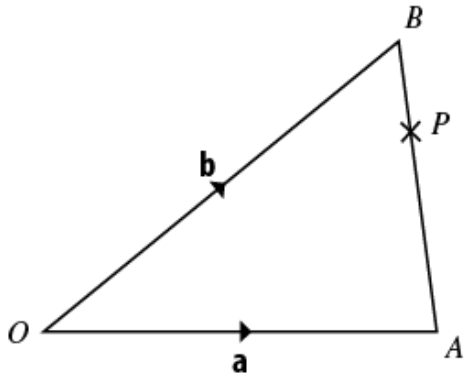


Diagram **NOT** accurately drawn

OAB is a triangle.

$\overrightarrow{OA} = \mathbf{a}$
$\overrightarrow{OB} = \mathbf{b}$

(a) Find $\overrightarrow{AB}$ in terms of **a** and **b**

---

(1)

P is the point on *AB* such that $AP : PB = 3 : 1$

(b) Find $\overrightarrow{OP}$ in terms of **a** and **b**.
Give your answer in its simplest form.

---

(3)

(4 marks)

Figure 11.21 – Source: GCSE Maths Takeaway, available at http://mrbartonmaths.com/students/legacy-gcse/gcse-maths-takeaway.html

My students could not do Part b); not because they did not understand vectors, but because they were not secure in dealing with the addition of whole numbers and negative fractions that is required to get the final answer. Had I instead thought carefully about all the skills involved in

mastering vectors – not just the obvious ones – I could have prevented this. I could have asked them a question like the following at the start of the lesson:

$$-\frac{3}{5} + 1 =$$

A $\frac{8}{5}$ B $-1\frac{3}{5}$ C $\frac{2}{5}$ D $-\frac{2}{5}$

Figure 11.22 – Source: Craig Barton for Diagnostic Questions

This would have revealed their weakness and I could have addressed it completely separately from our study of vectors, providing explanations, examples and practice as required. Instead, their confusion with this relatively small skill became tangled up in their wider perception of vectors, and hence vectors became yet another topic in maths that they could not do. Likewise, me subsequently trying to teach them this skill alongside navigating through a complex vectors problem likely imposed a significant burden on their working memories, leading to cognitive overload, with little learning of either fractions or vectors taking place. And this was 100% preventable.

I have found it beneficial when planning a lesson to make a list of all the prerequisite knowledge that students need to have acquired and automated in order to have the best chance of acquiring the new knowledge you want to teach them. It is very similar to the process suggested when we discussed Deliberate Practice in Chapter 8. Once again, this is a particularly useful thing to do as part of a departmental meeting when discussing an upcoming topic – everyone makes their own list, and then you compare. Our curse of knowledge makes this a difficult exercise, and so the more help the better. For example, without help from my department I would not have thought to include a question on placing decimals on a number line before teaching histograms (this skill is, of course, required in order to accurately plot frequency density).

b) To prime students' long-term memories

The assessing of baseline knowledge serves another purpose that is of direct benefit to the student – their long-term memory is primed for the acquisition of the related new knowledge you intend to teach them. Mccrea (2017) explains that 'we can increase the chances of learning happening by activating, or warming-up relevant areas of long-term memory before exploring a new topic'. This is a major benefit of asking good questions. They force the student to direct attention to the relevant topic, and essentially start the process of retrieving and organising relevant knowledge in their long-term memory that will make the acquisition of new knowledge so much easier. So, if I want to teach my Year 10s how to add algebraic fractions, some diagnostic questions at the start of the lesson on adding non-algebraic fractions, simplifying expressions and expanding double brackets will not only allow me to assess their understanding of these relevant skills and intervene accordingly, but also bring these skills to the forefront of students' minds so they are ready to be applied. I think of priming like an athlete doing stretches in preparation for a race – the correct areas are warmed up and the body (or, in our students' case, the mind) is more prepared to accept the challenge of what is to come.

c) To make use of the Testing Effect

The Testing Effect is related to priming, and will be covered in much more detail in the next chapter. But in essence, when students are compelled to retrieve knowledge, that very process modifies and enhances what is stored in long-term memory (*eg* see Karpicke and Blunt, 2011). This is, of course, good news if that knowledge is needed for the forthcoming lesson, but it can also be useful to ask questions on unrelated concepts that cover something students did last week, month or year. Given how quickly these diagnostic questions can be asked and answered, you can get through five or more in a matter of minutes. Setting up a schedule to ensure that all prior skills and concepts are revisited throughout the year via these questions is one of the best habits I have developed. The only downside is what to do if students get the question wrong. If it is not crucial to the content of the reset of the lesson, I tend to be honest with my students, reassuring them not to worry, that it is best we found that out now, and that we will come back to it some time soon. As discussed in the last section, I have seen and delivered too

many lessons that got derailed by an unrelated starter activity that did not go as planned. If I do not need to deal with it now, then I don't, and instead come back when I am more prepared.

d) To make use of the Pretest Effect
When I interviewed Robert and Elizabeth Bjork for my podcast, one of the most fascinating concepts they discussed was the Pretest Effect. This is where being tested on material you have not studied before actually enables you to better understand and remember the material when you do study it. Whereas priming refers to focusing students' attention on material they have seen before that is related to what they are about to learn, pretesting is for content that is brand new. We will discuss pretesting in more detail next chapter, but two studies merit a quick mention here. Richland *et al* (2009) found that generating a wrong answer increases our chances of learning the right answer, and Kornell *et al* (2009) found that as long as students are given appropriate feedback, you will still see a testing benefit even if they get the answer wrong on a pretest. Hence, as counterintuitive as it sounds, asking a diagnostic question at the start of the lesson on something students have not encountered before but will be doing this lesson, and having the subsequent discussion to ensure students are clear what the correct answer is, may well have two benefits. It will give you an insight into who knows what, which will better inform your teaching. But secondly, it may actually aid your students' ability to understand the new topic. More on this – and the potential pitfalls – next chapter!

2. Middle of the lesson
Imagine you are halfway through a particular lesson. The students are working on some carefully prepared sequence of questions (Chapter 7) or Purposeful Practice activity (Chapter 10). They are all at different points, working at their own pace. You are trying to make an informed decision whether or not now is the right time to move onto the next part of the lesson, maybe to introduce the next skill or concept. One of the best ways to do this is to ask a well-designed diagnostic question.

William (2011) describes diagnostic questions used in this ways as 'hinge questions' (Fun fact: see what happens if you visit hingequestions.com). He explains:

> The central idea here is that the teacher designs each lesson with at least one 'hinge' in the instructional sequence. The hinge is a point at which the teacher checks whether the class is ready to move on through the use of a diagnostic question. How the lesson proceeds depends on the level of understanding shown by the students, so the direction of the lesson hinges at this point.

Diagnostic questions presented in the way suggested earlier in this chapter are perfect for this as they are quick to ask, quick to get data back from, and do not simply tell you where certain students are going wrong, but why they are going wrong. This allows you to intervene where needed. Conversely, if everyone gets the question right, then onwards and upwards!

3. End of the lesson
As discussed in Section 3.10, I no longer end my lessons on a difficult question. For me, a key purpose of the end of the lesson is to inform the starting point of the next lesson. I want to get a quick and accurate sense of my students' understanding of what I have taught them that particular lesson so I know my starting point the next time I see them. Hence, I use a diagnostic question.

Of course, we need to bear in mind that the distinction between learning and performance warns us against inferring that students have definitely learned everything when we assess them immediately after instruction. However, wrong answers certainly tell us there are problems we need to address, and the specific nature of those problems. Likewise, if everyone gets the last question correct (and it is a good one!), then I have evidence to move on next lesson, but crucially making a note to revisit the concept at numerous points in the future.

A common practice amongst teachers is to issue an Exit Ticket at the end of the lesson. This is essentially a question that students complete in the final minutes of the lesson and hand to the teacher on the way out. The teacher then has a quick look over the students' responses and uses this to inform their planning. Diagnostic questions are perfect for this purpose. They are lightning fast, both for the students to complete and for the teacher to mark. I have started printing them out and adding a section for the student to

explain their answer. This enables me to glance at their choice of A, B, C and D, and then dig deeper into the reason behind this choice if I need to. It also allows the students to reap the benefits from self-explaining discussed in Chapter 5. An exit ticket might look like this:

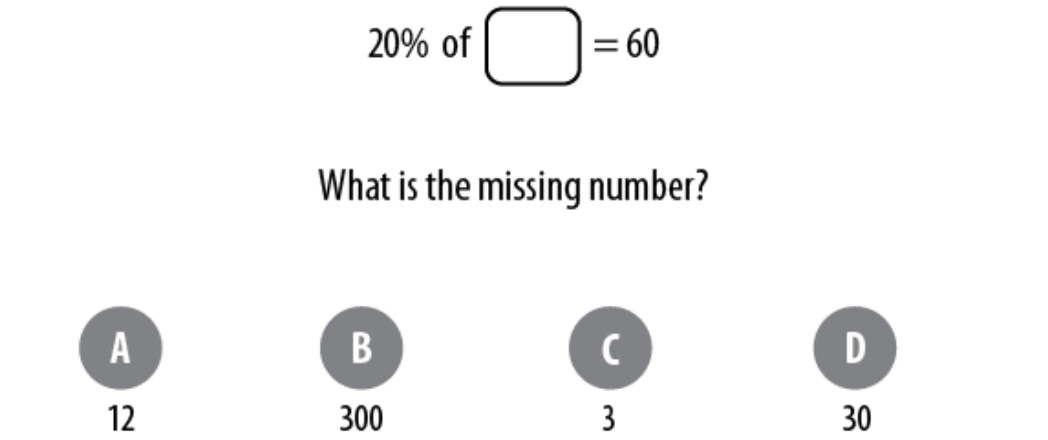


Figure 11.23 – Source: Craig Barton

Finally, the end of the lesson is also an opportunity to give students a question based on what they will be studying next lesson. This serves a dual purpose – it allows me to get a sense of their understanding of key concepts so I can adjust my planning accordingly, and it also provides an opportunity to take advantage of the Pretest Effect outlined above.

4. Homework

As we have seen, writing good diagnostic questions is flipping hard…which makes it the ideal activity to give to your students! Students must

demonstrate a sound understanding of the key concept or skill, together with an appreciation of where other students may go wrong.

A homework following a lesson on sharing in a ratio may look like this:

- *Write a diagnostic question on sharing in a ratio*
- *Write a model explanation for the correct answer*
- *Write a sentence explaining your choices for each incorrect answer*

The same homework can be set to all students, giving each the opportunity to demonstrate the depth of their knowledge of a concept. Such a homework also has the advantages of being relatively quick for me to mark, illuminating in terms of the insights I can learn, and may provide me with a new batch of diagnostic questions that I can use with my students. Everyone's a winner.

What I do now

Ciofalo and Wylie (2006) write: 'There is flexibility both in terms of when in a lesson a diagnostic item might be used and how a teacher might use it. The most important thing to consider is that the teacher is evoking information she needs to inform subsequent instruction or to provide feedback to students'.

The information I receive from my students' responses to diagnostic questions at various points in a lesson is vital. It allows me to respond in real-time to their ever-changing needs. The information is quick to collect and accurate. As such, the use of diagnostic questions has probably become the most important part of my teaching.

If only everyone agreed with me…

OceanofPDF.com

11.7. Seven common criticisms of multiple-choice questions

What I used to think

When I gave talks on formative assessment and the use of diagnostic questions, I used to be very nervous. This was not due to stage fright, but more because the mere mention of 'multiple-choice questions' tends to invoke a certain reaction in some people. Indeed, it is similar to the reaction I now get when I hear the phrase 'learning styles'.

Multiple-choice questions certainly have their critics, and I have met many teachers – both at my workshops and on Twitter – who simply refuse to use them. They have questioned why on earth I would be advocating such a harmful approach to assessment and teaching in general. I'll be honest – for many years I shared their concerns. But hours of reading, and speaking to people who know far more about the issues than me, have alleviated most of my fears.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Butler, A. C. and Roediger, H. L. (2008) 'Feedback enhances the positive effects and reduces the negative effects of multiple-choice testing', *Memory & Cognition* 36 (3) pp. 604-616.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.

- Little, J. L., Bjork, E. L., Bjork, R. A. and Angello, G. (2012) 'Multiple-choice tests exonerated, at least of some charges: fostering test-induced learning and avoiding test-induced forgetting', *Psychological Science* 23 (11) pp. 1337-1344.
- Marsh, E. J., Roediger, H. L., Bjork, R. A. and Bjork, E. L. (2007) 'The memorial consequences of multiple-choice testing', *Psychonomic Bulletin & Review* 14 (2) pp. 194-199.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

My takeaway

Here are seven criticisms of multiple-choice questions that I hear regularly, and my response to them.

1. Students have a one-in-four chance of getting the answer right by guessing
This is certainly true – I have done the maths. I have already discussed that William (2011) deals with this issue by giving students questions with more choices and more than one correct answer, and how for ease of data collection and consistency I prefer to keep all the questions I use in the same format. Instead, I counteract the problem of guessing in two ways. First, you will see that I ask students to provide their explanations as a means of illustrating their method for getting to the answer. My colleague, Colm, also asks students to vote with a closed fist to indicate that they really don't know. Second, I ask a follow up question. This immediately reduces the odds from being correct from pure guesswork from $\frac{1}{4}$ to $\frac{1}{16}$.

2. Students can work backwards to get the answer
This is also true. One obvious way to get round this is to ask the kind of questions where it is not possible to work backwards. Indeed, a good question to ask colleagues when showing them your questions is: *can someone work out the answer without doing what I need them to do?* But this is sometimes easier said than done. Solving equations is a classic example. If you ask a question where the four answers consist of solutions, it is always possible to figure out the answer by substituting each of the four possible

answers back into the original equation. Thus, the question becomes a test of a student's ability to substitute instead of solving equations. However, you can get around this problem by asking questions like this:

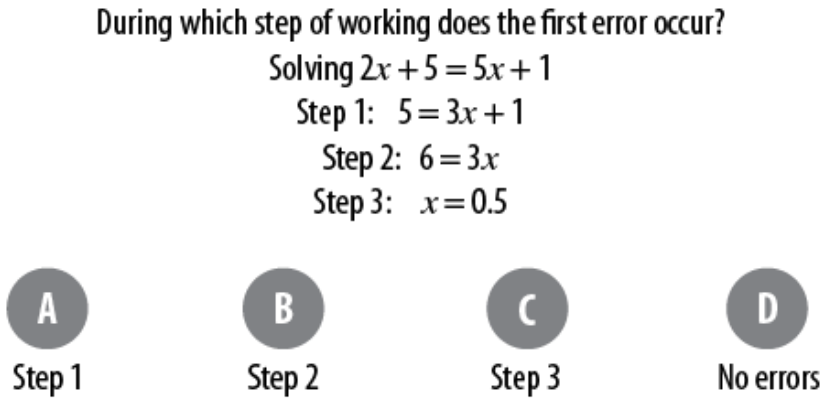


Figure 11.24 – Source: Craig Barton for Diagnostic Questions

For topics where this is not possible, asking students to provide their reasons can help alleviate this problem to some extent. I have also found that it is effective to remind students that these questions are tools for learning and not assessment. This is related to classroom culture, and also being honest with students – something that will become even more important in the next chapter. If students know that you are asking these questions so they can get better at a given skill and so that you can learn how best to help them, as opposed to merely being a way to record some more marks against their name, then they are far more likely to attempt the question in the way it is intended.

3. We do not have an opportunity to discover other answers students may have come up with

Once again, multiple-choice questions are guilty as charged. By limiting each question to four possible answers, we are imposing a constraint on the number of answers we can hear from students. If the question is a poorly designed one – with a common misconception excluded from the choices of wrong answers – then this could lead to big trouble, with students unable to share their thinking. In essence, we are preventing other errors from emerging.

There are ways around this. The first is to ensure the questions we use are good ones, with the most common student misconceptions carefully thought out, discussed, and built into the answers. This enables us to plan our responses to these misconceptions in a manner that is not as easy with other question types. In this sense the limited number of incorrect answers is an advantage, as we can *plan for error* – something discussed in the section below.

Secondly, after asking the question, voting and discussion of answers, we can ask the students, 'Did anybody think the correct answer was something else?'. Now, if this is a high-stakes test where students' responses are going to be recorded, then there is little incentive to be honest. But if students see this questioning process not as assessment but as a key part of their learning, and if the sharing of alternative answers is encouraged and praised, then students are more likely to be forthcoming. Any other answers that come to light can then be discussed as a class, and I may also choose to update future versions of this question with this answer.

However, I am not going to lie – this is an issue with all multiple-choice questions. But it is a price I am willing to pay for the many benefits that using diagnostic questions brings.

4. These questions are unfair – they are trick questions
Consider the following question:

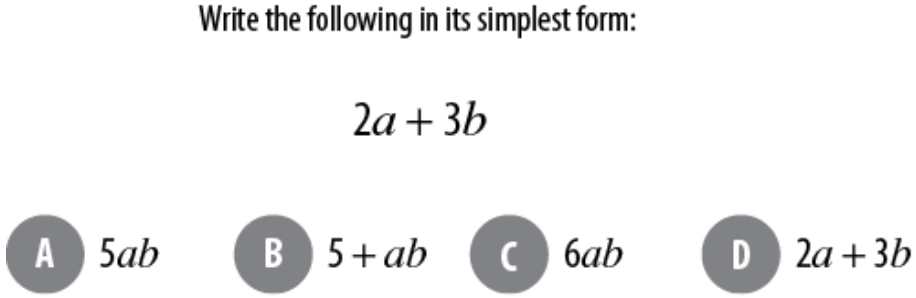


Figure 11.25 – Source: Craig Barton for Diagnostic Questions

Is this an unfair/trick question? I don't think so.

Wiliam (2011) discusses a similar version of the question, and shares the response he gets from teachers when they see it:

> The question is perceived as unfair because students 'know' that in answering test questions, you have to do some work, so it must be possible to simplify this expression; otherwise, the teacher wouldn't have asked the question – after all, you don't get points in a test for doing nothing. Such a question may well fall short of standards and fairness guidelines established for items in a high-stakes test. But for finding out whether students understand a key principle in algebra, it is a useful question. If a student can be tempted to simplify 2a + 5b, then the teacher should want to know that, because addressing this misunderstanding will be essential before the student can make progress in algebra. The fact that this item is seen as a trick question shows how deeply ingrained into our practice is the idea that assessment should allow us to sort, rank, and grade students, rather than inform the teacher what needs to be done next.

5. The exam is not multiple-choice, so what is the point in giving students multiple-choice questions to practise?

Related to the concern about trick questions is the more general point that the vast majority of questions on the exams students take are not multiple-choice, so what is the point of using them lots in the classroom? Surely it is better to give students questions more similar in form to the ones they will be tested on?

The first point to raise here was covered in Chapter 8 on Deliberate Practice – the nature of the practice does not (and, indeed, in many cases should not) need to match what is required in the final performance. Just as giving students problems to solve may not be the best way to help them become good problem-solvers, giving students exam questions may not be the best way to help them do well in exams. Diagnostic questions help you pinpoint specific misconceptions, and understand the reasons for these misconceptions, in ways that other types of questions cannot.

Secondly, Marsh *et al* (2007) find that the positive benefits of the Testing Effect (see next chapter) are still achieved via multiple-choice questions, and

that these benefits are not limited to simple definition or fact-recall multiple-choice questions, but extend to those that promote higher-order thinking. Indeed, there is evidence that the specific use of multiple-choice diagnostic questions may actually improve future performance on non-multiple-choice exams. Following my interview with Robert and Elizabeth Bjork, I contacted Robert for extra clarity on this point, as it seemed so important. He replied: 'In total, the positives of using good multiple-choice diagnostic questions are huge. In most of the experiments that have been carried out the final test is a cued recall test – without alternatives available – but initial multiple-choice testing has far greater benefits for that test than does initial testing that matches such cued recall testing'.

