Teaching Complex Ideas

How to Translate Your Expertise into Great Instruction

ARNOLD WENTZEL

ROUTLEDGE

Teaching Complex Ideas

Integrating insights from learning science with practical guidelines and stepwise approaches, *Teaching Complex Ideas* helps educators masterfully translate their expertise into easy-to-understand, interesting, and memorable instruction. Covering areas such as identifying the critical ideas within a complex topic, designing clear explanations, and making lectures useful and engaging, this resource brings together subjects and skills never before adequately addressed in a single book. Using real world examples and full of practical tips, this book guides college instructors to improve their understanding of their subjects, select the most valuable ideas to teach, and integrate those concepts with other aspects of teaching such as presentation design, technology, and assessment of understanding. This practical book helps professors at any stage in their career convert even the most complex ideas into great teaching.

Arnold Wentzel is Professor of Innovation Methodology, Management, and Research Writing in the Faculty of Industrial Engineering at the Universidad Antonio Nariño, Colombia.

This page intentionally left blank

Teaching Complex Ideas

How to Translate Your Expertise into Great Instruction

Arnold Wentzel

Routledge
Taylor & Francis Group
NEW YORK AND LONDON

First published 2019
by Routledge
52 Vanderbilt Avenue, New York, NY 10017

and by Routledge
2 Park Square, Milton Park, Abingdon, Oxon, OX14 4RN

Routledge is an imprint of the Taylor & Francis Group, an informa business

Library of Congress Cataloging-in-Publication Data
A catalog record for this title has been requested

Visit the eResources: www.routledge.com/9781138482371

ISBN: 978-1-138-48236-4 (hbk)
ISBN: 978-1-138-48237-1 (pbk)
ISBN: 978-1-351-05811-7 (ebk)

Typeset in Sabon
by Integra Software Services Pvt. Ltd.

Contents

Figures

Tables

Preface

THE STORY OF HOW I (ALMOST) FAILED AT TEACHING

Let me tell you my story. It is not because I am one of those 'super-teachers' whom we all admire – I cannot live up to that standard. No, I tell it because when I say that anyone who knows their subject can learn to teach well, it is because this is what happened to me.

At college, my fellow students used to say that I should never teach. I am socially awkward, and, back then, I would bore people with irrelevant details and make the simplest ideas sound complicated. And, they were right, as I found out twenty years ago.

One day, out of desperation, I applied for a teaching job at a private university for rich kids. I got the job but almost got fired within a month. My students were to be undergraduate Economics students who, I was told, were quite arrogant and entitled. But, "Don't worry" people told me, "Just remember you know more than them." Listening to them was my first mistake. Walking into my first class I was nervous but, fortunately, overconfident in the power of my content knowledge. After an introduction to the topic, I moved into the more complex material, where my supposed 'expertise' lost its power, and my audience too. As bored and puzzled expressions swept across the students' faces, I lost my confidence, my hands started shaking visibly, and, most embarrassingly of all, I became confused – making it look like I was not an expert after all.

The students' first course evaluations were brutal – on a scale of 1 (excellent) to 5 (awful), I scored worse than a 4 on average. Returning home that day, I wondered if I would always be below-average in most things. Fortunately, I didn't get fired only because I committed to improving dramatically by sitting in on other lecturers' classes, reading books on teaching, experimenting and practicing my lectures for hours. By the end of the semester, I was rated 2.5 by the same students – average, but good enough to stay on.

Strangely, reaching an average level inspired me to try harder. I tried my hand at teaching a wider variety of courses to as many audiences as possible: from undergraduate Arts students to MBAs and PhDs, in courses that included: Accounting, Economics, Ethics, Education, Finance, Management, Mathematics, and Research Writing. I failed at some of them too, but improved gradually.

Then, ten years ago, the constant improvement paid off when I was awarded the Teaching Excellence Award of the university for which I worked at the time. Eventually, I started teaching teachers how to teach, and this forced me to make my implicit knowledge as explicit as possible and distill what I learned over the years into simple and useful techniques. This book is the result.

THE PURPOSE OF THIS BOOK

This book has a unique unifying theme: *good teaching starts from ideas*.

I wrote this book because it is the kind of book that I was looking for when I almost got fired from my first teaching position. At that time, I found most books on teaching of limited use because they were either: too student-centered, lacking connection to the complex ideas I needed to teach; or they focused on content to the exclusion of how to connect it to students. I could find nothing that taught experts in complex ideas how to turn that into great instruction.

Great teachers of complex subjects are those who are idea-centered – who make ideas matter to their students. They are student-centered, not by explicitly aiming to be so, but as a natural result of the way in which they teach the content.

Content is fundamental to teaching, but, as my story shows, there is good reason why the emphasis on content knowledge in teaching has a bad reputation. People like Lee Shulman repaired and revitalized this emphasis with his broader concept of pedagogic content knowledge, which is the *combination* of the knowledge of the ideas in a subject *and* the knowledge of how to teach these ideas. Shulman (1986 p.9) summarized it as "…in a word, the ways of representing the subject that make it comprehensible to others".

Lack of pedagogic content knowledge can be hidden or forgiven in simpler courses and when students are younger. But as students become older and subjects become more complex, a teacher's pedagogic content knowledge becomes more important. The teachers that these students remember are often not those who understood the ideas of complex subjects, but who understood the ideas well enough to be able to explain them and make them interesting to *others*. Like Richard Feynman, they are the 'great explainers.'

If you are reading this book, you probably already understand your subject. But it is not enough. This book takes you further: it aims to show you how to transform your understanding and expertise into great teaching. It explains how experts should think in order to make their ideas teachable – that is, clear, interesting, memorable and able to be expanded into future learning.

It has long bothered me that instead of expanding, the effect of teaching wanes after a course. Usually, a professor conveys as many ideas as can be crammed into the limited time, but, within weeks after the course's completion, it is as if the course never happened. Students remember very little, if anything, and use even less. The result is something that starts wide, but ends up very narrow – the opposite of Figure 0.1.

However, by applying five core skills, learning can expand, as shown in Figure 0.1. The first skill is knowing what not to teach. Courses that contain complex ideas are among the most useful, but they can be overwhelming. Fortunately, for the teaching of complex ideas to increase in value, students do not need to know and remember everything. Quite the opposite. They just need to understand that which can help them to expand and revise their understanding indefinitely – and that is often much less than expected. Following the Pareto principle, I estimate that in most of these courses, 20% of the ideas give students access to 80% of the understanding, both now and in the future. These are the ideas that can expand into greater understanding with time, and the other four core skills are applied to them.

Some ideas will be difficult to understand, so knowing how to explain them will give students the sense of efficacy as they realize

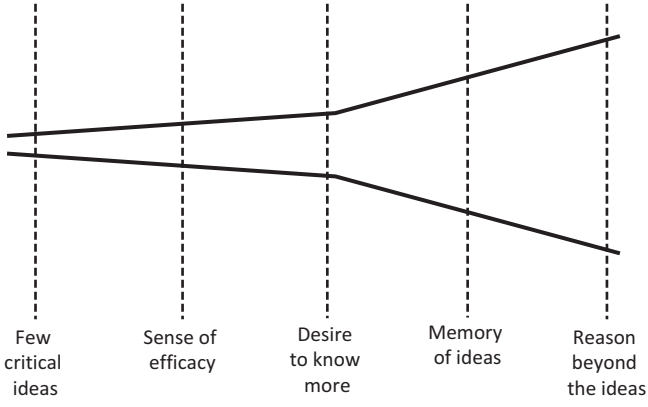

Figure 0.1 Learning That Expands Over Time

that they have the ability to understand. By making these ideas more interesting, the professor evokes a desire in his/her students to know more – either at that time or later in life. Enabling students to remember the important ideas then ensures that students have access to those ideas when they leave the course. Finally, teaching students how to reason with the important ideas will allow them to expand their understanding and go beyond what was taught in the course. This is because we do most of our learning outside the classroom anyway, and expandable learning makes this easier and faster.

Those who master the five core skills (as explained in Chapters 1-6) will find that their teaching is not in vain but will ripple through their students' lives for decades to come. In Chapters 7-10, the core skills are then connected to the more conventional skills discussed in other books, namely: assessment and the use of technology.

THE AVERAGE APPROACH TO BECOMING GREAT

I found that the most useful books on teaching are usually the non-theoretical books. While they were often inspiring, these books often made me feel rather inadequate as I tried to measure myself against the 'super-teachers' in such books. I admit that this still happens to me – especially while I was doing additional research for this book. Because my approach is different, this book will hopefully not do this to you.

My approach to improvement was, and still is, based on inversion: to become good, simply stop doing what is bad. Figure 0.2a shows an average teacher. Like everyone else, in some skills, he is rated above the average, in others, below the average. If he identifies his below average skills, and simply raises them to an average level, the result is Figure 0.2b. He is now well above average, and on his way to improving further.

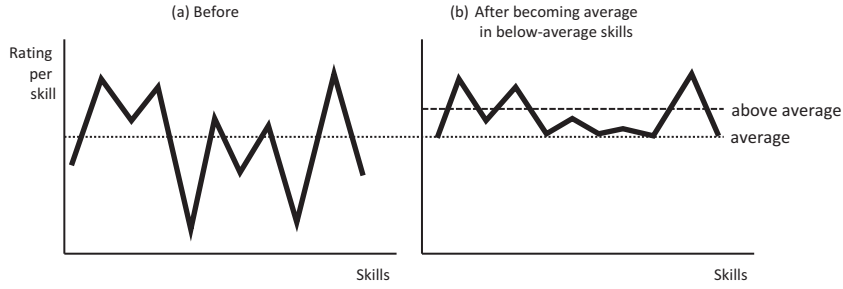

Figure 0.2 How to Become Above-Average

This works, because it is easy to be average, and you can learn even from average teachers how to improve. It is surprisingly easy to learn from them and you don't need a book for that. Also, in a world where most people cluster around the average, your students will notice your improvement and gradually respond more positively to your teaching. I have found this quite motivating. Once this motivation starts working in you, it will pull you up further, from above-average to good, to great.

Teaching is complex, with many skills to master, and inevitably most of us are below average in some of them. Fortunately, not all skills are critical. The critical skills are the five core skills discussed in Chapters 1–6. If you are below-average in any of them, the techniques in this book will help you raise them first to an average, and then to an above-average level. Since these skills are critical, understanding them will equip you with the insight to improve less critical skills, and raise your performance in any problematic ones to average or higher levels.

WHAT YOU WILL GET FROM THE BOOK

Teaching can be frustrating for experts, especially those of us who are professors and teachers. Our audience sometimes struggles to understand us, and despite our best efforts, they never quite share our excitement about the ideas. This book will be useful to such professor-experts or really all experts who want to, or need to, share their ideas.

In this book you will learn how to master the few skills that will make the biggest difference to well you translate your expertise into instruction. You will discover the following:

- How to teach less so that your audience understands more (Chapter 1)
- A step-by-step approach to creating clear explanations of complex ideas without dumbing them down (Chapters 2–3).
- Techniques for helping your audience to remember important ideas based on the science of forgetting and powerful mnemonic devices (Chapter 4).
- The unique F^3 approach to making even boring ideas more interesting (Chapter 5)
- How to develop students' reasoning ability without wasting time trying to teach them how to reason (Chapter 6).
- How to use a little-known century-old technique that transforms assessments into learning experiences and how to eliminate cheating without coercion (Chapter 7).

- Detailed advice on how to design assessments that accurately measure learning and how to use some unconventional assessments (Chapter 8).
- A systematic approach to resourceful teaching that works even when technology fails (Chapter 9).
- The simplest way to design presentations in a way that will significantly increase the value your audience gets from your lectures (Chapter 10).

Throughout the book, the insights from learning science are integrated with practical guidelines and stepwise approaches, so you will benefit from the science without having to wade through theory.

HOW TO USE THE BOOK

The obvious audience for this book is experts who need to teach, which includes, but is not limited to university professors, professors and students in teacher training programs, and teachers at the higher grades of secondary school (Grades 10–12). The guidance the book provides would be especially useful to those teaching courses that are idea-intensive (e.g. Economics, Social Studies, Geography, Biology etc.), and somewhat less relevant to those teaching purely skills-based courses (such as Mathematics, Accounting, Art etc.).

Regardless of whether you are reading this book in order to improve your teaching or as a textbook in a teaching methodology course, the critical chapters are the first six, and it is best to read them in the order in which they appear. The remaining chapters can be read in any order.

Because this is a practical book, every chapter contains exercises that ask you to apply the skills to your teaching. I recommend doing those exercises before moving on to the next chapter. Also, consult the Routledge website for this book for guidance on selected exercises and a supplementary chapter.

Throughout the book, I use boxes to refer you to more information on the web. Instead of giving you links that may be outdated by the time you read this, I recommend keyword searches in Google that will lead you to the right resources, as I do in the example below.

BOX EXAMPLE

For more information on this topic, enter the following in the Google search box: **cat AND (dog OR canine)**. You can also try **"the history of pets"** for more information from a different perspective.

Notice that I put the keywords that you should enter in a bold font. When I use AND, OR, brackets or double quotes, you should use them too, exactly as shown.

eRESOURCES

Additional resources for download can be found by visiting the book product page at www.routledge.com/9781138482371. Once there, click on the tab that reads "eResources" and then select the file(s) you need, which will download directly to your computer.

REFERENCE

Shulman, L.S. 1986. Those who understand: Knowledge growth in teaching. *Educational Researcher*, 15(2):4–31.

This page intentionally left blank

Acknowledgments

Firstly, thank you to all my students over the last thirty years who put up with my teaching – I know it wasn't always great, or good. I sometimes learned at your expense, the rewards of which were reaped by future students.

I learned about teaching from many people: my late grandfather, Charl de Beer; my late father; my mother; the handful of school teachers who helped me (Mrs Steenkamp, Mrs Kotze, Mr Keyser, and Mr Wright); university professors who inspired me (especially Roelof Botha, Christie Schoeman, and Lorraine Greyling); Annette Lombard (the one who didn't fire me); and many more. Special thanks must go to people who commented on early drafts, like Naiefa Rashied, and also Jacqueline Jones who again did excellent copy-editing, as she did on my previous books. I wrote the first draft of this book while on sabbatical in Colombia – so I want to thank Eleonora for making this possible.

This page intentionally left blank

Chapter One

Know What Not to Teach

If we want people to understand as *much* as possible, we should teach them as *little* as possible. Good teaching starts when we know what *not* to teach.

Only experts have the understanding to teach like this, but only those who are teachers and professors at heart, will care enough to do so. Even if you are both – what I call a professor-expert – you may realise this is easier said than done. This chapter will make it easier to do.

CAN YOU FIND THE 'LEAD'?

Nora Ephron, the screenwriter of some of the most iconic romantic comedies (such as *When Harry Met Sally*), learned this principle early. When she walked into Charles Simms's Journalism class in 1956, he explained how to write the lead for an article, that is, the first paragraph that contains the most important information. If people read only the lead, and nothing else, they should understand the meaning of the article.

Simms read a list of facts and asked them to write the lead for an article in the school newspaper. A more modern version of what Simms told them in class that day would be something like this: "The principal announced today that all teachers will travel by bus to Cape Town to attend a full-day education conference on Thursday. The conference will feature various world-famous education experts such as Ken Robinson (author of *Creative Schools*) and Carol Dweck (author of *Mindset*). Our teachers will learn about the latest educational approaches and introduce them into their classrooms."

Now, do what Simms asked of his class: In one short sentence, write the lead of the article. Is it "Teachers will be addressed by famous education experts" or rather "Teachers will soon be using new teaching methods"? Or something else?

The lead is neither of those, and most people miss it. The result of writing the wrong lead is that the readers of the article will miss the

critical information. And they will find this quite distressing, as you will realize when I reveal the lead that was hidden among the facts.

The same thing happens when we teach. When confronted with the many facts of a topic, many of us cannot identify the lead. So, when we teach these facts, our students miss the meaning of the facts and forget it all soon afterward. This chapter will help you find that 'lead' – not only for Charles Simms's article but for any topic you have to teach. Once you have the lead, you will know what is worth teaching. By implication, everything else is *not* worth teaching in the limited time you have. The proportion of content not worth our time is surprisingly large, as you will see.

LESS TEACHING EQUALS MORE UNDERSTANDING

Benjamin Franklin once wrote to a friend: "I apologize for the long letter. I had not the time to write a short one." This is true of teaching too. Good professors not only know what to teach; they know what *not* to teach. Their classes tend to be shorter and more focused. But when I was still a struggling novice professor, I really should have introduced most of my lectures by saying: "I apologize for the long lecture. I had not the understanding to prepare a short one."

For a person with little understanding, everything is important. In contrast, good professors reveal their understanding in their ability to cut to the heart of a topic (Bain, 2004) and to recognize the deeper principles (Ambrose, Bridges & DiPietro, 2010). Instead of covering as many ideas as possible, they explore a few critical ideas in depth from different perspectives (Bransford, Brown & Cocking, 2000).

Misha Gromov (2011) explains that understanding is the ability to see what information is redundant and knowing how to compress that information into a few valuable ideas or patterns. One cannot understand a telephone directory because there is very little that is redundant – every number is important. You can only memorize it and repeat it as a list. But one can understand physics, and most other subjects, because it is possible to reduce them to 'laws' or 'critical ideas' from which less important ideas are derived. Someone who understands a subject can see that most information contributes very little to understanding.

A friend of mine complains that his wife usually gives him the "six-hour version of a six-hour event". Someone who lacks understanding may do the same. If the topic in the textbook covers 50 pages, such a person would teach the 50-page version of the 50-page chapter, and will probably expect no more from his students.

Compare that to the 'great explainer', Nobel Prize-winning physicist, Richard Feynman, who made a point of helping his audience see the

critical ideas, and 'grew' his lectures only from those ideas. He often reiterated that the vast collection of scientific facts could be compressed into a small number of ideas or laws. Once he went as far as compressing all of physics into one sentence (Feynman, 2011, p.4):

> If … all of scientific knowledge were to be destroyed, and only one sentence passed on to the next generations … what statement would contain the most information in the fewest words? I believe it is … that all things are made of atoms – little particles that move around in perpetual motion, attracting each other when they are a little distance apart, but repelling upon being squeezed into one another.

He believed that by applying thinking and imagination to just a few critical ideas, one's learning could expand indefinitely. This is something we will discuss further in Chapter 6. For now, let us look at how to find the few critical ideas in any topic.

CONSIDER WHAT YOUR AUDIENCE IS BECOMING

You cannot identify the critical ideas unless you first know whom you are teaching, and what would be important to them (not to you) when they put their understanding to use. Don't think of your audience simply as who they are at the moment. Rather, think of them as who they will be in the future, after they are transformed by the understanding you want them to have. I learned this while working at a university whose teaching philosophy was 'learning to be' as opposed to 'learning about'.

If you teach to get students to 'learn about' a topic, you will be tempted to convey many ideas in an unselective manner. The implicit assumption here is that the audience remains unchanged, except that they will have collected more ideas. However, if you teach to get your students to learn 'to be' or 'to become' something, you will choose much more carefully which ideas you want them to engage with. You now aim to change the audience, so the audience matters.

Richard Feynman (2011) started his introductory physics lecture with exactly this philosophy: he assumed that every person was going to be a physicist. Of course, he did not mean that he wanted everyone to become a physics professor or a full-time physicist, rather he wanted them to think like physicists to better understand what was happening around them in the physical world.

So, spend some time thinking about this before doing any other preparation. Ask yourself: What is a realistic expectation for what

my students will be, or become, outside the course, as a result of their understanding? Sometimes it is obvious – when you teach education, you want most of your students to become teachers. Sometimes, it is less clear. When I teach economics, do I want them to become economists and policymakers? Perhaps only if I am teaching advanced courses. When teaching an introductory course, I simply want them to think like an economist when reading the news and for them to know how to use their understanding in investment or business decisions.

When you teach, you should engage with students as if they are, or are becoming, the kind of person you want them to become. This makes the critical ideas stand out. Instead of asking yourself the vague question: "What is important?" you will ask: "Why would a person, who is (being or becoming) a physicist/economist/etc., need to understand this topic?" What this does is shift the focus away from you onto the ideas and their relevance to your students.

Those who don't do this regularly make one big mistake. They teach as if they want students to become like them. This kind of subconscious narcissism eliminates empathy. It makes them think: "If it is important to me, it must be important to them", and this causes them to regard everything *they* know as important. Whatever you define as your audience, they should be *different* from you.

WORK BACKWARD TO WHO YOUR AUDIENCE IS NOW

Understanding the ideas that you teach should ultimately transform your students in some small way. But you need to be sure that this transformation is not too radical: it needs to be within the sphere of the 'adjacent possible'. Steven Johnson (2010) describe the adjacent possible as "a kind of shadow future, hovering on the edges of the present state of things, a map of all the ways in which the present can reinvent itself". The 'shadow future' is what you want your students to be, or become, and the 'present state of things' is their current understanding and interests. The seeds of students' transformation should already exist in their present state, and they 'reinvent' themselves as you connect the critical ideas to their present state.

Remember one thing about their present state: you are not teaching aspiring professors. They are usually young adults who (initially) can think of many other things they would rather be doing than sitting and listening to you. Compared to you, they have a very limited understanding of the ideas of the subject, and, as a result, have a much weaker

memory for the apparently simple ideas from last year's courses, or even last week's class.

There are many practical ways to make yourself familiar with your students' present situation, some mentioned in Box 1.1. As explained in the Preface, I use boxes to direct you to further information on interesting or useful topics. Instead of giving you links that may be outdated by the time you read this, I will recommend keyword searches in Google that take you to the right resources.

BOX 1.1 KNOWING WHO YOUR STUDENTS ARE NOW

To find strategies to help you to understand your students, try searches like: **"who your students are"** or **"understanding students"** (use the double quotes). One useful article is Blum (2016). Diagnostic assessments are good tools for determining how much your students know, and, for that, just enter **"diagnostic assessment"** in the search box. It is useful to get an idea of the misconceptions or false beliefs about a topic, and, to find these, simply enter: **misconceptions about [insert topic] or false beliefs about [insert topic]**. A more complicated method involves concept maps, and, for that, simply enter: **externalizing mental models**. Of course, nothing beats talking to your students informally, though these sources will help you to ask the right questions.

Once you notice the gap between students' present state and their shadow future, you should establish if the two states are 'adjacent' or easy to connect. Again, do not look at it from your perspective – the gap may seem minuscule to you, but may well feel like a wide, gaping canyon to your students. The way you will communicate the critical ideas, and how well students understand them, will depend on whether it is possible to build a bridge *within minutes*. If it takes longer, you are too far away from their present state.

A lack of awareness of the gap makes us suffer from the curse of knowledge (see Box 1.2). The more we know about something, the more difficult it is to imagine what it felt like to know a little. It will be useless if you find the right critical ideas, but students cannot understand your explanations of them. In Chapters 2 and 3, we'll get into the art of explanation. For now, forming an idea of the gap between who students are now, and who you want them to be, will help you choose critical ideas that are comprehensible.

BOX 1.2 CURSE OF KNOWLEDGE

For more information simply enter: **"curse of knowledge"**. There is an interesting experiment that illustrates it quite well, and to find that, enter: **"curse of knowledge" AND "tapping experiment"**. This experiment powerfully demonstrates how easy it is to regard information as obvious from our own perspective, while we are completely confusing those around us. If you enter: **cartoon "curse of knowledge" giving directions**, you should find a funny illustration of the idea in the form of a policeman giving directions to a confused man.

When you assume that your students know more than they really do, you will neglect to teach the critical ideas. You will most likely overwhelm them with details or teach them the ideas they need the least and neglect the ideas they need the most. Some of the critical ideas may seem obvious to you, but I have found it is dangerous to make such assumptions. Not teaching apparently obvious, but critical, ideas is a major cause of students losing interest in a course.

In short, knowing who you are teaching will give you a better awareness of what your audience is able to learn and which ideas they need to understand first. By itself, this does not solve the problem of knowing what to teach – which is what we turn to next.

WHY YOU ONLY NEED VERY LITTLE TO UNDERSTAND A LOT

A little knowledge may be a dangerous thing, as the saying goes, but a little *understanding* can be a powerful thing *if* it is of the right kind. You may have heard of the 80/20 rule or the Pareto principle: that 20% of your effort can produce 80% of the results. But this rule applies mainly to networked systems.

Fortunately, the understanding of a subject is a networked system. Consider the differences between facts and ideas, and knowledge and understanding. A fact is generally something we do not doubt, though this can vary across time, place and discipline, whereas as an idea is a connecting principle that binds different facts together. We *know* facts, but we *understand* ideas. Knowledge is the ability to list the facts that are connected to form an idea. One understands an idea when one can explain how and why facts are connected to form that idea and distinguish between what is critical and non-critical. Unlike facts, ideas can be

questioned, because not everyone will see the same connections made or select the same facts to connect.

Understanding is a network of connected facts and ideas, and the brain that enables this understanding is also a highly connected system. We should, therefore, find that that 20% (or less) of the ideas in most subjects give us access to 80% of the understanding of the subject.

However, the distinction between facts and ideas is rather fuzzy in practice. We don't find them in neatly separated packages; instead, they lie on a continuum of connectedness. The distinction depends on the perspective from which you view this continuum. For example, a number may be a fact in a telephone directory, but it may also be an idea if you view it as being the result of connected human observations and agreements, or when you start to see true patterns in the way numbers are constructed.

There are several points to take from this. Firstly, the recognition of facts and ideas will differ among disciplines. Secondly, if you understand an idea, you have already started to make connections. Thirdly, once we start taking these connections for granted, these ideas become facts. Lastly, we can connect ideas to form more ideas.

Coming back to the telephone directory example, it is quite useful to understand why the Pareto principle only applies to highly connected systems. Information is mostly a list of unconnected facts (telephone numbers), as represented in Figure 1.1a. Such information can perhaps be categorized, but one fact is not connected to any other fact. The information cannot be understood or compressed, only memorized. Every number stands by itself and knowing one number does not help you to know what the next number will be. To receive 5% of the benefit of a telephone directory, you need to memorize 5% of the numbers, and to receive 90% of the benefit, you need to memorize 90% of the numbers. There are no critical ideas, so 20% of the effort gives you 20% of the results.

When ideas are connected, as in Figure 1.1b, understanding becomes possible. Understanding one idea enables us to perceive the ideas that are connected to it by logic or experience. The well-connected ideas are the powerful ones. Take, for example, the idea right in the center of Figure 1.1b – it is connected to more ideas than the others. If this figure represented the knowledge of a topic, obviously you should learn this central idea first. Since this central idea gives us access to most of the other ideas, it would be possible to compress our knowledge into an understanding of the central idea and its connections. As you see in Figure 1.1c, this idea is a critical idea. If I had only a few minutes to teach the topic, I would teach this idea.

If you can find the 20% best-connected ideas, and only teach them well from different perspectives, students should gain around 80% of

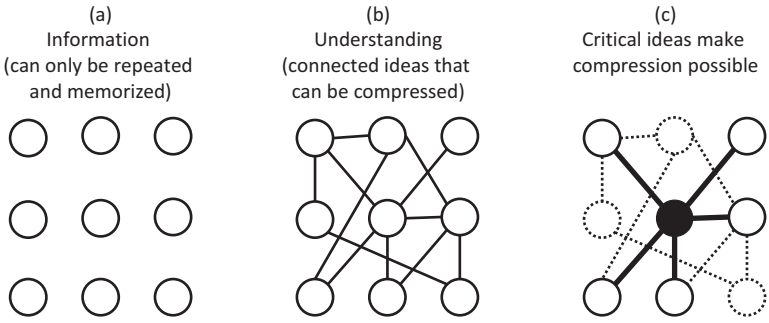

Figure 1.1 Information, Understanding, and Critical Ideas

the understanding. It follows that you need much less time teaching the non-critical ideas once students understand the critical ideas.

In network theory, this degree of connectedness is captured by the concept of centrality. In networks, especially ones that are scale-free (see Box 1.3), those nodes (circles in Figures 1.1 and 1.2) that have the most connections are the central nodes. One finds centrality in human knowledge too. There are many non-central ideas with a few connections and a much smaller number of central ideas with lots of connections. These central ideas are the critical ideas.

But don't take it too literally. As shown in Figure 1.2, if we map the connections between ideas in a topic, those with centrality will not necessarily be found in the center. Also, we may find that there is more than one central idea. In the figure, for example, there are two ideas that stand out as having the most connections.

Cognitive science has long recognized the networked nature of understanding, and educational theory is catching up. Applications of the Pareto principle are found in the literature on 'threshold concepts' (Meyer and Land, 2003) and 'big ideas' (Wiggins, 1989).

Threshold concepts are those that students find difficult, but once these concepts are understood, they have the ability to transform students' thinking about the subject because these concepts are so well connected. Understanding only a few threshold concepts makes it possible to understand many more ideas. Similar to threshold concepts, 'big ideas' are ideas that help us to make sense of a lot of information, and are useful for solving many problems. Grant Wiggins (2010) puts it like this: "An idea is 'big' if it helps us make sense of lots of confusing experiences and seemingly isolated facts. It's like the picture that connects the dots… in a complex field." Both threshold concepts and big

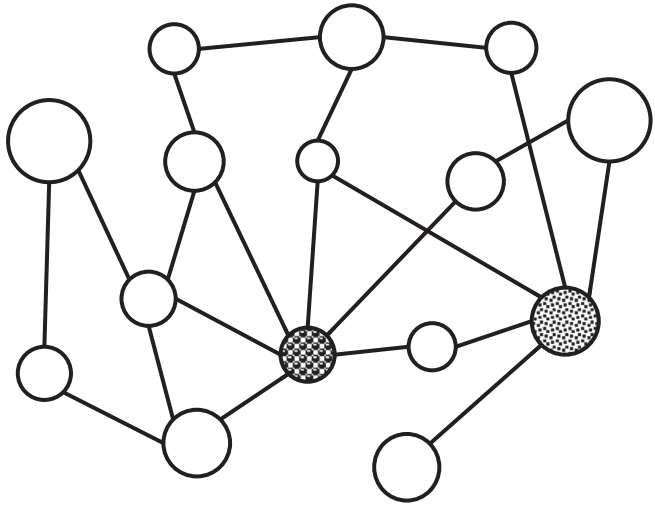

Figure 1.2 Central Ideas Are the Best-Connected Ideas

ideas acknowledge only a few critical ideas make most of the understanding possible.

BOX 1.3 THE BIG IMPACT OF JUST A FEW IDEAS

To better understand the 80/20 rule, search for **"Pareto principle"**. For a more technical explanation, search for **"scale free network" AND "power law" AND 80/20**. In scale-free networks, we find that a few nodes have the most connections, and this generates so-called power laws – very similar to what we find in networks of understanding. In a scale-free network, one can cut a large number of connections without significantly harming the performance of the network, so only a small number of nodes are really critical. This results in a reverse Pareto principle: one may remove 80% of the ideas and only lose 20% of the understanding. To get to the educational applications of these ideas, enter **"big ideas" AND Wiggins** and also try **"threshold concepts"**.

One can put the ideas in any topic or subject into three categories. At the core (see Figure 1.3), we have very few critical ideas. From them, we derive a bigger number of important ideas. The majority of ideas are nice to know, containing details that are useful to experts, but which do not significantly accelerate the understanding of a student. They are found at the periphery of Figure 1.3).

You should *not* teach nice-to-know ideas if students do not yet understand the important ideas, and you should *not* teach important ideas if students do not understand the critical ideas. Given the limited time, we should therefore *not* teach the majority of the ideas.

FINDING CRITICAL IDEAS AND PUTTING THEM IN SEQUENCE

To find the critical ideas requires insight. Even with insight, it is hard. Fortunately, there are techniques that make it easier and faster: the usefulness technique and the expanding elevator speech.

The Usefulness Technique

This technique is simple, but it requires imagination and experience. It starts from one of the most important questions a student can ask: "Why should I learn this?" or as Mathematics professor, Donald Saari, calls it "WGAD?" (who gives a damn?) (Bain, 2004, p.39). Since it is designed to focus on usefulness, it also makes you think how to bring the content to life.