Little *et al* (2012) provide an explanation for this. They found that multiple-choice questions actually have an advantage over different question types in that to get the question correct students must not only consider why their choice of answer is right, but also why the other answers are wrong. They explain:

> Our findings suggest that when multiple-choice tests are used as practice tests, they can provide a win-win situation: specifically, they can foster test-induced learning not only of previously tested information, but also of information pertaining to the initially incorrect alternatives. This latter advantage is especially important because, typically, few if any practice-test items are repeated verbatim on the subsequent real test.

We will consider this point further as we look at the next criticism.

6. They are easier than normal questions, so they do not challenge my students enough

To address this issue, we must again concede that poorly designed diagnostic questions may indeed be easier, as students can guess the answer or work backwards from the alternatives they are presented with. But good diagnostic questions can in fact be more challenging than their non multiple-choice counterparts for the very reason explained above – in order to get them right, students must consider why the alternative answers are wrong.

Consider the following question presented in two different ways:

Version 1:

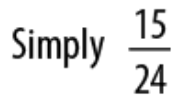


Figure 11.26 – Source: Craig Barton

Version 2:

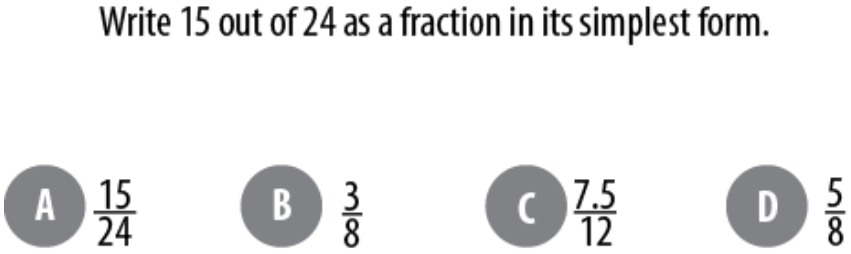


Figure 11.27 – Source: Craig Barton for Diagnostic Questions

Which is more challenging? No prizes for guessing that I believe it is Version 2. To get Version 2 correct, a student must first consider and rule out A (can it be simplified?), B (have I divided the numerator and denominator by the same amount?) and C (can I just halve the numerator and denominator?). This is arguably a greater challenge than simply getting the answer correct, and will leave students better placed to deal with variations of this skill when they encounter them.

Another example illustrates this point further. Consider the following versions of the same question:

Version 1:

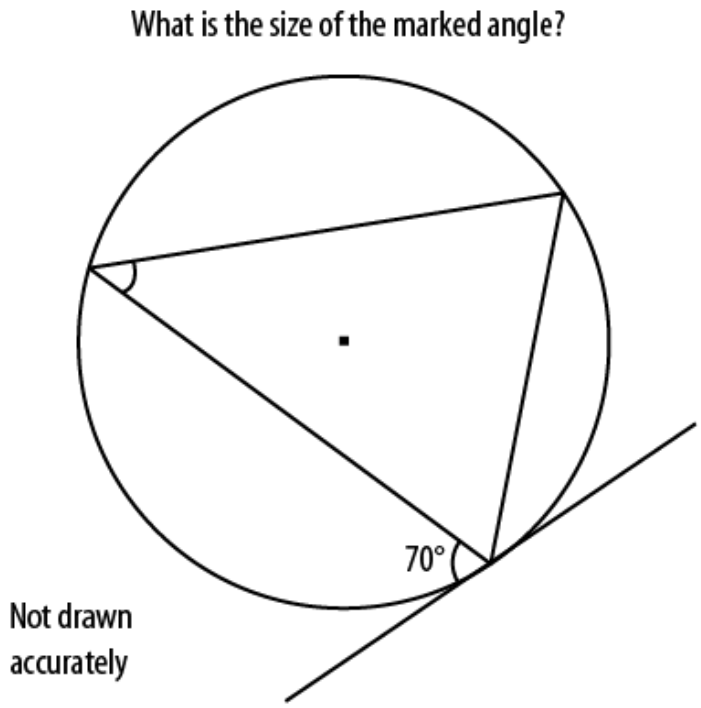


Figure 11.28 – Source: Craig Barton

Version 2:

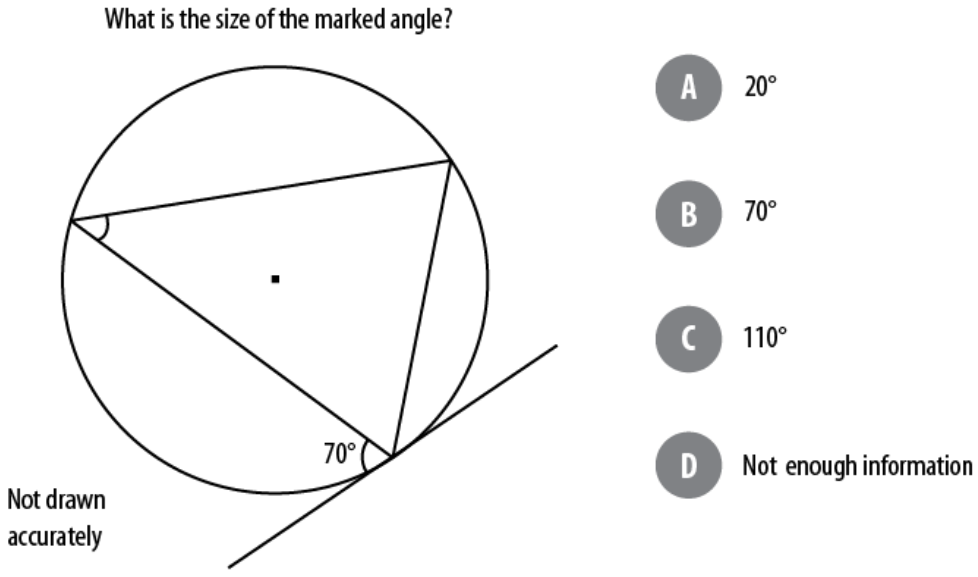


Figure 11.29 – Source: Craig Barton for Diagnostic Questions

Consider the thought process a student must go through to get Version 2 correct. They must consider and rule out A (is this something to do with a

tangent meeting a radius at 90°?), B (are those two lines parallel, so is this alternate angles? Or is it alternate segment theorem?), C (is this to do with opposite angles in cyclic quadrilaterals adding up to 180°?), and then fight against all their instincts that tell them *not enough information* is never the correct answer!

In addition, which version are students more likely to attempt? As I have already argued, I believe students are more likely to offer an answer to a multiple-choice question than one without options. If fewer students opt out, we learn more about their understanding from their answers. And we get these extra actionable insights without having to sacrifice challenge.

7. The distractors are dangerous – they can actually cause students to develop misconceptions

I have saved the big one until last. It is a fear shared by many teachers, and one that has caused me a few sleepless nights.

We have just seen some of the advantages of plausible incorrect answers (distractors) in that they make students think harder. But by showing students questions with these plausible alternatives, am I in fact helping students develop misconceptions that they would not have done otherwise? Indeed, is it the case that the more plausible the alternative answers – and hence the better the question – the more dangerous it is? Thus we are faced with a similar dilemma to the one discussed in Chapter 3 concerning cognitive conflict – is it better to simply show students correct answers and methods, instead of running the risk that students may get confused between what is right and what is wrong?

Marsh *et al* (2007) do find that multiple-choice distractors (or 'lures' as the authors refer to them) may become integrated into subjects' more general knowledge and lead to erroneous reasoning about concepts. However, they also find that misconceptions usually exist *before* seeing the multiple-choice question. They explain: 'Rarely did students select the correct answer on the initial test and then produce a lure on the final test. Nor were students likely to select Lure A on the first test and then produce Lure B on the final test'. So, it seems that the majority of the time diagnostic questions do not cause misconceptions, but in fact *highlight* misconceptions that were already

present. Moreover, diagnostic questions have the advantage of highlighting the specific nature of this misconception so it can be corrected.

Butler and Roediger (2008) offer more reassurance for fans of multiple-choice testing. They conducted an experiment where subjects studied passages and then received a multiple-choice test with either immediate feedback, delayed feedback, or no feedback. Feedback in this sense was informing the subjects if they were right or wrong, not the lengthy individualised process we teachers normally associate with the term. In comparison with the no-feedback condition, both immediate and delayed feedback increased the proportion of correct responses and reduced the proportion of intrusions (*ie* incorrect responses based on distractors from the initial multiple-choice test) on a delayed cued recall test. Of course, we must be very careful generalising these findings to the world of maths, but the implication for teachers from this study is very simple – if using multiple-choice questions, ensure students are told whether they are right or wrong in order to reduce any negative effects from exposure to plausible incorrect answers.

We will discuss the merits of immediate versus delayed feedback in the next chapter, but the key point here is that, following the asking and discussion of the diagnostic question, it should be crystal clear to students what the correct answer is. Lemov (2015) makes a similar point when discussing the concept of Own and Track, suggesting that teachers physically wipe any incorrect answers and suggestions from their boards following a discussion, lest they should lead to extra confusion as to what the correct answer was.

Indeed, we may be able to turn the potentially alluring nature of the alternative answers even further to our advantage. In Chapter 5 I discussed the Self-Explanation Effect, whereby learners who attempt to establish a rationale for the solution steps by pausing to explain the examples to themselves appear to learn more than those who did not. Similarly, in Section 8.4 we looked at the Hypercorrection Effect, whereby errors committed with high confidence are more likely to be corrected than low-confidence errors. If we only reveal the correct answer at the end of the discussion process, ask the students to pause and self-explain – maybe going so far as to write down their reasons as to why it is correct in their books.

We may end up with the best of all worlds, particularly if the student was initially confident in their choice of (incorrect) answer.

What I do now

No teaching tool is perfect, and diagnostic multiple-choice questions are no exception. Indeed, I would never dream of bringing my students up on a diet consisting solely of such questions. Students need practice answering a wide variety of questions, both short and long, open and closed.

However, I am less concerned about the criticisms of such multiple-choice questions, and more focused on their benefits. Ironically, their closed nature, combined with a limited number of possible responses, perhaps provides their greatest advantage over other forms of questioning – it means student answers can be quickly interpreted, and teacher responses and interventions can be carefully planned in advance. That point alone means they will always play a prominent role in my lessons.

OceanofPDF.com

11.8. Anticipating mistakes and misconceptions

What I used to think

I used to think that good subject knowledge was all about being able to do the mathematics myself, and that was enough to make me a good teacher.

Hence, when I walked into my first ever lesson as an NQT in September 2005, to teach a bunch of unsuspecting Year 8s the beauty of adding fractions, I thought somebody had set me up. I had never seen – nor considered – the kinds of mistakes these students were making, for the very reason that I had never made them myself. I kept saying things like, 'Don't you see, you do it like this', but of course the students did not see it that way, just as – and far more importantly – I did not see it their way.

Back then I was a good mathematician, and a terrible teacher.

Sources of inspiration

- Coe, R., Aloisi, C., Higgins, S. and Major, L. E. (2014) *What makes great teaching? Review of the underpinning research*. Available at: https://www.suttontrust.com/wp-content/uploads/2014/10/What-Makes-Great-Teaching-REPORT.pdf.
- Lemov, D. (2015) *Teach like a champion 2.0: 62 techniques that put students on the path to college*. Hoboken, NJ: John Wiley & Sons.
- National Research Council (2000) *How people learn: brain, mind, experience, and school: expanded edition*. Washington, DC: National Academies Press.
- Wiemann, C. (2007) 'The "curse of knowledge", or why intuition about teaching often fails', *APS News* 16 (10) (no pagination).

My takeaway

We have seen throughout this book how, as teachers, we may well be burdened by the curse of knowledge (*eg* Wiemann, 2007). Nowhere is this more evident than when it comes to anticipating the mistakes students will make and the misconceptions they hold. As a general rule, teachers are experts in their chosen domain. More often than not they have been in top sets when learning the subject, perhaps studied the subject at degree level, and if they did ever struggle with a topic it was so long ago that the precise nature of their struggles may well have been forgotten. This is the very reason why a report from the National Research Council (2000) suggests that expert learners may in fact be the worst people to teach novices.

In their review of what makes great teaching, Coe *et al* (2014) identify content knowledge as having strong evidence of impact on student outcomes. They explain that teachers need to have a strong, connected understanding of the material being taught, but that they must also understand the ways students think about the content, be able to evaluate the thinking behind non-standard methods, and identify typical misconceptions students have.

It is that final part that I struggled with throughout most of my career. Until you have taught a topic many times to many different students, how on earth are you supposed to anticipate the typical misconceptions students have when they are misconceptions you have never had yourself?

Even experienced teachers struggle. If you have heard me speak in the last four years the chances are I will have subjected you to my *Guess the Misconception* game (it is only a matter of time before Ant and Dec pay me millions for the TV rights). I present teachers with a diagnostic question that has been answered tens of thousands of times on my Diagnostic Questions website and challenge them to pick the most popular misconception. When I did this at MathsConf in Leeds, out of a room of around 200 teachers, only three teachers managed to select the most popular misconception from all four questions I presented.

You can experience some of the fun yourself by guessing the most popular incorrect answer on the following question:

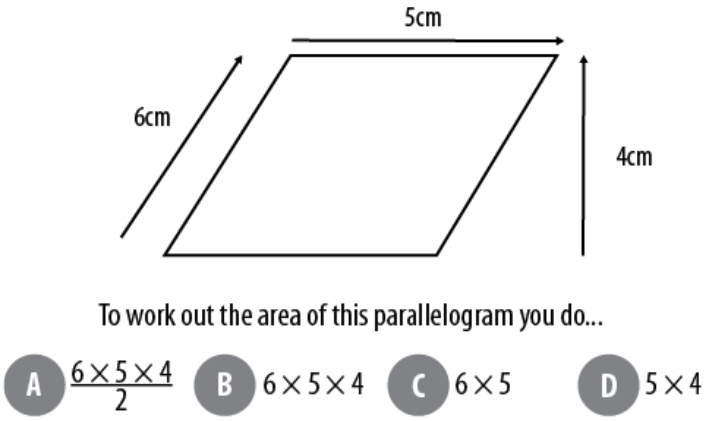


Figure 11.30 – Source: Craig Barton for Diagnostic Questions

The most popular misconception for all students is A (18%), with student explanations revealing that many are confused with the formula for working out the area of a triangle. However, for 11- and 12-year-old students, the most popular misconception is B, with students explaining that to work out area you just multiply things together (obviously). Interestingly, 11- and 12-year-olds perform better on this question than 15- and 16-year-olds, and I will leave you, dear reader, to speculate as to why that might be.

As teachers, the most prominent misconceptions we believe our students to hold may not be the same as the ones they do in fact hold. Does this matter? Well yes, it does. As I hope I have shown in this chapter, it is not the case the students are either right or wrong. They are right, or they are wrong for very different reasons, and it is the underlying cause behind their errors that dictates the specific type of support and intervention that they are likely to benefit from. Not everyone needs the same help. If we assume students will go wrong for one reason, and in fact they go wrong for another, then the help we offer may not be as effective as it could be.

So, how do we overcome this curse of knowledge, and apparent inability to anticipate where our students will go wrong? In short, we do what Lemov (2015) calls 'plan for error' – being aware of the incorrect answers students are likely to give in advance means you are better placed to deal with them. My favourite way of doing this (surprise, surprise) is to use diagnostic questions.

A huge advantage of using diagnostic questions is that you can plan for a finite set of errors in a way you cannot with non-multiple-choice questions. You can consider each incorrect response, why a student may select it, and how you intend to deal with it, preparing resources, demonstrations and explanations as needed. This means more of your important thinking can be done outside the pressures and heat of the classroom environment.

Of course, that does mean the questions need to be really well designed. There are two ways of increasing the likelihood of that happening.

The first is to write questions together with other teachers. I am going to discuss additional benefits to doing so in the next section, but it is enough to say here that writing a question and then challenging colleagues to list as many possible wrong answers as they can increases the chances of misconceptions being anticipated.

The second is to use my Diagnostic Questions website. Now, I promise I am not being a dodgy salesman here, because the website is 100% free to teachers and schools and always will be, but I use the website to help me plan every single lesson I teach. Say I am teaching expanding single brackets. I will search for all the expanding single bracket questions but order them by 'most misconceptions'. I will then study a few of these questions, looking at the proportion of students opting for each answer, and reading their explanations where needed. This enables me to get a sense of the kinds of questions students are likely to find difficult – this is related to the need to cover as much of the example space as possible, which we discussed in Chapter 7 – and where exactly students are likely to go wrong. Hence, I can anticipate the mistakes students may make and the misconceptions they are likely to hold before I teach the lesson.

And what do we do once we have anticipated a misconception? Well, we plan for it. We arm ourselves with explanations and relevant materials. This is not easy, but it is far easier to do in the comfort of our own homes, desks or the staff room, with colleagues on hand to help, than being alone in the heat of the moment, with 30 confused faces staring expectantly at us.

What I do now

When planning a lesson I ask myself: *Where do I want my students to get to, and what mistakes and misconceptions could stop them getting there?* I list as many misconceptions as possible, asking colleagues and using my Diagnostic Questions website. I then write or choose questions to assess these misconceptions, and I prepare to deal with the responses.

OceanofPDF.com

11.9. The benefits of teachers writing questions

What I used to think

I used to think the most effective way of teachers working together was to jointly plan lessons. Hence, each year when Year 11s and Year 13s had finished their exams and we had some 'gained time', us maths teachers would pair up to plan lessons and resources together for the next school year.

The problem was, many of the heated discussions that followed were rarely about pedagogy…

Sources of inspiration

- Coe, R., Aloisi, C., Higgins, S. and Major, L. E. (2014) *What makes great teaching? Review of the underpinning research*. Available at: https://www.suttontrust.com/wp-content/uploads/2014/10/What-Makes-Great-Teaching-REPORT.pdf.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

My takeaway

I am going to say something potentially controversial here – I think the joint planning of lessons is a bit of a waste of time. I find that it often leads to arguments and disagreements, but not necessarily about the right things. The problem is that teachers are very different:

- Some like group work, some don't.
- Some like worksheets, some don't.
- Some use PowerPoint, some don't.

- Some use animations, some don't.
- Some like Calibri, some like Comic Sans.

Squabbling over these matters is not the most effective way to use teachers' limited time. Instead, I believe teachers should write questions together.

We saw in the previous section how difficult it can be for teachers acting alone – whether they have many years of experience, or are just starting out on their careers – to anticipate the mistakes and misconceptions that students may have, and why it is so important to do so. Creating diagnostic questions together in the manner I will describe soon can help make this task a lot easier.

But that is not all. Here are some additional benefits to teachers writing questions together:

Good questions are transferable
Teachers often say to me, 'Have you got anything good for teaching (insert topic)?', and I think to myself: *Well, I have got something that works quite well for me, with my teaching style, with the knowledge of my students, the time of day I taught it, and so on, but it might not necessarily work for you*. Likewise, my colleague could teach the best lesson in the world, and in my hands, without significant adaptation, it could be a disaster. The problem is, a good lesson is not necessarily a good lesson for everyone, but a good question is, and once you have a good question it lasts a lifetime.

Good questions are not technology dependent
If you have teachers in your department who are whizzes on IT, then they can get out the iPads, voting devices or whatever else they like if that is how they choose to ask questions and collect the data. Likewise, if you have teachers who have been known to write on the expensive interactive whiteboard with a permanent marker, then they can print out or write up the question and get students voting with their fingers.

Good questions are not style dependent
As I hope I have made clear throughout this book, my teaching has undergone a significant transformation over the last few years. My lessons

are much more teacher-led, favouring a model of explicit instruction with a more careful use of rich tasks and inquiry-based lessons. But that does not mean I will benefit any more or any less than my colleague who loves projects, investigations, group work and so on. We can both ask the same questions, both get the same results, and deal with them as we choose. This is certainly not the case when asking teachers to all use the same activity.

Good questions help promote consistency in teaching and assessing
Related to all this is the fact that good questions ensure a certain amount of consistency in teaching and assessing. If you are in the kind of department where you have freedom how you plan and deliver your lessons, then ensuring all teachers are asking the same questions at certain stages of their lessons or throughout the topic unit ensures that everyone has a good idea of where students need to be heading.

Their creation promotes positive discussions between colleagues
I have already outlined some of the problems I have encountered when colleagues are asked to plan lessons together. This does not tend to happen with questions. Perhaps it is the fact that planning a question seems a smaller, more manageable task than planning a lesson, or perhaps there is simply less room for non-pedagogical considerations to get in the way. What is particularly nice is the fact that all teachers can get involved. I worked with one teacher who was vastly experienced, but for many reasons he felt a little out of place in our department, and as such rarely contributed to discussions in our weekly meetings. However, when asked to write and comment on questions, he came to life. He was able to call upon his 30-plus years of teaching experience, and share insights into where he had seen students go wrong in the past that us less-experienced teachers could simply not have come up with on our own. The discussion of shared experiences between teachers in a focused, positive manner is my favourite thing about teachers creating questions together. If I wasn't concerned how corny it sounded, I might be tempted to describe it as beautiful.

Their creation is particularly useful for less-experienced teachers
I return to the point I made at the beginning. As a novice teacher I simply had no idea of the types of mistakes students made and the misconceptions they held. It was only when talking to colleagues that I was able to learn

from their experience, and I believe the joint creation of questions in the way I am about to describe is the most effective way to tap into these benefits.

What I do now

There is a real danger that maths departmental meetings can get bogged down with administrative tasks. It is a crime to have such a wealth of knowledge in the room together and not use it effectively. Hence, we always free up time to plan and write diagnostic questions together.

The process goes like this:

1. We choose a topic that is coming up on the scheme of work for a particular year group. We tend to do this on a rota system, so Year 7 one week, Year 8 the next, and so on. This ensures that all staff have the opportunity to use the questions created, and allows for a wider range of topics to be covered.

2. One of us chooses a question from my Diagnostic Questions website on that particular topic and we play *Guess the Misconception*. Everyone decides what they think is the most popular incorrect answer, and we have a quick vote and a discussion. This gets everyone in the right frame of mind to start thinking about misconceptions. It is effectively *priming*, but for teachers.

3. Then everyone creates a diagnostic question on one piece of paper for that given topic, and their four answers (one correct and three incorrect) on another piece of paper. If we have been organised, the topic can be emailed out in advance of the meeting, and teachers can bring their completed question and answers along with them to save precious minutes.

4. People pair up and swap the piece of paper with the question on it with their partner. It is their partner's job to write down what four answers they would choose to go with this question.

5. The pair then discuss both questions and try to come up with an agreed set of four answers for each question that they are happy with.

6. Both questions and answers are then given to another pair who evaluate them alongside my Five Golden Rules from earlier in this chapter. The most important one to consider is always the last – is it possible to get the questions correct whilst still holding a key misconception?

7. The pair gives feedback on the questions and passes each back to the original author who may decide to make some tweaks.

8. The final version of the questions can either be scanned or turned into PowerPoints and shared with the rest of the department for use in lessons, homeworks or low-stakes quizzes.

Once teachers have done this a few times, the whole process takes around 10 to 15 minutes. I am ridiculously biased, but I believe there is no better way to keep the discussion positive, and focus it on the things that really matter for teaching and learning. There is simply no better way to spend precious time in departmental meetings.

OceanofPDF.com

11.10. If I only remember 3 things…

1. Without an effective formative assessment strategy, we are in danger of teaching blindly, being completely unresponsive to the needs of our students.
2. Diagnostic multiple-choice questions offer a way to allow teachers to identify and understand students' misconceptions in an efficient, accurate manner.
3. The joint creation of diagnostic questions can be a valuable, positive experience for all teachers, and is a great feature to build into departmental meetings.

OceanofPDF.com

12. LONG-TERM MEMORY AND DESIRABLE DIFFICULTIES

Have you ever taught something one lesson, been pretty sure the students have understood it, and then returned the next day to find it was as if the previous day never even happened? I have had a Year 10 student swear blind to me that they had never seen 'one of these weird wiggly graphs' before, despite this being their third lesson on cumulative frequency diagrams.

Such experiences are soul-destroying for all involved, and occurred all too regularly during the first ten years of my career.

Most of the issues tackled so far in this book have been concerned with the early knowledge acquisition phase of learning. We can use the principles of Cognitive Load Theory to present material in the most effective way. We can break down processes using the principles of Deliberate Practice. We can intelligently choose and present our examples and practice questions to students. And we can help them develop the ability to solve problems via the careful batching of related problems, combined with the power of Purposeful Practice.

That is all well and good. But if students cannot retain that knowledge, then we are wasting our time.

In order to tackle this issue, we are going to look at explicitly and deliberately making the action of thinking *difficult* for students. We will see exactly what such difficulties entail, why they are likely to be beneficial for learning, and consider how best to apply the principles both in and out of our maths lessons.

OceanofPDF.com

12.1. How long-term memory works

What I used to think

Back in Chapter 1, I made the rather worrying confession that for much of my career I had not really considered exactly how students think and learn. My ignorance in these fields was matched only by my ignorance when it came to my students' memories. I did not really have a clue how memory worked. I assumed some students simply had a better memory than others, and that there was little we could do to influence that. Likewise, despite years of frustration over students' apparent lack of ability to retain and recall information from one lesson to the next, I appeared powerless to be able to do anything to change that.

Thankfully, the Bjorks came into my life.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real world: essays illustrating fundamental contributions to society*. 2nd edition. New York, NY: pp. 59-68.
- Bjork, R. A. (1975) 'Retrieval as a memory modifier' in Solso, R. L. (ed.) *Information processing and cognition: the Loyola symposium*. Mahwah, NJ: L. Erlbaum Associates, pp 123-144.
- Bjork, R. A. (2011) 'On the symbiosis of remembering, forgetting, and learning' in Benjamin, A. S. (ed.) *Successful remembering and successful*

forgetting: a festschrift in honor of Robert A. Bjork. New York, NY: Psychology Press, pp. 1-22.

- Yan, V. (2016) 'Retrieval strength *vs* storage strength', *Learning Scientists Blog*. Available at: http://www.learningscientists.org/blog/2016/5/10-1

My takeaway

The new theory of disuse

In Chapter 1, I introduced the following very simple model of how students think and learn which has served us well throughout this book. Now, the eagle-eyed amongst you may have noticed something. Whilst there is a downward arrow coming from working memory – we have seen in Chapters 1 and 4 that working memory has a finite capacity, and once that capacity is exceeded, items being processed may be forced out – there is no such arrow from long-term memory. This is because Robert and Elizabeth Bjork's *New Theory of Disuse* (for example, see Bjork, 2011) is built on the assumption that instead of memories decaying and disappearing from long-term memory after prolonged periods of disuse, they remain in long-term memory but simply become less accessible.

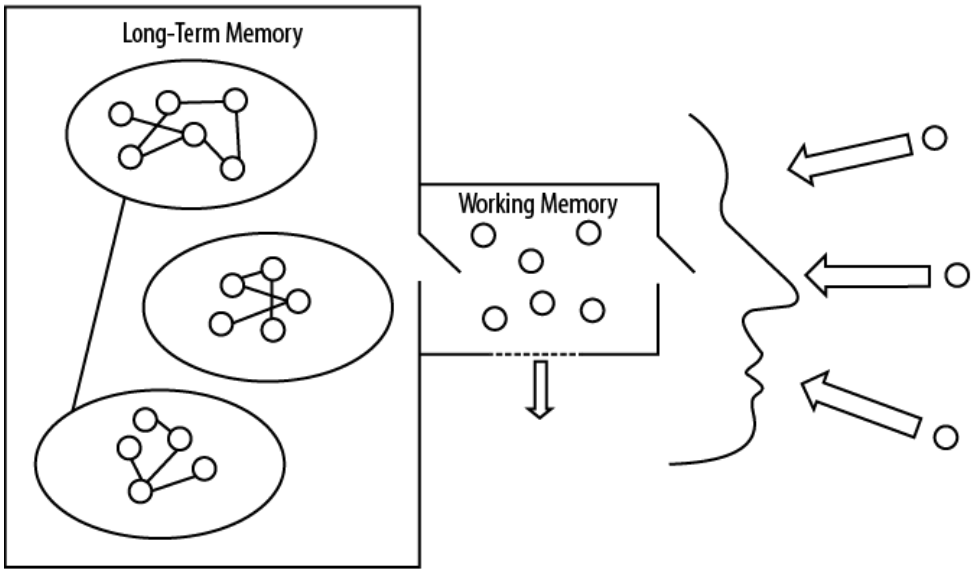


Figure 12.1 – Source: Craig Barton

This is a subtle, but important difference. It has immediate implications for our discussion of misconceptions in Chapters 2 and 11 – if the Bjorks'

theory is true, then we may never be able to fully rid students' memories of erroneous beliefs, and our only hope is to make them less accessible relative to the correct way of thinking. But it has much wider implications than that.

Retrieval and storage strength

In order to better understand the accessibility of the content of long-term memory, we need to consider the retrieval and storage strength of memories. Bjork and Bjork (2014) explain that each item in long-term memory has a retrieval and a storage strength.

Retrieval strength is a measure of how easily recalled something is currently – in other words, can it be recalled in the present moment? The retrieval strength of a given item can be high or low, and can fluctuate back and forth in between. Retrieval strength is measured by current performance, such as answering questions on a test or in class.

Storage strength is a measure of whether information is deeply embedded or well learned – in other words, can it be recalled later? Barring organic brain damage, the New Theory of Disuse is built on the assumption that storage strength cannot decrease, rather it is presumed to only accumulate. Storage strength cannot be directly measured, and so must be inferred. Is that information easily recalled at some point in the future? Or, if you forget that information, does it become faster to relearn it the next time?

The following diagram may help illustrate the difference:

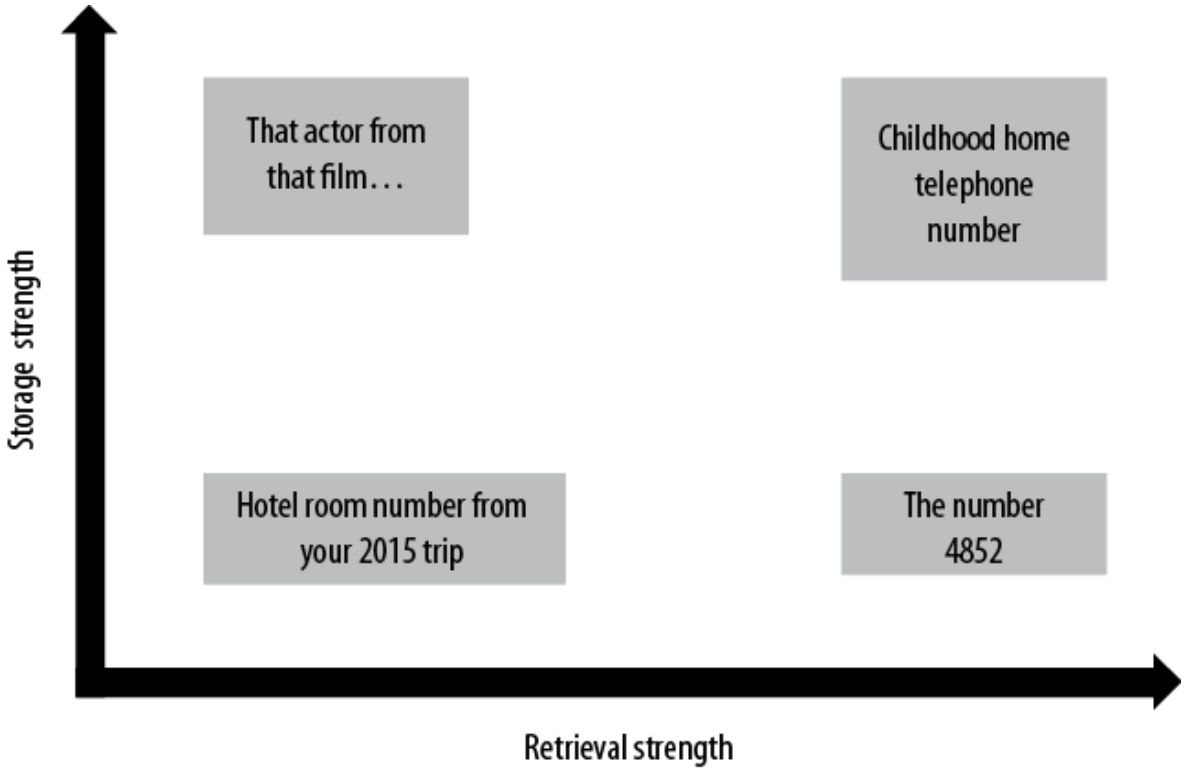


Figure 12.2 – Source: Craig Barton

Ideally, we want relevant and correct items in our students' long-term memories to have both high retrieval and high storage strengths, and that can be achieved by restudying the information.

When I interviewed Robert and Elizabeth Bjork for my podcast, I asked them what is the biggest misconception people have about memory. Robert explained that it was assuming memory worked like a recording device, say a hard-drive on a computer. When you retrieve some information from your computer's hard-drive, that information is not fundamentally changed in any way. Human memory does not work in the same manner. The retrieved information becomes more recallable than it would have been otherwise, and other information in competition with the retrieved information – in other words, information associated to the same retrieval cue or set of cues – becomes less accessible. Using our memories changes our memories as both retrieval and storage strength increase. Or, as Bjork (1975) put it, '**retrieval is a powerful memory modifier**'.

But there is a twist. Yan (2016) explains the dynamics of how retrieval strength and storage strength are related to each other with respect to learning:

- When we study information, both retrieval strength and storage strength increase.
- After studying that information, the higher the storage strength, the slower the loss of retrieval strength (*ie* slower forgetting).
- When revisiting information, how much storage strength increases is inversely related to the current retrieval strength: the lower the retrieval strength at that moment, the greater the boost in storage strength (*ie* long-term learning).

As we shall see, the time we revisit that information, how we revisit it, and what happens in between are crucial for the development of long-term learning.

The importance of forgetting

For the Bjorks, forgetting means items of memory are less accessible. In other words, the retrieval strength of a memory has decreased. This is natural. As time passes without the memory being accessed, its retrieval strength falls. Likewise, as we learn new information, procedures and skills, we create the potential for competition with related information, procedures and skills that already exist in long-term memory. Hence, access to that earlier knowledge can then be inhibited or blocked by related aspects of the newer, more accessible, knowledge.

So why is it important that we forget? Because by giving our memories a chance to suffer a decrease in retrieval strength, we can then enjoy a bigger boost in storage strength when we subsequently retrieve them. Therefore, the best way to remember in the long-term is to first start to forget.

But we can go too far. If the act of retrieval occurs too long after we have first learned the material, it may have become so inaccessible that we are unable to retrieve it at all. It is this delicate balancing act that is the key to making the most out of the wonderful workings of memory.

What I do now

I allow students to have time to start to forget the things I have taught following the techniques I will describe in this chapter, and I then induce retrieval. This act of retrieval could take the form of restudying the material in revision lesson, but as we shall see in Section 12.6, it is far more effective to test.

Finally, I make sure I tell my students what I am doing and why, because revisiting familiar material feels good, whereas forgetting does not!

OceanofPDF.com

12.2. The problem with performance

What I used to think

I used to think that student performance was a very good indicator of student learning, and that attempting to boost my students' performance could only be a good thing. Hence, I would create conditions (consciously or unconsciously) that were favourable to boosting performance.

Such conditions would include:

- Massed practice of a specific skill or concept so students get comfortable with it
- Sticking to one skill or concept at a time so students do not get confused
- Ensuring classroom conditions are as similar as possible to develop consistency
- Only giving students tests when I needed to
- Giving my students detailed feedback on exactly where they have gone wrong and how best to improve

The interesting/depressing thing was that students would often perform very well at the time, but then the very next day/week/month it was if they had never learned anything at all.

Sources of inspiration

- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real*

world: essays illustrating fundamental contributions to society. 2nd edition. New York, NY: pp. 59-68.

- Soderstrom, N. C. and Bjork, R. A. (2015) 'Learning versus performance: an integrative review', *Perspectives on Psychological Science* 10 (2) pp. 176-199.

My takeaway

It is no exaggeration to say that the Soderstrom and Bjork (2015) paper has changed my whole approach to teaching. The main premise is one of those things that sounds so obvious when you say it out loud, but I had been missing it for much of the 12 years of my career: *learning and performance are two different things*.

Performance is what we can observe and measure during teaching. Learning – what we are defining to be a change in long-term memory, and which is arguably the ultimate goal of teaching – is something we must try to infer. More often than not, current performance is the easiest way of doing this – a student answers a particular question correctly, or scores well on an end-of-topic test, and we conclude that they have learned what we needed them to.

That would all be fine, apart from the rather inconvenient and widely replicated finding that current performance can be a highly unreliable guide to whether learning has actually happened.

During lessons, teachers – consciously or otherwise – create conditions designed to boost performance. We provide cues and prompts to elicit correct answers to our questions. Students are also pretty good at mimicking, which allows them to say what they think the teacher wants to hear. We present topics in nice tidy blocks, keep class conditions consistent, don't give too many tests, and provide plenty of feedback. All of this is likely to keep performance flying high, but may have no impact on learning as nothing may have changed in long-term memory.

But it gets worse than that – when we design conditions to boost students' performance, we may actually hinder their learning. By not giving students an opportunity to begin to forget, we do not allow items in long-term

memory to gain the necessary boost in storage strength required for long-term learning. Likewise, conditions that appear to slow down performance or induce the most errors are often the conditions that lead to the most learning, precisely because students must work hard to retrieve items from long-term memory, thus fundamentally changing it.

Hence, as teachers we are at risk of being fooled by current performance. This can lead us to choose less effective conditions of learning – the ones that lead to a short-term boost in performance – over conditions more effective for learning. Moreover, it can lead students to prefer certain revision techniques, such as cramming, over others as they can see short-term improvements in performance.

Soderstrom and Bjork (2015) summarise this nicely:

> Given that the goal of instruction and practice – whether in the classroom or on the field – should be to facilitate learning, instructors and students need to appreciate the distinction between learning and performance and understand that expediting acquisition performance today does not necessarily translate into the type of learning that will be evident tomorrow. On the contrary, conditions that slow or induce more errors during instruction often lead to better long-term learning outcomes, and thus instructors and students, however disinclined to do so, should consider abandoning the path of least resistance with respect to their own teaching and study strategies. After all, educational interventions should be based on evidence, not on historical use or intuition.

What I do now

I purposefully create conditions designed to decrease performance in the short term, but increase learning in the long term. Therefore:

- Instead of massed practice, I use spacing.
- Instead of teaching skills in separate blocks, I interleave.
- Instead of keeping conditions in the classroom consistent, I vary where appropriate.

- Instead of testing only when I have to, I use a low-stakes quiz every single lesson.
- Instead of always giving lots of feedback, I delay or reduce feedback.