It involves the following steps:

1. Ask who you want the students to be, or become, outside of the teaching situation. If you teach Management, the answer may be an entrepreneur, a CEO or an employee; or, if you teach Economics, it could be a policymaker or an investor. Try to identify only one or two important roles.
2. Then, ask yourself why a person in this role would need to understand the topic. How exactly would it help them in practice? What would they be able to do if they understood the topic? Try to find one, two or three reasons and write them in short, full sentences. If you have more than three or four statements, you may be trying to do too much.
3. Put the statements in a sequence that makes sense; that is, either as a narrative or in a logical sequence, as shown in the example below.

4. Expand on the critical ideas using techniques, such as the job story, imaginary conversation or the expanding elevator speech, as explained later.

For example, suppose you are teaching about inflation. If the course is in Management, one group that would be interested is the CEOs of companies, or perhaps managers of different departments (such as those in human resources or marketing). Many students aspire to be in one of these roles. When they are there, why would they need to understand inflation? Perhaps it would help the company to determine by how much to raise its prices, and also what salary increase to offer employees. So, one might write something like the following sentence. "It is important for a CEO to understand inflation because: (1) a company should not increase their prices faster than the inflation rate, because it may lose customers; and (2) a company should increase their employees' salaries at more or less the same rate as the inflation rate in order to avoid employee dissatisfaction and strikes."

In this example, I identified two critical ideas. They are not the only ones, so there is an element of judgment involved based on the prior knowledge of the students and the objectives of the course. It is not necessary to look at every possible role. Once students know how to use their understanding in one or two roles, they will be able to extend it to other roles.

It is best to teach the critical ideas in a sequence that is easy to remember. Sequence is under-researched and underestimated in the education literature, but Ritter *et al.* (2007) contains a good overview. What they found is that an appropriate sequencing of ideas enables people to achieve understanding with 30% fewer facts, so this supports what we are trying to do by organizing teaching around critical ideas.

Using an appropriate sequence means putting the statements of the critical ideas in an order that makes them flow into each other with the aid of connecting sentences. Here is an example, with the connecting phrase in italics: "A company should increase their employees' wages at more or less the same rate as the inflation rate in order to avoid employee dissatisfaction and strikes *and these employees will have fewer customers to serve if* the company's prices rise faster than the inflation rate." If you struggle to find a connecting sentence, it could mean that: you are missing an important idea; that the ideas should be placed in a different order; or that you should try a narrative sequence.

Using a narrative sequence means embedding the ideas in a story. Stories need not be exciting to be remembered. In fact, even rather mundane stories are easy to remember, as long as it is possible to imagine them.

For example, in this case, it may the story of a CEO encountering a series of crises that can only be resolved by an understanding of inflation. Or it may be the story of a human resources manager who discovers that employees are dissatisfied while negotiating with a labor union, and then using her insight to help her friend, the marketing manager, to solve his problem with price-setting.

BOX 1.4 STORYTELLING

To find out more about the use of stories, simply search for **"using storytelling in teaching"**. You may be surprised how some apparently boring subjects, like Accounting, can benefit from it.

There is usually no single correct sequence. It must just simply make sense. What you will notice is that, when thinking like this, you will often not follow the sequence of the textbook.

The Expanding Elevator Speech

Sometimes, the usefulness of a topic is not immediately evident or easy to highlight. In such cases, the expanding elevator speech technique may be useful. An elevator speech is a very short statement (about the time it takes to ride an elevator) in which you try to get a point across clearly. A good elevator speech will provoke questions and lead to a longer conversation later. It is the seed from which a conversation can grow, so the critical ideas of the longer conversation should appear in the elevator speech.

Figure 1.3 depicts an expanding elevator speech. It starts with a few critical ideas in the center. Expanding only on the ideas in the center leads to a more detailed explanation in the second ring. Then, expanding only on the ideas in the second ring leads to the ideas in the outer ring. The critical ideas are the seeds from which everything grows.

The easiest way to generate an expanding elevator speech is to ask yourself the following questions in the following order:

1. If you had only 5% of the time you usually have to teach a class or a course, what would you teach? For example, if the class was scheduled to be 45 minutes, you would identify the ideas you would focus on if you only had 2¼ minutes (5% of 45). The same applies if you were considering a course, though the total time would be much longer, perhaps in the region of 30–40 teaching hours. I

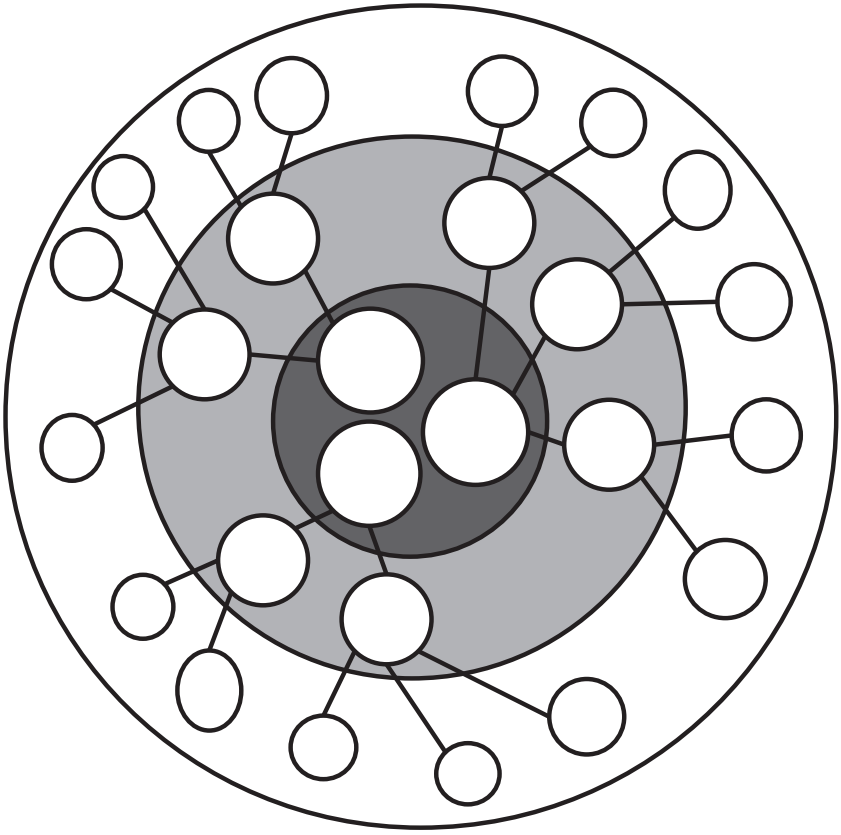

Figure 1.3 Expanding an Elevator Speech

actually prefer to start with a more extreme version: a one-sentence statement of what I want to argue in the lecture, which takes a few seconds. Here you find the critical ideas.

2. What is the best sequence for the critical ideas?
3. If you had only 20% of the time you usually have to teach the class, what would you teach? What you would teach in this time should follow from the critical ideas.
4. If you had 80% the time you expect to have to teach the class, what would you teach? You find this by expanding on the ideas in the 20% version. (To allow for flexibility in dealing with questions, longer discussions, and unforeseen events, it is best not to elaborate all the way to 100% of the available time.)

If you prepare from a chapter in a textbook, the final 80% version should not contain everything in the chapter, because it should focus only on the most important ideas derived from the 20% version, which, in turn, are derived from the critical ideas.

This technique can also be used to elaborate on the critical ideas you find with the usefulness technique. The usefulness technique is a different way to do steps 1 and 2, so you only need to apply steps 3 and 4 to help you elaborate on the critical ideas. It does require some intuition to create the 20% and 80% versions, so initially, you could try some of the elaboration techniques in the following sections.

An alternative to the expanding elevator speech is concept mapping (Box 1.5).

BOX 1.5 CONCEPT MAPPING

Any technique that maps knowledge as a network, such as concept mapping, would also be useful to find critical ideas. A concept map shows the possible relationships between concepts, and the relationships are indicated by arrows and linking words. In a concept map, the critical ideas are most likely those concepts with the most connections. For more information, search for: (**"concept mapping" OR "concept map"**) **AND examples**. For material created by the inventor of concept maps, search for: **"Joseph Novak" AND "concept maps"**.

ELABORATING ON CRITICAL IDEAS

Elaboration is not an excuse to introduce ideas that do not follow from the critical ideas. In other words, elaboration is *not* adding the ideas that you did not mention before. This will simply take you back to trying to teach everything. Instead, it should do two things: (1) *grow* the lecture from the critical ideas; and (2) *deepen* the students' understanding of these same critical ideas.

Growing a lecture is similar to growing a tree. Think of the topic as the trunk, the critical ideas as the few main branches, the important ideas as the twigs and the nice-to-know ideas as the leaves.

This leads you to teach in a different sequence to the textbook. Starting on the first page of a chapter and working all the way to the last page in the exact order it appears is not productive, since this does not help students to distinguish between what is critical, important and 'nice to know'. If the topic is complex, students will most likely become overwhelmed by all the details.

Mathematics teacher, Kalid Azad, compares it to revealing a photo. Following the textbook sequence is like unveiling the photo, starting on one side

Figure 1.4 Two Approaches to Elaboration

Photo credit: Ruan Klopper Photography

and working to the other end (Figure 1.4a). As a result, it takes the students a long time before they understand what they are looking at. By contrast, it is more useful to see the big picture from the start, but not in as much detail, and to fill in the details over time (Figure 1.4b). In this way, the students understand from the beginning and deepen their understanding gradually.

Elaboration by Growing: The Imaginary Conversation Technique

Earlier, I mentioned that I often start with a short statement of one sentence and grow a course from there. This is similar to what Richard Feynman did when teaching physics, except that it can also be done for one lecture at a time. This helps to structure the lecture like a conversation that gradually deepens using the "imaginary conversation technique" I describe in my book on argumentative research writing (Wentzel, 2018, p.59).

The first statement is not only designed to capture the essence of the message I want students to remember, but also to provoke questions. A good example is the sentence in the first paragraph of this chapter: "If we want students to understand as *much* as possible, we should teach as *little* as possible."

I then anticipate the questions that the first statement will provoke, and provide short answers to each one. After that, I anticipate the

questions that these answers will, in turn, incite, and so on… I then end up with something resembling a tree, with the initial statement being the trunk and the critical ideas growing directly out of the trunk: the further away from the trunk, the more detailed, and less critical, the ideas.

The process works like this:

1. Make the initial broad statement to an 'imaginary friend' who may be someone who represents who you want your students to become. This statement cannot, and should not, cover every detail. It simply has to be interesting enough to stimulate questions.
2. Imagine this friend responds with a question. The friend is only allowed to ask a question, and the question should be an open-ended question that encourages you to explain more. Such questions are usually 'how' or 'why' questions. The imaginary friend is not allowed to ask closed-ended questions, that is, those that can be answered with a single word or phrase. A closed-ended question with a single answer stops the conversation and the elaboration process.
3. You then respond to this question, and your response becomes the next statement that starts the process again. Return to step 1 until you have enough smaller statements to give you clear guidance on what to teach.

Taking one of my recent lectures as an example, this is how such a conversation might go:

Me: The progress in artificial intelligence makes us rethink creativity, so the role of the innovation manager must change.

Friend: Why does progress in artificial intelligence make it necessary to think differently about creativity? And how should the role of innovation managers change as a result?

Me: Artificial intelligence is developing so fast that it has revealed that 'shallow' forms of human creativity are merely algorithmic, except that humans are slower and more biased. As a result, innovation managers have to help companies identify: where 'shallow' human creativity can be replaced by artificial intelligence; how they can become more open to 'deep' human creativity; and how the two kinds of creativity can work together.

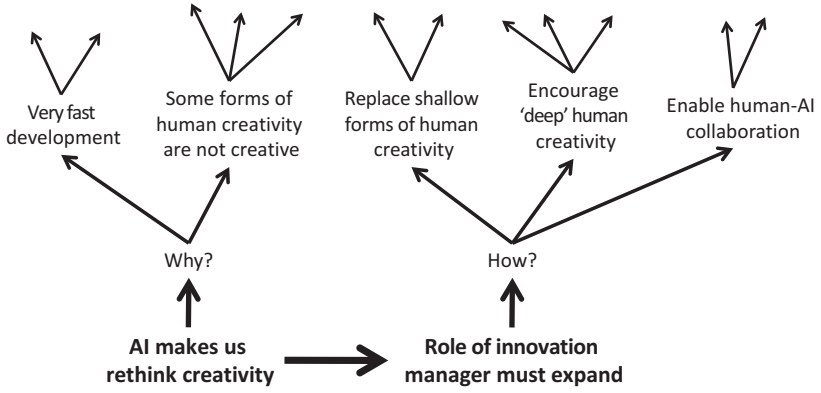

Figure 1.5 Elaborating on Critical Ideas through Conversation

This leads to a tree structure, as shown in Figure 1.5.

From the above, note that the imaginary friend does not always have to respond with a single question. Also, note that my answers are again in the form of short statements which suggest or provoke more questions from my imaginary friend. So, I will then have an imaginary conversation on each one of the branches and gradually grow the lecture outline. While constructing this tree structure, you should not consult the textbook, because then the temptation will be to include ideas that do not logically grow from the critical ideas.

Keep on expanding the structure until you believe you have sufficient material for the lecture. You will most likely have too many ideas, but with this elaboration technique, it is easy to see which ideas are most important (closer to the trunk) and which ones may be dropped (further from the trunk). Given that this kind of elaboration develops like a conversation, it is quite easy to run the class with the aid of methods that involve more active learning (such as discussions).

Once you reach what you believe to be the outer edges of the tree structure, you can consult the textbook. But this should only be to find the relevant information to explain the ideas on the branches. If you add ideas that do not fit on the branches or add new branches, then you are defeating the purpose of this technique.

As mentioned, we can elaborate not only by growing a lecture from the critical ideas but also by deepening understanding of these ideas. The next section looks at one way to do that.

Elaboration by Deepening: The Job Story

As mentioned before, good professors deepen understanding by exploring the critical ideas from different perspectives. It is better to teach a few critical ideas from many perspectives, than as many ideas as possible from a single perspective.

As you think about who you want your students to be, or become, you will discover that there are many different possibilities. Each possibility offers a different point-of-view on the critical ideas, and these different perspectives can be combined in a lecture.

One useful way of doing this was inspired by the 'jobs-to-be-done' approach of management guru, Clayton Christensen. What Christensen *et al*. (2016) argue is that people do not buy products for the sake of products, but for jobs they want to do. As a result, you can make a product more useful by adapting it to the job the customer wants. For example, people do not simply want a drill (product), they really want to hang something on the wall (the job they want to do with the drill). Likewise, they do not want a sandwich, they want to satisfy their hunger, and the sandwich is a means to this end.

To find the job-to-be-done, Alan Klement (2016, p.170) created the 'job story' and it is useful especially when teaching topics that have real-world applications. Klement suggests that a good job story follows a format such as: When ___ (a particular problem arises), I want ___ (a certain kind of solution), so that ___ (my life improves in a specific way). A product may have more than one job story depending on who you have in mind.

Applying this to teaching, we think of the audience members as the customers – not as who they are now, but rather as who we want them to be when using their understanding. If we think of a critical idea as a product, we can conclude that they too need the critical idea, not for its own sake, but in order to get a job done.

Suppose I am teaching basic combinatorics (a branch of mathematics dealing with enumerating possibilities). There are many different people who could use combinatorics, and each one's job story offers a different perspective on the topic. Combinatorics would be useful to people working on molecular biology, network security, insurance, gambling, and even some everyday situations. Every job story based on a critical idea offers a different perspective. It is quite possible that the same group has different job stories in different contexts. For example, in network security, the critical ideas would suggest one job story when working with file encryption and another when deciding on your own password.

If I am teaching the topic of basic combinatorics, my critical idea may be that, in our daily lives, we will spend our money more wisely if we can identify the number of possibilities in situations where there is some risk involved. Then I could write job stories for different problem situations, such as games of chance, selecting a password or choosing insurance. One job story could then be: When I am deciding whether to buy a lottery ticket (problem), I want to know my chances of winning (solution) so that I can know how to spend my money wisely (improvement). This is one of many scenarios I could have chosen.

Every scenario presents a different aspect or highlights different details of the critical idea. Combining them in a lecture deepens the students' understanding and increases the chances that their understanding will transfer to situations outside the class.

Elaboration by Active Learning: Guided Discussions

Elaboration does not always have to be the sole responsibility of the professor. In fact, less professor-talk is what you would expect from professors using critical ideas. By using more active learning approaches, like guided discussions, students can be part of the elaboration process. However, to ensure that discussions actually generate learning, advance preparation is necessary.

The first step is to create thought-provoking questions that make students think about the critical ideas. Grant Wiggins (2007), who did more than anyone to promote the notion of critical (or big) ideas, called them 'essential questions'. Good essential questions trigger debate because they do not have one correct answer and, in the course of trying to answer them, further questions are raised (see Box 1.6). In a similar vein, Bain (2004, p.37) found that many of the best college teachers taught by means of "big questions". These questions are ones that connect the critical ideas of the subject as a whole to the critical ideas in various topics and students' own experiences.

It is because essential questions raise even more questions that can be used to guide discussions, especially when used together with an instructional model like Socratic seminars (Estes & Mintz, 2016, Box 1.6). Such a seminar starts with you, before class, identifying, a broad, thought-provoking opening (essential) question and then, deriving from this, two or three basic questions that call for taking a position. For every basic question, develop around six to eight follow-up questions to guide discussion around each basic question. Not all of these questions will be asked, so think of them more as options. They are there when needed to reinvigorate the discussion or shift the direction. Before walking into class, one might have a map as seen in Figure 1.6.

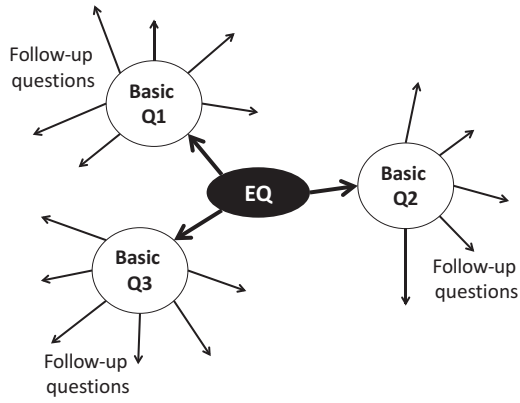

Figure 1.6 Discussion Map Derived from the Essential Question

For example, if the critical idea is: "Competition is good for business as a whole, but not always for individual firms," then the essential question might simply be: "Is competition always desirable?" The basic questions could then be designed to explore this question from different perspectives, such as: large and established companies, young entrepreneurs or start-ups and/or consumers. The follow-up questions are also designed to make students realize that the essential question and basic questions do not have simple or obvious answers. Figure 1.7 shows an example of possible follow-up questions organized around one of the basic questions.

One may introduce the topic with something that exemplifies the complexities of the essential question, such as an artifact, news article or video. Then the discussion is opened with the essential question. As the discussion progresses it can be shifted using the basic and follow-up questions. To ensure that learning happens, it is important to require students to respond to each other and provide reasons for their answers. At the end of the class, review the discussion and arrive at a conclusion in the form of the critical idea.

BOX 1.6 GUIDING DISCUSSIONS DERIVED FROM CRITICAL IDEAS

To find out more about essential questions, search for: **"essential questions" Wiggins**. For more information about the Socratic seminar enter: **"Socratic seminar" Estes Mintz**.

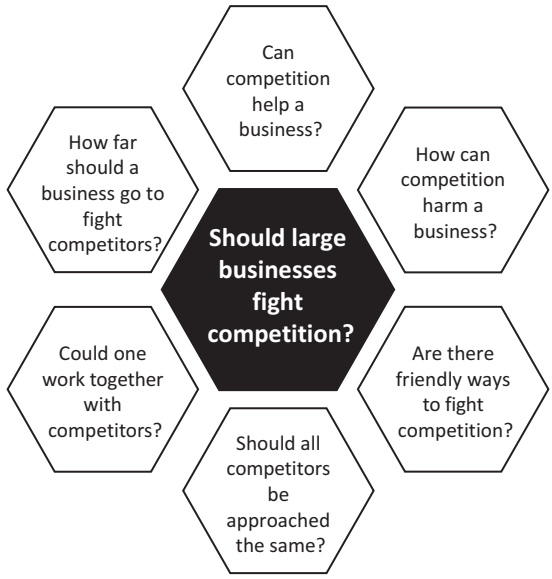

Figure 1.7 Follow-Up Questions

THE PHILOSOPHY OF SACRIFICE

The biggest problem in the elaboration stage is that non-critical ideas tend to sneak in. To counter this, it is best to adopt a philosophy of sacrifice when preparing lectures. It employs the Pareto principle, but in reverse.

This philosophy comes from my own discipline of economics. In the first class you attend, they teach you that resources are scarce, so you cannot do or get everything you want. From the beginning, you are taught to think in terms of sacrifices. For example, with limited money, you have to choose between buying an ice cream or a chocolate bar. Every choice means sacrificing one thing in order to get something else. As long as resources are scarce, the best you can do is to make the choice that involves the smallest sacrifice.

In a teaching situation, the scarcest resource is time. We simply cannot teach everything. The more time we spend helping students to understand one idea, the less time there is to support the understanding of other ideas. I used to think that I could cheat this choice simply by talking as fast as possible. Yes, indeed, if you have a nimble tongue and can cram enough content into your slides, perhaps you can *say* everything, but that is not the same as students *understanding* everything. Quite the opposite. All of us know that listening to a fast talker, with

slides that look like the small print of a contract, is guaranteed to leave no memory traces (at best) or, more likely, completely confuse us.

With limited time, we will always lose something. But we can choose what we lose. Do we talk fast and cover everything, but have students lose understanding of almost all the ideas? Or do we teach the ideas in the order they appear in the textbook, and hope that the least important ideas were the ones that we rushed through at the end? Or do we choose intelligently, in advance, which ideas to sacrifice so that the critical ones will stand out and be understood?

Unless they teach the simplest content, the best professors know they cannot cheat choice because complex topics require sacrifice. They sacrifice the teaching of some ideas so that they can effectively promote the understanding of the ideas that matter.

Without a philosophy of sacrifice, you will subconsciously employ an erroneous philosophy of perceived value. You will ask yourself the worst question possible when judging whether an idea is critical, which is: "Is idea *x* a valuable idea to teach?" Most ideas are probably valuable to an expert, so asking this question will cause you to include one idea after another, until you are back at trying to talk as fast as possible. While the techniques discussed earlier are designed to take your mind off of this question of perceived value, an explicit sacrifice philosophy is necessary to guard your limited time from the inclusion of lurking non-critical ideas.

It is much better to ask yourself, firstly: What other ideas will I *not* be able to teach properly if I teach idea *x*? Then, secondly, ask: Are these ideas that I sacrifice more valuable or less valuable to the audience than idea *x*? If they are less valuable, sacrifice them. If one of them is more valuable, sacrifice idea *x*.

So, when deciding whether to teach an idea, do not ask yourself what the students will gain from it. Rather, ask yourself what the students will *lose* if you do *not* teach it. Those ideas that cause them to lose the most if they don't understand them are the critical ideas.

With a sacrifice mentality, it is possible to categorize ideas, not on their value, but on the effect of *not* teaching them, as summarized in Figure 1.8:

- Subject-critical (SC) idea: An idea that, if it is not taught, will make it difficult for students to understand most topics in the subject as a whole.
- Topic/theme-critical (TC) idea: An idea that, if it is not taught, will make it difficult to understand most of the ideas in a particular topic.
- Nice-to-know (NTK) idea: An idea that, if it is not taught, will cause students' understanding to be less complex or precise compared to an expert.

- No-difference (ND) idea: An idea that, if it is not taught, will make little difference to students' understanding.
- Harmful (H) idea: An idea that, if not taught, will actually make the topic easier to understand for now.

IMPLICATIONS FOR PRACTICE

By now, it should be clear that critical ideas and the corresponding philosophy of sacrifice should have a dramatic effect on lecture preparation. Let's wrap up by considering how it should also inform other aspects of educational practice.

The Content of a Course

SC-ideas are usually taught in previous courses and form the foundation of more advanced courses. It is your responsibility to ensure that students possess an understanding of these ideas, because without this, students will struggle to understand new ideas in your course. One cannot always trust that they really learned these ideas in previous courses.

TC-ideas would be the threshold concepts or big ideas in a particular theme or topic in your course. These ideas are connected to many other ideas, so you should spend most of your time teaching these ideas.

NTK-ideas help the student to approach the detailed knowledge of an expert, but, since overall understanding does not depend on NTK-ideas, you should only teach them if there is time. ND-ideas should not be taught or can be assigned for self-study. H-ideas should not even be assigned for self-study, except perhaps for the most advanced students.

As shown in Figure 1.9, at the start of a course, a professor should spend most of his/her time teaching, revising and assessing SC-ideas from previous courses. As the current course progresses, more time should then be spent on TC-ideas. Towards the end of the course, there may be more time to teach NTK-ideas.

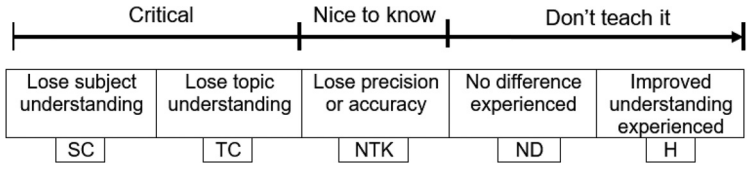

Figure 1.8 Degrees of Sacrifice

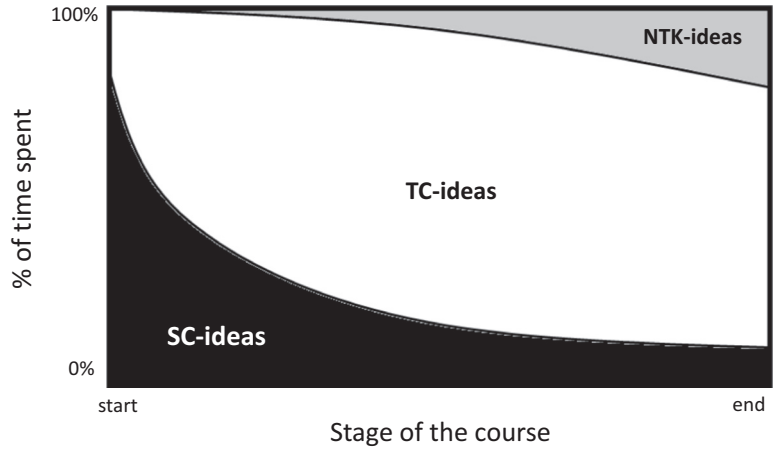

Figure 1.9 Time Spent on Different Ideas during a Course

Presenting the Content

If your lectures are prepared around critical ideas, the actual lectures should reflect that. Chapter 10 will explain, in detail, how to use critical ideas to design lecture slides. For now, I will just consider how it might influence the structure of a conventional lecture.

At the start of the lecture review, check students' understanding of the SC-ideas that will be necessary to understand the TC-ideas in the topic, even if those SC-ideas were taught in previous courses. Then introduce the TC-ideas of the current topic in an interesting manner (as explained in Chapter 5).

The body of the lecture should grow only from the critical ideas. If you are using a textbook, you will not cover everything, nor will it be covered in the same sequence. This is because your aim is not merely to *cover* the content, but rather to *reveal* it. Only if students gain a deep understanding of the most powerful ideas, will they be able to extend their understanding to other ideas and contexts. When ideas are skipped in the textbook because they are not critical, but may appear in an assessment, you should alert the students that these sections are for self-study.

At the conclusion of the lecture, return to the critical ideas. Summarize them and assign tasks to students to think about these ideas. Avoid anything that simply requires the reproduction of content from other sources. It would be ideal if these tasks could be completed in class. This is practical because your lectures will be shorter. Furthermore, it is desirable

because working on these tasks promotes understanding and reveals gaps in students' understanding, which can then be addressed in class.

An additional benefit of teaching with critical ideas is that it makes your teaching more flexible. In the military, there is a saying: "No battle plan ever survives a real battle." Similarly, no lecture plan ever survives a real lecture, especially when it is designed to encourage participation and understanding, but knowing the critical ideas tells you what you need to focus on when unexpected things happen. If for whatever reason, you have much less time than expected, you would teach the critical ideas first. Similarly, if you had more time, you would expand from there to the ideas that are important to know.

Assessment

Assessment of students' understanding will be discussed in detail in Chapters 7 and 8, so some brief comments will suffice for now. The understanding of students should be assessed regularly through formative assessments (assessments to get a sense of where they still need to develop but which do not count for marks), and at the end, through a summative assessment (an assessment that counts for marks and determines whether a student passes or fails a course).

Formative assessments should only focus on SC-ideas and TC-ideas. Even SC-ideas that were taught in previous courses should be part of every assessment in a course, especially if those ideas form the foundation of the current course's TC-ideas, and students should be warned about this.

The majority of the questions in a test or exam should test the ability of students to *apply* TC-ideas. Memorization is sometimes important and will be addressed in Chapter 4. If memorization is assessed, it should be focused on the building blocks of TC-ideas that cannot be logically derived from other ideas.

Assessment is a powerful tool for learning but there is limited time to assess students' understanding. The philosophy of sacrifice should apply here too. Time should not be wasted on assessing students' ability to memorize ND-ideas or H-ideas. The focus should be on the application of SC-ideas and TC-ideas. For advanced students, additional tasks can be designed that are based on NTK-ideas.

Designing Programs

A complaint, which never seems to get resolved, is that degree programs do not equip students for the real world. Universities respond to this dilemma by either adding new content to their programs or replacing

content. Over time, even though the content increases and/or gets updated, the complaint remains. The problem lies exactly in this response: *adding* new content can make programs less useful. To increase usefulness, content should be *reduced*.

It is important to distinguish between the content that students can only learn from experience, and the core content that will enable them to take advantage of such experience. Programs should focus on content that will enable graduates to make sense of as many new ideas and skills as possible. For example, employers don't need graduates who *know* everything – they want graduates who can *learn* everything.

Content that students will learn, or re-learn, in the real world should not be taught. But this is often the very content that gets added to programs only to be re-learned, forgotten or become outdated. Instead, it should be cut because it takes valuable time away from the teaching of the content that is critical regardless of changes or context. It would be more efficient to devote that time to the critical ideas and skills, and teach these from multiple perspectives. When such students join the real world, they will then learn faster because they will be able to easily expand their understanding.

When designing a program that will truly equip students for the real world, first define who students will become as a result of the degree. Spending time speaking with role models about this definition is more important than discussing content selection with them. From this definition: (1) identify ideas and skills that are common across all the roles that a student may eventually play; then (2) remove those that can only be learned from experience or that cannot be taught meaningfully in an academic context; and (3) remove those that are changing in an unpredictable way. What you are left with then is probably the critical content that should form the core of the program. The program should then be grown from this core only, resisting the temptation to add what sounds impressive or is simply about keeping up with other universities.

CONCLUSION

There is a Zen proverb: "It is the silence between the notes that makes the music." Music without pauses and silences is just incomprehensible noise with no rhythm. Music that uses all possible notes is cacophony. Similar to this proverb, one could say that: "It is what we don't teach that makes the learning happen." Good professors don't teach everything and that is what makes them good.

When we try to teach as much as possible, we are really trying to *end* the learning process, but when we teach as little as possible we are

trying to *start* the learning process. I know from my own teaching in the past that those who try to teach as much as possible have good intentions. They suffer from an irrational 'FOOMO' (fear of others missing out). Over time, I realized that my students actually miss out more when I teach too much, because they leave my lecture with information overload, which in turn, leads to loss of interest and forgetting.

One of the biggest distractions from critical ideas is the textbook. For this reason, I try to teach without a textbook if I can, relying instead on a variety of sources and/or students' own research. In fact, without exception, I have done my best teaching in courses without a prescribed textbook. If I do prescribe anything it is keyword searches, as I do in the various boxes of this book.

Whether you have to use a textbook or not, following the steps will be the first step in transforming your teaching. Let's summarize them:

1. Define your audience. Who do you want your students to be, or become, as a result of understanding what you teach?
2. Identify the critical ideas in the topic or theme you want to teach.
3. Grow your lecture outline from only these critical ideas.
4. Ensure that the philosophy of sacrifice also informs other aspects of your teaching such as time management and assessment.

BOX 1.7 MORE ON THE IMPORTANCE OF SACRIFICE

The philosophy of sacrifice is an ancient one, captured in the principle of *via negativa*. I particularly enjoyed Nicolas Taleb's explanation of it in his book, *Antifragile*. A web search for "***via negativa***" will also yield many interesting results. For a general discussion of how it applies to all aspects of life, I recommend the book, *Essentialism: The Disciplined Pursuit of Less*, by Greg McKeown. He gives many illustrative examples, including that, in recent times, the quality of a movie's editing (basically cutting) is the best predictor of whether it will win the Oscar for Best Picture. Even in innovation methodologies, such as TRIZ, trimming is one of the best ways to improve invention.

POSTSCRIPT: SO, WHAT WAS THE 'LEAD'?

Read the story of Charles Simms in the introduction again to refresh your memory. You really only need to follow the first two steps of the

usefulness technique to find the 'lead' or the single critical idea. If the facts will appear in the school newspaper, the audience is the school students. Why do you want them to know these facts? What do you want them to do or become as a result of reading it?

Because you want them to know that: "There is no school on Thursday", there is your lead, nothing more. The rest are just nice-to-know details.

When Simms revealed this, Ephron (2008, p.x) experienced it as an "electrifying moment". She realized that all the facts are "utterly meaningless if you haven't figured out the significance of the facts. What is the point? What does it mean? He planted those questions in my head." I hope this chapter planted those questions in your head too.

EXERCISES

1 Do the Feynman exercise:
   a) Suppose all knowledge about your subject will be wiped out soon, but you can pass just one sentence on to future generations – no longer than one 'tweet'. If you applied thinking and imagination to this sentence, it should be possible to derive many of the other ideas of the subject. What would be your sentence?
   b) If this sentence was your starting point, would you teach the subject differently?
2 Suppose you are offered an opportunity to teach one of the courses in your subject on educational television, but on one condition: you have to teach it in five lectures of fifteen minutes each.
   a) Who will you regard as your audience?
   b) Which topics would you teach?
   c) What will be the critical idea in each topic?
3 Select one topic in a course that tends to be overwhelming to students, or for which there is usually not enough time to cover all the content:
   a) What do you want your audience to be, or to become, as a result of understanding this topic?
   b) Use one or more of the techniques and identify no more than three critical ideas.
   c) Place the ideas in a logical sequence.
   d) Elaborate on each of the critical ideas.
4 Select another topic in a course, perhaps one that tends to be boring to students:
   a) What do you want your audience to be, or to become, as a result of understanding this topic?

b) Use one or more of the techniques and identify one or two critical ideas, and convert them into essential questions.
c) Create a map of questions that can be used to guide a class discussion.

REFERENCES

Ambrose, S.A, Bridges, M.W. & DiPietro, M. 2010. *How Learning Works: Seven Research-Based Principles for Smart Teaching*. San Francisco: Jossey Bass.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Blum, S.D. 2016. What It Means to Be a Student Today. In: *Handbook of Academic Integrity*, T. Bretag (ed.), Singapore: Springer, 384–406.

Bransford, J.D., Brown, A.L. & Cocking, R.R. (eds). 2000. *How People Learn: Brain, Mind, Experience, and School*. Washington DC: National Academy Press.

Christensen, C.M., Dillon, K., Hall, T., Duncan, D.S. 2016. *Competing Against Luck: The Story of Innovation and Customer Choice*. New York: Harper Business.

Ephron, N. 2008. *I Feel Bad About My Neck*. New York: Vintage.

Estes, T.H. & Mintz, S.L. 2016. *Instruction: A Models Approach*, 7th edition. Pearson.

Feynman. R.P. 2011. *Six Easy Pieces: Essentials of Physics Explained by Its Most Brilliant Teacher*. New York: Basic Books.

Gromov, M. 2011. Structures, Learning and Ergosystems [available at http://www.ihes.fr/~gromov/PDF/ergobrain.pdf, accessed 14 June, 2018].

Johnson, S. 2010. The Genius of the Tinkerer. *Wall Street Journal*, 25 June.

Klement, A. 2016. *When Coffee & Kale Compete*. NYC Publishing.

Meyer, J. & Land, R. 2003. Threshold Concepts and Troublesome Knowledge: Linkages to Ways of Thinking and Practising within the Disciplines. Enhancing Teaching-Learning Environments in Undergraduate Courses Project, Occasional Paper 4.

Ritter, F.E., Nerb, J., Lehtinen, E., & O'Shea, T.M. (eds). 2007. *In Order to Learn: How the Sequence of Topics Influences Learning*. New York: Oxford University Press.

Wentzel, A. 2018. *A Guide to Argumentative Research Writing and Thinking: Overcoming Challenges*. London: Routledge.

Wiggins, G. 1989. The Futility of Trying to Teach Everything of Importance. *Educational Leadership*, 47(3):44–48, 57–59.

Wiggins, G. 2007. What Is an Essential Question? [available at http://www.authenticeducation.org/ae_bigideas/article.lasso?artid=53, accessed 15 June 2018].

Wiggins, G. 2010. What is a Big Idea? [available at http://www.authenticeducation.org/ae_bigideas/article.lasso?artid=99, accessed 8 June 2018].

Chapter Two

What Makes an Explanation Great?

Explanation is one of the few critical skills every professor has to master, because, as Bain (2004, p.117) found in his study of the best college teachers: "Generally the most accomplished of the teachers had the best ways of explaining things." Unfortunately, there is very little information available to those who want to improve their ability to explain complex ideas (see Box 2.1). There is much literature on scientific explanations, which is about how scientists themselves make sense of the world and come to understand it. But this is different from creating pedagogic or teaching explanations, which is communication that aims to help *other* people understand the ideas that you *already* understand.

BOX 2.1 SCIENTIFIC AND TEACHING EXPLANATION

For a broad overview of scientific explanation, see: http://www.iep.utm.edu/explanat/or search for **"theory of explanation" or "scientific explanation"**. By comparison, the literature on how to create great pedagogic explanations is minuscule. Among them, the best practical works on teaching explanation are *Lecturing and Explaining* by George Brown (1978), and, even better, the *Art of Explanation* by Lee LeFever (2012).

The next two chapters will fill this gap. If we can diagnose where teaching explanations go wrong and figure out why some explanations are more effective than others, anybody can learn how to become a great explainer. This chapter will first look at the abstract principles that should guide you as you develop explanations, and this will then prepare you for the practical techniques outlined in Chapter 3 that can be used to create great explanations.

WHAT EINSTEIN NEVER SAID ABOUT EXPLANATION

Einstein never said: "If you can't explain it to a six-year-old, you don't understand it yourself." Even if he said it, you should not believe it. The fact that people believe this saying shows how little we understand about the art of explanation. By critically viewing this falsely attributed quote, we can gain much insight into explanation.

The reason this saying is so appealing is that it does contain a grain of truth. Richard Feynman was the one who captured the whole truth when he said, "I really can't do a good job… of explaining [it] in terms of something you're more familiar with, because I don't understand it in terms of anything else that you're more familiar with" (Gleick, 2011, p.393). For Feynman, when explaining, the challenge is to reformulate complex ideas, so that your *audience* can understand them *in terms of things that they already understand* at the time of the explanation.

In the case of most complex ideas, the audience is often university students or intelligent laypeople, not six-year-olds. For example, to explain special relativity theory, an audience needs to have a prior understanding of the concepts of variation, light, speed, time, and distance. The average six-year-old's grasp of these concepts is too far removed from the prior understanding that is needed to understand and remember a single explanation without difficulty. In other words, their understanding does not lie in the 'adjacent possible' (mentioned in Chapter 1). You may be able to summarize the effects of special relativity in a way most toddlers will find interesting. But their current understanding does not reach far enough and will first need to be extended by preparatory explanations if they are to understand why those effects will occur. However, if you cannot explain special relativity theory in terms of the things that a high-schooler is familiar with, then you probably do not understand it that well.

So, the saying can be more accurately rephrased as: "You only understand an idea if you can explain it in terms of what your audience already understands." This restated saying suggests that poor explainers are those who try to explain ideas in terms of things they themselves know, or in terms of things they believe their ideal audience ought to know. They narcissistically think about their audience in terms of themselves, so they either: overestimate how far the audience's current understanding can reach; or they connect to the wrong things in their audience's understanding. Contrast this to the great explainers, some of whom are mentioned in Box 2.2.

BOX 2.2 GREAT EXPLAINERS IN ACTION

The famous Feynman lectures can be found here: http://www.feynmanlectures.caltech.edu/I_01.html or simply search **"Feynman lectures"**. His explanations in these lectures were not always at the level that a layperson could understand. He improved with time, so to see an older Feynman in action, see https://www.youtube.com/watch?v=eqtuNXWT0mo or search for **"Richard Feynman" AND "fun to imagine"** on YouTube. The field of science is replete with other great explainers, often better than Feynman. Some of the well-known ones are Carl Sagan, Richard Dawkins, Bill Nye, Michio Kaku, and Neil deGrasse Tyson, to name a few.

If you watch some of the great explainers in action, you will notice three things. Firstly, they understand their audience – which may be six-year-olds, but in the case of complex topics are most likely older. More specifically, they know what their audience already understands *before* the explanation starts. Secondly, they understand the topic they want to explain so well that they can distinguish between the critical and peripheral ideas. Thirdly, they understand how to bring the two together, specifically how the critical ideas are connected to the ideas of the audience.

Together these three things capture quite well what Shulman (1986, p.9) called pedagogical content knowledge (PCK), simply defined as "the ways of representing and formulating the subject that make it comprehensible to others." Without PCK, we tend to rely too much on the textbook. Great explainers think carefully about how to translate the ideas in terms of the things that their audience already understands. This involves organizing the content around critical ideas, finding connections that make sense to the audience and putting the ideas in a sequence that strengthens those connections.

NO EXPLANATION WITHOUT UNDERSTANDING

Explanations aim to bring about understanding, so we can learn a lot by simply relating the two, as you will see in this section and the next.

In Chapter 1, I pointed out that human knowledge can be represented as a network, specifically a small-scale network where only a few nodes are critical. Understanding is not the ability to perfectly reproduce this network. Rather, it is the ability to: recognize what is critical and what is redundant; and then use this distinction to compress knowledge into a few connected ideas. If this understanding is correct, it will enable us

to figure out the remaining ideas and connections when "just a little imagination and thinking are applied" (Feynman, 2011, p.4).

For scientists, the perfect manifestation of understanding is an accurate 'theory of everything': a few critical ideas combined with a reliable method of inference – that is, a method for deriving new ideas from old. If found, a theory of everything would be extremely efficient: it would not contain any redundancies; and, given sufficient time, all the connected important and nice-to-know details could be inferred from it. If you had such a theory, you would not need to know everything, because you would already understand everything, *even things that you don't know* at that moment.

Something similar can be seen in language. A person who understands a language, such as Spanish, is not someone who knows every conceivable sentence and every conceivable response to every conceivable sentence. Such a person would know a lot and *appear* to understand Spanish. It would have to be some super-being with an almost infinite memory that was able to memorize every single possibility [This is similar to John Searle's (1980) famous Chinese room thought experiment, for which you can find a good summary on Wikipedia].

We would be deceived for a while into believing that this being understands Spanish. But one of two things would expose this: (a) change and (b) teaching. Even the smallest change (like a new word) that is not included in what was memorized will leave it not knowing how to respond. It might know every single connection at a given time, but would be unable to make *new connections* when its expectations do not match reality. If we asked this being to teach Spanish, its lack of understanding would be revealed even faster. Since this being does not understand, it would not explain anything to you; instead, it would only repeat one sentence after another and expect you to memorize them all. Nobody would want this super-being to teach them!

A person who understands a topic would notice patterns and redundant information and teach much more efficiently. For example, in any language, there are patterns (regularities, correlations, symmetries etc.) that connect thousands of words. By definition, these patterns repeat themselves, and something that repeats itself a thousand times need not be taught a thousand times. It can be compressed into one pattern, taught once and then applied in as many different combinations as possible as many times as possible.

Hence, a person who understands Spanish is a person who knows these patterns and can use them to generate every conceivable sentence that may come up in normal conversation, even ones that never happened before. She will be able to infer all those sentences using social cues, a relatively small number of words and connecting them all appropriately by using a handful of grammar rules. She would be able to speak and respond to sentences, *even if she never heard them before*, because she understands the

language. This is the person I would want to teach me. Since this person understands, she is more likely to help me to understand.

So, someone who teaches complex topics with understanding will not teach everything, but rather a few critical ideas, as well as the methods for inferring other ideas from them (see Chapter 6). If you understand a subject, you do not need to know everything about it. Understanding will enable you to construct a network of relatively few critical ideas, and gradually infer new ideas and connections as required.

The point of all this is that explanation is not merely simplification. A good explanation communicates understanding, and understanding uses patterns to compress knowledge, so it is clearly more than mere simplification. In fact, recent research in artificial intelligence and information theory (Tishby & Zaslavsky, 2015; Hoel, 2017) suggest that appropriate compression of complex ideas can increase the informational value of such ideas. The ideal understanding is a reduction of knowledge to a connected essence from which most other ideas can be inferred through reasoning or imagination. Explanation does not try to simply remove inconvenient complexity, it enfolds it into a form of understanding from which the complexity can gradually unfold.

If explanation reflects understanding, then, as understanding improves, so should the ability to explain. The next section explores this and shows that different levels of understanding correspond to different levels of explanation.

LEVELING-UP YOUR EXPLANATIONS

What is the highest level of understanding? If you believe what George Bernard Shaw said, that: "Those who can, do. Those who can't, teach", you would think that the highest level is being able to 'do' things with your knowledge. As Figure 2.1 shows, that is, in fact, the lowest level of understanding. Lawson (2006) demonstrated that being able to do something does not mean you understand it. Her study revealed that even in the case of simple objects that we use every day, most of us understand very little or misunderstand how these objects work. Doing gives us a 'feeling-of-knowing' experience (Hart, 1965), yet when someone interrogates our understanding, many gaps in our understanding are exposed.

Being able to explain ideas to experts requires some understanding. Anyone who has ever presented a paper at an academic conference knows that talking about one's work in front of experts can be intimidating, especially when they start picking on small details. But explaining one's ideas to such experts, especially ideas that they are familiar with, is not intellectually demanding. There are many things we can take

for granted, so we don't have to think carefully about distinguishing between critical ideas and nice-to-know details, and we can assume that an audience of experts can make most of the connections without our explanations.

Many experts have a huge volume of tacit knowledge that enables them to excel at what they do, but they are not sufficiently aware of their own understanding to convey it to others. For this reason, many need popularizers to make their tacit knowledge explicit so that others may understand.

Explaining to students or laypeople, who don't possess the prior understanding of experts, is significantly more demanding. Obviously, an explanation given to those with little or no expertise will be less accurate because it will contain fewer details, but that is exactly why it requires a higher level of understanding in order to produce these. It takes a high level of understanding to know what is critical, what can be simplified and which details can be left out without adversely affecting the learning process in the future. On top of that, one needs the insight to see how the critical ideas can connect to those ideas that the audience already knows.

It is this second-highest level of understanding (in Figure 2.1) that professors should strive to achieve. It is an ideal, and even the great explainers, like Feynman, miss the mark sometimes. Goodstein (1989, p.75) relates a story of a time when a friend asked Feynman to explain something in physics, and he responded: "I'll prepare a freshman lecture on it." Then, a few days later, Feynman returned and admitted: "You know, I couldn't do it. I couldn't reduce it to freshman level. That means we really don't understand it." I'm no Feynman, so I have failed many more times. Fortunately, as professors, we receive many chances to explain the same thing in our career, so we have the time and opportunity to improve.

The highest level of understanding, and an impossible ideal in the case of most complex topics, is the ability to explain to a six-year-old.

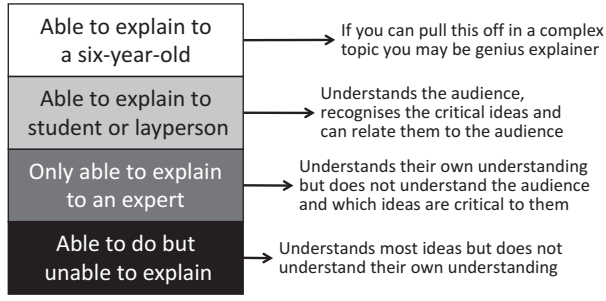

Figure 2.1 Levels of Explanation and Understanding

Anyone who can explain topics like discounted cash flow, continuous-time Markov processes, the citric acid cycle, simultaneous quadratic equations or special relativity to a toddler in a way that really makes them understand, would not just be a great explainer, but a genius explainer. Having said that, imagining how something could be explained to a toddler is sometimes a good way to find the critical ideas.

Now we can see why George Bernard Shaw was mistaken. Instead, Lee Shulman (1986, p.14) restated it much more accurately as: "With Aristotle, we declare that the ultimate test of understanding rests on the ability to *transform* one's knowledge into teaching. Those who can, do. *Those who understand, teach*" (emphasis added). It is simply impossible to be a great explainer without possessing great understanding.

The result of great explanations should be to get students to move up the levels of understanding as well. And students need not rely only on a professor's explanations to achieve this (see Box 2.3).

BOX 2.3 SELF-EXPLANATION TO ENHANCE UNDERSTANDING

Explaining to oneself is quite a powerful method for becoming aware of gaps in one's understanding, and it has been found to be effective in moving students up the levels of understanding. In fact, in my early years of teaching, I used it quite often myself when preparing my lectures. The seminal paper by Chi *et al.* (1989), and the excellent book, *Small Teaching* by James Lang (2016), raised the profile of this technique. You can find more information by searching for **"self-explanation" AND learning**.

HOW UNDERSTANDING CHANGES

You probably know from experience that there are different degrees of understanding, ranging from high (expert) to low (ignorant) and various states in between. However, it is not a smooth process and it does not start from zero (like the dashed line in Figure 2.2). It may appear a trivial point to make, but it is a dangerous misconception.

The view behind the dashed line is mistaken, because, firstly, as suggested in the first chapter, understanding takes place within a connected network of ideas, and even adding just one new connection can cause understanding to change suddenly. Secondly, no person ever starts with zero understanding. Knowledge and understanding emerge when ideas

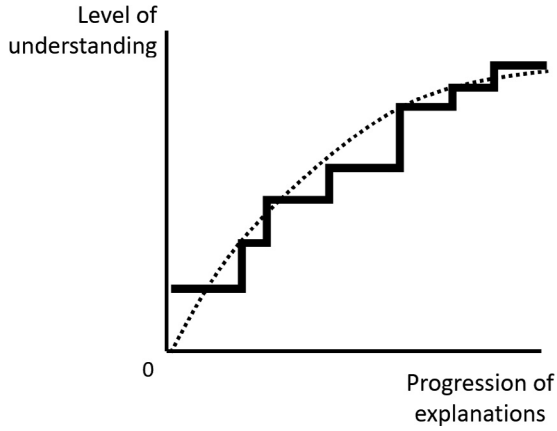

Figure 2.2 Understanding as a Series of Jumps Starting from a Base

are connected to other ideas, so we can only understand something new if we have some foundation of familiar ideas to which we can connect the new ideas. If there were no previous ideas to which new ideas could be connected, understanding would be impossible and explanation useless.

The solid line in Figure 2.2 is more accurate. It starts from a low, but non-zero, level and shows that understanding improves as a series of abrupt changes. Similar to the ratchet effect, once understanding is achieved, it is difficult to reverse. When we search for patterns or hidden objects, it is impossible to 'un-see' them once you find them – the same is true of understanding. This is something that professors can exploit because it means we can build on a solid base of prior understanding, but it can also make our life difficult if a student's understanding is, in fact, a misconception (see Box 2.4). This is almost unavoidable, because, as will be explained later, lower levels of understanding almost inevitably contain misconceptions.

BOX 2.4 DEALING WITH MISCONCEPTIONS

There is much interesting research on how to change students' misconceptions. You can search for it by simply entering **"conceptual change"** or for a more focused search **"conceptual change" AND misconceptions.** My approach is first to identify the most likely misconceptions by thinking about past lectures or searching the web for:

misconceptions about [topic]. When teaching, I make students aware of their misconceptions before teaching, and, in fact, I exploit these often to make lectures more interesting (see Chapter 5). Usually, misconceptions exist for a reason and there is often some truth in them. I then draw on such reasons and half-truths and show how they can be more useful when modified by correct ideas. Also, see diSessa (2014) for a good review of the topic, and references to studies that give a more detailed description of how to deal with misconceptions.

Figure 2.2 also suggests that there will be times when explanations will seem to have little impact. As a person is exposed to a progression of effective explanations, there will be periods during which the level of understanding appears to remain flat, punctuated by a series of sudden improvements. This is obviously because it takes time for students to internalize the network of ideas. This period can be lengthened unnecessarily by the teaching of non-critical ideas.

BLOCKS TO UNDERSTANDING

All students experience so-called 'blocks to understanding' at some point. They appear on the horizontal sections of the solid line in Figure 2.3, where they prevent the next abrupt transition to a higher level of understanding. If the student cannot leap across the block to the next level of understanding, understanding remains stagnant. These blocks may be caused by prior misconceptions (see Box 2.4) or by poor teaching.

The following are some of the common phrases that can be heard from students when teaching blocks understanding:

- "The ideas are too new to me." In this case, a professor does not use *familiar* ideas students already understand to which to connect the new ideas.
- "This is too much information." This suggests that the new ideas are not organized around a few *critical ideas*.
- "Everything is so complicated." The professor is not gradually adding and integrating *complexity in layers*.
- "I understand parts, but it still does not make sense." This usually means that the new ideas are not connected to each other in a *sequence* that strengthens the connections.

- "I get the explanation, but I forget it quickly thereafter." Students are not expected to *think* about the ideas, so their connections are short-lived.

While a professor may be at fault when each of the blocks above appears, it does not mean that the student has no role to play. In each case, it is not purely the professor's actions that cause the problem, but the kind of thinking (or lack thereof) that is triggered by his actions. Without thinking, we do not form connections in our brains. If an explanation does not make us think or makes us think in the wrong way, the result will be a lack of understanding.

Regardless of whether it is students' misconceptions or teaching that blocks understanding, such blocks have to be removed from both sides – from the student who comes from a low level of understanding, and from a professor who comes from a high level of understanding. In Figure 2.3, a professor cuts the size of the block down through explanation (the axe), and so reduces the amount of time and energy that students have to spend in order to achieve understanding. Yet, even though the block may have been removed, the student still has to leap over it through the act of thinking about the new ideas. Again, it is professors who set up the opportunities for thinking through meaningful tasks and authentic assessments.

An effective explanation is not simply one that makes the ideas appear easy. Rather, it aims to give students a sense of self-efficacy and so create the impression that thinking about the ideas is likely to yield positive results for them. Unless the explanation makes a student think about the ideas, it will have no long-term effect on the student's understanding.

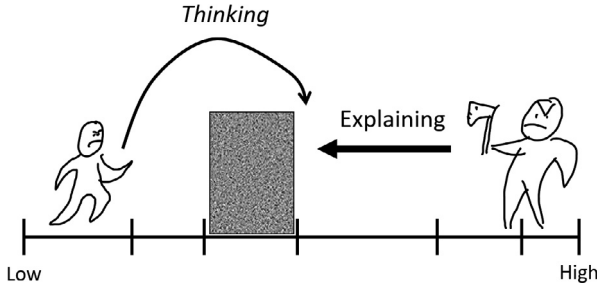

Figure 2.3 Understanding Is a Joint Effort

Table 2.1 Explanation Solutions to the Blocks to Understanding

Experience of the Block	Explanation Solution
"The ideas are too new to me"	Help students connect the new ideas to what is familiar to them
"This is too much information"	Organize the information around a few critical ideas
"Everything is so complicated"	Integrate complexity gradually and in layers
"I understand parts, but it still does not make sense"	Find an appropriate sequence in which the ideas can be connected
"I get the explanation, but I forget it quickly thereafter"	Make students think about the explanation

Fortunately, each one of the ways to experience a block to understanding can be transformed into positive principles to guide explainers, as summarized in Table 2.1.

If you are familiar with Vygotsky's zone of proximal development (ZPD), you may have realized that explanations are most effective in this zone. The ZPD is the area that connects what a student understands with what a professor understands. In that zone, we find the ideas and skills that a student does not understand yet, but can understand with the guidance of an expert.

Every solution in Table 2.1 represents the different elements found in good explanations that ensure that the potential of the ZPD is not wasted. Let's look at each of the five elements in turn.

FIRST ELEMENT: FAMILIARITY THAT BREEDS UNDERSTANDING

Without connection, there are just lists of facts. We, therefore, need to be aware of where some connection with the audience is possible. There are three places to search for potential connections between new ideas and the audience.

Firstly, if new ideas can be connected to what the audience already knows, usually through analogies or examples, then understanding can occur faster. This helps the audience to see that they already have more understanding than they realize and that the new ideas just expand that understanding.

Secondly, one can relate new ideas to what is happening in the world. This is very powerful because both the teacher and student share that knowledge, and so it is more likely that the teacher will use a language

familiar to students when connecting the new ideas to familiar real-world events.

Thirdly, to make it more likely that these connections will be longer-lasting, explanations need to evoke emotion. One needs to know what an audience cares about in order to find ways to show how new ideas help them understand the thing better and, by extension, how they can then act if necessary. This is something we will elaborate on in Chapter 5 as part of making ideas interesting.

SECOND ELEMENT: ORGANIZATION AROUND CRITICAL IDEAS

It is useless if you know your audience, but connect what they understand to the wrong ideas. Their future understanding will expand faster and more efficiently if their current understanding is connected to the few ideas that matter, that is, the critical ideas.

In textbooks, we find lots of facts, often neatly categorized in lists and sections, but rarely connected to each other and to critical ideas. But facts, without such connections, have zero meaning; they are lifeless. One can know all the world's facts, yet have no understanding – like the earlier example of the hypothetical super-being who memorized every single sentence and every response to every single sentence.

Humans cannot achieve such a feat, though I did try when I started with Accounting in eighth grade at school. For two years, I memorized the treatment and entry of every possible transaction. This enabled me to pass, but with a huge effort, and the effort required increased as we covered more and there were more different transactions.

Fortunately, for me and the rest of humankind, there is a more efficient and intelligent way of learning. As we look for connections between facts, we start seeing patterns, that is, commonalities and regularities. This enables us to eliminate what is redundant and identify the few critical ideas from which most other ideas can be derived. Critical ideas then become the seeds from which the other ideas can gradually grow, as was shown in Figure 1.3. In the case of Accounting, I discovered that there was a fistful of critical ideas that made almost all understanding possible: those captured by the accounting equation. Once I understood these ideas, I was able to continue with the subject up to my third year at university, without ever expending as much effort as I did in my final exam in ninth grade.

The same process occurs in teaching that is efficient, except that we do not derive the ideas (see Box 2.5), we explain them. We introduce new facts or ideas and then help the students see where each one fits

and what it means. Explanations reconstruct the process by which they were derived, but in the least messy way, without the dead ends and mistakes that would have happened in the real discovery of these ideas. Explanation, provided it focuses on critical ideas, can therefore further enhance the efficiency of the learning process.

BOX 2.5 LEARNING WITH OR WITHOUT EXPLANATIONS?

Instead of explaining, we can guide students to derive the ideas themselves, as in discovery learning, which is effective but can be rather slow. Search for **"discovery learning" OR "discovery-based learning"**. Related teaching methods you can look for are **"inquiry-based learning"** and also **"problem-based learning"**. These methods need not exclude explanations, in fact, explanations can make them more efficient, as you will find if you search for **"enhanced discovery learning"**.

Explanation that starts from the critical ideas and moves outward, is the fastest way to put a student on the path towards expertise. Studies agree that one of the biggest differences between novices and experts is in the way in which they organize their knowledge (see Bransford, Brown & Cocking, 2000, Chapter 2; Ambrose, Bridges & DiPietro., 2010, Chapter 2). Novices' knowledge often has little or no structure (as in Figure 2.4a), while experts' understanding is like a densely connected network (even more complex than Figure 2.4c). Experts acquire this structure as they discover, or are shown, patterns in disconnected facts. Learning is the process of absorbing this structure, and explanation speeds up this process.

Using explanation to move students from Figure 2.4a to 2.4c, without overwhelming them, is quite an art. You cannot just let an expert loose. For experts, any single fact or idea is connected to many others, so if we let them start anywhere and then simply follow the connections in any direction, students will be thoroughly confused at the end of a lecture. What we need is a structure that mediates between the expert and novice structure. This structure is what Hay, Kinchin and Lygo-Baker (2008) call a 'hub-and-spoke structure', as seen in Figure 2.4b.

The hub-and-spoke structure is most effective if it has the *critical idea/s as the hub* at the center. It helps the novice to see which of the facts and/or ideas in Figure 2.4a are most important. Furthermore, it is not densely

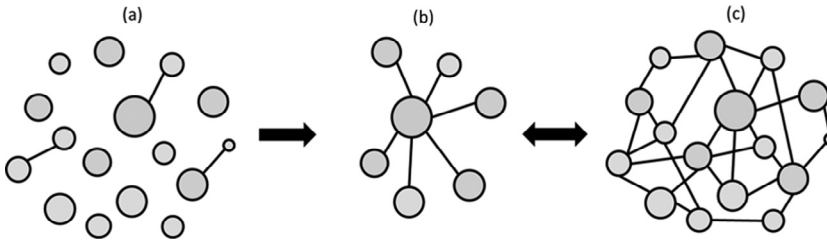

Figure 2.4 Something Mediates the Process of Becoming an Expert

connected, so it is easy for a novice to absorb. Once this structure is settled, it is then easy to complexify it and move toward Figure 2.4c.

This structure works well for a professor-expert too. Although the hub-and-spoke structure appears to be a simplification, it is derived from the expert structure. For that reason, it is flexible enough to accommodate new ideas in the future without major disruption, and therefore has the potential to become as complex as the knowledge structure of experts. With this structure, the conflict between simplification and accuracy is not as severe, because the simplifications are flexible enough to allow for more complex and accurate knowledge to grow from them (as seen in Figure 2.5).

So, a good explanation starts with critical ideas that are related to those with which the student is familiar. It then gradually adds details, while still remaining connected to the hub-and-spoke structure. This process of complexification is the next element.

THIRD ELEMENT: GRADUAL ADDITION OF COMPLEXITY THROUGH CONNECTIONS

For knowledge arranged in a hub-and-spoke structure to gradually grow into an expert structure, it needs to be made more complex gradually.

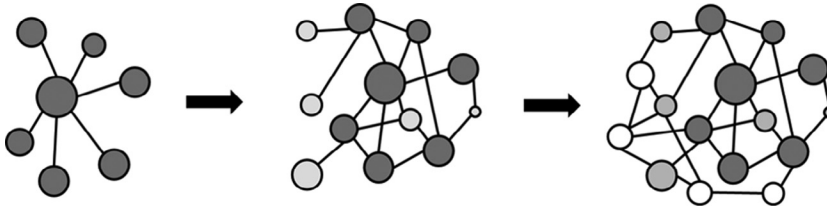

Figure 2.5 Growing an Expert Structure

This means connecting more facts or ideas, layer by layer. There are several ways in which this can be done:

- Add new connections between ideas (represented by lines) already in the hub-and-spoke structure, as seen in Figure 2.6a. The main ways in which to add connections are by showing: (a) cause and effect; (b) reason and action; (c) similarity and dissimilarity; and (d) parts and whole.
- Add new ideas (represented by the circles) by (a) inserting new ideas in between old ideas, as seen in Figure 2.6b. This happens, for example, when there is another factor that mediates a cause-effect, part-whole or reason-action relationship or some intermediate category; and (b) growing new ideas out of existing ideas, as seen in Figure 2.6c. This happens, for example, when we extend ideas by showing that there are more kinds, parts or effects of something or more reasons for something. Notice that ideas cannot be added without adding connections as well.

All three methods (Figure 2.6a, 2.6b, and 2.6c) of complexifying the hub-and-spoke structure work together to create an expert structure over time. But this complexification can only happen once the hub-and-spoke structure is firmly in place in the mind of a student. And it cannot happen all at once.