These are what Robert and Elizabeth Bjork call 'desirable difficulties'. They are conditions that make learning more difficult, but in a good way. They need to be distinguished from the various forms of extraneous load that we discussed in Chapter 4 – they too make learning more difficult, but in an entirely unhelpful way. Bjork and Bjork (2014) explain: 'desirable difficulties, *versus* the array of undesirable difficulties, are desirable because they trigger encoding and retrieval processes that support learning, comprehension, and remembering'. It is the very fact that retrieving information from long-term memory – particularly when it is not readily accessible – is effortful that makes these difficulties so beneficial to learning.

Just before we enter the world of desirable difficulties, a word of warning. All these techniques are designed to improve learning in the long-term, but at the necessary costs of a decrease in short-term performance. They feel more difficult and hence less appealing than their alternatives. This is fine, but that short-term lack of success may have a negative impact on students' motivation, as discussed in Chapter 2. Hence, with all the following techniques, I would advise caution.

First, they should only be used once students have acquired a certain amount of domain-specific knowledge using the techniques described so far in this book. After all, desirable difficulties work on the basis that they induce effortful retrieval that brings about a change in long-term memory. However, if the material is pitched at a level that is cognitively overwhelming, and thus students do not have a chance to transfer the knowledge from working memory to long-term memory in the first place, then there will be nothing to retrieve. Bjork and Bjork (2014) make the very same point, explaining that if 'the learner does not have the background knowledge or skills to respond to them successfully, they become undesirable difficulties'.

Second, I would strongly advise informing students exactly what you are doing and why you are doing it. That way, they may be better placed psychologically to cope with the inevitable short-term decreases in performance, especially if they see their fellow peers in other classes appearing to do better than them.

OceanofPDF.com

12.3. The Spacing Effect

What I used to think

I used to think that the best thing to do was to give students lots of practice on a topic all together so they could master it. Hence, I would teach topics in blocks, only returning to them during revision for exams.

I was missing out on one of the easiest imaginable ways to improve learning.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Benney, D. (2016) 'Optimal Time For Spacing Gaps', *mrbenney blog*. Available at: https://mrbenney.wordpress.com/2016/11/03/optimal-time-for-spacing-gaps/
- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real world: essays illustrating fundamental contributions to society*. 2nd edition. New York, NY: pp. 59-68.
- Carpenter, S. K., Cepeda, N. J., Rohrer, D., Kang, S. H. K. and Pashler, H. (2012) 'Using spacing to enhance diverse forms of learning: review of recent research and implications for instruction', *Educational Psychology Review* 24 (3) pp. 369-378.
- Cepeda, N. J., Vul, E., Rohrer, D., Wixted, J. T. and Pashler, H. (2008) 'Spacing effects in learning: a temporal ridgeline of optimal retention', *Psychological Science* 19 (11) pp. 1095-1102.

- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J. and Willingham, D. T. (2013) 'Improving students' learning with effective learning techniques: promising directions from cognitive and educational psychology', *Psychological Science in the Public Interest* 14 (1) pp. 4-58.
- Willingham, D. T. (2002) 'Ask the cognitive scientist. Allocating student study time: "massed" *versus* "distributed" practice', *American Educator* 26 (2) pp. 37-39.

My takeaway

What is the Spacing Effect?
The Spacing Effect refers to the finding that learning is better when two or more exposures to information are separated in time (*ie* spaced apart) compared to when the same number of exposures occur back-to-back in immediate succession. Given our work above, this should be of little surprise. The spacing of the sessions causes students to begin the process of forgetting the material, which lowers the retrieval strength. Therefore, when retrieval is induced at a later date, the material benefits from a larger boost in storage strength.

The Spacing Effect has been identified and replicated across many domains. Indeed, Bjork and Bjork (2014) describe it as 'one of the most general and robust effects from across the entire history of experimental research on learning and memory'. Willingham (2002) concurs, concluding: 'there is a mountain of evidence suggesting that spacing study time leads to better memory of the material; the effect applies to at least some of the types of learning students do – fact learning; and it seems to hold for school-age children'.

The concern is always the transferability from the laboratory to the classroom, and indeed from simple factual recall tests to the more complex processes involved in mathematics. Fortunately, the effect seems to generalise. In their review of effective learning techniques, Dunlosky *et al* (2013) conclude:

> On the basis of the available evidence, we rate distributed practice as having high utility: it works across students of different ages, with a wide variety of materials, on the majority of standard laboratory measures, and over long delays. It is easy to implement (although it may require some training) and has been used successfully in a number of classroom studies. Although less research has examined distributed-practice effects using complex materials, the existing classroom studies have suggested that distributed practice should work for complex materials as well.

What are the implications for teaching?
In a comprehensive review of research into spacing, Carpenter *et al* (2012) offer three pieces of practical advice for teachers wanting their students to take advantage of the Spacing Effect:

1. Teachers should dedicate part of each lesson to reviewing concepts learned several weeks earlier
The start of the lesson is usually the most convenient time to do this. In the past, I have been rather casual about my choice of topics – we haven't done dividing fractions for a while, so let's chuck one of those into the mix. But now I am more methodical. I make a list of all the skills students should know – a quick glance at previous years' schemes of work is a good way to discern this – and ensure I have them on a rotation schedule so nothing gets missed out.

In terms of the questions themselves, my go-to sources are either diagnostic questions as discussed in Chapter 11, or CorbettMaths's incredible '5-a-days'. These can be projected up onto the board as students enter the room, or printed out and left on their desks so they can crack right on. Alternatively – and more and more so these days – I might choose to cover previous material by means of a low-stakes quiz, which will be discussed later in this chapter.

2. Homework assignments should be used to re-expose students to important information they have learned previously
A couple of years ago our maths department completely rewrote all our homeworks. Each year group was to be provided with a lovely 30-mark

fortnightly homework, where 10 marks were dedicated to revision of previous topics, 10 marks were skill-based questions on the current topic, and 10 marks were application questions of the current topic. I will return to the issue of the application questions in the next section on Interleaving, but the mistake I now realised we made was to only have ⅓ of the homework testing prior knowledge. The revision section should have been the dominant section. Maybe 20 marks for revision and 10 marks for current would have been a better split.

The beauty is, such a change is not that hard to make. No new questions need to be written; it is merely a case of taking questions that used to be in the current topic section of previous homeworks and moving them to the revision sections of current homeworks. That way twice as many topics are being covered every two weeks. Students will inevitably find these homeworks harder and their performance will suffer, precisely because the act of retrieval is more effortful. But they will be learning more.

3. Teachers should give exams and quizzes that are cumulative
If we give our students topic-specific tests (as I have done for years!), then two things happen:

1. There is no opportunity to reap the benefits of the Spacing Effect. Students simply have no time to begin to forget the skills they have been taught, and hence the high retrieval strength during the test interferes with the more preferable development of storage strength.

2. It conveys the message that cramming is an effective revision strategy. If students know the relatively narrow range of concepts and skills that will be tested, they have an opportunity to cram their revision. This may involve pulling an all-nighter, or even merely glancing at their notes in the minutes before the test starts. This will boost the retrieval strength of the material that they are studying and, so long as the gap until the test starts is short enough, will probably be enough to make it recallable when needed. These students perform well on the test, assume they have learned the material well, and hence assume that their revision technique is effective. The problem is, as soon as they face a less predictable exam (such as a GCSE, A level or end-of-year test),

the lack of storage strength of the knowledge they need lets students down.

We can help students by ensuring that the tests they receive throughout the year are not topic-specific, but cumulative. They should test knowledge, not just from the last few weeks, but from the previous weeks, months and years. Sure, students will perform worse at them, but the effortful act of retrieving the information will be beneficial for their long-term learning. Just as importantly, it will highlight to students what they actually have learned and what revision strategies really are effective. Once again, I believe low-stakes quizzes are the key to this.

Is there an optimal spacing schedule?

This is obviously a key question to make all three strategies above as effective as possible. Unfortunately there is no straightforward answer. This is no surprise, given the number of variables involved, including the complexity of the material to be learned, the prior knowledge of the students, the conditions of learning, and so on.

However, there has been a body of interesting research conducted into optimal spacing schedules for a final test. Cepeda *et al* (2008) found that the longer the period between the initial learning episode and the time when you wish to recall the information, the longer the optimal spacing interval, assuming you only plan to revisit this material once. The authors go further, devising a formula to calculate an optimal spacing schedule. The research looked to trial a number of different gaps with a number of different retention intervals (RI). The gap + the RI = the total time between initial study and test.

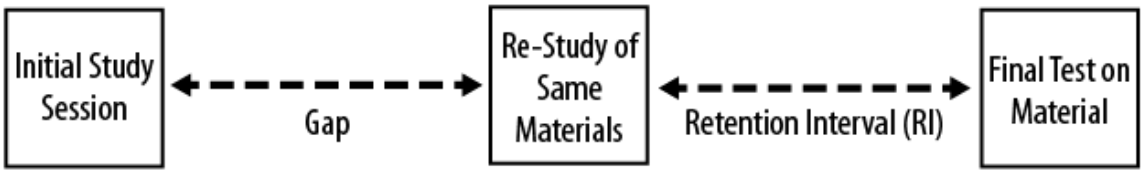


Figure 12.3 – Source: Cepeda et al (2008)

Cepeda *et al* (2008) did not find any magic ratio between the gap and the retention interval. As predicted, as the retention interval increased, so too

did the optimal gap for retesting, but not in a constant fashion. Specifically, as RI increases, the ratio of optimal gap to RI declines.

Benney (2016) has written a wonderful blog post with the following graph plotting optimal spacing intervals based on Capeda *et al*'s data, making it easy to discern some guidance.

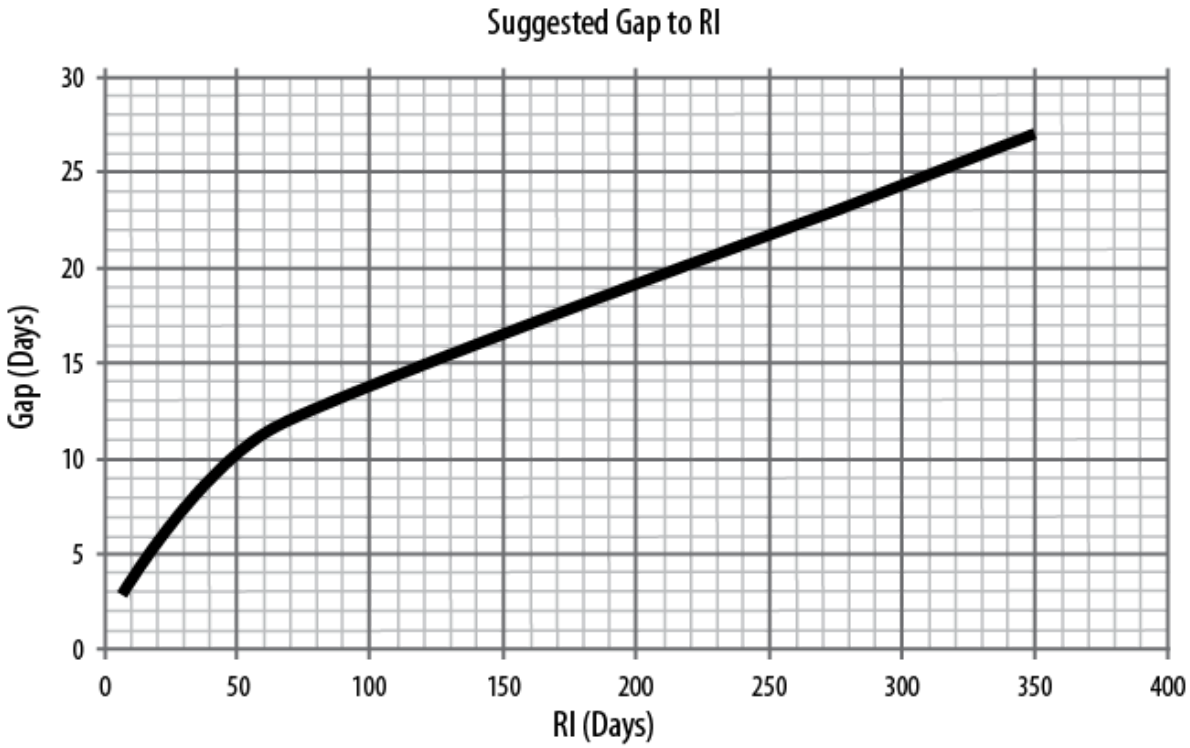


Figure 12.4 – Source: Damian Benney, available at https://mrbenney.wordpress.com/2016/11/03/optimal-time-for-spacing-gaps/

For example, if you finish a topic today and the test is in 60 days then (perhaps) your optimum gap before restudy would be 11 days, which leaves an RI of roughly 50 days (a 1:5 ratio or gap being 20% of RI). However, if your class finishes a topic today and the test is in 114 days then (perhaps) your optimum gap before restudy would be 14 days, which leaves an RI of 100 days (a 7:50 ratio or gap being 14% of RI).

This is not surprising if we assume that the optimal time to retest someone is when they are on the verge of forgetting. As Benney explains:

> It seems that it is preferable to have the restudy session within a shortish time of the original study despite the fact it gives a very long RI. This is surely because if the gap was any bigger too much of the original study material would have been forgotten and retrieval strength would be practically zero. It is better to keep a relatively short gap and trade off with a very long RI. I imagine a spacing session much later than the optimum gap would be more like a reteaching lesson rather than a restudy/recall lesson.

However, as Capeda *et al* (2008) point out, there is a danger in going too far the other way. The compression of learning into too short a period is likely to produce misleadingly high levels of immediate mastery (*ie* performance) that will not survive the passage of substantial periods of time.

Students need time to be on the verge of forgetting in order for the Spacing Effect to display its full power, but this can be tricky – if not impossible – to judge. In my podcast interview, I asked Robert and Elizabeth Bjork whether it was better to have too short or too long a spacing interval. They were quick to answer: too short always wins because of the danger of the knowledge becoming completely inaccessible. However, having read an early draft of this chapter, they clarified their point further:

> Too short would definitely be our answer if we were talking about the spacing of tests with no feedback. You don't want to make the first interval (gap) so long that the to-be-tested material has become completely inaccessible. When we are talking about re-study opportunities rather than testing, however, then it is probably better to err on the side of being too long. That is, even if most or all of the material cannot be recalled by the student, given that it is now being made available for re-study, that study opportunity could amount to a very powerful learning event.

A final point to make is that this study is based on the assumption that material will only be revisited once. In the classroom environment we can make much more use of the power of the Spacing Effect (and the related Interleaving Effect that follows) by regularly revisiting material via starters, homeworks and low-stakes quizzes.

What I do now

I follow the three practical suggestions from the Carpenter *et al* (2012) paper discussed above. However, I am not blind to the attraction – and potential benefits – of cramming. The day before a GCSE exam I will hold an exam clinic for my students, and if students are struggling on a particular topic, I will teach them and give them lots of practice. Likewise, when students are preparing for an exam, the days and hours before are likely to involve cramming. And this strategy can be effective, as it boosts the retrieval strength of the crammed information, which makes it easily accessible in the short-term.

However, the key point is that such a strategy is not beneficial for long-term learning. As Bjork and Bjork (2014) point out, 'little of what was recallable on the test will remain recallable over time'. As discussed above, if we keep giving students practice tests throughout the year where cramming can be an effective strategy to perform well, students will grow to depend on it. Cramming is fine as a one-off strategy before a high-stakes exam, but in the months and years before, students should be encouraged to recognise that a study schedule that spaces study sessions on a particular topic can produce both good exam performance and good long-term retention.

OceanofPDF.com

12.4. The Interleaving Effect

What I used to think

I used to think that topics should be taught, mastered and assessed in isolation from each other or students would get confused. Hence, I would do the following:

- teach topics in blocks.
- set topic-specific homeworks.
- use application questions related to the topic we had just studied.
- teach related concepts together.

Not for the first (or last) time in this book, I was quite astonishingly wrong.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real world: essays illustrating fundamental contributions to society*. 2nd edition. New York, NY: pp. 59-68.
- Emeny, W. (no date) 'The Elements: Interweaving', *Great Maths Teaching Ideas*. Available at: http://www.greatmathsteachingideas.com/the-elements-interweaving/

- Mayfield, K. H. and Chase, P. N. (2002) 'The effects of cumulative practice on mathematics problem solving', *Journal of Applied Behavior Analysis* 35 (2) pp. 105-123.
- Reddy, B. (2014) 'Design your own mastery curriculum', *Mr Reddy Blog*. Available at: http://mrreddy.com/blog/2014/03/design-your-own-mastery-curriculum-in-maths/
- Rohrer, D., Dedrick, R. F. and Burgess, K. (2014) 'The benefit of interleaved mathematics practice is not limited to superficially similar kinds of problems', *Psychonomic Bulletin & Review* 21 (5) pp. 1323-1330.
- Rohrer, D., Dedrick, R. F. and Stershic, S. (2015) 'Interleaved practice improves mathematics learning', *Journal of Educational Psychology* 107 (3) pp. 900-908.
- Yan, V. X., Soderstrom, N. C., Seneviratna, G. S., Bjork, E. L. and Bjork, R. A. (2017) 'How should exemplars be sequenced in inductive learning? Empirical evidence versus learners' opinions', *Journal of Experimental Psychology: Applied*. Advanced online publication.

My takeaway

What is the Interleaving Effect?
The Interleaving Effect contrasts a 'blocking' approach, whereby students study the same type of material over and over again before moving on to a different type of material, against an 'interleaving' approach, where students practise all of the problems in an order that is more random and less predictable. The latter approach has been found to enhance learning and transfer.

It is quite easy to get the concepts of interleaving and spacing mixed up (for me, anyway), and the two are inextricably linked as interleaving naturally produces spacing. A good way to think about it might be that spacing describes the scheduling of exposures to a single concept (A), and interleaving describes the scheduling of exposures to multiple concepts (A, B, and C).

It is relatively clear why spacing should enhance learning through the power of forgetting and subsequent retrieval discussed earlier in this chapter, but how about interleaving? Bjork and Bjork (2014) explain that one theory suggests that having to resolve the interference among the different things under study forces learners to notice similarities and differences among them, resulting in the encoding of higher-order representations, which then foster both retention and transfer. Another explanation suggests that interleaving forces learners to reload memories: if required to do A, then B, then C, and then A again, for example, the memory for how to do A must be reloaded a second time, whereas doing A and then A again does not involve the same kind of reloading. Such repeated reloadings are presumed to foster learning and transfer to the reloading that will be required when that knowledge is needed at a later time. Interleaving has an added advantage when it comes to mathematics – students get practice not just in how to carry out a strategy to solve a problem, but in also selecting the appropriate strategy in the first place.

Let's take a look at some of the research into the Interleaving Effect, and the practical implications for the classroom and beyond.

Content of lessons

Mayfield and Chase (2012) demonstrated the effects of interleaving within maths lessons. College students in need of mathematics intervention attended dozens of sessions over a period of several summer months in which they solved problems using five algebraic rules about the laws of indices. One group of students learned the rules through a procedure akin to blocking, in which each rule was learned and then practised extensively before moving on to the next rule. Another group learned the same rules through a procedure akin to interleaving, which involved continuous practice of previously learned skills. On subsequent tests, the interleaved practice group outscored the blocked practice group by factor of at least 1.3, both on skill-based questions, and crucially also on problem-solving.