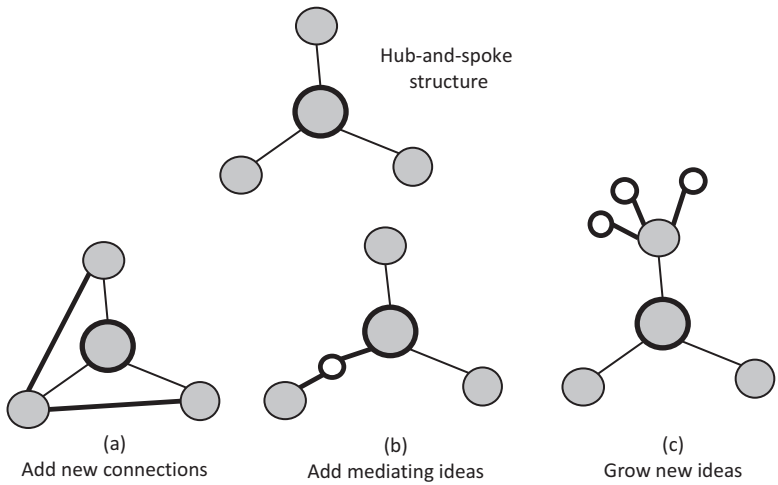

Figure 2.6 Ways of Adding Complexity to a Hub-and-Spoke Structure

It is best to first ensure that students understand the hub-and-spoke structure of critical ideas and then gradually add layers of complexity.

Good explanations, therefore, start with the critical ideas. Future explanations then return to these critical ideas and explain how these ideas are more complex than initially explained. And then further explanations do the same with these ideas, and so on, in an iterative process. Every layer adds details to the critical ideas of the previous layer. Figure 2.7 illustrates this.

Figure 2.7 does not show all the ways of adding complexity, but it illustrates the principle of first laying the foundation of critical ideas and their relationships, and then connecting a new layer of complexity on top of these ideas, and so on. At every layer, we find new critical ideas; but, if we collapse the layers, we should always return to the same critical ideas at the foundation. Ideas are created from connections between previous ideas. Any topic, therefore, is made up of ideas within ideas within ideas… like Russian dolls.

None of this complexification can happen without connection. Whether you add new ideas at same layer or add a new layer, it always has to be clear how it follows from the ideas before. The value of making connections should not be underestimated, because knowledge,

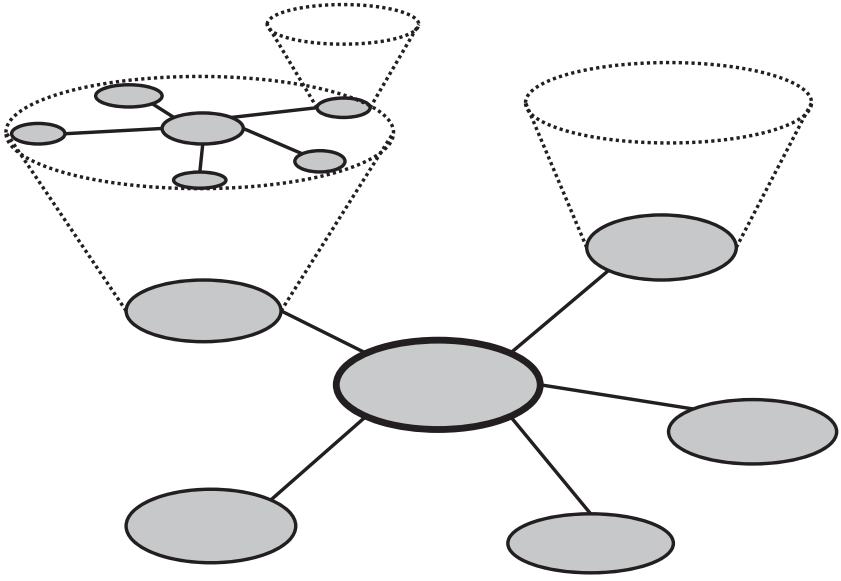

Figure 2.7 Good Explanations Gradually Add Complexity in Layers

and understanding of that knowledge, does not reside in the individual facts or ideas as such. Rather, knowledge and understanding reside in the connections between known facts and ideas.

Let's take an example from my field of Economics to illustrate. Some teachers may teach you the demand curve by showing you the equation $P_x = a - bQ_x$. If you have studied Economics, you already have all the connections to make sense of it, but if you have no prior knowledge, this idea will appear completely new to you.

Yet this apparently new idea can be explained by using only ideas that you already know, simply by making new connections between them. A good teacher will not start with the equation, but rather by reminding you of things you already know: (1) you already know what it means when a product has a **price**; (2) you already know what a **quantity** or number of products is; and (3) you already know what it means to **buy** something.

A demand curve is simply an idea that results from *connecting* these ideas of price, quantity, and buying. We **buy** (demand) a greater **quantity** if the **price** is lower, and less when the price is higher. When you see the connection between them, it may seem so obvious in hindsight that you may wonder why you need someone to teach it to you.

Figure 2.8 is a different way of showing this: it shows that someone will buy three products when the price is \$7 and eight products when the price falls to \$4. The demand curve simply connects all those points. Everything is derived from connecting already familiar ideas.

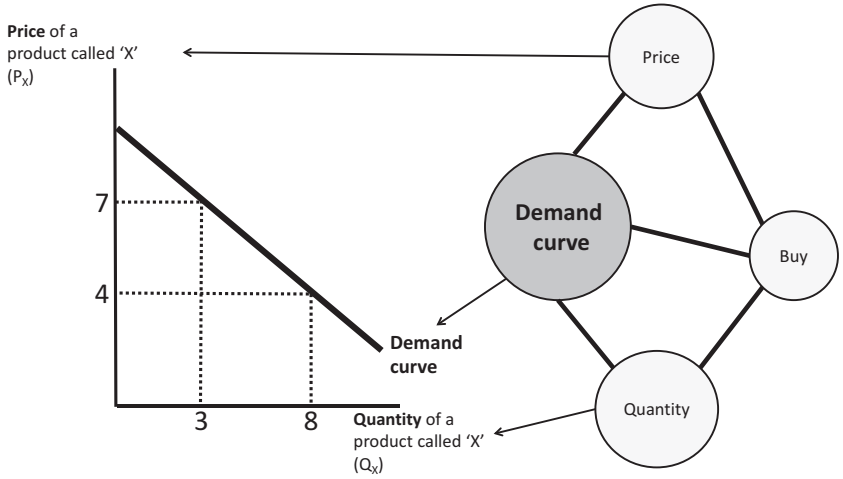

Figure 2.8 The Demand Curve as Connections between Ideas You Already Know

But why do we even need to understand this? If I were a business owner, I would need to understand this because it shows me that, if I raise my price too high, I will lose customers. It shows that when the price of a product goes up, people will buy less of that product.

The point of this example is simply to show the importance of making connections. Without connections, there can be no new knowledge and understanding, and it is only when we make connections between familiar ideas that we gain an understanding of new ideas.

Color is a simple metaphor that captures the relation between organization through critical ideas and gradual complexification. Critical ideas are the primary colors. We connect them by mixing them. Mixing yellow and blue results in a new color, green. Now we have three ideas (or colors) that can be mixed together to create an infinite number of colors with every new color at a higher level of complexity.

FOURTH ELEMENT: A LINEAR SEQUENCE THAT STRENGTHENS CONNECTIONS

Understanding is structured as a network, and even when it is simplified to a hub-and-spoke structure, it seems to require teaching many ideas at the same time (Figures 2.9a and 2.9b). The problem is that time is linear, so we can only teach one thing at a time (Figure 2.9c). Unfortunately, the right sequence is not always obvious. As mentioned in the first chapter, sequence matters for the quality of learning, so we cannot simply leave this to chance.

Even if you correctly derived the hub-and-spoke structure from the expert structure, but you teach the ideas in the wrong sequence, your explanation will still be confusing. Hyperlinks in websites, PowerPoint presentations and innovative presentation software (like Prezi) make it

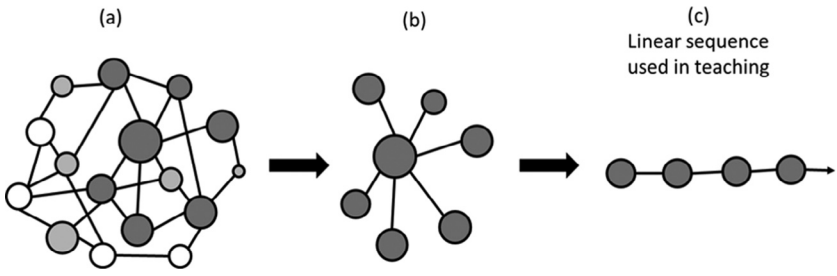

Figure 2.9 From Understanding to a Linearly Sequenced Explanation

easier to teach ideas as a network structure; but that still does not change the physical constraints and time limitations that force us to explain only one thing at a time.

The common solution to the sequencing problem that we find in many textbooks is to convert the network into a hierarchical structure. The problem is that, in hierarchical structures, more and more of the connections are lost as one moves to the lower, more detailed levels. Anyone who has read textbooks with sections, subsections, sub-subsections etc. can attest to this. So, while a hierarchical structure might work in texts where it is easy to page back as necessary, it is not ideal in verbal communication.

George Brown (1978), from whose book I learned much as a novice teacher, offers common sense advice for sequencing. It is based on the principle that explanations should start with something with which the student can relate. This principle is manifested in guidelines like the following: (1) start from the known and move to the unknown and unfamiliar; (2) start with the simple and progress to the complex; and (3) go from general and broad ideas to specific ideas. All of this is useful, if vague, and there are exceptions. The only rule is to start with something to which the audience can connect, and sometimes this may go in the opposite direction than these guidelines.

Such guidelines can be made more specific if derived from an understanding of the network structure of knowledge as explained thus far. To convert a network structure of knowledge into a linear teaching sequence, first convert it into a hub-and-spoke structure so that the critical ideas are clear. Then, when teaching, start at the hub and move outwards to the peripheral ideas, periodically returning to the hub before venturing further outwards. Figure 2.10 helps to make this more concrete.

In the middle is the topic of explanation from which three critical ideas follow (A, B and C). Each of them is connected to other important, but less critical, ideas (a1 up to c2). A good explainer would usually start with one of the critical ideas (maybe A) and then follow a logical path from A to a1 to a2, then returning to A before following the path from a3 to a4.

After this, the explainer would either return to A or connect to the next critical idea (perhaps B). This pattern is repeated further along the structure. For example, if we suppose that even more ideas were connected to a2, the explainer would follow those ideas and periodically return to a2.

This gives rise to the repetition of ideas, but the ideas that are repeated are always at the center or hub of the ideas being explained. This ensures that the central ideas are understood and integrated by the

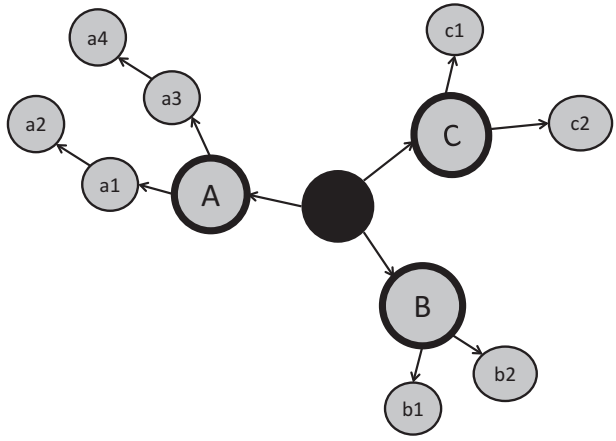

Figure 2.10 A Hub-and-Spoke Structure with No Clear Linear Sequence

students before the next layer of complexity is added. An explainer continuously integrates the details into the bigger and more important ideas in a sequence that resembles a spiral, as shown in the next two figures.

Figure 2.11 shows the most likely sequence of ideas in the explanation of critical idea A in Figure 2.10. Notice how the explanation starts from A, follows the flow of ideas from a1 to a2, then spirals back to A, before continuing with a3 to a4. At this point, the explainer might spiral back to A again to ensure that the details are integrated with the central

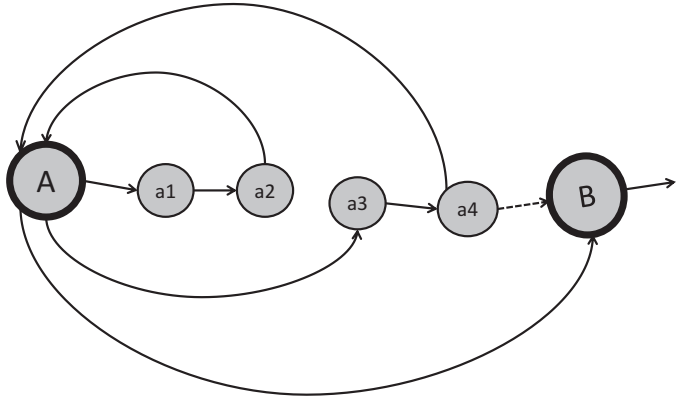

Figure 2.11 Sequence of Ideas Following the A-Branch of Figure 2.10

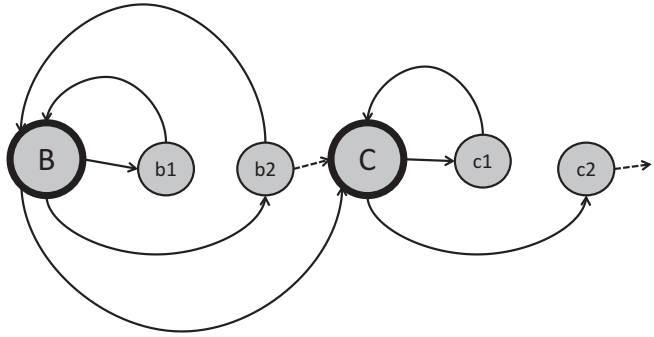

Figure 2.12 Sequence of Ideas Following the B- and C-Branches of Figure 2.10

idea (A). However, if there is a strong connection between a4 and the next critical idea (B), the explanation may proceed straight to B.

Figure 2.12 shows the spiral structure as it may turn out when explaining critical ideas B and C in Figure 2.10.

In Figure 2.12, the explainer spirals back more frequently. Starting from B, the explainer moves to b1, but if there is no clear direct connection to b2, he would first return to B and then again move outwards to b2. If there is a strong connection between b2 and C, the explainer might move to C, but if not, he would again spiral back to B to integrate the details and then connect to C. The same pattern is repeated for C.

An appropriate linear sequence, therefore, strengthens the connections and reinforces the structure through productive repetition. The same repetition would occur at higher levels of complexity. Every time an idea gets repeated, it is done with new knowledge in mind and so it offers a slightly different perspective each time.

FIFTH ELEMENT: PROVOCATION OF THOUGHT

It is one thing to make these connections in a lecture, but very different for these connections to also appear in the neural pathways of students. The only way to create neural pathways is through thought. Bain (2004) also emphasizes that the best college teachers know that students only really learn when they reason about what they have learned. As anyone who has ever passively watched an interesting documentary will know, while watching, it feels like you are learning (that 'feeling-of-knowing' mentioned earlier), but one or two days later you will find it difficult to remember much.

As noted in Figure 2.3, understanding is a joint effort involving teacher explanation and student thinking. The connections are made visible and available to the student by means of explanation. Yet, for the connection to be integrated into the existing understanding of the student, she or he has to think about the ideas. This can be done by making the ideas interesting (as explained in Chapter 5) through inciting actions that require thought (discussed in Chapter 6), or giving assessments that are authentic and exploit the testing effect (as explained in detail in Chapter 7).

PUTTING IT ALL TOGETHER

Before an explanation, the student is confronted with a loose collection of facts and ideas that do not make any sense. To make sense of these facts, the teacher organizes them around a few critical ideas, and then gradually makes those ideas more complex by adding new connections and introducing new ideas. Whenever new ideas are introduced, they have to be connected to the already familiar ideas in something that resembles a hub-and-spoke structure, and this then gives rise to even more new connections. The actual explanation when delivered to the students has to proceed in a sequence that starts from ideas to which the learner can relate, and then connects from there to the critical ideas, and then moves outwards along the hub-and-spoke structure. To ensure that the connections are made more permanent, the student has to think about the ideas and their connections.

This chapter was fairly abstract, but it lays the basic conceptual framework from which we can derive practical techniques for explaining. These techniques are explained in the next chapter.

EXERCISES

1. Analyze a chapter, or even just one explanation, in a textbook according to the five elements of explanation. Evaluate it accordingly:
   a) Does it connect new ideas to things with which students would be familiar?
   b) Is everything clearly organized around a few critical ideas?
   c) Is the knowledge structure gradually complexified?
   d) Is the sequence derived from a network structure? In other words, is there useful repetition to strengthen the connections and highlight how ideas are related or do connections get lost?

   e) Does it get students to think? Does it get students to think about the right things, that is, the critical ideas and connections between ideas?
2. Analyze a lecture or talk by a great explainer. For this, watch a video or read a transcript of a lecture that you found useful. If you don't know of any, find any of the physics lectures of Richard Feynman or an interesting TED talk.
   a) Does it connect new ideas to things with which students would be familiar?
   b) Is everything clearly organized around a few critical ideas?
   c) Is the knowledge structure gradually complexified?
   d) Is the sequence derived from a network structure? In other words, is there useful repetition to strengthen the connections and highlight how ideas are related or do connections get lost?
   e) Does it get students to think? Does it get students to think about the right things, that is, the critical ideas and connections between ideas?
3. Compare your results in exercises 1 and 2.
   a) What are the differences you notice in each element?
   b) What did you notice that was not discussed in this chapter?
4. Develop explanations for the topics or ideas you identified in Exercise 4 of Chapter 1. Analyze your explanations according to the five elements and improve any weaknesses you notice.

REFERENCES

Ambrose, S.A, Bridges, M.W. & DiPietro, M. 2010. *How Learning Works: Seven Research-Based Principles for Smart Teaching*. San Francisco: Jossey Bass.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Bransford, J.D., Brown, A.L. & Cocking, R.R. (eds). 2000. *How People Learn: Brain, Mind, Experience, and School*. Washington DC: National Academy Press.

Brown, G.A. 1978. *Lecturing and Explaining*. London: Methuen.

Chi, M.T.H., Bassok, M., Lewis. M.W., Reimann, P. & Glaser. R. 1989. Self-explanations: How students study and use examples in learning to solve problems. *Cognitive Science*, 13(2):145-182.

DiSessa, A.A. 2014. A history of conceptual change research: Threads and fault lines. In: Sawyer, R.K. (ed.), *The Cambridge Handbook of the Learning Sciences*, 2 nd edition, New York: Cambridge University Press.

Feynman, R.P. 2011. *Six Easy Pieces: Essentials of Physics Explained by Its Most Brilliant Teacher*. New York: Basic Books.

Gleick, J. 2011. *Genius: The Life and Science of Richard Feynman*. New York: Open Road.

Goodstein, D.L. 1989. Richard P. Feynman, Teacher. *Physics Today*, 42(2):70–75.

Hart, J. T. 1965. Memory and the feeling-of-knowing experience. *Journal of Educational Psychology*, 56(4):208–216.

Hay, D., Kinchin, I. and Lygo-Baker, S. 2008. Making learning visible: The role of concept mapping in higher education. *Studies in Higher Education*, 33(3):295–311.

Hoel, E.P. 2017. When the map is better than the territory. *Entropy*, 19(5):188.

Lang, J.M. 2016. *Small Teaching: Everyday Lessons from the Science of Learning*. San Francisco: Jossey-Bass.

Lawson, R. 2006. The science of cycology: Failures to understand how everyday objects work. *Memory & Cognition*, 34(8):1667–1675.

LeFever, L. 2012. *The Art of Explanation: Making your Ideas, Products, and Services Easier to Understand*. Hoboken: Wiley.

Shulman, L.S. 1986. Those who understand: Knowledge growth in teaching. *Educational Researcher*, 15(2):4–14.

Searle, J. 1980. Minds, brains and programs. *Behavioral and Brain Sciences*, 3(3):417–457.

Tishby, N. and Zaslavsky, N. 2015, April. Deep learning and the information bottleneck principle. In *Information Theory Workshop (ITW)*, 2015 IEEE (pp. 1–5). IEEE.

Chapter Three

A Step-by-Step Approach to Explain Complex Ideas Clearly

To help make sense of the abstract ideas of the previous chapter, imagine a caveman standing in front of you. He looks like an illustration brought to life, because, while hunting, he walked through a wormhole that transported him 80,000 years into the future.

He still has a club in his left hand. One thing changed in this journey through time: he picked up perfect English. Besides that, his knowledge has remained unchanged.

You look at each other in surprise. "I've got to put this on Twitter," you say to yourself. "What is Twitter?" the caveman asks you.

TRY THE TWITTER CHALLENGE

Fortunately, you happen to have Wikipedia open on your phone, so you give him the abbreviated version: "Twitter is an online news and…" You notice that you lost him already, but you continue nevertheless, "…social networking service on which users post and interact with messages known as tweets". His face is now contorted in utter confusion, yet you feel that you need to add something more: "Tweets are restricted to 280 characters. You have to be registered to post tweets, retweets and follow other users." Despite his excellent command of English, the caveman understood little.

You realize that you have forgotten everything from Chapter 2, and you tell him: "I am sorry, let me try to explain again." The caveman is relieved that he is finally able to understand another sentence, so he sits down and looks attentively at you.

Self-doubt fills you. You can assume the caveman understands English, but nothing more – he knows nothing about wires, wireless, electricity, text, characters, signals, social networks, mobile phones, computers, screens, buttons or anything you take for granted. He has never been exposed to anything you are trying to explain, so he knows less than a recently potty-trained toddler about technology. Can you do it?

Do the Twitter challenge now, and if you are not familiar with Twitter, consult the Wikipedia entry. Create an explanation of Twitter in the form of a short voice recording of around three minutes or a written description of fewer than two pages before proceeding.

Make notes of how you approached it. You will compare this explanation to the explanation that I created to get a sense of where your explanation did not work. Or your explanation may be better than mine, as some of my students have done over the years. If that happens, you will try to understand why your explanation was better. This is important because we learn much from recognizing what we do well.

The reason you may struggle with this challenge is that the second chapter was quite abstract. In this chapter, you will see how those abstract principles are translated into a step-by-step process for constructing good explanations. At the end of this chapter, you can try again by following the steps and analyze the difference between your attempts.

STEP 1: THE CRITICAL IDEAS FOR THE AUDIENCE

After getting to know your audience, you need to find the critical idea/s. With the caveman, I would choose the usefulness technique: asking why a caveman would need to know about Twitter or how it would be useful to him. The result: Twitter is the easiest way to share your experiences with others in your group.

The critical idea should only refer to things that are familiar to the caveman. Also, note that because the critical idea leaves out many details, it is not completely accurate – it is coarse-grained, like a photo at a very low resolution. This 'coarse-grained' explanation is, however, necessary initially as the caveman does not yet have sufficient background to understand all the details. But a good critical idea will help him, later on, to see how the details fit together.

Critical ideas, like the one for the caveman, tend to be rather broad and general. They help us to eliminate irrelevant information; but, because they are somewhat vague, they can be made more useful by exploring them in more detail. And there is much to explore. There are huge volumes of information embedded in good critical ideas. If we could release this information, we could explain critical ideas properly and increase their usefulness. The first step in releasing this information is breaking the critical idea up into more basic units: ideas that are more familiar, simple and concrete. Once we have this, we have the building blocks from which to construct an effective explanation.

The techniques mentioned in Chapter 1, like the imaginary conversation technique, can be used to release the information. But, here, I suggest using a more systematic technique, called 'un-chunking', which works well with the hub-and-spoke knowledge structure. It is based on the concept of chunking discovered in studies about memory and expert thinking (Miller 1956; Chase & Simon, 1973), but in reverse. To explain 'un-chunking', we need to take a quick detour through the idea of 'chunking'.

DETOUR: CHUNKING AND UN-CHUNKING

Novices look at situations or phenomena and see lots of disconnected facts, so they initially find it difficult to make sense of everything. Over time, they start seeing that not all information is relevant, and also notice that many of relevant facts are similar or connected. These groups of disconnected facts are then grouped together into so-called 'chunks.'

For example, the novice biologist may look at nature and only see lions, rabbits, eagles, tigers, antelope, fish, wolves, mice, and snakes. But the expert notices patterns. She sees that certain animals eat other animals, and so reduces the list of animals to two chunks: 'predators' and 'prey'. She then notices that these two interact in predictable ways and then combines this relationship to create a chunk called 'predator-prey population cycle'. Whatever does not fit into one of the chunks is regarded as irrelevant information that the expert ignores (for now); while the novice, by contrast, gets entangled in the details.

Chunks allow experts to generate information about items, even those that are unknown to them. For example, if I have never heard about something called a 'lulo', but I know that it belongs in a chunk called 'fruits sold in supermarkets,' I can still tell you a lot about a lulo, even if I have never seen one. I know that it will be edible and contain seeds. It is most likely sweet or pleasantly sour and probably has a protective skin and is fleshy inside.

Chunks also make it easier to remember new facts that can fit into one or more of the chunks. For example, we often forget items on a long shopping list, but when we group items into chunks (fruits, vegetables, dairy products), the list is much easier to remember. This is because chunking integrates ideas and reduces the number of ideas we have to keep in mind, thereby reducing our cognitive load.

But chunks can also increase ideas in a fruitful manner. Once we have chunks, we may recognize that they form patterns or groups. These groups are categories that divide the chunk into smaller parts again. For example, biologists noticed that there are different kinds of predators,

and broke the 'predator' chunk into those that are 'carnivorous', 'herbivorous', 'parasitic' and 'mutual'. This additional complexity is worthwhile if it makes the chunks more useful. However, even after such categorization, there will still be *vastly fewer chunks than facts*.

BOX 3.1 CHUNKING

For more information on chunking, you can simply enter **chunking** in Google. Or for more detailed information, combine it with other search terms, for example, **chunking AND memory, chunking AND learning, chunking AND teaching, chunking AND chess**, and so on.

If you think that you have seen chunks in the previous chapters, you would be correct. In Chapter 1, you saw that ideas are chunks that emerge as we connect facts and ideas; and in Chapter 2, we saw that chunks can consist of smaller chunks, as happens when we gradually complexify a knowledge structure (see Figure 2.7). It is when we notice patterns, that we find these chunks.

While chunking is a very efficient way to become an expert, it causes problems when experts try to explain things to novices. As experts advance, they learn to give every chunk a name or a label, and then start to use the labels in explanations as if these chunks are now simple, single units with obvious meanings. They forget that every label is, in fact, a complex set of many different connected facts and ideas.

No wonder the caveman got confused at the very first words: "Twitter is an online news and social networking service". He was perplexed because the explainer did not realize that 'online' is the label of a chunk that involves many other complex chunks such as 'computer', 'communication network' and 'electronic information', all of which are unfamiliar to the caveman. The same happens with 'news' and 'social networking'. Like many experts, the explainer fell prey to the curse of knowledge or the "expert blind spot" (Nathan & Petrosino, 2003, p.905).

Many of us have the same experience as the caveman when a computer salesman tells us that: "This computer has a dual-core 2.4 GHz processor, 16GB RAM, 256GB hard drive, 802.11ac wireless with WPA2 support and a 23 inch widescreen LCD monitor." It even happens that experts themselves 'curse' other experts with their knowledge, as evident from the jargon-filled title of Alan Sokal's spoof article: "Transgressing the boundaries: Towards a transformative hermeneutics of quantum gravity" (Sokal, 1996) (see Box 3.2).

BOX 3.2 EXPERTS FOOLING EXPERTS

For more information on this fascinating case of how experts have fooled experts, just search for **"Sokal affair"**. It is easy to laugh at this, but it happens to all of us when we pretend to understand a mechanic or computer salesman. This gives you a sense of what students feel like when they struggle to admit that they don't understand.

Chunking compresses facts into far fewer chunks. As a result, many details, connections, and ideas gradually become invisible as they are absorbed into these chunks. Compared to novices, experts will, therefore, experience a subject as becoming simpler, even as they acquire more knowledge.

For an expert, more and more knowledge gets packed into the same chunks. So, an expert who wants to become a great explainer needs to first reverse the chunking process and unpack the chunks *until he gets to the chunks that are familiar to the audience* (going from left to right in Figure 3.1). This process is what I call 'un-chunking': finding the ideas from which the chunks are formed. We do this until we find the ideas that can connect to the audience's current understanding and build explanations with these. In other words, un-chunking leads us to the building blocks for explanations.

In Figure 3.1, suppose the really big circle is the chunk that refers to the ecosystem. As you can see, it is also made up of other chunks.

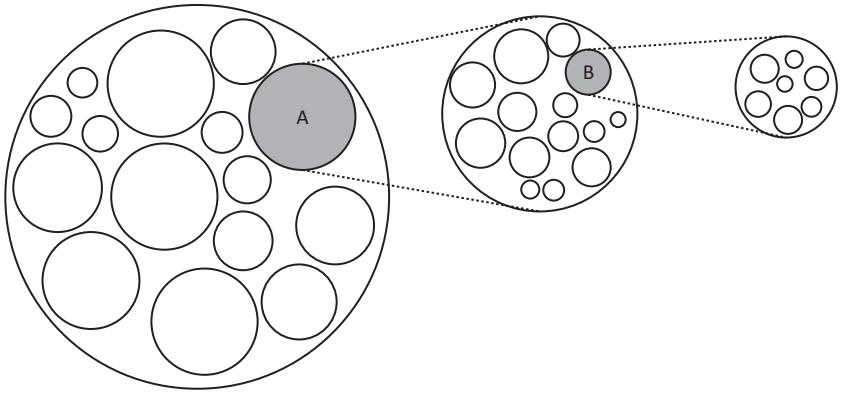

Figure 3.1 Un-Chunking – Unpacking Expertise Chunk-by-Chunk

One of those chunks (labeled A) is the predator-prey population cycle. Within that too, we find many chunks, including the chunk labeled B, which encompasses the concept of predators. If we unpack that further, eventually, we get to lions, sharks, eagles, and wolves, and, if we were teaching six-year-olds, we might start here, but with older students we may start at B. Un-chunking is this process of deriving more basic chunks, especially those that are hidden in our expert blind spots.

STEP 2: GENERATE THE BUILDING BLOCKS OF THE EXPLANATION

Un-chunking helps us to derive the simple and familiar ideas with which to build an explanation that the audience can understand. In an explanation, we will then connect these building blocks, and so make the hub-and-spoke structure more complex.

Since we will be connecting the building blocks, the different kinds of connections tell us how exactly to un-chunk. As mentioned in the previous chapter, when connecting the building blocks, we can make any one of these four kinds of connections: (a) cause and effect; (b) reason and action; (c) similarity and dissimilarity; and (d) parts and whole. Each one of the connections is also a way to unpack or un-chunk the idea in the center, as follows:

- Part-whole: Un-chunk up (is this idea part of something bigger?) and un-chunk down (what is it made up of?)
- Similarity-dissimilarity: Un-chunk across by asking what the idea can be compared to (an analogy) or contrasted to, and what would be good examples or non-examples.
- Cause-effect: Un-chunk up diagonally (what can cause it?) and down diagonally (what are the effects or what does it lead to?)
- Reason-action: Un-chunk up diagonally (what is the reason for it? Why do we need it or use it?) and down diagonally (how does it work? How do we use it?)

This is summarized in Figure 3.2. In the figure, the star in the middle is the idea that you want to explain and elaborate on, and then you ask the un-chunking questions and note the results.