On his *Great Maths Teaching Ideas* website, Will Emeny describes one application of interleaving in the content of maths lessons as 'interweaving'. He explains: 'Interweaving is my own term for building previously learned skills into questions on the current topic being studied. For example, if your

class have previously studied fractions and are now studying perimeter, make sure that they work on some perimeter questions featuring fractional lengths'. So, as opposed to completely chopping and changing between topics – one lesson on quadratic formula, the next on angles in parallel lines – previous learned concepts are skilfully integrated into the current concept being studied. Interleaving in this manner has the added advantage of showing students that maths is an integrated and connected subject, as opposed to the fragmented subject it can sometimes appear to be.

Curriculum design clearly plays an important role in the successful implementation of interleaving (or interweaving) within lessons. Schemes of work that are set up in such a way that previous topics or concepts can be seamlessly integrated into the practice of a new topic or concept lend themselves very well to this interleaving approach. An example would be solving linear equations. So long as negative numbers, fractions, and decimals have all been encountered before, these topics can be revisited by including them as coefficients or solutions to the linear equations. The very best mastery schemes are set up this way, and free resources such as Mr Taylor's 'Increasingly Difficult Questions' (available at http://taylorda01.weebly.com/increasingly-difficult-questions.html) are built on this principle (see Figure 12.5).

Likewise, one of my favourite sources of interleaved questions are those from the annual United Kingdom Mathematics Trust (UKMT) maths challenge. Consider Figure 12.6.

This question could be used at the end of a teaching unit on volume in order to revisit fractions of an amount as well as plans and elevations. These questions are available for free at www.ukmt.org.uk. They cover all age ranges and are brilliant to use once the core concepts have been covered.

But is there a danger we can interleave too much? After all, surely knowledge needs to be successfully transferred from working memory to long-term memory in order to benefit at all from the forgetting-retrieval process? And surely carefully introducing skills and concepts via worked examples and the careful selection of practice questions (as discussed in Chapters 6 and 7) is the best way to achieve this? We have seen in Chapters 4

and 9 that it is the very fact that such knowledge is organised, accessible and automated in long-term memory that allows experts to solve the kind of complex problems that novice learners cannot. Hence, surely a blocking approach – albeit it, a carefully planned form of blocking – is most effective in the initial knowledge acquisition phase?

Take the example of solving linear equations. Students would need sufficient conceptual understanding and procedural fluency of the basics of solving linear equations before interleaving previously studied concepts such as negative numbers, fractions and decimals. Otherwise there is a very real danger that our novice learner's fragile working memory will become overloaded and no learning will take place.

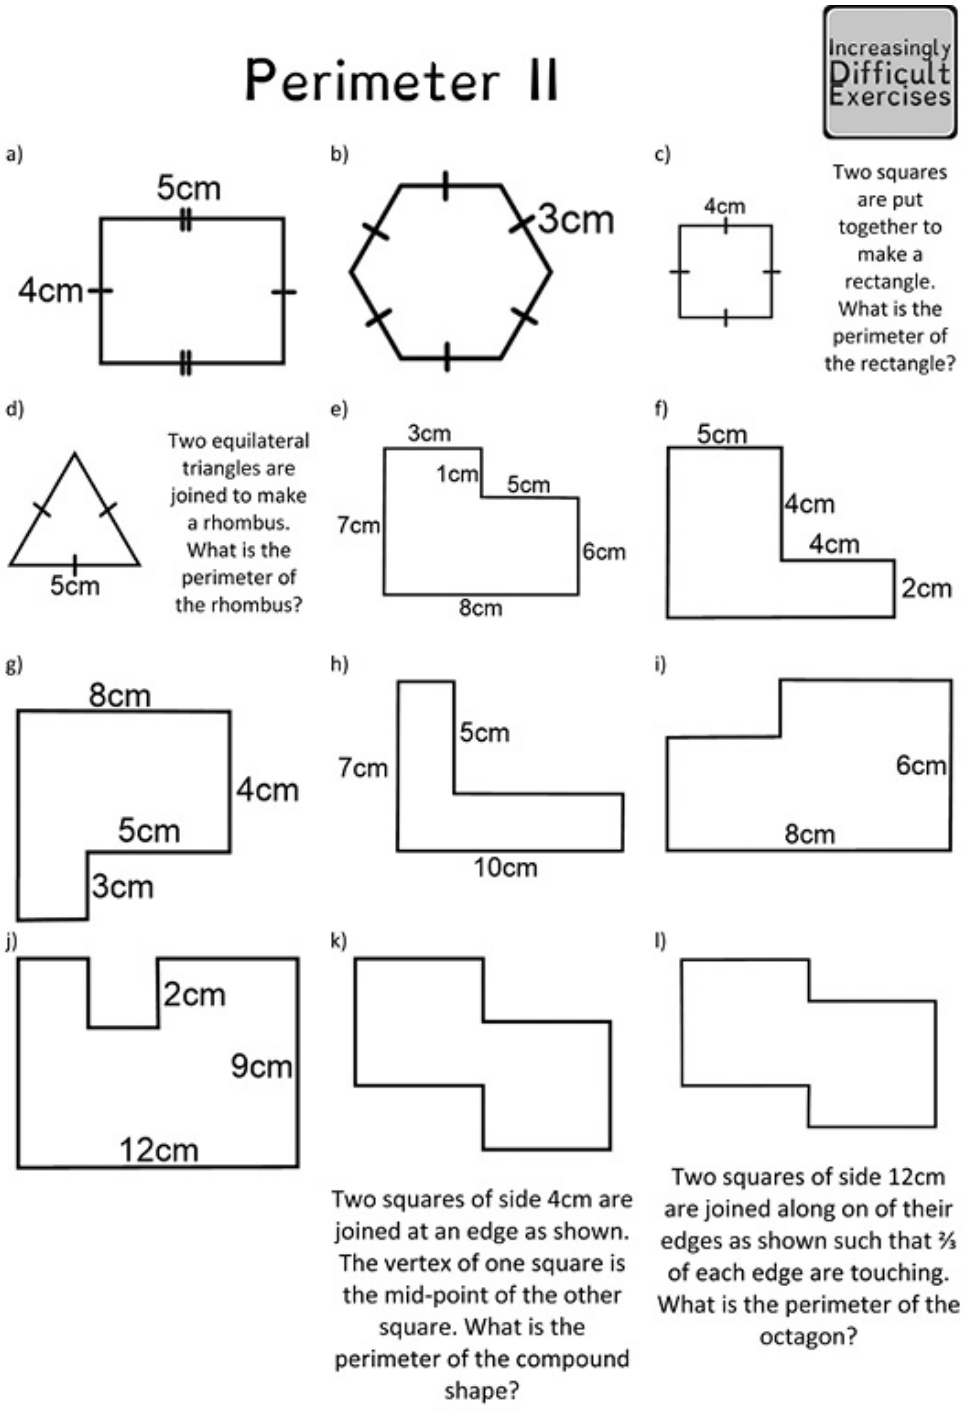


Figure 12.5 – Source: Mr Taylor, available at http://taylorda01.weebly.com/increasingly-difficult-questions.html

Yan *et al* (2017) lend support to this view. They compared pure blocked and pure interleaved approaches with three hybrid schedules – blocked-to-

interleaved, interleaved-to-blocked, and mini-blocks. They found that the blocked schedule led to the worst performance, but that mini-blocks and blocked-to-interleaved schedules were statistically as effective (but not better than) pure interleaving. Equally interesting was the perception of the participants themselves. The authors explain: 'Although participants demonstrated some metacognitive sophistication with respect to the relative benefits of blocked and interleaved study, pure interleaving was the least popular schedule, despite its being one of the most, effective schedules for learning'.

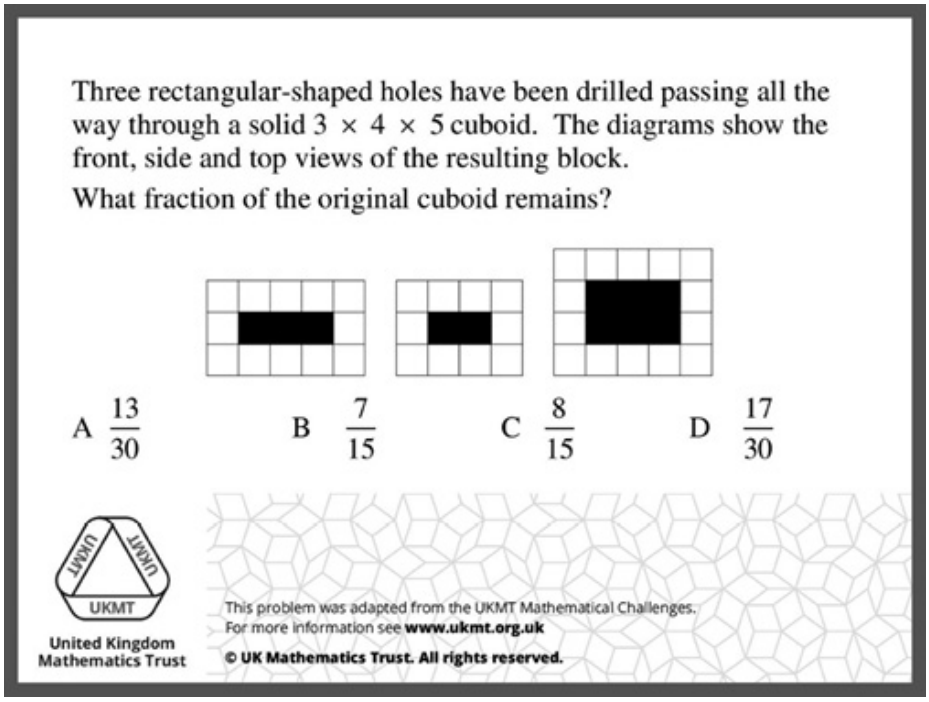


Figure 12.6 – Source: UKMT for Diagnostic Questions

In Chapter 2 we discussed the importance of students' perception of their ability to succeed. It seems that a hybrid approach – beginning with the careful selection and presentation of worked examples, followed by intelligent and purposeful practice – might give us the best of both worlds. It may be just as beneficial to long-term learning as pure interleaving, but with the added advantage that students taste the success needed to keep them working hard. And once you are satisfied that students are comfortable with the basics – the techniques involving diagnostic questions discussed in the previous chapter may be a good way to assess this – then interleaving other

elements can only be a positive thing in that students get a wider variety of practice of the current concept, as well as benefiting from the power of retrieval when accessing concepts they may not have encountered for quite some time.

One final point: as I discussed back in Chapter 3, I do not believe that interleaving should be used to interfere with the learning of a new skill or concept. I used the example of using measuring angles as a means of allowing students to discover circle theorems. Such an approach is fantastic for improving students' ability to measure angles – a skill that they may not have had to retrieve for sometime – but if students struggle with this then it interferes with their development and understanding of the new concept. I feel a better approach is to achieve a certain level of success or mastery in the new skill or concept first, before interleaving prior skills and concepts into the mix.

Content of homeworks

Rohrer *et al* (2015) suggest that we can reap some of the benefits of interleaving without radically changing our schemes of work, or even the content of our lessons.

Seventh-grade students saw their teachers' usual maths lessons and received regular homework assignments. Every student received the exact same problems, but the scheduling of the problems was altered so that students received blocked or interleaved practice. Blocked homeworks were topic-based, whereas the interleaved ones contained a mixture of questions across topics. Later, students received a review of all the content, followed 1 or 30 days later by an unannounced test. Students following the interleaving program performed significantly better on both tests, with a greater difference in performance found on the test taken 30 days later, suggesting a positive impact on retention. Working on such interleaved activities gives students experience in selecting mathematical strategies, as well as executing them.

This is a relatively easy win, but one which I have been slow to pick up on and take advantage of. As already discussed in the section on the Spacing Effect, we used to give our students predominantly topic-based homeworks.

The simple practice of moving away from topic-based homeworks, to more mixed homeworks is likely to be of great benefit, with students needing to distinguish between and retrieve several different areas of knowledge from long-term memory, instead of just the one. And this does not mean throwing out all your old homeworks – it just means spreading out the topic-specific questions out.

Again, this will feel difficult – students marks will suffer in the short-term – but it is all for the greater good.

Application questions

This is a massive one that I discussed in Section 7.4 with regard to a special type of non-example, and in Section 9.3 in my approach to problem-solving. A common approach to sequencing lessons and presenting homeworks is to cover skill-based questions first, and then move on to applications of that skill. So, lesson 1 you may cover the basics of something like Pythagoras, and lesson 2 move on to applications of Pythagoras. Likewise, a Pythagoras homework or assessment is likely to follow the same format, with skill-based questions first, followed by application questions.

Rohrer *et al* (2014) argue that the correct solution to most mathematical problems involves two steps: identify the strategy needed to solve the problem, and then successfully carry out that strategy. But by following a lesson on Pythagoras lesson with a series of application problems, all of which require the use of Pythagoras, I am denying my students the opportunity to practisce that first step – identifying the strategy. Hence, all I was really testing when giving my students contextual problems was whether or not they could carry out Pythagoras's theorem, not whether they could recognise when Pythagoras's theorem was (and crucially was not) needed, *and then* apply their skills.

To go back to my earlier point, there may be some advantages to blocked practice in lessons in the first instance. But there is plenty of scope for improvements in the application section of homeworks and assessments. The easiest fix is to simply mix up the application questions so that students are required to distinguish between problems. Again, this does not require writing new questions. Questions can simply be taken from the application

sections of homeworks students have studied in the past. So, there may be contextual questions on simultaneous equations, estimating the mean and Pythagoras on the same homework.

As discussed in Chapters 7 and 9, perhaps the ideal solution is to present students with problems that have similar surface structures and different deep structures (or 'SSDD' – I am deluding myself that the more I say it, the more likely it is to catch on). In Section 7.4 I gave the example of a series of questions all involving isosceles triangles, one which was angle-based, one which involved area, another was a linear equation, and then one on Pythagoras. Likewise in Section 9.3 I suggested that questions all involving the surface structure of a rectangle could be used to cover concepts such as area, loci and transformations. These types of questions really do force students to focus not just on the similarities between questions, but crucially on the differences, which may well be vital for better transfer to different situations. Creating a collection of such problems that can be used in lessons, or on homeworks and assessments, is one of my goals for the future. I'm dreading telling my wife.

Related concepts

Finally, we have the issue of when to teach related concepts. Should area and perimeter be taught together, or apart? How about factors and multiples? Mean and median? The list goes on, and the answer is by no means straightforward.

After analysing the data on over 20 million answers from students around the world, I can safely say that the two biggest causes of incorrect answers on my Diagnostic Questions website are:

1. Misremembering or misapplying rules
2. Confusing related concepts

For example, whilst 23% of students opted for answer B in the question below, a combined 30% of students went for answer C (the mean) and D (the range).

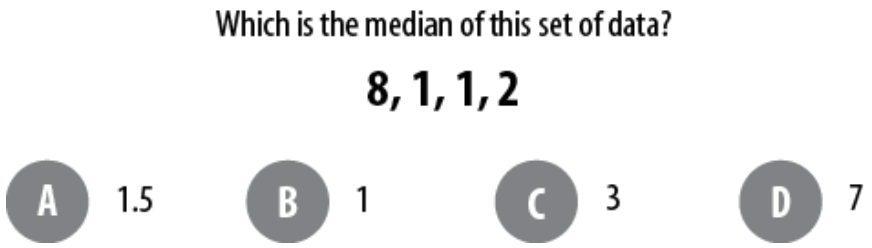


Figure 12.7 – Source: Craig Barton for Diagnostic Questions

Perhaps by teaching such concepts apart from each other, we can ensure that this confusion does not happen. Indeed, this is the approach favoured by Bruno Reddy, who, in his 2014 post, 'Design your own mastery curriculum', explains that one of the aims was to separate minimally different concepts. So, Bruno presents a scheme of work where in Year 7 area and perimeter are taught separately, as are mean and median.

However, there is a problem with this approach. Let's briefly return to Rohrer *et al* (2014) and their point that the correct solution to most mathematical problems involves two steps: identify the strategy needed to solve the problem, and then successfully carry out that strategy. Presenting these related concepts together forces students to distinguish between them, and hence benefit from interleaving and the power of non-examples. Otherwise, if something like perimeter is taught together in a block separate from area, then students may learn that all they have to do on every question is to add lengths together, never having the opportunity to make the distinction between questions related to perimeter and those related to area.

So what is the solution? In short, I think separating minimally different concepts during the early knowledge acquisition phase is the way to go. So, perimeter should be presented separately from area. The Interleaving Effect can still be used to great effect by ensuring things such as fractions, decimals and even algebra find their way into perimeter questions (once the key concept of perimeter is secure, of course!). Likewise, teach area separately. But then, they need bringing together. This can happen in homeworks and assessment. Or it can happen in lessons via the use of Purposeful Practice tasks such as *find the areas of all the different rectangles you can make that have a perimeter of 20m*. But it should only happen once conceptual

understanding of the initial two topics is secure. It does not even need to happen in the same year. Separate the concepts in Year 7, and then you have another four years to bring them together.

What I do now

Much like spacing, interleaving can have a powerful effect on students' long-term learning. However, it has an advantage over spacing in that it does not take any extra time. I have discussed how homeworks and assessments can be tweaked to take advantage of the Interleaving Effect, with no new content needed, and this is an incredibly worthwhile exercise that our department undertook as a team.

Likewise, interleaving within lessons should mean that instead of spending, say, two weeks on fractions in Year 7, you spend one week, and then spread that extra week out across the year, dropping fractions into topics such as probability, solving linear equations, ratio, perimeter, and so on. This takes a lot of planning, both in terms of the order of the curriculum and the design of lessons themselves, but the benefits should be well worth it.

OceanofPDF.com

12.5. The Variation Effect

What I used to think

I used to think that consistency was the key to good learning. Keep learning conditions the same wherever possible – same room, same seats, *etc*. I felt this served two main purposes. First, it was great from a behaviour-management perspective – once I had a seating plan that worked, and all the key characters were sat exactly where I needed them to be, then it meant I could concentrate my efforts on my teaching. Second, it was less for the students to think about, so they could just get on with learning.

What could possibly be wrong with that?

Sources of inspiration

- Barton, C. (2016) 'Tom Bennett', *Mr Barton Maths Podcast*.
- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real world: essays illustrating fundamental contributions to society*. 2nd edition. New York, NY: pp. 59-68.
- Smith, S. M., Glenberg, A. and Bjork, R. A. (1978) 'Environmental context and human memory', *Memory & Cognition* 6 (4) pp. 342-353.

My takeaway

Now, neither of the points I made above are invalid. When I interviewed Tom Bennett – one of the country's leading experts on behaviour – for my podcast, he explained that without good behaviour, everything else falls apart. Behaviour comes first. So, if a consistent seating plan in the same

classroom is needed to maintain this behaviour, then I would not change it. Likewise, we have seen throughout this book how fragile students' working memories are. In Chapter 4 we considered the Redundancy Effect, whereby even irrelevant information has to be processed by students' working memories. Just as this applies to classroom displays and the teacher moving around the room, it seems sensible to assume it also applies to changing classroom circumstances. A child suddenly moved to a new seat or sitting in a new room will have plenty of other things to occupy their working memory than the key concept of the lesson you are trying to teach them.

But there is a counter-argument to all this that is worth discussing.

Another of Bjork and Bjork's (2014) desirable difficulties is varying the conditions of practice. In short, when instruction occurs under conditions that are constrained and predictable, learning tends to become contextualised. Material is easily retrieved in that context, but not so if tested in a different context. In contrast, varying conditions of practice – even varying the environmental setting in which study sessions take place – can enhance recall on a later test. Indeed, Smith *et al* (1978) found that studying the same material in two different rooms, rather than twice in the same room, led to increased recall of that material.

This has become known as the Variation Effect and since Smith *et al*'s 1978 study there have been many others which have replicated these findings. Why does varying conditions work? Well, it may enhance long-term learning because the material becomes associated with a greater range of memory cues that serve to facilitate access to that material later.

I find the Variation Effect fascinating. Creating cues to trigger memories can be viewed as one of the hallmarks of successful teaching. It is how useful devices such as analogies and mnemonics work. But the key is that we do not want students' memories to be tied to a specific cue, for as soon as that cue is no longer available, the memory may no longer be accessible. It is why I argued against the permanency of even the most useful classroom displays in Section 4.7 – students need to ensure they can retrieve the information when it is not present on the wall.

Whilst it is clear that many experiments in this field are concerned with the recall of words – and so we must be very careful in making any wild claims about solving mathematical problems – the fact that recall in mathematics plays an important part – in terms of definitions and rules – suggests that we should, at the very least, take note of these results.

What I do now

So, how do we get the balance right between the benefits of cues and consistency, whilst also tapping into the power of the Variation Effect?

Variation by location

For me, the first few weeks of meeting a new class are all about getting to know the students, and establishing my expectations and routines for behaviour. During this time, I would not consider moving seats or classroom. I am looking for consistency.

However, as I get to know a class, and feel I have a good handle of them, then I will experiment with variation. I find it hard to justify breaking up pairs of students that work well together, so more often than not I will move pairs of students around the room to different locations. However, to ensure that their success is not entirely dependent on the person they are sat next to – both explicitly in terms of copying, and implicitly in terms of the comfort and familiarity that person brings – I will occasionally split hard-working students up. I find random seating plan generators are a good way to do this so the students know the process is as fair as it can be, and hence are less likely to kick up a fuss.

But just to reiterate, if this negatively affects the behaviour of the class, and thus impacts my ability to teach them well, then I will stop. And with some classes, I may never even start (Year *10 – 2013*, you know who you are). Likewise, I'm reluctant to do it in the early knowledge acquisition phase – a time when my students' working memories may be pushed to the limit.

A much easier win is to try to ensure students experience the place they are going to be sitting any high-stakes exams before they sit them. This could be as simple as booking the exam hall for a lesson, or running a low-stakes quiz or a mock exam in there. I have taught too many students who could

perform in the classroom, but who then went to pieces in the exam itself. As we discussed in Chapter 2, much of this could be due to maths anxiety, but the unfamiliar location could also play a role.

I also advise my students not to do all their revision and practice in the same environment – be it their room, in front of the telly, *etc*. Consistent environmental conditions during revision could be just as susceptible to cue-dependent recall as what happens in school.

Other forms of variation

But the Variation Effect does not need to be limited to seating plans and locations. Anything that changes the conditions of study can help break cue-dependent recall. But this must always be balanced against the potential cost of a fall in consistency – something that I feel needs carefully managing especially during the initial knowledge acquisition phase. We need students attending to things that positively impact on their learning, not on changing classroom conditions.

Here are some ideas I have experimented with. Some have worked wonders with certain classes and students, and been a bit of a disaster with others. As you read them, consider which you think may work with your students:

- Changing the structure of lessons. I always used to start my Year 11 lessons with one of Corbett Maths's 5-a-days, but I tried using one in the middle of a lesson every now and again
- Changing the format of questions in low-stakes quizzes – see Section 12.7
- Asking students to work for some parts of the lesson in books, some on file paper and some on mini whiteboards
- Asking students to write in a different pen
- Asking them to use a different calculator
- Working alone, in pairs or in groups
- Using a video to introduce or illustrate a concept

- Swapping teacher – so I teach the parallel class in Year 9 for a lesson, and *vice-versa*

As with all the techniques described in this chapter, these strategies are designed to make learning more difficult. As such, I would only use them once I was happy that students were relatively secure in the knowledge I was testing. Also, explaining to the students exactly why you are doing these things, and the benefits you expect them to bring, is likely to make acceptance of the process a lot smoother.

OceanofPDF.com

12.6. The Testing Effect

What I used to think

I used to think that the purpose of tests was solely to give me information about my students' understanding. They were merely a vehicle of assessment. So, I just used to give my students tests when I had to, and be rather apologetic about it.

As a result, neither my students nor I came to enjoy or even value tests very much. We both preferred to get back to things that might actually improve their learning.

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast.*
- Bjork, E. L. and Bjork, R. A. (2014) 'Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning' in Gernsbacher, M. A. and Pomerantz, J. (eds) *Psychology and the real world: essays illustrating fundamental contributions to society*. 2nd edition. New York, NY: pp. 59-68.
- Dunlosky, J., Rawson, K. A., Marsh, E. J., Nathan, M. J. and Willingham, D. T. (2013) 'Improving students' learning with effective learning techniques: promising directions from cognitive and educational psychology', *Psychological Science in the Public Interest* 14 (1) pp. 4-58.
- Karpicke, J. D. and Roediger, H. L. (2008) 'The critical importance of retrieval for learning', *Science* 319 (5865) pp. 966-968.

- Roediger, H. L., Putnam, A. L. and Smith, M. A. (2011) 'Ten benefits of testing and their applications to educational practice' in Mestre, J. P. and Ross, B. H. (eds) *Psychology of learning and motivation: cognition in education, Vol 55*. San Diego, CA: Elsevier Academic Press, pp. 1-36.

My takeaway

I know I have made some big claims already throughout this book, and here comes another one: *viewing tests not just as tool of assessment but as tools of learning has been the single biggest change to my teaching over the last two years*.

The Testing Effect (also known as the Retrieval Effect) simply blows my mind. It is the widely replicated finding that far from being neutral and uninfluential, the retrieval of information produces learning, and is usually found to be significantly more effective on learning than restudying. Tests act not only as passive assessments of what is stored in long-term memory (as is often the traditional perspective in education) but also as vehicles that modify what is stored in long-term memory.

Indeed, in their review of effective learning techniques, Dunlosky *et al* (2013) conclude: 'On the basis of the evidence described above, we rate practice testing as having high utility. Testing Effects have been demonstrated across an impressive range of practice-test formats, kinds of material, learner ages, outcome measures, and retention intervals'. Bjork and Bjork (2014) explain that much laboratory research has demonstrated the power of tests as learning events, and in fact a test or retrieval attempt, even when no corrective feedback is given, can be considerably more effective in the long-term than reading material over and over.

Like many of the other desirable difficulties we have discussed in this chapter, the power of the Testing Effect is probably due to the cognitive strain experienced when trying to reconstruct knowledge as opposed to restudying it, something that is related to the concept of germane load (Section 4.9).

That is all well and good, but isn't that what we do in schools – teach students something, and then test them later? Well, to a certain extent; but

in the past I may have been guilty of too much teaching and not enough testing.

Let me show you what I mean. On our scheme of work our students have a test every half term of content largely studied in the previous six weeks (mistake #1 – see Spacing Effect above). Let's say it is just after Christmas and the topics I am teaching my Year 8s that half term according to our scheme of work are: percentages, solving linear equations and angles in polygons. I would teach each topic as well as I could, set them topic-specific homeworks (mistake #2 – see Interleaving Effect), do a couple of revision lessons before the test, and tell the students to revise for the test for homework (mistake #3 – promoting the benefits of cramming). What could possibly go wrong?

Well, notably absent from any of this is some actual test of retrieval. Sure, students may do homeworks under test conditions, but with Google, Corbett Maths, parents, siblings and WhatsApp to hand, there is no guarantee that is the case. My revision lessons will be full of my conscious and unconscious cues and hints designed to improve performance. Finally, more often than not, students' revision consists of a glance over their notes, where they are met with that reassuring sense of familiarity.

The conditions I am creating are essentially the same as those set up by a fascinating study by Roediger and Karpicke (2006). They compared different learning sequences in preparation for a free-recall vocabulary test. One set of students followed a pattern of study-study-study-test (similar to my Year 8s), whereas another followed study-test-test-test. When tested a week later, even though the latter group had much less exposure to the material, only having been taught it once, they significantly outperformed the group who had studied the material on three separate occasions. Testing led to better long-term retention. This finding is supported by Karpicke and Roediger (2008) who conclude: 'Repeated studying after learning had no effect on delayed recall, but repeated testing produced a large positive effect'.

Once again, as these findings lie outside the realms of the maths classroom, we must be careful in generalising. But I have seen this so many times myself. I teach students something, and they seem fine. We then revise it

again in class, and with a little help from me they seem fine again. Students feel confident about the topic, so they just glance over their notes. Then it comes to their end-of-term tests, and it all goes to pot. Retrieval leads to learning and retention in a way that repeated studying simply does not.

What I do now

Quite simply, I test my students every single lesson in the way I will describe later in this chapter.

OceanofPDF.com

12.7. The many, many other benefits of tests

What I used to think

Having learned about the benefits of retrieval on learning, I was becoming rather partial to tests. It was just a shame, I mused, that tests did not have any other benefits…

Sources of inspiration

- Roediger, H. L., Putnam, A. L. and Smith, M. A. (2011) 'Ten benefits of testing and their applications to educational practice' in Mestre, J. P. and Ross, B. H. (eds) *Psychology of learning and motivation: cognition in education, Vol 55*. San Diego, CA: Elsevier Academic Press, pp. 1-36.
- Rosenshine, B. (2012) 'Principles of instruction: research-based strategies that all teachers should know', *American Educator* 36 (1) pp. 12-39.
- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

My takeaway

Two of Rosenshine's (2012) 'Principles of Instruction' concern testing:

1. Begin a lesson with a short review of previous learning.

1. Engage students in weekly and monthly review.

The reason testing takes up two places on that rather illustrious list is probably because the benefits of testing are not just confined to the direct effect they have on learning that was described in the previous section. In fact, in terms of looking at the benefits of testing, we are just getting started.

Roediger *et al* (2011) summarise the additional benefits of testing, together with relevant research, in a paper that I would encourage everyone to read, entitled 'Ten Benefits of Testing and their Applications to Educational Practice'. Here, I will just pick out a few.

Testing Provides Feedback to Instructors
This is an obvious, but nonetheless important point. Testing can provide teachers with valuable feedback about what students do and do not know. Teachers may overestimate their students' knowledge, and testing provides one way to improve that. We have seen many times throughout this book that the problem of the *curse of knowledge* permeates education. That is, teachers (especially those with relatively little experience) can fail to realise the state of knowledge of their students and hence pitch their lessons at too high a level. Testing – along with Chapter 11's diagnostic questions – provides a more objective and accurate measure of just where this state of knowledge lies.

Testing Identifies Gaps in Knowledge
This benefit is not just for teachers. Taking a test permits students to assess what they know and what they do not know far more accurately than other studying strategies, such as rereading notes or highlighting. Such strategies tend to give students a nice, comforting feeling of familiarity, and hence lead them to make poor decisions when allocating study time. A student may read their notes on fractions and think, 'Yep, I can nail them'. But such familiarity can be misleading. There is no hiding when doing a test. Students can see in black and white (or maybe red) what they can and cannot do. This means students can concentrate study efforts on areas in which their knowledge is deficient. This helps answer the eternal question: 'What should I revise, sir?'. In short – as I tell my students – if you could not retrieve that information in a test, then you haven't learned it fully, and you need to work on it some more.

Frequent Testing Encourages Students to Study
Identifying gaps in knowledge is all well and good, but what then do students do with that? After all, Wiliam (2011) makes the point that as teachers we sometimes assume that most of student learning happens within our lessons, whereas given that we only see our students for a few hours a

week, it is likely that what happens outside of lessons is of greater importance. Fortunately Roediger *et al* (2011) found that having frequent quizzes, tests, or assignments seems to motivate students to study. We have all taught students who procrastinate throughout the year, giving off a calm air often accompanied by reassuring phrases like, 'Relax, sir, I have got it covered', and often do not study properly until the night before a test. More frequent testing across the year encourages students to study more, and hence take full advantage of the Spacing Effect. And of course, the fact that students are better-informed as to what to focus their study on means that effective learning outside of our lessons is more likely.

Testing Produces Better Organisation of Knowledge

We have seen how the act of retrieval can change the long-term memory, and because tests are primarily designed to induce retrieval, the benefit can be significant. Retrieval can improve the conceptual organisation of practised materials, especially on tests that are relatively open-ended. Retrieval causes students to organise knowledge more than does reading. As students actively recall material, they are more likely to notice important details and weave them into a cohesive structure in their long-term memories.

Testing can Facilitate Retrieval of Material that was not Tested

Retrieval practice does not simply enhance retention of the individual items retrieved during the initial test: taking a test can also produce retrieval-induced facilitation – a phenomenon whereby testing also improves retention of non-tested but related material. In Chapter 11 we discussed that one of the advantages of using diagnostic multiple-choice questions over other forms of questions was that students were not only forced to consider why one particular answer is correct, but also why the others are wrong. This effectively means students are getting more bang for their buck when answering a question, thus making this benefit even more powerful.

Three reasons to show students the correct answers

Do students need to know if they are right? I know this seems like an obvious question to ask, but according to the Testing Effect, all the benefits of tests come from the act of retrieval itself. No mention is made of what

happens following this retrieval – in particular, does it matter if the student is wrong and never discovers that fact?

Back in Section 8.4 we looked at three reasons to ensure students had access to the answers to classwork. They were:

1. Practice makes permanent
2. The Self-Explanation Effect
3. The Hypercorrection Effect

All of these reasons are just as relevant to testing, and can only be achieved if students have access to the answers. Likewise, I would argue that the positive benefits are likely to be larger if those answers are given relatively quickly following the completion of the test. The student who waits a week to get the results to their test back may have practised incorrectly in between, and is less likely to remember the reasons for the answers they gave and thus less likely to benefit from self-explaining or hypercorrecting.

What I do now

I am far more aware of the many, many benefits of regular testing. Likewise, I make sure my students are in no doubt whatsoever that the single best way they can revise in mathematics is to continually test themselves throughout the year on the full range of the specification they are working towards. Put the highlighters and the revision guides away.

Testing is one of the most powerful tools that both teachers and students can have at their disposal. As such, I give my students a test every single lesson. And exactly how I do that is the subject of the next section.

OceanofPDF.com

12.8. Low-stakes quizzes

What I used to think

If you had asked me even as recently as 12 months ago whether I could afford to spend five to ten minutes each lesson on a test, I would have asked you if you were having a laugh. Have you seen the typical maths scheme of work? Then there was the time out of lesson. After all, for a test to be useful, I have to mark it, collect in the scores, assign some kind of level or grade, and ideally write reams of personalised feedback. With my marriage in continuous crisis between Sept and June each year as it is, I simply could not spare the time.

So, I would only give tests when I absolutely had to, and my students, my wife and I were more than happy with that arrangement.

Sources of inspiration

- Ashcraft, M. H. (2002) 'Math anxiety: Personal, educational, and cognitive consequences', *Current Directions in Psychological Science* 11 (5) pp. 181-185.
- Barton, C. (2017a) 'Dani Quinn – Part 1', *Mr Barton Maths Podcast*.
- Barton, C. (2017b) 'Nick Rose', *Mr Barton Maths Podcast*.
- Christodoulou, D. (2017) *Making good progress?*. Oxford: Oxford University Press.
- Hendrick, C. and Macpherson, R. (2017) *What does this look like in the classroom? Bridging the gap between research and practice*. Woodbridge: John Catt Educational Ltd.

- Karpicke, J. D. and Blunt, J. R. (2011) 'Retrieval practice produces more learning than elaborative studying with concept mapping', *Science* 331 (6018) pp. 772-775.
- Karpicke, J. D. and Roediger, H. L. (2008) 'The critical importance of retrieval for learning', *Science* 319 (5865) pp. 966-968.
- Little, J. L., Bjork, E. L., Bjork, R. A. and Angello, G. (2012) 'Multiple-choice tests exonerated, at least of some charges: fostering test-induced learning and avoiding test-induced forgetting', *Psychological Science* 23 (11) pp. 1337-1344.
- Lovell, O. (2017) 'ERRR #009. Andrew Martin, Load Reduction Instruction, Motivation and Engagement', *ollielovell.com*. Available at: http://www.ollielovell.com/errrpodcast/errr-009-andrew-martin-load-reduction-instruction-motivation-engagement/
- Marsh, E. J., Roediger, H. L., Bjork, R. A. and Bjork, E. L. (2007) 'The memorial consequences of multiple-choice testing', *Psychonomic Bulletin & Review* 14 (2) pp. 194-199.
- McDermott, K. B., Agarwal, P. K., D'Antonio, L., Roediger, H. L. and McDaniel, M. A. (2014) 'Both multiple-choice and short-answer quizzes enhance later exam performance in middle and high school classes', *Journal of Experimental Psychology: Applied* 20 (1) pp. 3-21.
- Putwain, D. W. (2007) 'Test anxiety in UK schoolchildren: prevalence and demographic patterns', *British Journal of Educational Psychology* 77 (3) pp. 579-593.
- Sadler, P. M. and Good, E. (2006) 'The impact of self- and peer-grading on student learning', *Educational Assessment* 11 (1) pp. 1-31. Siegler, R. S. (2002) 'Microgenetic studies of self-explanation' in Granott, N. and Parziale, J. (eds) *Microdevelopment: transition processes in development and learning*. Cambridge: Cambridge University Press, pp. 31-58.

My takeaway

Why don't students like tests?

I think there are two main reasons why students may not like tests:

1. They see tests as high-stakes tools of assessment
If students are only used to doing tests at the end of every half term, under strict timed conditions, and with stakes such as set-changes, school reports, parental involvement and potential embarrassment in front of their peers up for grabs, then it is little wonder they do not enjoy them. Indeed, Putwain (2007) found that tests cause anxiety in students. Back in Chapter 1 we looked at the debilitating role that maths anxiety can play on a student's enjoyment and learning in mathematics. The 'worry' component of anxiety takes up valuable working memory space, which makes the processing of the kind of complex tasks you are likely to find on a high-stakes exam more difficult, thus inhibiting performance. For Ashcraft (2002), timed tests are a key source of this anxiety.

2. They don't appreciate the benefit
A recurring theme throughout this chapter on desirable difficulties is that the approaches feel harder, performance dips in the short term, and as such students do not appreciate the benefit to their long-term learning. This is no truer than in the case of tests. Karpicke and Roediger (2008) found that students' predictions of their performance were uncorrelated with actual performance. Indeed, the students in the study were not aware of the benefits of practising recall, even after they had successfully done it! As with all the desirable difficulties we have looked at in this chapter, this is most likely because retrieving information is hard, whereas other forms of revision, such as reading notes and highlighting, feel easier, give a sense of familiarity, and hence seem more effective.

I believe both of these can be overcome by regularly using low-stakes tests – or, to be more precise, *low-stakes quizzes*.

Low-stakes quizzes
Very few of the benefits of testing identified in the previous two sections require the tests to be marked, graded and recorded. And yet, many of the negative aspects of tests stem from these processes. If students can come to see testing as a regular part of their lesson – and crucially, as a regular part of their *learning* – where their performance will not lead to negative

consequences, then they may start to reap some of the significant benefits identified in this chapter.