Not all ideas can be un-chunked in all directions, and that is fine. The purpose is only to find possible building blocks that may be part of an

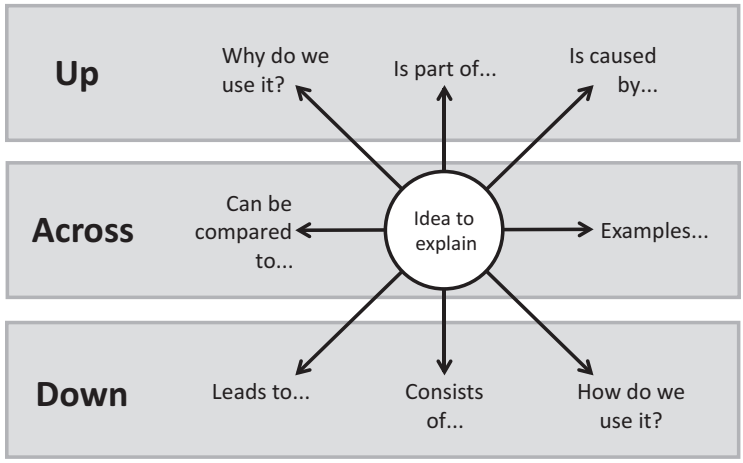

Figure 3.2 Un-Chunking to Find the Building Blocks of an Explanation

explanation, and not all of them will be useful. If, while using this method, other chunks come to mind, write them down too. At this point, you are really searching, in a structured way, for possible components that you can use in constructing an explanation.

Returning to the caveman, un-chunking will help us find possible ideas to use in a more comprehensible explanation of Twitter. As summarized in Figure 3.3, the result is:

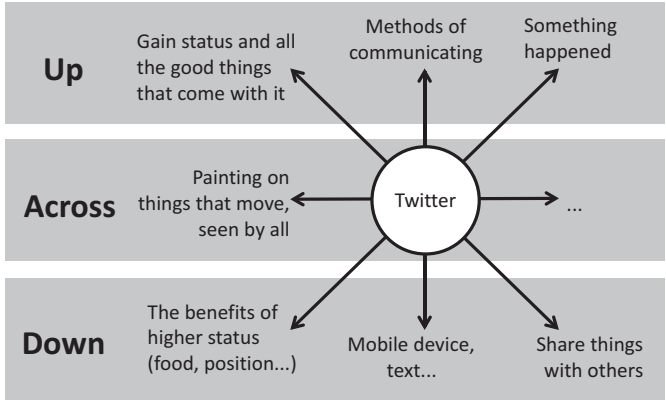

Figure 3.3 The First Round of Un-Chunking the Idea of Twitter

- Part-whole: Twitter is part of a larger category of methods of communication within a community, and methods by which members of a community can interact with each other. Twitter itself is made up of things like phones, cellular phone networks, characters, retweets etc., most of which will be unfamiliar to him.
- Similarity-dissimilarity: I can compare Twitter to other methods of communication that the caveman uses. Some of these are: shouting to his fellows; meeting around a fire; or painting things that happened to him on a cave wall. I can also contrast Twitter to these methods, especially if I want to explain how Twitter overcomes some of the problems with methods like cave paintings; maybe by pointing out that Twitter is like making cave paintings but on things that can move around. I can show him examples of Twitter, but it is not certain that he will understand what he sees.
- Reason-action: Why would a caveman want to use something like Twitter? Perhaps because something happened to him that he wants to share with others, and by sharing it he can gain status. How would he use it? I have to be careful here because the caveman does not know anything about mobile devices, but he understands that objects can move and can understand the use of symbols like pictures, spoken words or hand signals to explain to others what happened to him.
- Cause-effect: What might cause a caveman to want to share things with others? Maybe he thinks that some important event occurred that could give him more status if others knew about it. Obviously, the effect of this would be that he would gain more status, and, by extension, be afforded the benefits that come with more status and recognition from other tribe members.

Un-chunking for explanation is an iterative process. That means that un-chunking one idea will generate more basic ideas. If any of them are unfamiliar to the audience, they can be un-chunked to generate even more basic ideas, and so on.

The same with the Twitter challenge. Figure 3.3 shows only the first round of un-chunking. You can then take one of the resulting ideas that the caveman would find difficult (e.g. text characters or mobile device) and apply un-chunking to them. Doing this will introduce additional layers to your explanation (as in Figure 2.7). For the purpose of this example, we will just construct an explanation for the first layer of understanding.

By using this technique, all the content you generate for an explanation is derived from the critical ideas and will ensure that you remain focused on the critical ideas. Combining the three ways to complexify a hub-and-spoke

structure from Chapter 2 with un-chunking gives rise to more ways to generate the raw content of the explanation (see Appendix 3.1).

BOX 3.3 THE ORIGINS OF UN-CHUNKING

The technique was derived from a combination of sources. As explained, chunking comes from studies about how experts think, and also from a neuro-linguistic programming (NLP) method called 'chunking', as explained on the ChunkMaps website (http://www.chunkmaps.com). The kinds of un-chunking for explanations are inspired by the explanation of 'conceptual understanding' on the Instructional-Design Theories site (http://www.indiana.edu/~idtheory/methods/m6c.html). The procedure of un-chunking and re-chunking closely follows the creative process behind argumentative research writing (see Wentzel, 2018, Chapter 4).

PIT STOP: WE UN-CHUNK SO THAT WE CAN RE-CHUNK

Now you may wonder, if understanding is about compression, why are we increasing the number of ideas through un-chunking?

Firstly, remember that un-chunking is about finding the possible building blocks of an explanation, and you will not use all of them. You will not be doing un-chunking with your students, so they will not see the increase in the number of ideas. You will do it on your own so that you, as the expert, become aware of the ideas that are needed for an explanation, especially the basic ideas that are familiar to students, and which you may be taking for granted.

Secondly, when constructing an explanation from these building blocks, you will be making connections and so put them back together again, so the ideas will decrease in number as a result of the explanation. In the explanation, you will simulate the process of gaining understanding by showing how the basic ideas combine to create greater understanding, as set out in the next few steps. In a way, the remaining steps explain re-chunking.

STEP 3: WHERE TO START?

With the building blocks ready, we now need to choose: which ones to use; which ones to avoid for now; and which ones may be potential starting points. Review the ideas that emerged during un-chunking and categorize them as shown in Table 3.1. Place every idea that the

Table 3.1 Choosing Building Blocks

Connection	Avoid (for now)	Can be used
Part-whole	Mobile phones, social networks, characters, retweets, electronic signals, electricity, Internet etc.	Methods of communication in a community without technology, such as conversations, facial expressions, hand signals, paintings, songs, dance etc.
Similarity-dissimilarity	He might not understand examples of Twitter without some explanation first.	An adaptation of the above methods of communication that create something similar to Twitter e.g. moving cave wall paintings or people shouting while running (to help understand the notion of 'mobile').
Reason-action	Do not mention concepts like text, characters, letters, words, buttons, phone etc.	The information he would share with others e.g. danger, successful hunts, family events etc. that would benefit him (e.g. increase his status). The adaptations mentioned above could be useful e.g. mobile paintings.
Cause-effect	None came up during un-chunking.	The events mentioned above that could be beneficial to share e.g. births, hunts, predators approaching etc. The benefits of higher status and recognition – such as admiration, leadership or more food.

caveman would not be familiar with in the 'avoid' column, and perhaps take note of any that may be important soon, but will need some explanation first. Everything else goes into the 'can be used' column. If any of them are a bit abstract, then spend some time making them more concrete, as I did with the methods of communication.

The most fundamental rule of starting an explanation is: start with something relevant to the topic that is also familiar to the audience. Any one of the building blocks in the third column of Table 3.1 meets that rule.

Though this narrows down the possibilities, there will still be too many ways to start the explanation. To choose among them, consider which ones are most likely to engage the audience. Perhaps, while doing un-chunking, you already had some ideas.

Some of the most common ways to start are simply to: state the critical idea; raise a provocative question based on the critical idea; tell a story that illustrates the critical idea or talk about the benefits of understanding the critical idea. With the caveman, I would be inclined to start with a story that involves him, which shows the potential

benefits of understanding Twitter. Box 3.4 suggests more ideas for starting an explanation.

BOX 3.4 INTRODUCING A TOPIC

There are so many ways to start an explanation, but you will find little with a direct search. To get a sense of the possibilities, you have to focus the search on classes and presentations. Search for one or more of the following: **"starting a lecture", "lesson hooks", how to begin a lesson, how to introduce a lesson topic in class, how to open a presentation**, or a combination of these search terms. For example: **"how to" AND (start OR begin OR open OR introduce) AND (class OR lecture OR presentation OR lesson)**.

STEP 4: CONNECT THE BUILDING BLOCKS

If a new idea cannot be rooted in things we already know, it will be impossible to understand. Some ideas can only be understood after several layers of explanation, but even these ideas have, at their foundation, things that we know. If all ideas are ultimately connected chunks made from things that we already know, there may be two reasons for not understanding something: (1) the new idea is too strange; or (2) the new idea is invisible.

If the idea is too strange, it may be because it is completely new to us or we have come across it before but simply forgotten it. We have not integrated it into our network of understanding, because we don't see how it is connected or built up from anything that is already in this network. The solution is to make the connections clear. The most effective way to make an idea appear familiar is through metaphors and analogies. If the situation allows, another way is to simply show how the idea is applied in a familiar context, and let the students see how it interacts with a reality they already know. If this is not possible, having students imagine an authentic situation can also work.

Richard Feynman was a master of making analogies. His explanations of even the most complex ideas are filled with the simplest of analogies: rubber bands explaining magnetism; bouncing tennis balls explaining steam pressure; and many more. Some authors (Hofstadter & Sander, 2013) have gone as far as arguing that all thinking is analogical. Others (Lakoff & Núñez, 2000) demonstrate that almost all mathematical ideas are built on two metaphors: the journey between two points and a container.

By grasping that new ideas can be used, such ideas become more familiar to the student, so, in a way, applications are analogical too;

because, after an application, a professor will draw parallels between what happened and the new idea. It involves showing how the ideas may work in authentic situations to explain why things happened, solve problems or predict consequences. Such situations are created to get the student to think beyond the actual situation. Again, Feynman constantly did this, regularly relating the ideas he explained to real-life situations or getting his audience to imagine scenarios in which they could use the ideas to make sense of things. Ideally, students should be involved in applications; but, if this is not feasible, narratives in the form of true stories, news events, case studies or short movies can also be useful.

Often, if the idea is invisible, while not yet aware of it, we already understand the idea intuitively. However, we cannot think about an idea and connect it when it is below our conscious awareness. This is often the case with abstract ideas. In this case the solution is to make the new idea itself stand out. Abstract ideas can be made more visible by means of demonstrations and the instantiation of ideas.

Instantiation is simply giving examples that show that students already know the idea or have seen it before. However, examples can be confusing if they are not signaled as such. Before giving an example, always precede it with a phrase like 'for example…' and make it clear how the example helps us to see a broader or more abstract idea.

Demonstrations are similar to applications, except that they show the existence and nature of new ideas through situations created especially for use in the classroom. We are all familiar with the demonstrations and experiments that science teachers do in class, but they don't have to be physical and can include games, role plays, and simulations. One of my favorites is a simulation of an old-fashioned open-outcry market that shows students how much they already know about supply and demand, even if they never heard of the terms before. (Many such simulations are available from the Council for Economic Education website.)

Looking at the un-chunking questions, these methods of connecting the building blocks are confirmed. The similarity-dissimilarity connection implies the use of analogies, metaphors, and examples; the cause-effect connection implies the use of narratives; the reason-action implies using demonstrations and applications on the 'action' side, and reasoning on the 'reason' side (more on reasoning in Chapter 6.)

An often-under-estimated technique for making connections is implied by the part-whole connection. There is no better way of showing this than a diagram. In fact, diagrams are also useful for analogies, reason-action, and cause-effect connections. I believe that it is very useful to develop visual literacy skills, or at the very least, to become familiar

with the use of the most fundamental types of diagrams (see Box 3.5). Feynman's lectures and writings contain many rather simple diagrams that capture the essence of complex ideas. The simple-looking Feynman diagrams in theoretical physics became famous, not only because they replaced complex equations, but also because they produced deeper insight than these equations. Visuals are especially important when making lecture presentations, as explained in Chapter 10.

BOX 3.5 VISUAL TEACHING

A good book on the basic visual archetypes that are useful to explain most ideas is Roam (2008). A simple search for **"visual teaching"** or **"visual teaching strategies"** will also deliver useful results. This has nothing to do with so-called 'learning styles' – it is a concept that has been scientifically discredited (see Chapter 7). The fact is that the visual sense is well developed in almost all humans, and one of the easiest ways to quickly understand complex ideas.

Let's apply these different ways of connecting ideas to the Twitter challenge. If you did un-chunking correctly, some ideas may already have come to mind. It is possible, but not necessary, to use all of the methods to connect building blocks in an explanation of Twitter. Although not shown here, I believe the best method to use in this case would be diagrams.

- Instantiation: There are many examples of people communicating with other group members in order to gain status. So, in my explanation, I might use some of them, e.g. paintings of hunting experiences on cave walls, or telling stories about the hunt as the tribe gathers around the fire to cook the meat.
- Analogy: I could use some of these examples and extend them. I could say that a mobile device is like anything that the caveman carries around – like an axe, a spear or a club. Except there is one difference: this is a mobile device on which you can also paint pictures – in the same way you would paint pictures on a cave wall.
- Demonstration: Once the caveman understands my most basic explanation, it may be useful to demonstrate to him how Twitter works using some mobile phones.
- Application: I could discuss how the caveman might use Twitter (maybe after a successful hunting expedition) or ask him to imagine situations where something like Twitter would be useful to him.

After this step, you have the starting point of your explanation and you have figured out different ways to connect the building blocks. Now you have to work out an appropriate sequence.

STEP 5: PUTTING IDEAS IN SEQUENCE

As established in step 3, an explanation should start with something familiar and build on that. Ideally, it should also be something that attracts the attention of the audience. In the case of the caveman, he knows and cares about hunting and his position in the group, so that would be a good starting point.

Because the critical idea I want to convey is that he can gain status by using Twitter, I would remind him of the ways in which he currently tries to gain status (painting about his hunt on cave walls). So now I have a good analogy for Twitter – which I can use to compare to, but also contrast with, Twitter. I will explain to him that he cannot gain as much status by painting on his cave walls, because it cannot be seen by people in other caves. Now I can start to introduce him to the basic idea of Twitter as a way to share his experience with people in other places (though I probably should warn him it is going to sound like magic to him).

At this point, he would have a very basic idea of Twitter that is not completely accurate, but it would have laid the foundation for me to add further details. I can now add a new layer of complexity by explaining that there are things like mobile devices by showing that these are similar to objects that he carries around, and that the paintings he makes can also appear on these things. Once the caveman understands this, I can add yet another layer of complexity by showing him an actual phone and how I can draw pictures on my phone and send it to other phones.

The final product may be something like this:

> Sometimes when you have a great hunt, you want to tell others about it because they weren't there to share the experience with you. So, you usually paint it on your cave wall. But still, it can only be seen by a few people.
>
> But, in our society, there is a way to share it with many more people – people living in other caves far away. We call this way of sharing 'Twitter'.
>
> Twitter even allows people to follow what happens on your hunts in the future. And as more people follow you, your reputation can spread so that you can become known as a great hunter!
>
> Now, I warn you, the way Twitter works is going to sound like magic. Imagine this: next time you paint a picture of your hunt it

appears on the cave walls of other caves, even those far away. This is the magic of Twitter.

Actually, the magic of Twitter goes further. Your pictures need not only appear on things that stay in one place (like cave walls), but also on things that move around and the things that you carry with you (like your club or spear). So, with something like Twitter, your paintings can appear on things people carry with them. What Twitter does is this: every time you paint a picture on your cave wall, it also appears on other people's clubs or spears. With Twitter, your paintings no longer stay in one place – rather, you can share them with everyone who wants to know about your experiences.

Here, let me show you [now I physically demonstrate it, but only once I am sure he understands the ideas so far]. This small thing is like my club, but I call it a 'phone' [I point to my phone]. On this phone, I paint pictures. I know these pictures look strange, but that is just our way of doing it. We call it letters. When we put letters together, it makes people see pictures in their head.

Now see how it happens. I 'draw' a picture on my phone. Then I can make other people see it too when they look at their phones. [I get a friend with a phone and I send a tweet]. Look what just happened on my friend's phone… the same picture is there now [showing that the same thing just appeared on my friend's phone].

Notice how my explanation gradually added new ideas and connections, moving from the critical idea outwards. A rough mapping of how the explanation proceeded can be seen in Figures 3.4 to 3.6 and will also help you to see how there was a gradual complexification.

You also notice how my actual explanation kept on spiraling back (as explained in the second chapter) to the critical idea of sharing but in a slightly different way every time. This ensured that the new ideas are regularly integrated with the critical idea so that the caveman's understanding can become more solid.

Figures 3.4 to 3.6 show the gradual complexification of the hub-and-spoke structure, and how more ideas get added and connected, at higher levels of complexity. Over time, it approaches an expert structure, which is based on a hub-and-spoke structure (see Figure 2.5).

At this point, there are still many things I have not explained – like social networks, retweets, followers and the 280-character limit. But this would be too much information for the caveman and these ideas (while important) are not necessary for him to understand the critical idea. I will only get to these ideas in a future explanation once I am sure he understands my initial explanation.

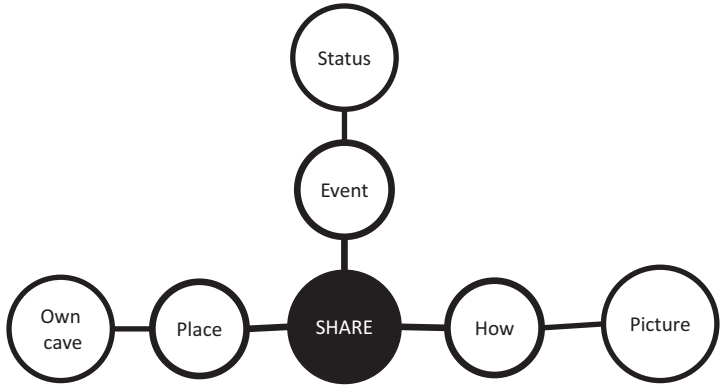

Figure 3.4 First Stage of the Twitter Explanation

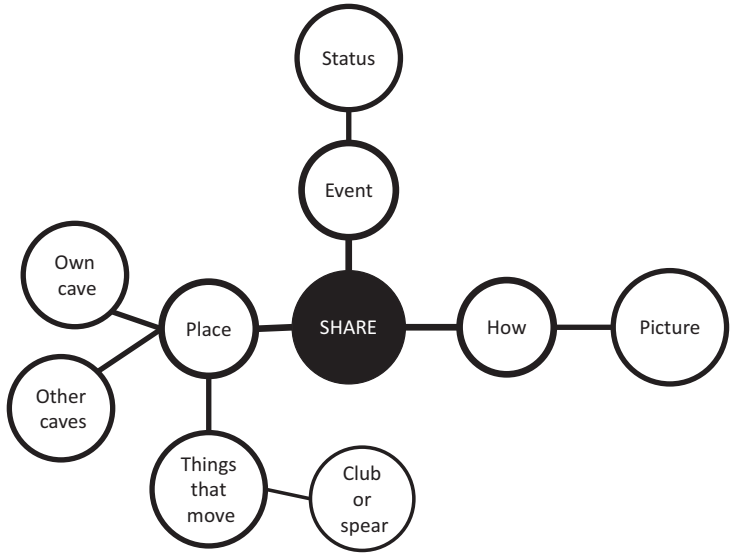

Figure 3.5 Next Stage of the Twitter Explanation

Also, my explanation of Twitter is not fully accurate yet, so I will keep on adding more complexity in layers. Perhaps the next layer would be adding the idea of followers, the idea of retweets and the idea of character limits. Only much later would I consider explaining the idea of typing/writing and the technology behind Twitter. But, even without

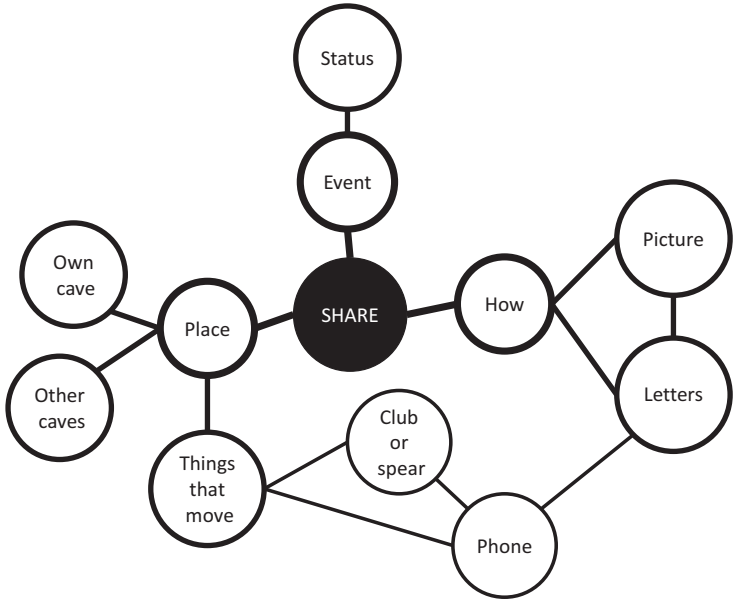

Figure 3.6 Final Stage of the Twitter Explanation

these ideas, the caveman should understand the central idea of Twitter as the easiest way to share experiences (and thoughts) with others. A more detailed analysis of the Twitter explanation, and the methods used to connect the ideas, can be found in Appendix 3.2.

REFLECTION: HOW ACCURATE SHOULD AN EXPLANATION BE?

Great explanations of complex ideas will never be 100% accurate at first. They will be coarse-grained, omitting complications, subtleties and exceptions. Sometimes, they may contain ideas that are not completely accurate, as in the case with of analogies (all analogies break down at some point).

One sees great explainers struggling with this conflict. On the one hand, they want to promote understanding, and, for that, they have to employ simplifications that other experts will criticize. However, on which side of the conflict you choose to stand depends on your audience. When teaching, the audience is comprised of students and novices, so our purpose is to help them understand. We have to start at the center

of the network of ideas, where the ideas are not as detailed, and work our way outward towards the details and complexities. If students do not grasp the critical ideas on which they can hang the details, they will simply forget the details or cram them into their short-term memory for an exam. It is for this reason that I agree with LeFever (2012) when he argues that one should initially trade accuracy for understanding.

In fact, the great explainer Richard Feynman was known to do this. Notice how Hillis (1989, p. 82) reminisces about one of Feynman's explanations:

> This was a typical Richard Feynman explanation. On the one hand, it infuriated the experts who had worked on the problem because it neglected to even mention all of the clever problems that they had solved. On the other hand, it delighted the listeners since they could walk away from it with a real understanding of the phenomenon and how it was connected to physical reality.

This shows how the greatest experts can be great explainers if *they know what to neglect* initially, with the aim of helping their *audience* understand. Feynman, like all of us, nevertheless struggled with this dilemma, but he always prioritized understanding now in order to achieve accuracy later. Consider this observation by Feynman (1999, p.3):

> Now, what should we teach first? Should we teach the correct but unfamiliar law…? Or should we first teach the simple … law, which … does not involve such difficult ideas? The first is more exciting … but the second is … a first step to a real understanding of the first idea. This point arises again and again in teaching … but at each stage it is worth learning what is now known, how accurate it is, how it fits into everything else, and how it may be changed when we learn more.

He basically summarizes many of the arguments in this book so far. Reading his introductory lecture "Atoms in Motion", you will notice that he is not afraid to simplify (or rather compress) ideas, and then return to them later and explain that what he said before is not entirely accurate. He does it six times in this lecture alone! His purpose is to get the audience to first understand the idea in broad terms, and then come back and gradually, layer by layer, fit more details into the idea.

Furthermore, I hinted in Chapter 2 that this conflict between simplification and accuracy may be a false dilemma. Understanding is the

compression of knowledge to an essence from which many more ideas can be inferred through reasoning and/or imagination. This compression, when communicated in an explanation, will appear to be a simplification, though it is better to think of it as 'generative simplification'. In a generative simplification, knowledge is not eliminated or twisted. Simple explanations are more like seeds that contain the potential for the details to grow from them. The details only grow when the understanding of the simple explanation awakens the growth of these seeds. The dilemma is false because it is not possible to achieve understanding and perfect accuracy simultaneously – rather, we use understanding to provoke accuracy in an endless cycle of growing expertise.

BOX 3.6 DON'T 'DUMB THINGS DOWN'

There is a difference between generative simplification and oversimplification (or 'dumbing down'). The link is a multi-part article that helps you to distinguish between the two and avoid oversimplification. (https://wordpress.com/view/argumentativeresearchwriting.wordpress.com)

The caveman does not know, and does not care, that I was simplifying. He wants to understand Twitter first, because this enables him to grow his understanding by exploring, for example, how a phone is different from a moving club, or how letters are different from pictures.

STEP 6: PROVOKE THINKING

Finally, I can now improve the caveman's basic understanding by asking him to think about other situations where he could use Twitter, or maybe ask him if he would like it if his children had something like Twitter too. By getting him to think, I can also stimulate his asking the kinds of questions that will reveal whether he is ready for another layer of complexification.

HOW TO FOLLOW THE PROCESS WITH, OR WITHOUT, A TEXTBOOK

Finally, let us summarize. Figure 3.7 captures the six steps in constructing an explanation.

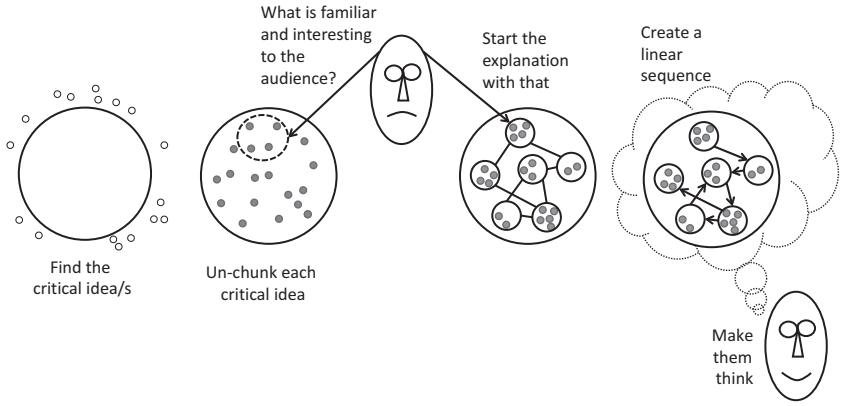

Figure 3.7 The Process of Constructing an Explanation

This process is not difficult when you teach a course or explain a topic from various sources, including your own understanding. It becomes more difficult when you rely on a textbook because a textbook can paralyze even the best explainer if it is not used appropriately. Assuming that you already know your audience, let us apply and modify the steps in cases where you must use a textbook:

1. You have to teach from a place of understanding. If not, open the textbook and study the topic. You may return to this step periodically because the next steps will most likely make you realize that there are parts you don't understand.
2. *Close* the textbook and keep it closed in this step. Find the critical ideas of the topic. You will not find the critical ideas by reading the textbook, in fact, doing so will only cause you to get lost in the details. *Only* consult the textbook if this step makes you realize you don't understand something, and then try again to find the critical ideas with the textbook closed.
3. Use the technique of un-chunking to generate the most important content that will help you to explain the idea. Ideally, you should not read the textbook during this step unless you are looking for something specific, like an example or useful fact. Again, it is through your understanding that you generate this content. If you read the textbook, you will simply repeat the details from the textbook and not think clearly about how to elaborate on the critical ideas.

4. From the ideas generated in step 3, identify which ones are familiar to the students. These are potential starting points, and, from these, identify which ones are likely to lead to engaging introductions. If the explanation builds on a previous explanation, part of the introduction should also include the connection between the explanations. The textbook remains closed unless you are looking for something interesting.
5. Identify places where the ideas may be too abstract or where students may not have sufficient or strong connections. Think of analogies or design applications that will make abstract ideas clearer, and think of illustrative examples or design simple demonstrations to help explain ideas that may be too unfamiliar to students. Also, consider other methods like diagrams and narratives. The textbook remains closed unless you are looking for something useful.
6. Without referring to the textbook, find a good linear sequence to explain the ideas, starting from what you selected in step 5. Employ the spiraling strategy, as described in the second chapter, to ensure that critical ideas are integrated. Sometimes you may find yourself alternating between steps 5 and 6 because it may be easier to first find a good sequence and then determine places where the connections need to be strengthened.
7. Develop the complete explanation and present it (more on this in Chapter 10). Also, ensure that you create opportunities for students to think about the ideas. Consider how to make the explanations more interesting (see Chapter 5) and where it would be helpful to get students to memorize something and assist them in doing so (see Chapter 4). Use the textbook only if you need to refine or look up something.
8. Return to step 1 if this explanation leads to another explanation and make the connection between them. Through this process, gradually make students' knowledge structure more complex.

EXERCISES

1. Analyze the talk or lecture of a great explainer.
   a) Who is the audience? Does the explainer talk as if he or she is aware of who they are?
   b) Does the explainer organize the talk around one or more critical ideas?
   c) Once you are sure you understand the talk, identify the critical idea/s and un-chunk them. From this, do you recognize

the content of the talk? Are there ideas that you can see the explainer did not talk about? Why do you think the explainer neglected some ideas?

d) How do they start the talk? Is the starting point derived from the critical idea? Is the starting point something that is familiar and engaging to the audience?
e) Does the explainer make connections between ideas?
f) Does the sequence make sense to you? Why? Does the explainer sometimes spiral back to previous ideas? Does this improve the listeners' understanding?
g) Does the explainer sometimes simplify in order to promote understanding? How? Are the simplifications helpful or generative?
h) Does the explainer get the audience to think? How?

2. Take another everyday phenomenon or something that is in the news (such as nuclear bombs, traffic lights or black holes) and try to explain it to a caveman, using the steps in the chapter.
3. Identify one idea (not a whole topic) that your students find difficult to understand or a topic in which your explanations don't work well. It should ideally be something that usually takes around 5–10 minutes to explain. Construct an explanation using the steps in this chapter.
4. Take a complex idea in your subject, one that requires a series of explanations. After constructing the individual explanations, put them in a sequence that strengthens the understanding of the critical ideas.

APPENDIX 3.1: WAYS TO MAKE A KNOWLEDGE STRUCTURE MORE COMPLEX

In the second chapter, I highlighted three ways to complexify a hub-and-spoke structure: (1) adding connections between ideas; and adding new ideas, either by (2) inserting new ideas in between old ideas or (3) growing new ideas out of existing ideas. I also identified four kinds of connections that one can make: (a) cause and effect; (b) reason and action; (c) similarity and dissimilarity; and (d) parts and whole. By combining them, we obtain different ways of elaborating on ideas, as shown in Table 3.2.

Each cell described in Table 3.3, numbered one to twelve, represents one possible way to make a knowledge network more complex.

Table 3.2 Ways to Elaborate

Ways to elaborate Type of connection	Add connections between existing ideas	Insert new ideas in-between existing ideas	Grow new ideas from existing ideas
Cause-effect	1	2	3
Reason-action	4	5	6
Dis/similarity	7	8	9
Part-whole	10	11	12

Table 3.3 Examples of Elaboration

	Description	Example
1	A known cause or effect is also related to other existing causes and effects.	The economic problems in Germany after World War 1 were caused by the rise of populist leaders such as Hitler.
2	Show that an intermediate cause or effect complicates an existing cause-and-effect relationship.	The economic problems in Germany *led to widespread poverty* and this helped the rise of Hitler.
3	Show that more causes or effects exist than were previously indicated.	Besides the rise of Hitler, *there were other causes of World War 2, such as…*
4	Connect a known reason (why?) for an existing idea or known action (how?) to existing ideas.	Companies raise their prices *because* they want to protect their profits.
5	Show that an intermediate reason or action complicates an existing reason-action relationship.	Companies raise prices because they *want to cover their higher costs,* and they want to do this in order to protect their profits.
6	Show that more causes or effects exist than were previously indicated.	Companies raise their prices because they want to protect their profits, but there are *other reasons to do this, such as…*
7	Show that two previously unrelated ideas are more similar than previously thought, or that they are opposites.	A church is like a production process, with inputs (sinners), a process of change (religious conversion) and outputs (saints).
8	Insert an intermediate analogy or a simpler example that leads to a more complex example.	A simple example of a production process is a church and *with that, we can understand the example* of car manufacturing plant.
9	Introduce a new relation of similarity (example, analogy etc.), or a relation of dissimilarity (non-example).	The manufacturing of cars is an example of a production process, but the distribution of those cars is not.
10	Show that two existing ideas stand in a part-whole relationship: one is part of	Animals fall into the *category* of multicellular eukaryotic organisms. Animals can be *divided into* smaller

	Description	Example
	the other, or, one is a larger category into which the other one falls.	categories, like bilaterian and non-bilaterian organisms.
11	Show that there is an intermediate category between two ideas that stand in a part-whole relationship.	Crabs are animals, but crabs *also form part of invertebrates, which is a sub-category* within the category of animals.
12	Show that there are more kinds or categories of an existing idea.	Dogs can fall into *many different kinds* of categories. They are pets but they can also be categorized as mammals, carnivores or predators.