Hence, each lesson I teach now consists of a low-stakes quiz. Now, just like with the presentation of worked examples described in Chapter 6 and Chapter 11's diagnostic questions, I am very particular about the specifics of these quizzes. It is a format I have tweaked, transformed, and tweaked some more over the last year working with my students, and it now looks like this:

- I call them *quizzes* as opposed to tests. Tests are scary; quizzes are fun.
- I print the quizzes out. Even though this is a strain on the photocopying bill, it is worth it. Printing out allows me to ask a wider variety of questions, such as transformations, angles, *etc*. It also helps encourage good practice, such as labelling diagrams and filling in missing information. Projecting questions up on the board has the classic issue of students at the back struggling to see, and printing out the quizzes ensures all eyes are facing down, which can improve concentration. Finally, if you are going to adopt the Quiz-Homework-Quiz Combo that I will describe at the end of the section, students need to have a clear record of the question and the answer so they can revise properly. Fitting the quizzes onto one piece of A4 (back-to-back if needed) keeps me in the relative good books with regard to the maths department budget.
- Students do the quizzes in silence, and on their own. The single most important thing is that these quizzes induce each student to retrieve information from their own long-term memories. This is not a collaborative exercise, and likewise I do not want the processing of noise and other distractions taking up valuable space in their fragile working memories. Once students get used to this, and realise that these quizzes are tools of learning and not assessment, the negative effects of these potentially intimidating conditions soon diminish.
- These are not open-book quizzes. I used to allow students to have their notes with them whilst they took certain in-class quizzes. My logic was that it would reduce the stakes and pressure of the quiz, whilst also

showing students the benefits of taking neat, well-presented, detailed notes in class. But I return to my previous point – these quizzes must induce retrieval in order to benefit from the Testing Effect. I want students searching for answers in their long-term memories, not in their books.

- There is no time pressure. Whilst I may have in mind that these quizzes will last around 5 to 10 minutes, this is flexible and rarely told to the students. Instead I will make a judgement call, based my observation of how the students seem to be getting on. I ask my students to turn over their quizzes when they have finished. When half my students have done so, I will calmly announce that there is one minute to go. Now, that is not to say that I am opposed to timed quizzes by any means. Students need practice of completing work under timed conditions in order to gain valuable experience pacing their work for when they sit high-stakes timed tests. Also, I believe answering questions under timed conditions can lead to the development of fluency and automaticity in key knowledge. But certainly when I am getting classes used to the concept of regular low-stakes quizzes, I tend to avoid explicit time constraints.

- For students who finish early, there is an extension question. More often than not – as described below – this is a UKMT problem that combines several areas of maths in an unconventional way. I do not expect every student to get to this problem, and it is certainly not the focus of the quiz.

- Just before I project up the answers, I give students 30 seconds to look back over their work and place a score out of 10 next to each question to indicate how confident they are in their answer. This serves two purposes. The first is that I find it a far more effective way of getting students to check their work than simply saying 'check your work' (if there is a phrase students ignore more routinely, I am yet to hear it). But this confidence score has an even greater purpose when we go through the answers.

- I project the answers up at the end, and most of the time I ask students to mark their own work (see below for a discussion about self- *versus* peer-marking of these quizzes). There is no talking at all at this stage – I want all of my students' limited working memory capacity focused on the answers on the board and the answers in front of them. Of course, I want students to pause and consider why they got *any* of the questions wrong in order to take advantage of Chapter 5's Self-Explanation Effect, but in particular I want them to focus on questions they got wrong that they were confident were correct. This is my attempt to take advantage of the Hypercorrect Effect discussed in Section 8.4. Hence, I ask my students to look at the confidence scores next to each question they got wrong, and place an asterisk next to those they gave a high confidence score to. The class then has a few minutes of silent contemplation – each student beginning with the asterisk questions – to see if they can figure out where they went wrong. This technique is designed to counteract two classic tendencies I have witnessed in many students over the years. The first is to simply settle for a score of, say, 8/10 – merely glancing at the correct answers for the couple of questions they got wrong, but not with any real focus or attention. The second is the tendency to divert most attention towards the questions they struggled with – those questions they didn't really understand and took a bit of a guess at. So a student may have really struggled on Question 8, be not at all confident about their answer, and hence be determined to find out how to do it properly. This seems intuitively the most sensible thing to do – that is, until you really think about it. The Hypercorrect Effect tells us that the most significant gains to learning come from errors created with high confidence. Whilst the student is focusing their attention on troublesome Question 8, they may skip over the fact they also got Question 2 wrong and yet were certain about their answer. By explicitly drawing students' attention to these high-confidence errors via their confidence score, we are reducing the possibility that they fail to pay sufficient attention to them. And of course, high-confidence errors not attended to are more likely to be repeated in the future as we have the wiring of paths of production that Anderson's model of thinking (discussed in Section 1.1) explains are so difficult to change. Obviously

we want students to reflect and correct all errors, but it is the high-confidence ones that must take priority.

- Following this period of quiet contemplation, I ask students if there are any questions they still do not understand. If there are issues that can be resolved quickly, I will do it there and then, often involving a student who got a question right and is confident about their answer. If there are issues that cannot be resolved quickly, I will not tackle them in that moment. Instead, I will tell the students in question not to panic, and that I am going to go away and plan the best way to help them understand the concept better. I will then address it in a future lesson, or if it is a small group of students then I may cover it at break. The only exception to this is when – as discussed in Chapter 11 – the question is testing baseline knowledge necessary to acquire the new skill or concept I wish to teach. In those cases, I will take as long as I need to ensure students are comfortable with the material.

- Students answer these quizzes on sheets of paper, which they then place in their Low-Stakes Quiz Folders. Stuck to the inside of each folder is a Personal Record Sheet where they fill in their marks. The folder removes the hassle and untidiness of sheets stuck in books, as well as the tendency for them to find their ways into trouser pockets or scrunched up in the bottom of school bags.

- I will not ask students for their marks, nor will the marks be shared with anyone else. In a 2017 interview with Oliver Lovell, psychologist Andrew Martin explains that competing with yourself is more motivating than competing with others. Hence, the Personal Record Sheet has the advantage of spurring students on to improve, whilst avoiding any anxiety that may come from having to share their score. Students know that I will look over their quizzes and Personal Record Sheet when I occasionally take in their folders (see the point about no-stakes tests below), but I make it very clear that this is not to record any scores, or tell their parents, or anything like that. I make sure they know it is so I can get an idea of what I need to do to help the class learn better. This is all part of the shift from viewing tests as tools of assessment to tools of learning.

- The content of these quizzes always consists of things students have encountered before, and they are never topic-specific. I want to ensure my students reap all the benefits of the Spacing and Interleaving Effects. The only exception to this might be when I wish to tap into the Pretest Effect, which will be discussed towards the end of this chapter, but in such cases these questions will only make up a small proportion of the quiz.
- Finally – and perhaps most important of all – these low-stakes quizzes are not a small aside to the lesson that we only do if we have time – they are a key part of the lesson. As such, I do the quiz at the start of each lesson to ensure I do not run out of time doing something else and hence have to sacrifice the quiz. As soon as students see you are willing to forgo doing the quizzes, they immediately start to lose importance in the students' eyes.

What format should the low-stakes quizzes be?
Seeing as I am setting one low-stakes quiz every day, I want to be able to quickly get my hands on top-quality sets of questions and answers. Bearing in mind the importance of the Variation Effect identified in Section 12.5, I like to mix up the format of the questions to make the process of retrieval more effortful, and hence the effect on learning more durable.

In Chapters 8 and 11 we discussed how the format of the practice does not (and perhaps should not) match the format of the final performance. Most relevant here is the finding by McDermott *et al* (2014) that the format of the quiz (multiple-choice or short-answer) did not need to match the format of the critical test (*eg* end of unit exam) for the benefits of testing and retrieval to emerge.

My low-stakes quizzes are comprised of a combination of the following four things. I build them myself based on what my students have studied in the past. It can be tricky keeping track of the topics that have and have not been covered, together with the topics that students have struggled with and hence need revisiting, but a good old-fashioned Excel spreadsheet does the job.

Short-answer questions
These are the bedrock of my low-stakes quizzes, and I have three favourite sources to get the questions, each of which also provide the answers and are (at the time of writing) completely free:

1. Jonathan Hall's *Maths Bot Test Maker* (available at mathsbot.com/testMaker), which allows you to build low-stakes quizzes, selecting questions from a whole host of topics, which can be projected on the board or (as I prefer) made into a printable sheet.

2. Pete Mattock's *Brockington Homework* collection (freely available on TES at tes.com/teaching-resource/brockington-college-maths-homework-booklets-11223661), which has thousands of questions for all year groups, all in Word and so completely editable, and you can tweet Pete for the answers.

3. John Corbett's *5-a-day* (available at corbettmaths.com/5-a-day), which now stretches from primary right up to the top end of GCSE, ensuring I will always find questions to suit the needs of my class.

Diagnostic Questions
Throughout Chapter 11 I extolled the virtues of diagnostic multiple-choice questions, so it will be no surprise that many of my low-stakes quizzes make use of these. They are super-fast to mark, and as Little *et al* (2012) explain, they have an extra advantage in that, to get the questions correct, students must consider why a particular answer is correct as well as why the alternative answers are wrong. Printing these questions out means that students are encouraged to write their working and the reason for the correct answer, thus benefiting from the Self-Explanation Effect (see Chapter 5).

Multi-mark questions
I may include one of these at the end of the low-stakes quiz, but only if students have studied all the relevant content. Having discussed in Section 9.4 the possibility of novice learners struggling on problems and not actually learning anything, I obviously want to ensure that that does not happen on these low-stakes quizzes. Predominantly, they are tools to induce retrieval,

not push novice learners' memories to the point of cognitive overload. As we get closer to a high-stakes exam, I will include more of these multi-mark questions as I look to hone students' exam technique.

UKMT questions
I will always have a question from the United Kingdom Mathematics Trust (UKMT) competitions (available at: ukmt.org.uk) up my sleeve for students who have finished. These are in multiple-choice form, and hence have the advantage of being quick to mark, but they tend to require more thought and time to solve. The beauty of the UKMT questions is that often they require students to use their knowledge in novel ways, or combine two different areas of knowledge and hence take advantage of the Interleaving Effect.

Should we use exam papers?
It would seem we have the perfect set of low-stakes tests readily available to us in the form of past exam papers. After all, these comprise short- and long-answer questions – heck, AQA's GCSE papers even have diagnostic questions – and cover the relevant syllabus. So long as we are careful to run them as low-stakes quizzes with no pressure on performance as described above, will they be even more effective than the formats described above?

Potentially, no they won't. We need to remember what exam papers are designed to do. Exam papers are designed to assess students' understanding at the end of a particular course or unit. They are designed to cater to a wide variety of abilities. They may ask questions in unconventional ways, and often contain unstructured multi-mark questions. Indeed, Christodoulou (2017) makes the point that often these types of exam questions are more like project-based teaching than assessment of key skills.

When judging whether the Testing Effect will apply or not, the important factor is not the complexity of the material, but whether a task involves retrieval or not. It is the act of retrieving knowledge from long-term memory that strengthens the memory. If content is not in long-term memory to begin with, then pupils cannot retrieve it. Exam papers do not always encourage retrieval from long-term memory, but instead test elements of problem-solving. We have seen in Chapter 9 that problem-

solving does not always lead to learning, and hence students may work through exam papers for extended periods of time and not actually learn anything. We have also seen in Chapter 8 how exams should perhaps be viewed as the final performance, and practising in such a manner makes it very difficult to pinpoint the necessary areas in which improvement is needed.

I am as guilty of this as anyone. My students have been in their seats no more than a few seconds at the start of Year 11 before I present them with their weekly past paper schedule that they are to complete at home over the course of the next 9 months. And what happens? Students routinely score around 60%, answering the questions that they can on the topics they like, struggling or leaving out the same ones over and over again, getting varying amounts of help along the way. I would much rather have more control over the content and types of questions students answer, and the environment in which these tests are sat, so I can ensure they are a true test of retrieval that all concerned will benefit from.

Carefully selecting individual questions from exam papers to add to low-stakes quizzes can be an excellent idea. However, I feel presenting students with full exam papers – whether in class or for homework – should be left as late in the process as possible – no exam papers until after Christmas of Year 11 seems a sensible policy.

Self-assessment v Peer-assessment

Who should mark these low-stakes quizzes? We have already discussed how very few of the benefits of testing require me, their teacher, to mark their tests. But is it better for students to mark their own work, or for a peer to do so? Sadler and Good (2006) investigated this in Grade 4 Science lessons.

Their study reported three key findings:

1. For either peer- or self-assessment to work students must be familiar with how to mark accurately. Strategies such as exposing students to mark schemes and showing them examples of other students' work and then discussing them are likely to be successful in achieving this aim. The authors sum this up with the following: 'For optimal student-

grading, we suggest training, blind grading, incentives for accuracy, and checks on accuracy compared to teacher grades'. Now, so long as the questions in the low-stakes quizzes are short-form, this will not be an issue. But if including multi-step, multi-mark questions, then clearly a student-friendly mark scheme or, better still, a worked solution carefully annotated, may be called for.

2. There are indications of bias in the students' marks: when grading others, students awarded lower grades to the best-performing students than their teacher did, but when grading themselves, lower-performing students tended to inflate their own low scores. This is a problem that should hopefully be overcome over time. My experience has been that as students realise over the course of the year that the marks are not the important thing, but the act of retrieval itself, then this is much less of an issue.

3. Most interesting to me is the impact on learning. When students were given a follow-up test based on the same content that they had been marking, students who graded their peers' tests did not gain significantly more than a control group of students, but those students who corrected their own tests improved dramatically. This implies that self-assessment is more beneficial to learning than peer-assessment.

Hence, the majority of the time I allow students to mark their own low-stakes quizzes. It is much quicker – no time is lost swapping books or arguing about what they meant to write. My students also seem to prefer this, especially the first few times they do quizzes. Having someone else mark your work feels like an assessment, whereas marking it yourself feels like learning.

Dylan Wiliam agrees, which is always a good sign. Writing in Hendrick and Macpherson's (2017) *What does it look like in the classroom?*, he explains: 'The benefits of testing come from retrieval practice and hypercorrection and that's why the best person to mark a test is the person who just took it. So more testing and less marking'.

However, we need to keep in mind Siegler's (2002) two key findings on self-explanations from Section 5.2: children who were asked to explain the experimenter's reasoning learned more than children who explained their own reasoning; and explaining why correct answers are correct and why incorrect answers are incorrect yields greater learning than only explaining why correct answers are correct. Likewise, we have seen in Section 6.4 the benefits of self-explaining steps in Supercharged Worked Examples. Hence, I sometimes choose to complete a low-stakes quiz myself, making a few key mistakes along the way, and challenge the students to mark and correct it.

Likewise, there may be occasions where I want students to mark each other's work. But this only comes much later in the process, once students are both confident and fully on board with the fact that these quizzes are tools of learning. As well as the benefits related to self-explaining, seeing other students' work and approaches can be very beneficial and illuminating. But if this comes at the price of increasing students' anxiety or preventing them taking risks for fear of being wrong (or being seen publicly to be right), then I simply do not do it.

Low-stakes *v* no-stakes

What about low-stakes versus no-stakes? When I interviewed author and TeachFirst researcher Nick Rose for my podcast, he explained that he favoured *some* stakes to the quizzes students take. Whilst it is clear that the mechanism of retrieval is the key to tapping into the benefits of the Testing Effect, we do need students to put the required effort in to bring about that retrieval. If they know the outcome does not matter at all, then there is a danger that some students will simply not bother thinking hard and hence not learn. Something as simple as insisting students record their scores on a Personal Record Sheet, reminding them that you will be looking through their folders, or occasionally engaging in peer-marking might be a way of increasing the stakes slightly above zero.

'But my students keep doing really badly at these quizzes'

This is something I hear a lot when I talk about how I run low-stakes quizzes. Teachers explain to me that they like the idea, and they tried them with their students, but the students couldn't get any of the questions right. Teachers in this position tend to do one of two things – either give up, or

change the delivery of the quizzes, maybe allowing students to work together or making them open-book. Whilst I appreciate the frustration, my advice is to stick to your guns and persist exactly as I have described above. Any modification risks watering down the act of retrieval from long-term memory, which is the main driver of the increase in learning. And if students seem unable to retain and access the knowledge needed to do well in these quizzes, then I would argue that denying them the opportunity to practise the very act of retrieval that will enable them to retain and access such knowledge will only do them harm in the long term.

But – as with all these desirable difficulties – I know it can be tough in the short term until students start to see the benefits for themselves. I have seen it with my own students, especially the ones who have a troubled history with mathematics. Therefore, I would advise reminding students that they are not being assessed, no marks are being recorded, and that these are purely tools of learning that mountains of research suggest will work. Secondly, there is nothing at all wrong with including a few easier questions in order to give students the taste of success that Chapter 2 argued was a key determinant of motivation. But do not give up, or compromise the act of retrieval. I promise it will be worth it.

The Quiz-Homework-Quiz Combo
Let's end with something controversial!

When I interviewed Dani Quinn for my podcast, she explained how Michaela Community School run weekly low-stakes quizzes in mathematics. This has heavily influenced the policy our maths department has adopted for the current academic year with our Year 11s.

I do my usual low-stakes quizzes every lesson as described above. However, on Friday and Monday, every Year 11 maths class in our school receives a more formal low-stakes quiz. Friday's quiz (Quiz 1) contains 15 questions, printed out, all taken from the Brockington Homework booklets discussed above. Usually 5 of these questions will be on the most recent topic being studied, 5 will be from the current term, and 5 from topics studied in previous terms or years. All students receive the same quiz, regardless of set

– although as we get closer to exams there will be a Foundation and Higher version to account for the different content.

Students take Quiz 1 at the start of Friday's lesson, in silence. As described above, the timings are pretty flexible, and will vary between sets. As discussed in Chapter 7, I favour differentiation by time spent on activities rather than giving students different activities. When half the students have finished, the teacher announces that there is one minute to go.

The rest of Friday's lesson is dedicated to going through Quiz 1. The teacher projects up the answers, students mark their own, and are encouraged to ask if they are stuck. No marks are recorded or collected in by the teacher, and all the time the emphasis is that this is a learning tool, not an assessment tool. The time needed to go through these will again vary by set, but also the complexity and recency of the topics covered. Sometimes it may take all lesson, and is likely to involve plenty of discussion, as well as the principles of modelling and examples detailed in this book. Teachers are encouraged to have extra questions available on the topics covered in the quiz in case students need extra practice, and that is where Corbett Maths or classics such as Ten Ticks worksheets come into play. Any time left over is spent using UKMT questions, or Purposeful Practice tasks as described in Chapter 10. If students are still stuck at the end of the lesson, they are encouraged to attend a maths clinic at lunch.

Homework over the weekend is to prepare for Monday's follow-up quiz – Quiz 2.

Quiz 2 contains the exact same questions as Quiz 1, in the exact same order, but with the values changed. Quiz 2 happens at the start of Monday's lesson, again in silence. This time students' papers are collected in, and then it is back to the scheme of work for the remainder of the lesson. Quiz 2 is marked by the teacher, scores are recorded on a central spreadsheet that all maths teachers have access to, and papers are returned to students the next day with ticks, crosses and a score. No personalised feedback is given – see Section 12.10 on delaying feedback.

Here comes the controversial, but incredibly important bit. Each class has a benchmark score that will be set by the teacher based on their knowledge of the class. It is always above 80%. If students fall below this benchmark on Quiz 2, then parents are contacted, and students are placed in detention.

I know this sounds ridiculously harsh, but here is the rationale. Students know exactly what is coming up in Quiz 2. If they listen and are active during the review process on the previous Friday, seek help if they are still stuck, and then put in the effort at home, then there are no excuses. Quiz 1 is a test of ability, initial understanding and prior knowledge, so it would be unfair to impose a performance-based sanction here. Quiz 1 is a learning tool. Quiz 2 is fundamentally different. Performance in Quiz 2 is determined by effort, not ability. Why should students not be able to get 90%+ on these tests? Of course, everyone makes daft mistakes, hence my reluctance to set the benchmark at 100%, but if we have high expectations for our students, and give them every opportunity to meet these expectations, then they will rise to meet them.

It is early days, but already we are seeing a marked improvement on basic skills, effort outside of the classroom, and an increased awareness amongst students of the distinction between what they think they can do and what they actually can do. For the Quiz-Homework-Quiz combo to work, everyone needs to buy into the idea. Consistency is needed amongst staff, and support is needed from parents. We have found discussing the research and rationale behind the idea to be the key to this. Sure, there will be teething problems. Sure, there will be exceptional circumstances that call for a certain amount of flexibility. But I am convinced this will be the single most effective policy we have ever introduced.

What I do now

For students to see tests as the tools of learning that they can be – and hence to reap all the benefits discussed in the previous section – counter-intuitively they need more testing, but of the lower-stakes variety.

I used to say to myself: *I can't afford the time to do a test every lesson*. Now, I realise I can't afford not to.

12.9. The Pretest Effect

What I used to think

I used to think that there was no point testing students on things that they didn't know. Hence, I would save all my testing until the end of the topic.

Surely I was right this time?

Sources of inspiration

- Barton, C. (2017) 'Robert and Elizabeth Bjork', *Mr Barton Maths Podcast*.
- Kornell, N., Hays, M. J. and Bjork, R. A. (2009) 'Unsuccessful retrieval attempts enhance subsequent learning', *Journal of Experimental Psychology: Learning, Memory, and Cognition* 35 (4) pp. 989-998.
- Richland, L. E., Kornell, N. and Kao, L. S. (2009) 'The pretesting effect: do unsuccessful retrieval attempts enhance learning?', *Journal of Experimental Psychology: Applied* 15 (3) pp. 243-257.
- Roediger, H. L., Putnam, A. L. and Smith, M. A. (2011) 'Ten benefits of testing and their applications to educational practice' in Mestre, J. P. and Ross, B. H. (eds) *Psychology of learning and motivation: cognition in education, Vol 55*. San Diego, CA: Elsevier Academic Press, pp. 1-36.

My takeaway

Giving students a test on some material that you are going to teach them at some point in the future clearly has a benefit to you, their teacher. It is likely to identify any existing knowledge that you can subsequently build upon, as well as indicating any misconceptions that you can plan how to resolve. As I discussed in detail in Chapter 11, diagnostic multiple-choice questions are well suited to both of these aims.

However, as I teased back in Chapter 11, there may be another benefit to giving students a test on unfamiliar material – the Pretest Effect. This is where being tested on material you have not studied before actually enables

you to better understand and remember the material when you subsequently do study it.

I first came across this during my podcast interview with Robert and Elizabeth Bjork. They described a series of small-scale experiments they had run with students in science lectures. Some students were given a test on material that they could have no hope of getting right as they had never studied it before, and unsurprisingly performed poorly. They then sat through lectures on that material as normal. In the test at the end of the series of lectures, they significantly outperformed students who had not experienced such a pretest.

The Bjorks explained that the Pretest Effect is a relatively new phenomenon in the world of research. It is hypothesised to work by making students better prepared to attend to and learn the critical conceptual features of the targeted concepts during the subsequent instruction. This may be related to the concept of *priming* discussed in Section 11.6, although there we specifically looked at including retrieval of previously acquired knowledge.

There are several studies that seem to confirm the Pretest Effect.

One of Roediger *et al's* (2011) 'Ten Benefits of Testing' is: 'Testing causes students to learn more from the next study episode', explaining that when students take a test and then re-study material, they learn more from the presentation than they would if they restudied without taking a test.

Richland *et al* (2009) found that generating a wrong answer increases our chances of learning the right answer. In this study, one group of students was given the text on which they would be tested – passages with key facts marked – while a second group was given the opportunity to memorise the questions they would be asked, a third group was given an extended study period, and a fourth group was given a pretest. Even though they got almost every pretest question wrong, students in the pretesting group outperformed all other groups on a final test, including those students who had been allowed to memorise the test questions. It would seem that the act of unsuccessfully attempting to answer questions has a greater effect on learning than studying the questions on which you are to be tested.

Of all the desirable difficulties we have looked at in this chapter, the Pretest Effect feels the most counter-intuitive to me. We have the issue that none of these studies took place in the domain of mathematics, although the Bjorks' description of the science lecture series is not a million miles away. But how could it help students' learning, for example, to give them a test on simultaneous equations before they even know what simultaneous equations are and the process for solving them? Indeed, surely it will cause students to rely on ineffective strategies (such as trial and improvement), which are the very strategies that I argued in Section 9.4 may not lead to learning? Then again, perhaps the mechanisms of the Pretest Effect work by not only priming students to attend to the relevant areas of long-term memory when the new material is subsequently taught, but also by providing a purpose for learning that material as described in Section 2.5. Students have attempted this strange question, struggled, and now want to know how to solve it.

But what about the impact the inevitable low scores will have on students' morale? We have seen how the desirable difficulties we have looked at so far lower performance, but surely no more so than a test on something a student has never seen before. Indeed, the concept of *productive failure* discussed in Section 9.4 is related to the idea of a pretest, and a key argument against productive failure was that any benefits could be outweighed by the negative effect on student morale. This point is articulated really well by Richland *et al* (2009):

> When a learner makes an unsuccessful attempt to answer a question, both learners and educators often view the test as a failure, and assume that poor test performance is a signal that learning is not progressing. Thus, compared with presenting information to students, which is not associated with poor performance, tests can seem counterproductive. Tests are rarely thought of as learning events, and most educators would probably assume that giving students a test on material before they had learned it would have little impact on student learning beyond providing teachers with insight into their students' knowledge base. In terms of long-term learning, however, unsuccessful tests fall into the same category as a number of other effective learning phenomena – providing challenges for learners leads to low initial test

> performance, thereby alienating learners and educators, while simultaneously enhancing long-term learning. The current research suggests that tests can be valuable learning events, even if learners cannot answer test questions correctly, as long as the tested material has educational value and is followed by instruction that provides answers to the tested questions.

If we are going to do a pretest, students need to know that their performance does not matter, and that the benefits to their learning may not be immediately apparent. The duration of the struggle during a pretest is not likely to be as long as attempting to solve long-form, unstructured problems, especially if the pretest is made up of quick-fire, short-form questions, interspersed with questions on subjects students have been taught. Hence, my reservations are not as strong compared to those I have with productive failure as described in Section 9.4.

What I do now

Before we start a new topic unit I will always ensure that some of the questions in the daily low-stakes quiz cover prerequisite baseline knowledge that students need in order to access the new content I am about to teach them, and I will further assess this using diagnostic questions at the start of the lesson as described in Section 11.6. This has the dual benefits of enabling me to assess my students' understanding as well as priming their long-term memories ready for the related new knowledge. That seems a no-brainer.

However, I have also started experimenting with the Pretest Effect. So, I may also include a few questions on what I am about to teach. I do not concern myself too much with students' answers to questions on this new material. The questions are there to take advantage of the strange workings of the Pretest Effect. However, I may learn some valuable insight from my students' incorrect answers, in particular from their choices of wrong answers to diagnostic questions.

As ever, I tell my students exactly why we are doing a pretest, and not to worry!

OceanofPDF.com

12.10. Delaying and reducing feedback

What I used to think

I used to think – in fact, I used to *know* – two things about feedback:

1. The more the better
2. The quicker the better

Thus, when students handed homework in on a Friday, many a Sunday afternoon was spent marking students' tests and homeworks, never daring to leave a cross or a question-mark on the page unless it was accompanied by a explanation, personalised comment, worked example and follow-up question, switching frantically between my red and green pens. I used to dread the weekend.

Then Monday's lesson would come around, and it was time for DIRT time. Students would be given their homework back, complete with reams of beautiful feedback, and within about three minutes I would wish I had never bothered. Students either told me they had finished everything, or claimed they didn't have a clue what to do.

Sources of inspiration

- Barton, C. (2016) 'Dylan Wiliam', *Mr Barton Maths Podcast*.
- Hattie, J. and Timperley, H. (2007) 'The power of feedback', *Review of Educational Research* 77 (1) pp. 81-112.
- Soderstrom, N. C. and Bjork, R. A. (2015) 'Learning versus performance: an integrative review', *Perspectives on Psychological Science* 10 (2) pp. 176-199.

- Wiliam, D. (2011) *Embedded formative assessment*. Bloomington, IN: Solution Tree Press.

My takeaway

A full discussion of what makes effective marking and feedback is beyond the scope of this book. For me, the bible on this matter is Dylan Wiliam's *Embedded Formative Assessment*. However, I would like to briefly discuss one fascinating finding that has had a significant and positive impact on my teaching and (I think more importantly) on my work-life balance.

Hattie and Timperley (2007) explain that feedback is one of the most powerful influences on learning and achievement, but this impact can be either positive or negative. Given that the insistence by schools on written feedback is one of the key contributing factors to the workload crisis that is driving teachers out of our profession, we need to ensure that any time we do spend giving feedback is as beneficial to our students as possible.

When I interviewed Dylan Wiliam for my podcast, he uttered the immortal line: '**The only good feedback is that which makes students think**' – a definite contender for tattoo number three. Despite the hours I put into my marking, my feedback did not make my students think. Ironically (or perhaps, unfortunately) the more detailed the feedback in terms of the information it provides the student, the less they need to think. By fully explaining where students have gone wrong and then providing scaffolded support in the follow-up question, students can pretty much breeze through DIRT time on autopilot.

One possibility to remedy this – and, indeed, it is the final of Bjork's desirable difficulties that we will consider – is delaying and reducing feedback.

Soderstrom and Bjork (2015) present several studies that suggest that delaying or even reducing feedback can have a long-term benefit to students' learning. Why? Well, because regular, immediate feedback can cause learners to become overly dependent upon it, perceiving it as a crutch to their learning. To relate this to themes covered in this chapter, immediate feedback prevents students from thinking hard and having the opportunity

to start to forget. Like everything else in this chapter, delaying or reducing such feedback is likely to have a detrimental effect on short-term performance, but a positive effect on long-term learning.

But if we don't identify student errors quickly, and provide our students with the means to understand and correct them, surely there is a danger that mistakes and misconceptions may become embedded?

Perhaps the important thing here is not necessarily the time delay between the answer and the feedback, but what happens in between. If students answer a question incorrectly, do not know they are incorrect, and then subsequently answer another 100 questions in that manner, then that misconception is likely to be difficult to remove – after all, practice makes permanent. However, if students answer just one question on a particular topic (and with my movement away from topic-based homework as described in this chapter, this is ever more likely), and are then compelled to reflect on their work following reduced feedback as described below, we may get the best of both worlds. Any negative effects may be negated, with the added bonus that students need to think more. This is pure speculation, but it makes sense combining the findings of several papers.

What I do now

When initially marking a piece of homework (or Monday's low-stakes quiz), I now give no feedback whatsoever – I simply put ticks and crosses. Wiliam (2011) takes this a stage further, suggesting that instead of, say, putting ten ticks and five crosses, we instead simply inform the students that there are five mistakes for them to find and correct. We thus make feedback into detective work. I then give this work back to the students and see if they can identify the source of their errors and correct them. When I next take the books in, this is when I will give more detailed, task-focused feedback, again following the advice of Wiliam (2011).

I am trying to distinguish between three key reasons students may have got a homework question wrong:

1. They do not understand the concept.

2. They made a careless mistake.

3. They could not be arsed.

I would argue that the kind of detailed feedback I was writing will only help in the first case. Students who have made a careless mistake may be alerted to that fact by the presence of a cross and be able to correct it; and students who could not be arsed will be no more or less arsed no matter what I write. Moreover, students who did not understand the question at the time have an opportunity to try again, reflect, self-explain or seek help from a peer during DIRT time. Either way, students are compelled to think.

There is another advantage to reducing and delaying such detailed feedback. Have you ever got half-way through a set of books, only to find you have been writing pretty much the same piece of (lengthy) feedback on most of your students' attempts to a particular question? Good God, I know I have, and it is infuriating. By simply ticking and crossing I can identify patterns and trends much more quickly, and hence deal with them more effectively. For example, if most of my students have got a particular question wrong, I will probably choose to address it using example-problem pairs, Intelligent Practice, and so on, as opposed to writing the same thing in 30 books.

Directly related to this is a suggestion that Dylan Wiliam makes in Hendrick and Macpherson's *What Does This Look Like in the Classroom?* (2017) which has revolutionised how I go through exam papers, particularly mocks. Instead of spending hours marking a set of papers, or spending a full lesson going through a paper – which may greatly benefit a student who has got 20% but not one who has got 90% – I now simply give students their papers back unmarked, sit them in small groups, give them a blank paper, and challenge them to come up with the best composite paper they can. As Wiliam explains, this technique 'gives students retrieval practice when they actually do the test, takes advantage of the hypercorrection effect when students find out their answers were wrong, and also provides an opportunity for peer-tutoring (which, the research shows, benefits both those who receive, and those who give, help)'. It also helps us poor teachers have something resembling a life outside the classroom.

OceanofPDF.com

12.11. If I only remember 3 things…

- As teachers we can only observe performance, and performance may be a poor indicator of learning.
- At times, we need to carefully make learning desirably difficult in order make retrieval more effortful so students can reap the rewards of a boost in a memory's storage strength.
- Low-stakes quizzes may be one of the most powerful tools of learning at a teacher's disposal.

OceanofPDF.com

CONCLUSION

So, that was how I wish I'd taught maths.

Writing this book has at times felt like therapy, and there have been occasions during the writing process when I have nearly fallen off the wagon. Just two weeks ago, I came across a five-star, cut-and-stick trigonometry discovery task that was just crying out to be used with my Year 9s. I need to stay strong.

Writing this book has also been both humbling and challenging. It has meant facing up to 12 years of mistakes, false beliefs and missed opportunities. It has meant confronting, revising, and completely ditching quite a few ideas and practices that I have held dear for many, many years. It has meant thinking back to things I have said to colleagues and students with an air of regret.

The irony is that despite reading over 200 books and research papers, recording over 100 hours speaking to the world's leading educational experts, and writing a relatively lengthy book, I genuinely feel I know less now about how to teach students mathematics effectively than I did five years ago. Back then, teaching was easy. I did what I had always done, and if it didn't work then it was probably the fault of my students. Back then I was Outstanding. Now I realise I was clueless, and I am only slightly less clueless now. My only comfort comes from the Dunning-Kruger Effect – as we learn we become more aware of our deficiencies, and I certainly still have plenty.

At the start of this book I presented ten things I used to believe that I no longer do. So, it seems only fitting to end with ten things I now believe that I

wished I knew when I first started teaching.

1. Students remember what they are thinking about, or attending to.
2. Planning for achievement is more important than planning for motivation.
3. Practice does not make perfect, practice makes permanent.
4. Students do not think and learn differently due to their learning styles, but because of their amount of domain-specific knowledge.
5. My choice of examples and exercises are the single most important part of my planning.
6. It often makes sense to teach the How before the Why.
7. Students can be struggling but not learning.
8. Effective differentiation is best achieved in terms of the time students spend on a task, not by giving different students different tasks to do.
9. Retrieval, predominantly through frequent low-stakes quizzing, is the key to long-term learning.
10. Perhaps, above all, the best thing I can do to help my students become the independent problem-solvers I want them to be is to carefully and explicitly teach them.

And now comes the far more difficult part – putting all of this into practice, and doing so consistently. Not only will this be a struggle for myself, but given that some of the strategies outlined in the book necessitate a decrease in short-term performance, it will also be a challenge for my students and any colleagues who come along for the ride.

Teaching is flipping hard. The challenges are summed up nicely by one of my favourite quotes (which is perhaps a bit long for a tattoo):

> After 30 years of doing such work, I have concluded that classroom teaching … is perhaps the most complex, most challenging, and most

> demanding, subtle, nuanced, and frightening activity that our species has ever invented … The only time a physician could possibly encounter a situation of comparable complexity would be in the emergency room of a hospital during or after a natural disaster.

Shulman and Wilson, *The Wisdom of Practice*, 2004.

But that is also why teaching is the best job in the world. There is never a boring day, nor an easy day. And the prize for getting it right – or as close to right as we possibly can – is that we get to change the lives of the hundreds of students we teach.

I'll take that.

OceanofPDF.com

ACKNOWLEDGEMENTS

The danger of thanking individuals is that you inevitably leave someone out who will never speak to you again. However, there have been a number of people who have helped me in so many different ways that I have to take that risk.

Firstly, my colleagues. Geraldine, my NQT mentor (I am pretty sure I still owe you some paperwork from 2004), and Debbie my first Head of Department – you supported me, gave me massive opportunities, and didn't shout at me too much for my shoddy marking. Lorna and Anne, without you I would not have become an AST and had the privilege to learn from all the teachers I got to watch and work with. No one ever tells you how hard moving school is, and without the support of Karen, Alison and Andrea in those early Thornleigh days, I wouldn't have made it to Christmas. Colm, it was an honour to be your NQT mentor, and to learn from you, although it is slightly annoying that you are now my boss. Simon, for helping me realise my dream in creating Diagnostic Questions; and Ben for helping us avoid going bankrupt. And to the maths departments of Range High School and Thornleigh Salesian College, past and present, for indulging my new ideas with enthusiasm (most of the time), and allowing me to learn from you – I have been very lucky indeed.

I would like to thank each and every one of the students who I have taught over the last 12 years. I am sorry for all the times you told me you didn't get it, and I assumed it was your fault (although, sometimes, it definitely was). Teaching is the best job in the world because of you, and I am so proud of all you have achieved.

To my podcast guests. You are my heroes, and it is you that have started me on this journey that has transformed the way I teach. And to the people who listen to my podcast on the way to work, whilst walking the dog, paving the driveway or even enjoying a barbecue – it means so much to know that you are enjoying the conversations and learning as much as I am. Long may it continue.

To the authors of the following books that have significantly changed the way I teach, and which I would encourage every teacher to read:

- *Why don't students like school?* – Daniel Willingham
- *The seven myths of education* – Daisy Christodoulou
- *Embedded formative assessment* – Dylan Wiliam
- *What every teacher needs to know about … psychology* – David Didau and Nick Rose
- *What does this look like in the classroom? Bridging the gap between Research and Practice* – Carl Hendrick and Robin Macpherson
- *Teach like a Champion 2.0* – Doug Lemov
- *What if everything you thought you knew about education was wrong?* – David Didau
- *Memorable Teaching* – Peps Mccrea

To the bloggers who inspire me and challenge my thinking all the time, sharing their expertise with such generosity. I will definitely cause offence by leaving people out here, but the blogs that have had the most influence on my recent thinking are:

- *Filling the pail* – Greg Ashman
- *Science and Education* – Daniel Willingham
- *To the real* – Kris Boulton

- *The learning spy* – David Didau
- *Until I know better* – Dani Quinn
- *Resourceaholic* – Jo Morgan
- *Median* – Don Steward

To Alex and the team at John Catt Educational for approaching me about this book and putting their faith in me. And my undying gratitude and respect to Kris Boulton, not only for inspiring me on my podcast, at conferences and via his blog, but for being the first person to read this book. It felt like I was handing in a piece of homework I was a little unsure of to a teacher I so desperately wanted to impress, but Kris corrected errors with patience, pointed out omissions with kindness, and hopefully helped turn incoherent ramblings into something that makes sense.

Finally, the two special ladies in my life. Mum, thank you for believing in me and providing me with all the opportunities I had, even though it was pretty tough at times. And Kate – when we met on that date seven years ago, I don't think you realised what you were letting yourself in for by getting together with a workaholic teacher who, for example, would choose to write a book about teaching maths whilst on a romantic holiday in Kos. I know being married to me can be an absolute pain, but I really, really could not have done any of it without you. You are my world.

OceanofPDF.com