APPENDIX 3.2: ANALYSIS OF THE TWITTER EXPLANATION

Table 3.4 analyses the explanation by indicating what method of strengthening the connection is being used. The abbreviations are: Ins = Instantiation; Ana = Analogy; Dem = Demonstration; App = Application.

Table 3.4 Analysis of an Explanation

Explanation	Ins	Ana	Dem	App
Sometimes when you have a great hunt, you want to tell others about it… So, you usually paint it on your cave wall, but still, it can only be seen by a few people.	x			
But in our society, there is a way to share it with many more people – people living in other caves far away. We call this way of sharing 'Twitter'… as more people follow you, your reputation can spread so that you can become known as a great hunter!	x	x		
Now, I warn you, the way Twitter works is going to sound like magic. Imagine this: next time you paint a picture of your hunt, it appears on the cave walls of other caves, even those far away. This is the magic of Twitter.	x			x
Actually, the magic of Twitter goes further. Your pictures need not only appear on things that stay in one place (like cave walls), but also on things that move around and the things that you carry with you … your paintings no longer stay in one place – rather, you can share them with everyone who wants to know about your experiences.		x		x
Here, let me show you [now I physically demonstrate it, but only once I am sure he understands the ideas so far]. This small thing is like my club, but I call it a 'phone' [I point to my phone]. On this phone, I paint pictures …		x	x	

Explanation	Ins	Ana	Dem	App
When we put letters together, it makes people see pictures in their head.				
Now see how it happens. I 'draw' a picture on my phone. Then I can make other people see it too when they look at their phones… Look what just happened on my friend's phone … the same picture is there now.			x	

REFERENCES

Chase, W.G. & Simon, H.A. 1973. Perception in chess. *Cognitive Psychology*, 4:55–81.

Feynman, R.P. 1999. *The Pleasure of Finding Things Out: The Best Short Works of Richard P. Feynman*. Cambridge: Helix Books.

Hillis, W.D. 1989. *Richard Feynman and the Connection Machine*. Physics Today, 42(2):78–83.

Hofstadter, D.R. & Sander, E. 2013. *Surfaces and Essences: Analogy as the Fuel and Fire of Thinking*. New York: Basic Books.

Lakoff, G. & Núñez, R. 2000. *Where Mathematics Comes From*. New York: Basic Books.

LeFever, L. 2012. *The Art of Explanation: Making your Ideas, Products, and Services Easier to Understand*. Hoboken: Wiley.

Miller, G.A. 1956. The magical number seven, plus or minus two: Some limits on our capacity for processing information. *The Psychological Review*, 63:81–97.

Nathan, M.J. & Petrosino, A. 2003. Expert blind spot among preservice teachers. *American Educational Research Journal*, 40(4):905–928.

Roam, D. 2008. *The Back of the Napkin: Solving Problems and Selling Ideas with Pictures*. New York: Portfolio.

Sokal, A.D. 1996. Transgressing the boundaries: Towards a transformative hermeneutics of quantum gravity. *Social Text*, 46/47:217–252.

Wentzel, A. 2018. *A Guide to Argumentative Research Writing and Thinking: Overcoming Challenges*. London: Routledge.

Chapter Four

Help Students to Remember Creatively and Forget as Much as Possible

Kim Peek would have been a star student in many professors' classes. Not only could he speedread books, the left page with his left eye and the right page with his right eye, but he could remember most of what he read. According to his obituary in the New York Times (Weber, 2009), he could recall the contents of 12,000 books on a wide range of topics such as music, history, geography and literature.

There are a few other examples of these superlative memorizers; the most famous of which may be Solomon Shereshevsky. According to the neuropsychologist, Alexander Luria, Shereshevsky could remember long lists of random numbers and words after only hearing them once. And apparently, when tested fifteen years later, he could still repeat them accurately. Other remarkable cases are: Stephen 'the human camera' Wiltshire, who can accurately sketch any scene he has seen, even if only for a few seconds; and Jill Price, a woman who can recall details of every day of her life since childhood.

Some people call this kind of memory 'photographic memory' and most people, including many students and professors, associate it with success and intelligence. No wonder then, that a whole industry has been created around promises to help people improve their memories.

Before continuing, think for a moment about how well people like Kim, Solomon, Stephen and Jill would do if they were able to apply their memory abilities in your courses. Would they pass, fail or be top of the class?

Your answer will reveal the extent to which you need to read this chapter. This chapter will remove the misunderstandings about memory that are most damaging to teaching and replace them with ideas that are closer to what scientists know about memory. You will see that memory is not about adding and accumulating facts, but rather about forgetting

as many facts as possible. Also, memory is, in fact, quite creative, with imagination playing a big role. Furthermore, you see that memory, if used appropriately, is intertwined with understanding.

MEMORY IS CREATION AND RECREATION, NOT STORAGE AND RETRIEVAL

In my early days as a lecturer, Kim, Solomon and the others would certainly have passed, and, in fact, may even have been my best students. This was a sign that my early courses did not promote understanding and that I misunderstood the purpose of memory. This misunderstanding is, according to Bain (2004, p.83), common among mediocre professors "who think of memory as a storage unit and intelligence as the capacity to use the information in that tank". They think of teaching as an act of filling up those tanks, and look upon assessments as testing the level of students' tanks and how precisely their contents conform to the original facts learned from lectures and textbooks.

Thinking of memory as the storage of information is completely wrong. Firstly, the evidence is mounting that photographic memory, in the sense of perfect long-term storage of events or situations, does not exist, including each of the cases mentioned above. More importantly, research shows that human memories are neither stored nor retrieved (O'Loughlin, 2017). Rather, to commit a fact to memory, a memory has to be created around this fact by connecting it to known ideas, and when trying to remember the same fact, the memory has to be recreated.

Let me explain. To memorize anything successfully, we cannot merely store it as a loose unit in a single container. Instead, a fact can only stay in your memory if you connect it to already familiar ideas and to the associations you make in the current situation, as shown in Figure 4.1a. The fact changes to the degree that is attached to other things; it can no longer remain a mere fact, but becomes part of an idea. If it does not get changed in this way, it is forgotten. For example, the word 'serendipity' stayed in my memory only once I connected that word to the story of the three princes of Serendip and the person who used it in our conversation, and how it applied to the topic of innovation we were discussing at the time. I probably heard that word many times before, but never remembered it because I never connected it to other ideas I knew. This process of memorization ensures that, while you and I may agree on the definition of serendipity, it will not have identical meanings for us due to the different ways we connected it to our knowledge network.

Also, to successfully remember something, the brain does not simply find the memory in one place, but literally *re*-members it, that is, brings

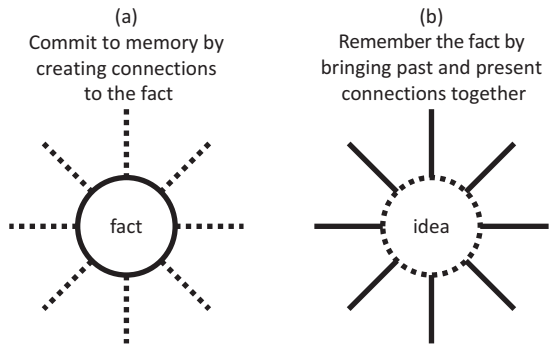

Figure 4.1 Constructing and Reconstructing Memories

all the members together. The original fact I try to remember is no longer an isolated fact, but an integral part of one or more ideas. When I remembered the meaning of serendipity sometime later, I recreated it from the original connections I made, as well as from the associations I have now made in my current situation. The memory is brought back to life from old and new connections, as seen in Figure 4.1b. This suggests that remembering something changes it, as shown in several classic and modern studies with the aid of stories and pictures (Bartlett, 1920; Carmichael, Hogan & Walter, 1932; Hemmer & Steyvers, 2009). For example, when shown two circles connected by a line, some people later remembered it as a pair of glasses while others remembered a barbell. The original fact changed as a result the connections made before and during the act of remembering.

If a person's memory was simply a loose bunch of facts, he/she should not be able to pass a course that requires understanding. As explained before, understanding is the ability to compress a vast collection of facts into a *network* of a few critical ideas from which further ideas can be generated. To create and expand this network, facts and ideas are connected. Hence, someone who can flawlessly recall vast numbers of facts, but cannot see the connections between them, should fail most university-level courses.

Kim Peek was like that: he held a massive stockpile of musical, historical, sports and geographical facts, but struggled when understanding was required. He had an IQ that was well below average because he was born with brain damage. Amazingly, he did manage to gain understanding, especially in music, towards the end of his life. This only happened as he started to make connections between musical facts. When recognizing a piece of

music, he became able to link it to other facts – about the composer, the musical era, style of composition etc. – and eventually applied the same method of understanding when he learned to play music. His comprehension increased as he gradually transformed his storehouse of facts into a knowledge network. No wonder, then, that he was the inspiration for the celebrated movie, *Rain Man*.

When memorization happens through connection, it is very hard to distinguish it from thinking. In fact, cognitive scientist, Daniel Willingham (2009, p.41), explains that, "Memory is the residue of thought". Connections are made by thinking, so we remember the things we think about the most. As shown in Table 4.1, this way of thinking leads to a very different kind of teaching and assessment than that of the storage view of memory.

The connectedness view of memory is interwoven with the ability to understand. Memory should support understanding, so, before expecting students to memorize a fact, you should first recall the network organized around the critical ideas of your lecture, and then ask: "Which idea in this network does this fact help students to understand?" If it is not clearly connectable, it is not worth remembering, in class or in an assessment. This would exclude much of the content in textbooks, especially lists of advantages, disadvantages, factors, types, and so on.

We also saw that facts change as they are memorized and remembered, so it is quite *un*natural to expect students to reproduce perfect replicas of facts in assessments. Life is not a trivia competition, rather, we memorize facts in order to expand our understanding, and assessments should reflect that.

As explained before, understanding enables one to make sense of previously unseen situations, so it is better to assess how well students can

Table 4.1 Contrasting Views on Memory

	Memory as storage	Memory as connectedness
Memorizing	Storing loose facts in specific containers in the brain	Connecting facts to what is already known and integrating them with ideas
Remembering	Retrieving previously stored facts from their separate containers	Recreating facts from their past and current connections
Teaching	Teach as many facts as possible in the given time	Guide students in connecting new facts to an existing knowledge network
Assessment	Measure how many facts have been stored	Determine how well facts can be connected to situations not seen before

recognize and use the connections they made to the same facts in new situations. So, for example, instead of asking students: "What is the definition of 'inflation rate'?" I prefer to give them some economic data for a foreign country and ask them to work out the inflation rate. To do this, they have to understand what an inflation rate is, even if they cannot reproduce the textbook definition exactly. As a result, I assess not whether they can store definitions, but rather whether they are *becoming* economists.

FORGETTING, NOT REMEMBERING, IS WHAT MAKES MEMORY USEFUL

By now we know that not everything is worth knowing, and in the previous section I argued that not everything is worth remembering either. If something is not worth remembering, it should be forgotten, and it is this forgetting that makes memory valuable.

The assertion that forgetting makes memory useful may sound strange. To understand this, consider what happens when someone just accumulates memories without ever forgetting any of them. Firstly, it would be overwhelming and make it difficult to pay attention to anything else. When you are reminded of a date in your past, it probably stops there, because you have forgotten most of what happened on that day. But for Jill Price, mentioned at the start, it does not just stop there. Each date comes with a flood of all the detailed memories connected to that date. In Sloman and Fernbach (2017, pp. 38–40), she describes it like this:

> Whenever I see a date flash on the television.… I automatically go back to that day and remember where I was, what I was doing, what day it fell on and on and on and on and on.… It is nonstop, uncontrollable and totally exhausting.… Most have called it a gift, but I call it a burden.

Alexander Luria wrote the same about Solomon Shereshevsky. Apparently, he could not think of a simple thing like a restaurant without being overwhelmed by the many details from other restaurants he visited. It got so bad that Shereshevsky tried in vain to get rid of memories by writing them on scraps of paper and burning them, and later turned to alcohol. Although the facts do attest to the veracity of this, it should be noted that this could be somewhat of an exaggeration, as Luria was prone to (Johnson, 2017).

To understand, our minds need to be free to pay attention to creating new connections, and not get stuck in a swamp of old connections.

Memory supports understanding because it lets us forget unimportant connections and facts that are wrong, irrelevant and redundant.

There is a second, more important way in which forgetting promotes understanding: forgetting enables us to reduce billions of observations and facts to far fewer concepts. To understand why this is important, consider again what happens if one cannot forget details. Jorge Borges (1964) wrote a short story, inspired by Luria's work, that explains this well. The story is of a man called Funes, whose perfect memory of every possible detail only allows him to see the world as a collection of *details*. Because he cannot forget details, he is unable to see *patterns*. As explained before, patterns are regularities: they repeat themselves, but with small differences every time. It is only when we can ignore the differences that we can see patterns.

But what if, like Funes, we cannot forget the differences and other small details? We would not be able to see the bigger picture, find the overall critical ideas, recognize concepts or identify deeper principles. Funes, for example, had trouble with simple things like dogs. Every dog was different in a very detailed way, so for Funes, each one was a different animal that needed a different name. For him, even the same dog seen from the front, from the side and from the back were different entities, therefore requiring different names. There were no categories or chunks of information for him. But, for us normal ('forgetful') people, this is not a problem. After seeing just a few dogs, we notice similarities and forget the differences. This allows us to create just one chunk called 'dog' and with this we can recognize dogs, regardless of size, breed or color, even dogs that we have never seen before.

Figure 4.2 explains this in the case of shapes. The first four objects all look like circles, but not to the person who is unable to forget differences. Because each shape is different, for someone with a perfect memory, each one would need a separate name (sircal, cyrcil, curkl and sahrcle). The rest of us will simply look at all four together, notice the similarity and forget the small differences. Forgetting like this then enables us to generalize, remembering only the essence, and giving it just one name: 'circle'. Later, when trying to remember what you saw in Figure 4.2, you will most likely forget the details and remember only that you saw several circles, albeit imperfect ones.

According to the traditional storage view of memory, any kind of forgetting is a memory 'imperfection'. But, as Schacter (2001) argued, every one of the seven major memory imperfections was designed to enable humans to reduce excess information to its essence or gist, and so recognize patterns in a changing multi-dimensional world. Without such an ability to generalize from detailed facts to broad ideas, we would not be

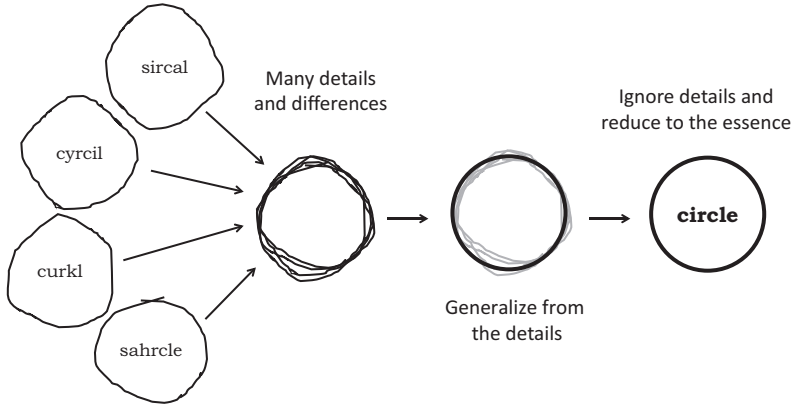

Figure 4.2 Memory as a Process of Forgetting

able to make sense of anything. This has now been confirmed by an ongoing stream of research in the cognitive sciences (Fernández, 2015, Nørby, 2015, Richards & Frankland,2017). A recent review of neurobiology research (Davis & Zhong, 2017) goes even further, suggesting that, without forgetting, we would not be able to learn.

It also explains why the renowned psychologist, William James (2007, p.679), wrote: "In the practical use of our intellect, forgetting is as important as remembering." We have to forget in order to see patterns, connections, chunks of information or critical ideas. Remember that understanding is the ability to *compress* a vast collection of facts into a network around a few critical ideas. If we could not forget what is non-critical, we would not see what is critical. And without critical ideas, we would not have something around which to grow a network of related ideas.

Table 4.2 summarizes the two contrasting views, and shows that the correct view (memory as selective forgetting) leads to a very different approach to teaching. Again, the implications of the memory-as-forgetting view are very similar to the ideas about understanding delineated in the previous chapters.

If forgetting is critical, then teaching should help students to forget in a way that reveals what is critical and assessment should be done to see how well they are able to use these critical ideas to generate information in new situations. For example, if students understand the deeper principles of patent circumvention, they should be able to generate information about patents they have never seen before, specifically the degree to which each one is vulnerable to legal action.

Table 4.2 More Contrasting Views on Memory

	Memory as perfect recall	Memory as selective forgetting
Memorizing	Copying information into memory folders in the brain without changing them	Recognizing which facts do not fit into the knowledge structure and are thus not worth connecting
Remembering	Retrieving information without losing details and without any inaccuracies	Identifying the essence from many facts and inferring new information from it
Teaching	Repeating as many different facts as possible in the given time and emphasizing details and exceptions	Teaching facts with the aim of getting students to see patterns or deeper principles
Assessment	Measuring how accurately stored facts can be reproduced	Determining how well new information is generated from old facts

HOW UNDERSTANDING AND MEMORY ARE INTERWOVEN

Experts not only understand their subjects much better than novices, they also have a much stronger memory capacity for the incorporation of new facts and details. The reason lies in something that has been recurring throughout the book so far: the way experts organize knowledge. While novices often attempt to learn facts and ideas in a disconnected or linear way, experts build a connected network structure of ideas in their subjects and organize this network around a few critical ideas. Every idea in this network is not just one fact. Rather, ideas themselves are made up of chunks – which are, in turn, made up of other ideas and facts that are connected.

It is much easier to remember something if it is connected to something that you already understand and fits into a chunk you already recognize. The more chunks you have, and the better connected they are, the easier it is to remember something that is new to you.

The classic illustration of this was shown by De Groot (1965) and elaborated on by Chase and Simon (1973). When, for between two and fifteen seconds, De Groot showed positions from various chess games to weak players, he found that they were unable to reconstruct the positions from memory. In contrast, the expert players could reconstruct every position almost exactly. The reason for this is that the weaker players saw many individual pieces, while the experts recognized patterns in the placement of the pieces and perceived connections between the pieces. Instead of having to memorize around 25 individual pieces, experts only had to remember four or five chunks.

The mathematician, George Boole, who had a remarkable memory, captured this well when he explained in a letter to Charles Babbage that his ability does not result so much from strength of memory as from "the power of arrangement, which provides its proper place in the mind for every fact and idea and thus enables me to find at once what I want" (University College Cork, n.d.). The 'arrangement' can be seen as a connected set of ideas in which one can easily find the 'proper place' for a fact or idea. With a well-organized knowledge network, memorizing and remembering related facts becomes almost effortless. It is not surprising then that Bain (2004) found that the best professors emphasized the importance that students understand knowledge structures rather than merely storing information.

Under conditions of rising complexity, understanding can and should replace memory, as with what happened to me when I started to study Accounting in Grade 8. Initially, I tried to memorize every single ledger and journal entry for every single transaction. While the content was simple, this worked well; but, after two years, it was too hard for me to keep up. However, once I understood the simple idea of the accounting equation and how it was connected to the recording of transactions, I could very often work out the relevant details of how to record *any* transaction. I eventually majored in, and even taught, Accounting at university.

One way, therefore, to know the details is to memorize every one. That is hard work. Another way is to generate the details, as needed, from an understanding of the connected critical ideas. And that is working smarter, not harder. In a world that changes and presents new situations all the time and given the limited capabilities of humans, it is clearly much more efficient to understand than to try to perfectly recall facts. Prior understanding also enables one to recognize differences and exceptions faster and to have a better memory for them. Understanding is really just a more flexible and sustainable form of memory.

We should not teach and assess students' knowledge as if understanding and memorization are separate intellectual actions. Each one is only effective and makes sense when supported by the other. Memory filters information through forgetting, and so supports compression, making understanding possible. Something can only be remembered by connecting it to other ideas and is made possible by understanding. Furthermore, although understanding may be a more efficient way to handle details, there are cases where memorizing the details first actually makes understanding more efficient. This is what we turn to next.

SOMETIMES 'STORAGE' WITHOUT UNDERSTANDING MAY BE NECESSARY

While, as highlighted in this chapter, the human brain is not just a bunch of containers for perfect memory storage, there are times when memories have to be retained without distortion. This is true in the case of information necessary for survival or when variation causes information to lose its value (such as dates in history, locations in geography, certain values in physics, and so on).

Sometimes memorization may be the only option. For example, when learning the vocabulary of a foreign language, the rules of a game or the terminology of a new subject, memorization is, in most cases, what is required. For example, the word for 'chair' in Spanish is '*silla*', and there is no way that you can work that out from understanding chairs; you just have to memorize it.

In instances where working things out from understanding takes too much time or energy, it is more efficient to memorize without the aid of understanding. When survival depends on action – for example, when one encounters a predator or what may be a poisonous plant – thinking about everything would take up precious time. Even in more conventional situations, having things in your memory allows you to act faster or devote your attention to more complex tasks. It is for this reason that you can solve more complex arithmetic problems faster if you have memorized the multiplication tables.

Mathematician, Alfred North Whitehead (1948, p.42), even went so far as saying that: "Civilization advances by extending the number of important operations which we can perform without thinking about them." The more things we can do automatically and without thinking, the more we free up our cognitive resources to be spent on thinking about new and complex ideas.

WHEN IS MEMORIZATION WITHOUT UNDERSTANDING APPROPRIATE?

It can therefore make sense to memorize without understanding. Being able to distinguish between what information should just be memorized in a course, and what not, is a sign of understanding. It includes the following:

- **Foundational facts that cannot be understood.** Some facts are a matter of convention, so one cannot work them out. For example, you cannot understand that '+' means to add

together, you just have to know it. You can read about the history of the plus sign, but there is no logical explanation as to why a plus sign must necessarily mean 'add together' – someone decided that and now everyone uses it. These facts are foundational in the sense that they come first because they cannot be derived from other ideas. Such facts include signs, rules, vocabulary, definitions of basic concepts and basic categorizations in a course. Complex ideas form as foundational facts are connected. So, if you don't know the foundational ideas, you will not be able to understand more complex ideas.

- **Key details and exceptions**. Critical ideas, as explained before, are not always 100% accurate because they leave out some details. Sometimes there could be an exception or some detail that is required for an important connection, and it would hamper understanding if students were not aware of it. These need to be memorized temporarily. A time should come later, when the student's knowledge structure is sufficiently complex, that the exception or detail will make sense.
- **Important, but excessively complex ideas that cannot yet be fully understood**. Some complex ideas are critical for an important connection, but would take too long to explain. In order to speed up students' learning, one may initially treat these ideas as facts and expect students simply to memorize them. As they start to use these complex ideas, they should eventually gain understanding, or one could return to them later and explain them in more detail once students have a more complex knowledge structure. For example, mathematical volumes have been written about the concept of infinity. But, in an arithmetic class, it is not necessary for students to grasp it completely in order to use it. Explaining infinity would delay the learning of more critical ideas, so it would be better to memorize the meaning of infinity at first, as if it were a fact.

In each of the three cases above, notice that memorization serves understanding, so there is still no justification for courses that can be passed by memorization alone. Memorization merely *accelerates* understanding by: (1) making communication and complexification possible (in the case of foundational ideas); (2) saving time (in the case of key details); and, (3) reducing the number of ideas that students need to keep in mind at a given time (in the case of over-complex ideas).

Memorization, in every case, creates placeholders in the knowledge structure of the student until understanding arrives. Figure 4.3 illustrates the process. Initially, a fact or idea is memorized, but the student cannot,

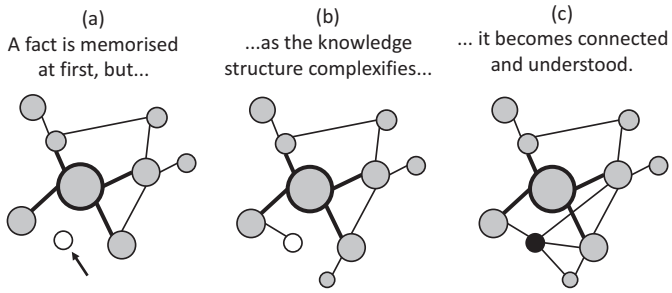

Figure 4.3 Memory as a Placeholder for Future Understanding

for whatever reason, understand it fully yet, so it is initially disconnected. Through effective teaching, the initially memorized information becomes connected as other ideas grow into and around them. The knowledge structure gradually complexifies until the student finally has sufficient background to make sense of it. At that point, memorization is no longer needed and the idea can be safely released from the inflexible 'holding place' of memory. For example, when learning a new language, you will often start by memorizing words and simple grammar rules. You then start using these words and grammar, and very soon you start getting a feel for the language. As you start understanding how the language works, memorization becomes easier and you will, in fact, start relying less and less on memorization of words and phrases.

Caine and Caine (1991) explain this with the metaphor of getting to know a new city. Initially, one memorizes routes, but this is done in order to form a map where all the routes are connected. Once you have such a mental map of the city, you can rely much less on memorizing routes and remember new ones faster, because you understand how different parts are connected.

Memorization without understanding should occur much less frequently than is seen in most courses. A foundation takes up much less volume than the structure built on it, so foundational facts are relatively few. Facts that do not have a clear connection to the critical ideas should not be memorized. Besides the three exceptional instances above, memorization just for the sake of memorization is pointless. I am embarrassed to say that my early students spent far too much energy on pointless memorization; energy that I should have redirected towards gaining understanding. Expecting students to memorize facts, for the sole purpose of making up marks in an assessment, is a waste of the power of memorization.

TECHNIQUES TO HELP STUDENTS MEMORIZE

When memorization is necessary, it does not have to be difficult. Obviously, when it is combined with understanding and interest, memorization is almost effortless. However, when it has to be done without understanding, mnemonic devices can make it relatively effortless too, and it would help our students if we exposed them to such techniques. (The word 'mnemonic' comes from the ancient Greek word '*mnemonikos*', which means 'relating to memory').

Mnemonic devices exploit the fact that the strongest memories are created through connection. Although they enable one to memorize without understanding, they still involve creating connections, albeit artificial ones, between what you want to remember and what you already know. Every mnemonic device is based on this principle.

Different people have different memory styles. This is not to be confused with the idea of learning styles, for which there is no credible evidence (see Chapter 7). It is, in fact, different *courses or subjects* (not people) that demand different ways of *learning*. Different *people*, on the other hand, prefer different ways of *remembering*: some remember better visually, while others remember better kinaesthetically or by auditory means. For example, you cannot learn to solve equations by singing about them – you have to solve them logically – but you can remember the rules used in mathematical proofs by making a song about it.

Having different mnemonic techniques helps you in that you can use the one that best fits an individual's memory style. This section will introduce you to some, and Box 4.1 will direct you to many more.

BOX 4.1 MNEMONICS FOR EFFORTLESS MEMORIZATION

There is a huge amount of useful information on the web about mnemonics and learning as you will see if you search for (**mnemonics OR "memory techniques"**) **AND learning**. The Wikipedia entry for 'mnemonic' is a good place to start. In her memory experiments and book, Lynne Kelly explains some methods that are not so well known, which you will find if you search for **"memory code" AND "memory experiments"**. Remembering definitions and vocabulary can sometimes be challenging, so search for **"mnemonic dictionary"** for ideas about remembering them with the aid of mnemonics.

What follows are not the only techniques, but the ones I tend to employ most often in my own studying and teaching. All techniques have one thing in common – they require you to first extract keywords or very short key phrases that describe what you want to remember. Sometimes the keywords are obvious or already given, but often you have to find them yourself.

Analogies

With analogical memory techniques, you remember one set of ideas by comparing it to another set of ideas that already has a pre-existing structure, like the rooms of a house, the components of an object, the parts of your body, your circle of friends and many more.

Suppose I want to remember the five components of the market environment of a business. In the textbook, it appears as follows:

- The people that buy the firm's products;
- Intermediaries through which the firm sells its products;
- Other firms that supply the firm with inputs;
- Competitors with which the firm competes;
- The greater society within which the firm operates.

My first step is to find a keyword that describes every point in a way that I will remember what the point is about. The keywords could be words that already appear in the points, or new words that better describe what the point is about for me. The result may be as follows:

Table 4.3 Extract Keywords

Point to remember	Keyword
The people that buy the firm's products.	Customers
Intermediaries through which the firm sells its products	Intermediaries
Other firms that supply the firm with inputs	Suppliers
Competitors with which the firm competes	Competitors
The greater society within which the firm operates	Society

The next step is to choose an existing phenomenon that has a structure to which the keywords can be related. Ancient Greek and Roman orators used the rooms in a building (see Box 4.2), but I'll use my circle of acquaintances and friends, in such a way that when I think of them they will remind me of the five components.

Table 4.4 Link the Keywords to People

Keyword	The person in my life
Customers	My students whom I teach and so provide them with a service
Intermediaries	Peter, who is like an intermediary friend in my life because through him I meet many other people
Suppliers	Welile, a friend from whom I buy products that I can't normally find in the shops
Competitors	Stephen, an old guy who once tried (unsuccessfully) to convince my girlfriend to go on a date with him
Society	My neighbors in the building where I stay

I will choose five people I know who play roles in my life that are similar to the five components. It is important that I can clearly explain, in a way that makes sense, why I connect each person to each component, otherwise I will not form a strong connection.

If I want to remember the components of the market environment, I simply think of the people in the environment around me, and they remind me of the components.

The problem with this approach is that it is very personal and is unique to every student. If you want to help all students simultaneously, you could try a more impersonal technique of relating it to something that everyone knows – like body parts. I can choose five parts that are in some way similar roles to the five components.

Now if I want to remember the components of the market environment, I think of my own body, and it helps remind me of the components. I have used this technique with many different things, like everyday objects: my car, a pen, a tree, and many more. The more

Table 4.5 Link the Keywords to Body Parts

Keyword	Body parts
Customers	Parts that have needs I have to satisfy every day, like my mouth or my stomach
Intermediaries	My hands that act as intermediaries between me and other people, and between me and objects
Suppliers	My eyes and my ears that supply me with new information all day long
Competitors	My head and my heart, which are always in competition
Society	The clothes that I wear – because they are all around my body like society is all around me

similar the thing is to what you want to try to remember, the more effective this technique is.

BOX 4.2 SPATIAL MEMORY METHODS

Spatial memory methods are part of the category of mnemonic techniques described in this section. It is an ancient technique, going back millennia (see Kelly, 2017). I mention it in a box because Maguire *et al*. (2003 p.90) found that, "Superior memory was not driven by exceptional intellectual ability or structural brain differences. Rather, we found that superior memorizers used a spatial learning strategy." For more information, search for **"memory palace"** or **"method of loci"**.

Categorization

Sometimes there is already a structure that exists in the ideas that you want to remember, but the textbook does not make the structure clear. If you can find some model or categorization hidden in the ideas, you can make it easier to remember because it reduces the load on your short-term memory.

For example, let's say someone gives you this shopping list to remember: carrots, milk, lettuce, apples, yogurt, banana, cucumber, cheese, and peaches. You will probably forget most of the nine products by the time you get to the shop. But you can reduce the number of items you have to remember to three, and this is much easier. In the list, there are three categories: vegetables (carrots, lettuce and cucumbers); fruit (apples, bananas, peaches); and dairy products (milk, yogurt and cheese). Because of this categorization, there are fewer items to remember at first (only three categories instead of nine products). And now the products are also easier to remember, because each category reminds you of the products.

Here is another example. Suppose, in a textbook, the authors list six kinds of strategies a business can follow:

- Innovation: the development of new products;
- Concentration: focusing on a single product to market;
- Market development: existing markets or products are developed more intensively;
- Rationalization: reducing costs by getting rid of products, assets or labor;
- Divestiture: selling off parts of the business;
- Liquidation: closing down the business

Table 4.6 Categorize to Remember

Category	Strategy
Strategies that EXPAND business activity	Innovation: the development of new products; Market development: existing markets or products are developed more intensively
Strategies that REDUCE the scope of the business	Concentration: focusing on a single product to market; Rationalization: reducing costs by getting rid of products, assets or labor
Strategies that ELIMINATE business activity	Divestiture: selling off parts of the business; Liquidation: closing down the business

From this, I could see three categories, with two strategies each. It is much easier to remember three categories or two items than having to remember six items. What I saw is shown in Table 4.6.

Not all ideas belong to a single inherent category, so I don't use categorization as often as the other techniques. I prefer visual techniques, which is what I turn to next.

Images

Regardless of memory style, most people have a near-perfect memory for images that can last for days or weeks. This has been known since Standing (1973) and was recently confirmed by Konkle *et al.* (2010). Julian de Freitas (2012) explains that they showed:

> participants a stream of three thousand images… Then, participants were shown two hundred pairs of images – an old one they had seen in the first task, and a completely new one – and asked to indicate which was the old one. Participants were remarkably accurate at spotting differences between the new and old images – 96 percent… despite needing to remember nearly 3,000 images, they still performed almost perfectly.

Visual memory techniques exploit this natural and effortless ability by connecting facts to images.

I will demonstrate one image-based mnemonic device that has never failed me since my father taught me this technique more than thirty years ago. It involves connecting ideas with related images, and then creating a visual unit from those images.

Suppose I want to remember the elements of a firm's internal control system, which are as follows.

- Control environment: how the business is managed e.g. philosophy, style and values;
- Information and communications: systems and reports related to operations and compliance- with internal and external stakeholders;
- Control activities: policies and procedures to execute an action plan and address risks;
- Risk assessment: identification and analysis of risk that could prevent the achievement of objectives;
- Monitoring: checking the performance of the system through monitoring activities or audits.

As usual, the first step is to extract descriptive keywords. There is no right or wrong here – the only thing that matters is whether the keyword reminds you of what the point is about. For me, the result is shown in Table 4.7.

Secondly, I attach a symbol, or a *very simple* picture that is easy to draw, to each keyword. This step does not require drawing ability, simply the ability to come up with a symbol or a really simple picture that makes sense to you. You need to be able to draw it fast and nobody needs to see your pictures, so do not create an artwork here. Again, there is no right or wrong here – the image should simply be able to remind you what the keyword is, or what the point is about. For me, the result was as shown in Table 4.8.

Thirdly, integrate all the symbols or images into one image. To integrate the images, they have to be made part of each other and combined so they form one image. Simply putting a border around the images or

Table 4.7 Extracting Keywords

Elements to remember	Keywords
Control environment: how the business is managed e.g. philosophy, style and values	Control
Information and communications: systems and reports related to operations and compliance; with internal and external stakeholders	Information
Control activities: policies and procedures to execute an action plan and address risks	Policies and procedures
Risk assessment: identification and analysis of risk that could prevent the achievement of objectives	Risk
Monitoring: checking the performance of the system through monitoring activities or audits	Monitoring

Table 4.8 Attach Images to Keywords

Control	Information	Policies and procedures	Risk	Monitoring
	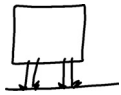			
An *angry face* of a person trying to control me	A *billboard* on which people communicate information	A *thick rulebook* in which you find lots of procedures	An *explosion* that destroys everything is a big risk	You can monitor what is happening better if you have good *glasses*

connecting the images with arrows will not work, because the images are still separate, and many images are more difficult to remember than a single image. The images have to be *integrated with each other into a single image* in a way that you can explain to yourself. It is not necessary for the integrated image to make sense or even look like a nice picture because the mind best remembers pictures when they are *strange*. The important thing is that the images are all integrated into one image. For me, the result was Figure 4.4.

When I have to recall the ideas, I simply draw the integrated image, and the different parts of the image remind me of the ideas I have to remember. There is the angry face (control) with glasses on it (monitoring). Because he is angry, an explosion (risk) comes out of his head. His body is the billboard (information), but this is a billboard that has a book (policies and procedures) as the screen of the billboard. There is no other image like this and it is a bit strange, but what is important is that I can explain how I created it. Combine this strangeness with humans' strong visual memory, and you have a powerful memory technique. I used to have 20-30 images per subject that I studied in school and university, and I never forgot any of them in an exam situation.

Rhymes

Rhyming (often combined with rhythm) is an age-old mnemonic device that was used to pass ideas from generation to generation in the absence of writing. With the rhyming technique, you create a structure containing words that will remind you of what you want to remember. Rhymes

Figure 4.4 My Integrated Image

are easier to remember, because if you forget something, you only need to find some words that rhyme with the other words, and that will probably prompt you to remember. The rhyme does not have to be great, it just needs to rhyme and make sense. Again, it starts by extracting keywords, and then putting them into a rhyme that uses these keywords. I created the following rhyme for me to remember the five components of the market environment:

For *customers* I care
But *competitors* I scare
And *suppliers* they pair
With *intermediaries* to get products there
While *society* is everywhere

The rhyme makes sense because a business should care for customers and its competitors would be scared of it. Also, suppliers help the business, often by working together with intermediaries. And society is all around the business. The rhyme is silly, but effective. If you are not good at finding words that rhyme, there are many free rhyming dictionaries on the web that can help you to find rhyming words.

If you can add some rhythm to a rhyme (like a rap song), it becomes even easier to remember, because the rhythm makes the memory more automatic. Another possibility is to put the words of a rhyme into a song that the students know very well – I know of many teachers who use this technique very effectively.

Stories

With this device, you create stories that serve as a framework that you use to link ideas together in an interesting way. Humans' memory for stories is almost unlimited. You can hear a story once and be able to retell it to someone days later. And we remember thousands of stories without effort because stories naturally make connections. This technique uses this ability.

Start by extracting keywords, as always. Let's take another example, a rather boring and intimidating list of the characteristics of the business environment, with the keyword already extracted, as shown in Table 4.9.

Next, I create a story in which I use every keyword. My story eventually was:

> I used to work in an office with a guy who had three heads. We were never sure what he was saying, because when he spoke all heads would talk and interrupt each other and sometimes they would even say completely different things. For a short time, we could bear it, but, after a while, we had to do something. We put him in a separate closed office where his talking heads could talk as much as they liked for as long as he wanted.

Table 4.9 Extract Keywords

Characteristics of the business environment	Keyword
It is a combination of many factors at different levels that need to be studied separately and together.	Multi-faceted
The components are interdependent. When one changes, the others are affected.	Inter-dependent
The environment is uncertain and can change very fast.	Uncertain
To cope with changes, the business needs to change too, either proactively or reactively.	Action
Businesses are affected differently by the environment – even departments in the same business will experience this.	Differences
Changes in the environment have both a short-term and long-term effect on the business.	Short- and long-term
Environmental factors have an ongoing effect that is relentless and harsh.	Ongoing

If you cannot see how this helps, here is the story again, but this time notice how different parts of the story remind me of different keywords:

> I used to work in an office with a guy who had three heads (multifaceted). We were never sure what he was saying (uncertainty), because when he spoke all heads would talk and interrupt each other (interdependence) and sometimes they would even say completely different things (differences). For a short time (short-term), we could bear it, but, after a while (long-term), we had to do something (action). We put him in a separate closed office where his talking heads could talk as much as they liked for as long as he wanted (ongoing).

Sometimes there are cause-and-effect relationships in a list of facts which you can use to create a story, but, if not, creating a vivid and strange story works just as well.

Associations

Using mnemonic devices that involve creating associations between the ideas is popular among memory champions and those who specialize in teaching memory improvement. Such devices are based on the principle that we remember things that are vivid (colorful, noisy, full of movement, emotional) and very different from the ordinary. When using associations to remember, you use your imagination to create vivid and absurd

connections between ideas. If the connection is vivid and strange enough, one idea will automatically make you think of the idea to which it is connected, and that idea will make you think of the next one, and so on.

The most common device here is the 'link' method. To illustrate it, let us try to remember the four types of unemployment. The keywords are simple to extract this time because they are simply the types of unemployment: structural, cyclical, seasonal and frictional unemployment. Next, I come up with a mental image that I associate with each keyword. The result for me was Table 4.10, but every person will have different images.

Finally, I associate one word with the next word in *as vivid and strange a way as possible*. This is the key. There has to be color, movement, sound and funny things happening in each association – if not, you will forget it. My associations are shown in Table 4.11.

Table 4.10 Keywords Attached to Mental Images

Keyword	Mental image
Structural	The Eiffel Tower
Cyclical	A bicycle
Seasonal	Sun and clouds
Frictional	Rubbing things together

Table 4.11 Linking the Mental Images

Keyword	Mental image
Structural	The Eiffel Tower
Link between Eiffel Tower and bicycle	*I imagine that the Eiffel Tower has hands and feet and it gets on to a giant bicycle and rides it, but all the time it falls off and so crashes into the buildings of Paris, sending people screaming and running away.*
Cyclical	A bicycle
Link between bicycle and sun and clouds	*I imagine a cloud riding on a bicycle, but this bicycle's wheels are two suns; and so, as the cloud rides the bicycle through the streets, people collapse from heat and shield their eyes as the bicycle passes them.*
Seasonal	Sun and clouds
Link between sun and clouds and rubbing	*I imagine the sun grabbing two unruly clouds who are running across the sky and then hitting and rubbing them together to discipline them, and this makes the clouds cry loudly and their tears fall like rain.*
Frictional	Rubbing things together

Now, if I want to remember the four types, I just need a way to link unemployment to the first keyword (structure). I could either think about how unemployed people need to build structures to have jobs (but this is not vivid enough, so I will probably forget it). Much better would be to imagine lots of unemployed structures or buildings, which, instead of standing upright, are lying down on the streets and snoring because they are unemployed. So, when I think of unemployment, I think of sleeping buildings (structural unemployment), and when I then think of structures I think of the Eiffel Tower on a bicycle. The bicycle makes me think of cyclical unemployment, which reminds me of the cloud on a bicycle with sun-wheels and this makes me think of seasons and seasonal unemployment… and so on. Through vivid associations, one idea makes me think of the next idea. Because it draws so much on imagination, this method can be quite engaging, and it was the secret behind Solomon Shereshevsky's astounding memory (Johnson, 2017).

BOX 4.3 MEMORY PEGS

A popular variation on the link method is a mnemonic device where the images, and their sequence, are pre-determined. It is extensively used in books like *Master Your Memory* by Tony Buzan (1998). To learn it and see some great examples, search for **"memory pegs"** or **"memory peg system"**.

Memory through Repetition

One of my teachers used to say that, "Repetition is the mother of learning." He implied that memorization is learning and that the most effective way to remember is to repeat. As explained in this chapter, this is a simplistic view is contradicted by the research. Millions of students who spent nights rote-learning and repeating lists of facts, and then immediately forgetting them, can attest to the ineffectiveness of mere repetition and the misery it causes.

Yet, there is a kind of repetition that is quite effective, and surprisingly effortless. It is called 'spaced repetition' and also requires a degree of forgetting in order to work well. Since it is related to the so-called 'testing effect', it will only be discussed in Chapter 7, where I look at the principles of assessment.

This completes a rapid tour of some highly effective memory techniques. Lastly, I just want to identify some of the things of which you need to be careful.

Take Care!

When using the memory techniques above with information that will have to study more than one time, you should make notes about the technique you used and how you used it. That way you can very quickly revise the work with less effort. For example, if you used the image technique, do like I did above; put the relevant image next to each idea and note down a short explanation of the reason for the image, and then draw the integrated image below the ideas with a short explanation. This way, you can quickly remember the ideas even if it was months or years since the last time that you looked at the image.

There are also some mnemonic devices that I would advise you to avoid or use as little as possible, because they make very weak and lifeless connections. These are techniques that involve text, and because the connections involved are not vivid or strange enough, you will often forget them under pressure. For example, if you want to remember the five components of the market environment (customers, competitors, intermediaries, suppliers, society), you can take the first letter of each word and get 'CCISS'. This looks simple and fast to do, but it is easy to forget. If you have more than two or three of these letter-words, you start forgetting them. And worse, with too many you will also start forgetting what the letters stand for. The same applies to using the letters to make sentences like this one: "Cute Clerks Inhaled Smoke Sadly". The letter-sentence suffers from the same problem – if you have more than two or three of them, you will forget what some letters stand for because humans don't form natural connections to text.

MEMORY AS DISTRIBUTED COGNITION

Except for the making of connections, the other feature of mnemonic devices is that they do not limit memory only to the brain. These devices spread the cognitive work involved in memory to physical objects, as in analogical methods, or imagined phenomena (like pictures or stories). This is also known as 'distributed cognition' and explains why mnemonic devices make remembering so effortless once they are mastered (more about this in Chapter 9.)

In a fascinating book, Lynne Kelly (2017) explains how pre-literate cultures used distributed memory to achieve extraordinary memory feats. Just one example she mentions is that of the Matsés people in South America. When their memorized knowledge only about medicinal plants was documented, the result was a 500-page document! Every one of the mnemonic devices in this chapter (and many more) has been used

extensively in pre-literate cultures, forming the foundation of a densely interconnected and multi-modal knowledge of plants, animals, climate, genealogy, and more.

What her book shows wonderfully is how distributed memory, when it serves understanding, can be used to build flexible knowledge structures that are easy to complexify. It also accelerates learning in any area because it makes it easier for the next generation to build on the previous generation's knowledge. Richard Feynman would have agreed. He (1963) was vehemently opposed to memorization without appropriate connections, but he acknowledged that memory plays a critical 'time-binding' function for humans (Feynman, 1999, p.184).

CONCLUSION

There are few people that have anything close to a photographic memory. Many do have astounding memories, but in the majority of these cases, they developed their memories using mnemonic devices. This chapter showed that these devices can be demonstrated and taught to students. Doing so will vastly reduce the effort required for memorization. However, this is not an excuse to expect students to memorize the many useless categorizations and bullet point lists in textbooks. Memory is tightly coupled with understanding, and should be subservient to it. Even in the rare cases when memorization *is* necessary, it should be like tape that holds something in place until the glue of understanding can make it stick.

BOX 4.4 STICKY TEACHING

There is another kind of way of making ideas both memorable and interesting that straddles the techniques in this and the previous chapters. This is best explained by Chip and Dan Heath (2007) in their very useful book, *Made to Stick*, which I can highly recommend to teachers. Also see their free resources, which you can find by searching **Heath AND "teaching that sticks"**; and their methodology called "anchor and twist", which you can find by searching for **Heath AND "anchor and twist"**.

Changing one's view of memory, from store-and-recall to connect-and-forget, has a dramatic impact on one's teaching, and especially assessment (see Chapters 7 and 8). Memorization should be tested only for those ideas that will promote future understanding. If someone with a photographic memory can pass your assessment, critically consider your view of memory.

EXERCISES

1. Analyze a chapter in a textbook for which you already identified the critical ideas.
   a) Are there any foundational ideas? What are they, and how can you help students to remember them? To which critical ideas will you connect them? (If this is not clear, reconsider your answer.)
   b) Are there any key details that need to be memorized before they are understood? If any, what are they, and how can you help students to remember them? To which critical ideas will you connect them? (If this is not clear, reconsider your answer)
   c) Are there any over-complex ideas that need to be treated like facts for now? If any, what are they, and how can you help students to remember them? To which critical ideas will you connect them? (If this is not clear, reconsider your answer.)
   d) Which ideas are worth forgetting? This may be a substantial part of the chapter.
   e) If your assessment should include some memorization, which memorized facts are the most appropriate to assess?
2. Identify at least three lists of ideas or terminologies that students need to memorize, but find difficult to remember.
   a) Apply a different mnemonic device to each one.
   b) How would you teach and demonstrate each device to students?
3. Consult Box 4.1. Which other mnemonic techniques not explained in this chapter do you think will be effective in your course? Look specifically for techniques that can help students remember the definitions of key terms in your subject.

REFERENCES

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Bartlett, F.C. 1920. Some experiments on the reproduction of folk-stories. *Folklore*, 31(1):30–47.

Borges, J.L. 1964. Funes the Memorious. In: *Labyrinths: Selected Stories and Other Writings*. Translated by J.E. Irby. D.A. Yates & J.E. Irby (eds). New York: New Directions.

Buzan, T. 1998. *Master Your Memory*. London: BBC.

Caine, R.N. & Caine, G. 1991. *Making Connections: Teaching and the Human Brain*. Alexandria: Association for Supervision and Curriculum.

Carmichael, L.C., Hogan, H.P. & Walter, A.A. (1932). An experimental study of the effect of language on the reproduction of visually perceived form. *Journal of Experimental Psychology*, 15:73–86.

Chase, W.G. & Simon, H.A. 1973. Perception in chess. *Cognitive Psychology*, 4:55–81.

Davis, R.L. & Zhong, Y. 2017. The biology of forgetting – a perspective. *Neuron*, 95:490–503.

De Freitas, J. 2012. Why is memory so good and so bad? Explaining the memory paradox. *Scientific American*, 29 May [available at http://www.scientificamerican.com/article/why-memory-so-good-bad, accessed 18 July, 2018].

De Groot, A.D. 1965. *Thought and Choice in Chess*. The Hague: Mouton.

Fernández, J. 2015. What are the benefits of memory distortion? *Consciousness and Cognition*, 33:536–547.

Feynman, R.P. 1963. The Problem of Teaching Physics in Latin America. First Inter-American Conference on Physics Education, Rio de Janeiro [available at http://calteches.library.caltech.edu/46/2/LatinAmerica.htm, accessed 24 July, 2018].

Feynman, R.P. 1999. *The Pleasure of Finding Things Out: The Best Short Works of Richard P. Feynman*. Cambridge: Helix Books.

Heath, C. & Heath. D. 2007. *Made to Stick*. Random House.

Hemmer, P. & Steyvers, M. 2009. A Bayesian account of reconstructive memory. *Topics in Cognitive Science*, 1:189–202.

James, W. 2007. *The Principles of Psychology*, Volume 1. New York: Cosimo Books.

Johnson, R. 2017. The mystery of S., the man with an impossible memory. *New Yorker*, 12 August 2017.

Kelly, L. 2017. *The Memory Code: The Secrets of Stonehenge, Easter Island and Other Ancient Monuments*. New York: Pegasus.

Konkle, T., Brady, T.F., Alvarez, G.A. & Aude, O. 2010. Scene memory is more detailed than you think: The role of categories in visual long-term memory. *Psychological Science*, 21(11):1551–1556.

Maguire, E.A, Valentine, E.R., Wilding, J.M. and Kapur, N. 2003. Routes to remembering: The brains behind superior memory. *Nature Neuroscience*, 6(1):90–95.

Nørby, S. 2015. Why forget? On the adaptive value of memory loss. *Perspectives on Psychological Science*, 10(5):551–578.

O'Loughlin, I. 2017. Learning without storing: Wittgenstein's cognitive science of learning and memory. In: M. A. Peters and J. Stickney (eds.). *A Companion*

to Wittgenstein on Education: Pedagogical Investigations, pp. 601–614. Singapore: Springer.

Richards, B.A. & Frankland, P.W. 2017. The persistence and transience of memory. *Neuron*, 94:1071–1084.

Schacter, D.L. 2001. *The Seven Sins of Memory: How the Mind Forgets and Remembers*. New York: Houghton Mifflin.

Sloman, S. & Fernbach, P. 2017. *The Knowledge Illusion: Why We Never Think Alone*. New York: Riverhead Books.

Standing, L. 1973. Learning 10000 pictures. *Quarterly Journal of Experimental Psychology*, 25(2):207–222.

University College Cork n.d. George Boole: Computer Science [available at https://georgeboole200.ucc.ie/boole/legacy/computerscience, accessed 22 July, 2018].

Weber, B. 2009. Kim Peek, inspiration for 'Rain Man', dies at 58. *New York Times*, 26 December 2009.

Whitehead, A.N. 1948. *An Introduction to Mathematics*. Oxford University Press.

Willingham, D. 2009. *Why Don't Students Like School?* San Francisco: Jossey-Bass.

Chapter Five

How to Make Boring and Complex Ideas Interesting

Almost fifty years ago, Elizabeth Loftus and John Palmer showed a movie to some students and changed the way we understand memory. The movie was just a few seconds long and showed a car accident. After the movie, the students, all of whom had seen the same movie at the same time, were divided into groups and asked how fast the cars were going. The differences between their answers were quite big. For example, one group estimated it was 51 km/h, while another thought it was 66 km/h (Loftus and Palmer, 1974). How could people remember the same event so differently?

The answer supports what the previous chapter said about memory: that we recreate memories from the associations we make at the moment the memory is formed and at the moment that we remember. Since the students had all seen the same movie, the associations they made during the movie could not fully explain the big difference. So, what was it?

There were also the associations made at the moment of remembering. This occurred when students were asked about the speed of the moving car. Maybe they were asked the same question but in different words?

That is indeed what happened. The researchers changed just one or two words when they posed the question to every group. But is it really possible to manipulate a recent memory so dramatically simply by changing one word? Do you think that I would be able to change your memory of something as easy as that?

Well, yes, it seems so. One group was asked, "How fast were the cars moving when they hit each other?" For another group, the word 'hit' was replaced with 'collided with', for another, it was replaced by 'smashed into', for another, it was 'contacted', and, for another, 'bumped into'. Think for a moment which group guessed the car was going at 51 km/h and which thought it was 66 km/h.

The groups who heard the question with the words 'collided with', 'bumped into' and 'smashed into' estimated faster speeds, while those who heard 'hit' and 'contacted' estimated slower speeds. The highest

estimate came from 'smashed' and the lowest from 'contacted'. Clearly, the associations students made with the words influenced how they remembered the event; sometimes even remembering things that were not there.

After that, we could no longer think the same way about eyewitness testimony and the way lawyers ask questions in court. Elizabeth Loftus, in fact, became an expert witness in criminal cases. Her first case was one in which a woman had killed her abusive boyfriend, where witnesses could not agree on how much time had elapsed from the time she picked up the gun until she fired it. The difference was crucial because it would determine if it was murder or self-defense, but witnesses could not agree on whether it was seconds or several minutes. Loftus testified to the unreliability of memory, and the woman was acquitted. Her work has led to more acquittals, making her many enemies who believed that she helped guilty people go free. But her response is, "I haven't had a situation where someone was acquitted because of my testimony and then went on to commit some awful crime" (Costandi, 2013, p.269).

The possibility that innocent people may be imprisoned or receive the death penalty weighs more heavily on her, so she has been campaigning for legal reforms for decades. As a result, in some courts in the USA, jurors now have to be informed of the unreliability of eyewitness testimony. It makes you wonder how many of your own memories have been subtly manipulated by others perhaps even whether or not you can control your students' minds simply by using different words.

Hopefully, some parts of this story were interesting to most of you. This may be because, having read the previous chapter on memory, it is relevant to you. In fact, think back to Chapter 4 and identify any other things about memory that struck you as interesting. My guess is that most of you will choose the idea that forgetting is critical to a good memory or perhaps the statement that there is no such thing as photographic memory. But before continuing, consider these ideas, and ask yourself: "What made them interesting?"

The story of Elizabeth Loftus, and whatever other ideas you identified, have a couple of things in common that make them interesting. In this chapter, we will identify them by analyzing what makes ideas interesting and applying it to teaching.

WHY SHOULD IDEAS EVEN BE INTERESTING?

On the surface, it seems that we want to make ideas interesting to ensure that students pay attention. Ideas to which nobody pays attention literally do not exist. As professors, we do not want this to happen to us

because it means not only that our students will forget the ideas we taught, but also that they will forget us. But this is a poor reason to make things interesting.

As with memory, interest serves understanding by enhancing memory and motivation. If something is interesting, we think about it. The act of thinking creates new connections and strengthens existing connections, and it is only the ideas we think about that we remember (Willingham, 2009). In addition, interest often generates an emotive response, and, if thinking is accompanied by emotion, it further enhances the memory of the ideas involved (Tyng *et al*., 2017).

Of course, by inducing thinking, interest also makes understanding possible. Understanding is not achieved by merely storing more facts; rather it involves connecting facts and ideas, thereby changing the knowledge network structure (also called 'conceptual change'). The problem is that conceptual change is always disruptive. All new ideas need to be connected to existing ideas, which include ones they acquired during the course but also preconceptions that students brought with them to the course. Whenever we introduce new ideas, it disrupts the structure and organization of these existing ideas (DiSessa, 2014).

Disruption is often resisted because it takes energy and suggests that we are at least partly wrong. But interest can remove much of the resistance and make us more willing to incorporate new ideas into our knowledge structure. For this reason, a person who is interested in a topic will not wait for a professor to introduce disruptions, but will independently search for new patterns to make sense of incoming information. In sum, then, when something is interesting, it is more likely that students will want to broaden their understanding.

INTERESTINGNESS DESCRIBED IN THREE 'F-WORDS'

It is difficult to completely define what makes an idea interesting because it is multi-dimensional. The three dimensions that I found to be important in teaching can be captured in three words: 'fascination', 'fun', and 'fumbling'.

The most obvious definition of 'interesting' is that of something that holds attention. This is best captured by the idea of 'fascination' that Sally Hogshead (2010) writes about. She explains that word comes from the Latin word '*fascinare*', which means 'to bewitch'. When you are fascinated by something, your attention locks onto it, and you are unable to move onto to something else. If we can make something relevant and attractive to the students, they are likely to find it fascinating.

Fascination is an important first step in generating interest because our students will not be motivated to learn what we explain to them

unless they are motivated to pay attention. But, by itself, fascination is of limited use in learning and teaching situations. We have all watched television documentaries or attended flashy presentations that captured our attention, but, one day later, we remember little. This is because simply paying attention does not translate into learning unless we actively use that attention – unless we think.

This is where the game designer, Raph Koster (2005), made a great contribution with his 'theory of fun', which adds the second dimension. He defines 'fun' as that which we experience when we find and master new patterns of thought. This is exactly what happens in games. To master a game, we have to figure out certain patterns, and once we have done this, the game is no longer as much fun. However, if we cannot see any patterns, the game is too difficult and definitely not fun. With a new game, we initially struggle as we try to figure out the game and how to play it. Over time, we figure out the patterns and eventually master the game. Games tend to be the most fun as long as they are pushing us closer to the edge of our capabilities, while simultaneously giving us a sense that mastery is a real possibility.

Like games, to gain an understanding of a subject or topic, we need to find patterns. If we are fascinated with the topic, we will have fun looking for patterns. Finding these patterns helps us to see connections and chunks, compress the information and try out the ideas to see how well we can generate new information. Once we find all the patterns and become competent in using them, we have mastered the subject and it becomes easy.

However, there is a problem. Once we have mastered a game, we keep on playing simply to experience the pleasure of winning, even if we are no longer learning new patterns. Similarly, in a subject, the patterns become sedimented once they are automatic and part of our long-term memory. At this point, we no longer need to think much and no new learning occurs.

It is difficult to switch to learning a new game or a new set of ideas, so our brains resist it, preferring the easy path of just repeating old patterns. As Willingham (2009) argues, the human brain prefers following old patterns over thinking, even though we enjoy thinking once we get into it. In a learning situation, something has to force us to be open to a new idea by upsetting our old ideas and making us fumble. This means that old ideas must be disrupted so that mental cracks can appear through which new ideas can slip in.

Murray Davis (1971) elaborates on the third dimension of interestingness as disruption. He recognized that ideas are interesting when they are relevant and not too obvious to your *audience* (even if they are

obvious to you). Ideas that an audience can use are relevant to them. But if the ideas are relevant but completely obvious, they are boring. If an idea disrupts us too much, we call those ideas 'absurd'. Obvious ideas are patterns that we already figured out (boring) and absurd ideas are ideas where the patterns do not make sense to us (also boring). An idea is interesting when it is useful and lies somewhere between being obvious and absurd. The 'sweet spot' of interestingness is shown in Figure 5.1.

In a dynamic learning situation, we have to push back against the obvious and disrupt students' understanding by upsetting what the students thought they knew. Davis makes the point that something is interesting if it denies or contradicts at least one thing that the audience assumes is true. Note the use of the word 'audience' – students will pay attention again if an idea disrupts something that *they* knew. If it disrupts your knowledge (or that of other experts), it will not necessarily be interesting to students, so it is important to know your audience before trying to disrupt them.

In summary, interestingness is a process that arises from the interaction of three things: **fascination** (getting attention), **fun** (using attention), and **fumbling** (disrupting attention). The three processes are interdependent, so if any one of them is deficient, it will harm the effects of the other two. This interaction can be expressed by multiplication:

$$\text{Interestingness} = \textbf{F}\text{un} \times \textbf{F}\text{ascination} \times \textbf{F}\text{umbling} = \text{F}^3$$

The three F-words give us the F^3 (F-cubed) approach to giving ideas an interesting quality. As Figure 5.2 shows, all three support each other and lead to an ever-upward spiral of learning if used together.

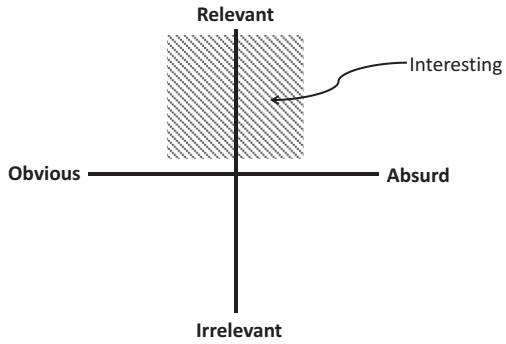

Figure 5.1 The 'Sweet Spot' of Interestingness

Source: Derived from Davis (1971)

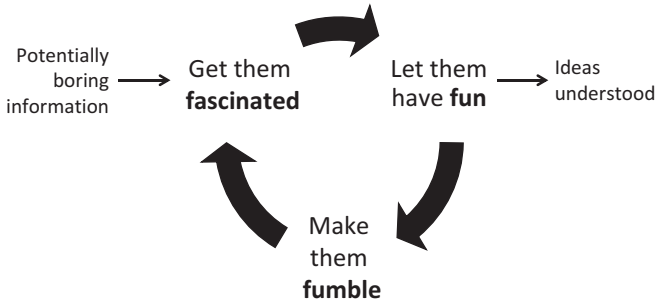

Figure 5.2 The Process of Interestingness

Figure 5.3 further shows that generating interest is a continuous process. At first, fascination draws the student in to pay attention to an idea, but this does not require a large investment of cognitive energy. The student then engages with the idea, which requires a lot more cognitive energy. As the student masters the idea, she uses less cognitive energy, and a disruption is needed if the knowledge structure is to be further complexified and refined. To deal with such a disruption takes more cognitive energy, but it prepares the student to jump to a higher level of understanding.

Returning to the story at the start of this chapter, it is possible to explain why (or not) it may have been interesting to you. Having just finished a chapter about memory, you will hopefully have recognized

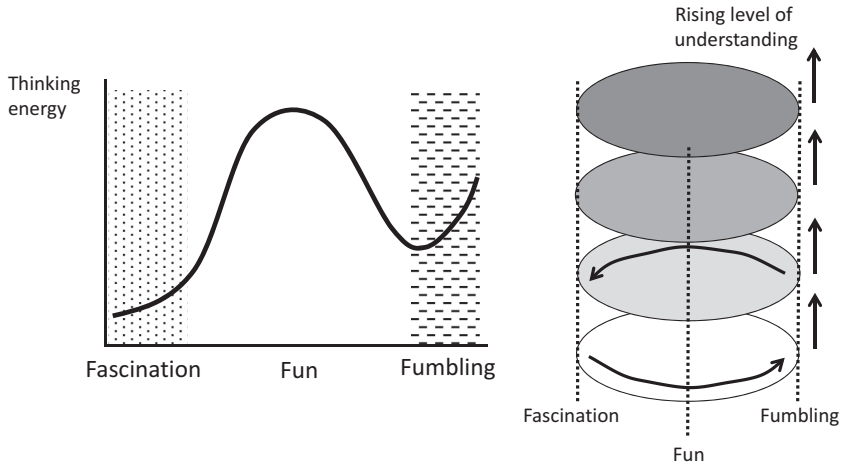

Figure 5.3 Interestingness as a Continuous Process

how relevant the topic is to you and your practice as a teacher. This should have generated at least some degree of fascination. If you were ever part of a court case or found yourself controlled by someone else, then even more so. The study was presented as a puzzle – both in terms of the cause of the different views and in terms of how it may be useful – and this required finding a pattern. I asked some rhetorical questions to encourage you to look for patterns yourself. If the pattern I revealed was not obvious or known to you, then reading the story would have been fun. If you understood the previous chapter, the ideas would not have surprised you so much that you would think of them as absurd. Lastly, I wanted to make you fumble a bit. Hopefully, it upset your ideas about memory by illustrating how easy it is to manipulate it. If I chose to continue with how one can go about fixing this problem, or how you can use this to your advantage, you would have been more open to it than before.

USING THE THREE F-WORDS TO GENERATE INTEREST

Taking each F-word in turn, let us see how to use them individually before considering how they work together.

Fascination with Patterns

When I first started job hunting after university, I briefly considered a profession that specializes in fascination: sales. During this period, and later, I was exposed to some really good salesmen. They taught me something simple about fascination – that if you want to get people's attention when selling something, you have to "sell the sizzle, not the steak".

When someone wants you to eat in their steakhouse, they won't entice you by showing you the raw steak and explaining to you how nutritious it is, or by telling you from what kind of quality cow it is, or by showing you how juicy it is. That is more likely to put you off, even if you are not a vegetarian. Rather, this person will simply throw the steak on the griddle, so that you hear the sizzling sound and have the smell of grilled meat reach your nose as you suddenly realize how hungry you are.

If you want to sell something, do not sell the thing itself. It is the benefit of the thing that gets attention, or even better, the benefits of the benefit. When we teach, we are also selling our ideas in exchange for students 'paying' attention. We need this 'payment' for the learning transaction to begin, and by offering opportunities for fascination, we make students willing to pay attention.

Fascination does not happen by talking about the steak or the idea itself, no matter how fascinating it is to us. It happens when we are able to connect the idea to our audience in a way that *they* can see that it is relevant to *them*. This means that you need to get to know your audience, what their goals are and what they care about.

For example, if I teach someone about investing in the stock market, I will not fascinate people by telling them what shares are, how to buy and sell them, about financial statement analysis or the role of brokers – they will lose interest. I am more likely to succeed if I talk about the benefits of investing, such as increasing one's wealth. But I am even more likely to succeed if I connect to something deeper inside of them, by showing them the benefit of the benefit (or the meta-benefit).

The benefit is gaining wealth, but what is the benefit of gaining more wealth? Maybe, if you gain wealth, you can increase your status or help your friends and family. Maybe it simply means you can finally be independent, doing what you want when you want and never having to ask anyone for anything. The deeper the idea connects to something that matters to a person, the more fascinating it will be to them.

Sally Hogshead (2010) explains that there are seven meta-benefits, which she calls 'triggers' of fascination and they are summarized in Table 5.1.

You pull a trigger by making students feel the emotion described by that trigger and showing how the idea you want to explain can help to relieve or intensify the emotion. Knowing your audience's values or unquestioned assumptions about life will help you to select the best trigger. If you are addressing a group of people with different values, it means that you often need to pull more than one trigger.

For example, when I used to teach Economics to privileged students in South Africa and I got to the topic of poverty, I could pull any one of the above seven triggers and get them fascinated with varying degrees of

Table 5.1 Seven Triggers of Fascination

Trigger	Description
Pleasure	Good feelings, sensual experiences, anticipation of pleasure
Mystique	Puzzling, unanswered questions, being part of a secret
Alarm	Fear, loss of possibilities, respond now
Prestige	Achievement, getting ahead of others, respect, admiration
Power	Control, command over others or over the environment
Vice	Rebel against rules, being different
Trust	Comfort, certainty, predictability

Source: Derived from Hogshead (2010)

success. I summarize them below (rather crudely) in order of effectiveness, as I perceived it, for this particular audience.

- *Alarm*: Poverty is rising and has led to revolutions worldwide where wealthy people have become targets and lost their possessions and position. This is likely to happen here unless we start doing something, and for that, we need to understand the dynamics of poverty.
- *Trust*: We don't need to worry about poverty as long as we know how to address it.
- *Power*: The best way to gain power is to appeal to the majority, and the majority in the country is poor. If you understand poverty better, you have a better chance of gaining political power.
- *Prestige*: People admire those who don't just look out for themselves, but who are seen to care about issues of human suffering, as many celebrities appear to. To know how to do this, you need to understand how the poor live and what they really need.
- *Vice*: Your parents want you to live the same life as theirs by getting you to think as they do and remain cocooned in your world. But you need to learn to think for yourself too – there is a bigger world out there where people are not rich, and where you can live a more meaningful life by getting involved with real issues like poverty alleviation.
- *Mystique*: Any number of puzzling questions from the news at the time relating to poverty was useful here. For example, why has the poverty rate risen even though the country became increasingly democratic? Or why did poverty rates fall in some South Asian countries that clamped down on democratic rights?
- *Pleasure*: Though I never pulled this trigger, I could probably have done it by means of a pleasant field trip to a really poor neighborhood and have them experience poverty in a limited way.

If my audience had been a different one, I would have pulled different triggers in a different way. Not all of the triggers I pulled were socially desirable (the power trigger perhaps), but my purpose was to attract their attention so that I could get them thinking beyond the triggers. The other two Fs helped with that.

Looking at the story at the beginning of the chapter, you may see how I tried to pull a few triggers. Initially, by not revealing the reason for the different answers immediately, I used the mystique trigger. By making you think about the possibility of this knowledge

being abused, I also pulled the alarm trigger; and then, by mentioning that you might employ this knowledge to your own benefit, I pulled the power trigger.

Fun Figuring Out Patterns

Ideas are patterns that need to be figured out, but fortunately, the human brain is naturally good at seeking and detecting patterns (Caine & Caine, 1991). The experience of fun is evolution's way to get us to enjoy making sense of apparently random incoming information and compress it into something more useful.

Fascination draws us into this process and motivates us to engage in figuring out patterns. Part of the fun is figuring out the idea itself, and another part is figuring out how to use it in order to make sense of a reality that seems to be random, noise or a bunch of disconnected facts. People find history less interesting if it is "just one damn thing after another" and more interesting if they can see a story, a repeating pattern or a conspiracy theory that holds it all together. Similarly, once you understand a theory, such as elite theory in political science and sociology, the news is no longer just a series of random events but will start to cohere around an organizing principle.

Each idea is a new pattern that allows us to see the world differently and organize our understanding differently. For example, before I learned about the market mechanism (demand and supply) in economics, I was aware of price changes. But, until then, it seemed to me to be just a lot of things happening with no clear meaning. Learning about the market mechanism allowed me to see patterns in these changes, understand what caused them and even make some accurate predictions. With this knowledge, I could even understand things with which I had no direct experience, like the price changes of crude oil, gold and currencies.

However, as with games, all of this is only fun for those engaged in the process, not for passive observers. This implies that one needs to encourage participation. This can be done inside a lecture (e.g. through discussions, quizzes) or outside (e.g. performing authentic tasks). It can be done in groups (e.g. through various kinds of cooperative learning), by talking, individual writing or even just by thinking (e.g. through using advance organizers or simply pausing for a few seconds after asking a question). It is not surprising then that Bain (2004) found that the best professors do not only rely solely on lectures, but encourage active learning as well. Box 5.1 contains some information about active learning methods.

BOX 5.1 ACTIVE LEARNING

To find more information about active learning, simply enter in the search box **"active learning methods" AND classroom**, or, to be more specific **"active learning" AND classroom** or even **"active learning" AND "college classroom"**. An extreme method is the flipped classroom, in which there is no lecturing during lectures, only activities. For information, just search for **"flipped classroom"**.

Since ideas change the way we look at and experience the world, they are not always easy to figure out, which is why we need explanations that enable us to think about, and with, these ideas. Explanation speeds up the mastering of a pattern as long as it leads to new insights. Ideas, such as Pavlovian conditioning or Bayesian probability, are fun as long as you discover new ways of using them to explain things. Eventually, they become obvious to you, and using them will be automatic and much less fun.

For something to remain fun, it has to push you to the limits of your ability. This is similar to Mihaly Csikszentmihalyi's (1990) idea of 'flow'. We are in a state of flow when we are fully engaged with what we are doing. It is a dynamic process that can only be maintained if the level of challenge rises as we become more skilled (see Figure 5.4). We remain in the flow channel as long as our skill level and the challenge we perceive are more or less in balance. When our skill level exceeds the perceived degree of challenge, we get bored, and when the challenge exceeds our skill, we become anxious or frustrated and give up.

So, once people are fascinated by an idea, we can make it fun by presenting a series of varied challenges that match their ability. These

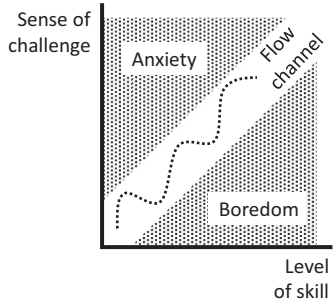

Figure 5.4 Keeping It Fun

Source: Adapted from Csikszentmihalyi (1990, p74)

challenges ideally involve active learning and may be as simple making sense of recent news events or an everyday phenomenon. For this, authentic scenarios or case studies work well. It is important to start out with challenges that are easy – maybe simple scenarios where the application of the idea is obvious, and then progress to more challenging tasks – like generating new applications of the ideas themselves.

For example, when I teach Behavioral Finance (psychology applied to financial markets) I present ideas separately with simple applications to financial markets and progress to more surprising applications like religion and romantic relationships. Then, I integrate the ideas by showing them examples from my own life where someone fooled me by those using the ideas in clever combination. Below is such an example:

> The other day I went to do some shopping when an attractive young lady walked up to me and asked if I knew about the Dead Sea. She talked about how the Dead Sea is associated with all sorts of health benefits. She then proceeded to sell me beauty products made of minerals found in the Dead Sea, even though I had no use for them. This is how she did it.
>
> As I answered her questions about whether I have a wife, girlfriend or mother, she gently took my hand and started to demonstrate the nail care product. She was done in three minutes and the result was a surprisingly smooth and shiny thumbnail. I was still admiring my nail when she started showing me all the products in the nail care kit, and every time she showed it to me, she asked me to hold it. Very soon I was holding all the different products in the kit. She told me that the kit would last me for eighteen months.
>
> I was informed that this kit would sell in a nearby shopping mall for the equivalent of $100, but that she was selling it at a special price of $37. As she showed me her invoice book I could see that most of the pages were filled out by previous customers. She pointed out how the majority of them actually bought more than one nail kit and added that I would get a bottle of some lotion for free with every kit I bought.
>
> While she was filling out my invoice for the nail kit, she kept on repeating further reasons for why I made a good decision. She asked who she should make out the invoice to: "Mr. Handsome, Sexy or Gorgeous?" and made me feel like I was clearly a well-informed buyer (which I obviously was not).
>
> But if I thought she was done with me, I was mistaken. She simply assumed that I would pay by credit card (this way she could sell me even more products) and started demonstrating yet another

> product. If it were not for the fact that I was late for an appointment, I probably would have bought that too. It took much effort for me to directly contradict her several times and point out that she was making assumptions that were not correct. After a few minutes, I extricated myself, paid the $37 cash and left. I felt I did quite well, but later realized I could have bought the same nail care kit for less than half the price elsewhere. I gave the nail kit to my girlfriend and it didn't even last one day.

I ask my students to spot all the behavioral biases to which I fell prey (quite a few). Finally, I ask them to find such examples in their own life and elsewhere (for example, by investigating companies' annual shareholder reports). Their own examples then become case studies in the following years' courses.

While I will address this in more detail in later chapters, I want to touch on the use of formative assessments in making and maintaining interest, because it is one of the easiest ways to keep students within the flow channel. Formative assessments consist of tasks that enable students to develop their understanding – such tasks can be informal in-class activities or challenging homework assignments, and they count very little toward students' grades. In contrast, summative assessments aim to evaluate students' performance and are usually high stakes tests or exams that determine if a student passes a course or not. Since fun requires relaxation, summative assessments are much less fun.

An authentic task is a problematic real-life scenario with enough detail so students can imagine it vividly, and in which they have to use their understanding to achieve something. The outcome is presented (not necessarily in writing) as some kind of performance or product. Authentic tasks encourage students to apply the patterns they are learning and see the relevance of the ideas, and are commonly used by good professors (Ambrose, Bridges & DiPietro., 2010, Bain, 2004). Compare the non-authentic and authentic tasks in the next table.

While students might not like to do any of the tasks, they will find the authentic one more interesting and think more deeply about the ideas. Also, it is easier to add layers of increasing difficulty to authentic tasks as students' abilities improve. The non-authentic tasks (with the exception of the calculation question) do not require thinking – only mentally copying and pasting from the textbook.

My own experience over the years has been that students learn far more from authentic tasks than from my explanations. To construct an authentic task, a simple procedure is to imagine a scenario with the five elements, as shown in Table 5.3 (see Wiggins & McTighe, 2011,

Table 5.2 Two Kinds of Tasks

Non-authentic tasks	Authentic task
Define inflation. Where do we find information on the inflation rate? If the CPI in 2015 was 105 and the CPI in 2016 is 116, calculate the inflation rate in 2016? Explain why the expected inflation rate is important to companies. List and discuss the negative effects of inflation.	It is that time again – wage negotiations for ACME Chairs, a company in Cape Town specializing in making chairs (using mainly unskilled labor). Last year, the company's wage negotiators really messed up – they offered a wage increase that was too low. As a result, the company suffered a damaging month-long strike. This year, they don't want to make the same mistake again, so they called you in to offer expert advice. Find out what is the current and expected inflation rate in South Africa and make a recommendation to the negotiators. Your recommendation should be written in less than one page. In it, you should suggest the proper wage increase and explain two arguments in support that the negotiators can use when negotiating with the labor union.

Table 5.3 Elements of an Authentic Task

Elements	From the example (Table 5.2)
Topic	Inflation
A detailed real-life scenario	Details about the company and its mistakes
Problem for someone	Impending strike due to bad economic advice
Role for the student	Economic advisor to make a good recommendation
Product or performance	In the form of a written recommendation

Module G), or to take an actual event or adapt it so that it contains the five elements.

In the example, my topic was inflation, where I imagined a company that had a problem because they didn't understand inflation. In this scenario, the student plays the role of an advisor who does understand inflation and presents his advice in the form of a report. However, a report is not the only way to have presented it – it could have been a role play, case study, verbal presentation, and many more.

With enough challenges, one's skills will grow, but unfortunately, we tend to prefer the boredom in the zone of mastery and may even resist new ideas that challenge us further. However, in a changing environment,

resisting new ideas is bound to cause stagnation. A student's knowledge structure needs to become more complex, so change is required. However, for students to get, and remain, sufficiently interested in such change, something has to regularly overcome their resistance, and this is where fumbling comes in.

Fumbling Toward New Patterns

To continue learning, students have to periodically open up to new ideas, which happens when their existing ideas are disrupted. Disruption makes students realize that their old patterns do not work as well as expected, and this causes them to fumble (in the sense of struggling to use old ideas and trying to clumsily reach for anything that may help). Active learning helps here too, especially if students are required to respond to provocative questions or make predictions that make them realize they need better ideas. This creates the kind of doubt that Feynman believed was the essence of learning (Gleick, 2011).

I really understood the value of getting people to fumble when I encountered Stephen Thaler's (1997) 'imagination engines' – computer programs that are creative. Thaler's computer program is a type of neural network, so called because it mimics the workings of a human brain. A neural network can learn new skills because, like the human brain, it has brain cells (neurons) and constantly adjusts the connections (or 'weights') between these neurons.

A neural network resembles a simple expert knowledge structure with many nodes and connections between them. As students learn more, their initial hub-and-spoke structure gradually comes closer to a simple expert structure (see Figure 2.5). But, as the structure becomes gradually more complex, new ideas and connections are needed for further learning; and this involves the personal creation of new knowledge. Unless there is some disruption, this will not occur, as Stephen Thaler demonstrated.

Twenty years ago, Thaler wondered what would happen if he disrupted one of his neural networks. First, he taught it several well-known Christmas carols. Then he started cutting some of its connections. To his surprise, at first, the network somehow reconstructed memories along other neural pathways to repeat the carols perfectly. But after cutting more connections, it sang something quite creative: "All men go to good earth in one eternal silent night" (Thaler, 2013, p.452). The disruptions caused it to create new carols it never sang before. However, there came a point where the disruptions were too severe and it no longer made sense.

Thaler realized that he had developed a truly creative machine, and with further refinements, his neural networks generated several inventions that

were sufficiently original to be patented. What he realized was that there needed to be an optimal level of disruption. If he disrupted the neural network too little, it would simply repeat the old sedimented patterns (the Christmas carols in this case) because it did not really have to 'think' (find new neural pathways). But if he disrupted it too much, the neural network sang complete nonsense and the result was random utterances.

There was a small range of disruption that led to new thinking by the neural network, and the same is true of humans. If an idea disrupts us too little, we think it is obvious and hardly pay attention. Disrupt us too much and we think the idea is crazy and we resist it. However, with just enough disruption, we become interested in changing. We start to think, and even if we don't agree with the idea, our minds have been permanently stretched.

Davis (1971) explained that the easiest way to disrupt an audience is to identify something they take for granted, and then to show that it is, in fact, the opposite. For example, in neuroscience, we used to think that emotion made us less rational; but, due to the work of Antonio Damasio, we now know that without emotion, we become irrational. Or we used to think that time was constant; but, because of Einstein, we now know that it changes depending on our movement through space.

Davis's strategy is a little too blunt for actual teaching situations. In lectures, I see disruptions as being subtler and occurring in different degrees of 'spiciness' – like Indian food, where:

- *Mild* disruptions show students that there is something important that they never considered or thought about;
- *Medium* disruptions show students that what they thought up to this point is incomplete or wrong;
- *Hot* disruptions show students that things are, in fact, the exact opposite of what they previously thought.

It does not matter if, in the end, my audience agrees with me or not, or even if I agree with my own disruption. What matters is that my disruption is so well-reasoned that the audience is forced to think and argue about it. At that point, they are fumbling and willing to consider the new ideas I want to them to engage with.

There are usually three situations in which disruption becomes necessary.

- Complexification: Making a knowledge structure more complex involves some degree of disruption. This is because it involves broadening ideas, breaking them up, adding new connections or breaking connections. For example, suppose I teach the history

of World War 2. I will, at some point, have to complicate matters by explaining a new cause – perhaps that the Allied powers were not completely innocent and contributed to the events leading up the war. Or, when teaching about government, I might, at first, use the metaphor of a family, but later, I have to show that this metaphor can be misleading.

- Change of context: When applying an idea to a different context, there is usually some disruption if some ideas need to be adapted. For example, when teaching financial literacy, I would explain that you should not be dependent on debt and not borrow money to pay expenses if you have emergency savings available. But when I explain the kind of financial literacy needed to survive under extreme poverty, I would explain that in this context it does not make sense to be debt-free and that you should, in fact, borrow money rather than draw on your emergency savings. This leads to a broader idea about what it means to be financially literate and is definitely disruptive.
- Common beliefs or preconceptions: When teaching new ideas, one often runs up against preconceived ideas that block further learning and have to be disrupted so that learning can progress. For example, if you teach Biology, you need to disrupt the widely held belief that everything needs a designer; or in Engineering, you need to disrupt the idea that models need to be realistic; or in Economics, I cannot teach international trade unless I disrupt the idea that imports are bad for a country.

My own experience suggests that whether you use mild, medium or hot disruptions depends on the situation, as Table 5.4 shows.

Most gradual complexifications of the knowledge structure do not change the structure dramatically, so the disruption is usually mild or medium. A change of context has the potential for more severe disruptions; while overcoming preconceptions very often requires the hot version of disruption.

Table 5.4 Choosing the Level of Disruption

	Mild	Medium	Hot
Complexification	X	X	
Change of context		X	
Common sense		X	X

One has to take care when disrupting students' knowledge structures. Disruption is useless if done for its own sake. As always, you should focus on only a few ideas: those that are the next step in continuing the learning process. Too much disruption will not allow students to consolidate changes in their understanding or think deeply about them. Once a person has been disrupted and they are fumbling, they are open to being fascinated by a new idea, and the cycle starts again. Let's put it all together now.

PUTTING IT ALL TOGETHER

There is no recipe for using the F^3 approach – it depends on the audience and the presenter's own style. Though ideal, it might not always be possible to use all three Fs, and their actual use may only be a small part of a presentation. The best I can offer is possible steps, as laid out in the exercises at the end of this chapter, and some guidelines derived from my own experience.

I sometimes teach Economics to MBA students. This is challenging because it is often a mixed audience, from people who know virtually nothing to people who keep up with the daily economic and financial news. I find fumbling most useful here in two ways. Firstly, those who know a great deal come to my classes ready to be bored, and they can quickly become disruptive unless I disrupt them first. Because they keep up with the news, they have many more unquestioned assumptions about Economics than those who know little. As a result, I have no problem disrupting them and keeping them engaged. Secondly, this helps me to deal with the mixed audience: I can now explain the basic ideas to those who know little, while the knowledgeable listeners are kept attentive with the disruptions I regularly direct at them. When I then introduce the 'fun' component in the form of an authentic task, I find that both groups enjoy it and learn from interacting with each other.

I employ fascination as early as possible when teaching skills and ideas that appear useless and boring on the surface – especially when I teach research writing skills to (post)graduate students. In the first lecture, I usually pull the alarm, mystique, power, vice, and prestige triggers. When pulling the alarm trigger, I usually throw in a bit of disruption in as well. I explain to them, with examples, how almost everything they have learned in their undergraduate studies has set them up to fail when researching and writing their thesis. Obviously, I then present my course as the solution. With the mystique trigger, I appeal to their innate desire to find out more about the things they care about; and a touch of vice may enter as I show them that, in research, you finally have some freedom to question authority. It is not difficult to pull the power and prestige triggers as I talk about how

it feels to see your ideas have an impact and how it advances your reputation and career. By the time I am done, hopefully, none of them leave thinking research writing is boring. I follow this up throughout the semester with a series of cumulative authentic research activities that leave them with a sense of achievement at the end.

The aspect of fun is indispensable – unless there is an active engagement with the ideas, the work of fumbling or fascination is wasted. So, fun is usually my follow-up after I have addressed fumbling or fascination. Fun is sometimes a way to make an apparently threatening topic (perhaps due to perceived difficulty) less so. When teaching Finance, I may sometimes start with a simple investment game that shows my students that they already know some useful ideas and so build up their confidence. During the course of the semester, the game becomes more challenging and I use their experiences in the game to set them up for disruptions.

In conclusion, let me take you on a quick tour of an actual lecture I give on the business cycle to business executives, so you can see how I apply the F^3 approach in totality.

I walk into the room with many skeptical listeners who think I have nothing to teach them, a few who are nervous because they know very little, and some in-between. I have a range of introductions designed to soothe the fears of the less knowledgeable and disrupt the rest enough for them to pay attention. One of my favorites is to ask them why they think a useless subject like Economics is even taught at this level. I would explain that it is useless because just about everything in the subject is obvious if you get beyond the jargon – and if it is not obvious it is usually of no practical consequence to a business. It does not matter where this discussion takes us except to conclude by saying that they should have an answer by the end of the session.

Since my topic is the business cycle, I explain to them that there are few concepts in economics that are so useful in timing the exploitation of business opportunities. I spend a minute or two pulling the power and mystique triggers and may even throw in some true and alarming stories of firms going bankrupt by ignoring the business cycle. I then launch into a brief explanation of the business cycle for those who do not have sufficient prior knowledge so they can benefit from the lecture. I keep it short and simple, focusing on a few critical ideas. At this point, I trade accuracy in the details for broad understanding, especially since the details will only become meaningful later.

Now they are ready to be disrupted, and just in time… before the knowledgeable ones become bored. I ask them: "Which part of the business cycle is the best for a business?" Most will answer the upswing phase when the economy is expanding. I then tell them this is mistaken and the reason some firms went bankrupt, and that the best part in the cycle for a

successful business is, in fact, the downswing, when the economy is stagnating. At this point, I no longer need to push my explanations onto them; they will pull everything out of me through their arguments and questions. This is the longest part of the lecture, but sometimes it feels like the shortest.

Finally, I end the lecture with an authentic task where, given the expected direction of the business cycle, they criticize or commend their own company's existing strategies and devise new or adapted strategies. We may even look at the cases of bankrupted firms I mentioned and see if they could have done things differently.

At this point, there is no need to return to the question of why a useless subject like Economics is taught. They can see that, even though the ideas I taught were obvious in hindsight, the application of these ideas is not always obvious when we are blinded by unquestioned assumptions.

CONCLUSION

Fun, fumbling and fascination are not entirely separate concepts, and they overlap quite a bit. By extension, interestingness is not so much a product as a continuous process, as shown in Figure 5.2. The process is summarized in Table 5.5.

To end, I return to the 'great explainer'. It is easy to regard teaching as a boring but necessary activity that takes time away from research. But Feynman showed that teaching is sometimes the most interesting thing a professor can do. For him, teaching was a way to test and improve his own understanding. More importantly, teaching disrupts us. He believed that because "teaching is an interruption" (Feynman & Leighton, 1984, p.5) it presents us with opportunities to question the assumptions of our discipline, gain new perspectives and find new questions for research.

Table 5.5 Summary of the F^3 Process

	Purpose	Action
Fascination	Attracts us to some ideas…	…by pulling one of our fascination triggers…
Fun	makes us want to take up the challenge…	…of finding and mastering new ideas…
Fumbling	…but once we master it, it becomes boring, so something has to create the need for new ideas…	…by disrupting our old ideas…

EXERCISES

1. Identify a topic that students find boring. Follow this three-step process to create fascination:
   a) Consider each of the seven triggers of fascination and write down as many ideas as you can about how you can use them to get students' attention.
   b) Given what you know about your students, which trigger/s are they most likely to find relevant? Eliminate the others for now.
   c) Which of the triggers are easiest to connect to the critical idea/s of the topic? Find at least one trigger for the topic, and, if appropriate, one for every critical idea
2. Identify an idea that students find boring. Follow this process to create opportunities for fun:
   a) Work out how you will get students fascinated with the idea, or at least the topic of which it is part.
   b) How will you reveal the pattern of connections that create the idea? Consider active learning methods (see Box 5.1). Or will it simply be an explanation accompanied by a series of thought-provoking rhetorical questions with pauses? The point is not to talk non-stop, but to give students the chance to participate, even if they do so just mentally, in trying to figure out the pattern.
   c) Design an authentic task based on the idea or overall topic, that can be done individually or in groups.
3. Identify a topic that students find boring. Follow this three-step process to create fun:
   a) Identify the points at which the knowledge structure needs to change due to complexification, preconceptions or changing context.
   b) Choose one of these for which to introduce a disruption. It is best not to disrupt too much, so eliminate potential disruptions that are not related to critical ideas. Students should fumble only to open their minds to important changes in their knowledge structure.
   c) What degree of disruption will you use (mild, medium or hot)? Work out how you will do it.

REFERENCES

Ambrose, S.A, Bridges, M.W. & DiPietro, M. 2010. *How Learning Works: Seven Research-Based Principles for Smart Teaching*. San Francisco: Jossey Bass.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Caine, R.N. & Caine, G. 1991. *Making Connections: Teaching and the Human Brain*. Alexandria: Association for Supervision and Curriculum.

Costandi, M. 2013. Evidence-based justice: Corrupted memory. *Nature*, 500:268–270.

Csikszentmihalyi, M. 1990. *Flow: The Psychology of Optimal Experience*. New York: HarperCollins.

Davis, M.S. 1971. That's interesting! *Philosophy of the Social Sciences*, 1(4):309–344.

diSessa, A.A. 2014. A history of conceptual change research: Threads and fault lines, 2nd edition. In: R. K. Sawyer (ed.), *The Cambridge Handbook of the Learning Sciences*, 88–108. New York: Cambridge University Press.

Feynman, R.P. & Leighton, R. 1984. The dignified professor. *Engineering & Science*, November:4–10.

Gleick, J. 2011. *Genius: The Life and Science of Richard Feynman*. New York: Open Road.

Hogshead, S. 2010. *Fascinate*. New York: HarperCollins.

Koster, R. 2005. *A theory of fun for game design*. Scottsdale: Paraglyph Press.

Loftus, E.F. & Palmer, J.C. 1974. Reconstruction of automobile destruction: An example of the interaction between language and memory. *Journal of Verbal Learning and Verbal Behavior*, 13(5):585–589.

Thaler, S.L. 1997. A quantitative model of seminal cognition: The creativity machine paradigm (US Patent 5,659,666). Paper available at www.imagination-engines.com, accessed 24 April, 2007.

Thaler, S. 2013. Creativity Machine® Paradigm. In: Carayannis, E.G. (ed.), *Encyclopedia of Creativity, Invention, Innovation, and Entrepreneurship*. New York: Springer-Verlag, pp. 447–456.

Tyng, C.M., Amin, H.U., Saad, M.N.M. & Malik, A.S. 2017. The influences of emotion on learning and Memory. *Frontiers in Psychology*, 8:1454.

Wiggins. G. and McTighe, J. 2011. *The Understanding by Design Guide to Creating High-Quality Units*. Alexandria: Association for Supervision and Curriculum.

Willingham, D. 2009. *Why Don't Students Like School?* San Francisco: Jossey-Bass.

Chapter Six

If You Want Students to Reason Like Experts, Don't Teach Them How to Reason

What does it mean when you give a group of professors and some bright high school students the same quiz and some of the students outperform the professors? Does it mean that the professors are not really experts and have little to teach the students?

This is exactly what Sam Wineburg (1991) did when he gave some History professors and students a test on the American Revolution. Some of the professors did not specialize in this topic, and a few students did better than them on the quiz. But for Wineburg, this was neither surprising nor embarrassing, because he did not accept the conventional definition of 'expert'.

What is an expert really? In a subject like History, is it a person who can answer the most questions about history, or is it a person who knows how to be a historian?

Being a historian, sociologist, geographer or botanist is not a quiz contest. Instead, experts are people who know how to make sense of conflicting, incomplete, and vague ideas and use them to improve not only their own knowledge but also create knowledge in the discipline. They know how *to be* an expert, as mentioned in Chapter 1.

This is why, after the quiz, Wineburg gave them the real test. Students and professors were given the task of drawing conclusions from real historical documents, which contained, as one would expect, many gaps, inconsistencies, and conflicting views. For example, one of the tasks was, after having been given a variety of source documents, to select from a group of paintings, the one that best reflected the actual Battle of Lexington.

Here is where the professors excelled. They thought that the task was very difficult and it pushed them to their limits. Wineburg observed how they: tried to interpret the evidence in different ways; endeavored to make sense of the conflicts and the gaps; went back-and-forth; generated

alternative explanations, with provisions and qualifications; and finally identified the illustration they thought was the least wrong, and hence most trustworthy. They did this by going back to the first principles of studying history in order to distinguish between the critical and non-critical information and reasoning from that toward a tentative answer that approximated the likely truth. In other words, they really struggled.

So, how did this show that the professors were experts? Consider what the students did. For them, the task was much easier, and they treated it like a multiple-choice quiz. They thought that there had to be one right answer and that their job was simply to select the one that was pre-determined as correct. There was no attempt to construct knowledge or arrive at a sophisticated understanding, because they thought of knowledge as given and constant, not as something that is imperfect, changeable, contested and constructed.

If knowledge were something that is given and constant, the difference between those who understand and those who do not, would not be so clear. It only becomes obvious when experts and novices are thrown into situations with no well-defined questions or answers, where knowledge has to be reasoned out. Referring to one of the professors in the study, Wineburg (1991, p.84) put it like this: "Her expertise lay not in what she knew, but in what she was able to do when she did not know."

This confirms everything from the previous chapters. To remind you: understanding is the ability to compress information into a much smaller number of connected ideas; and then, by using *reasoning* and a bit of imagination, generate new information, even in situations not experienced before. In the study, the professors had compressed their knowledge into a few deep principles. Through reasoning, they unpacked those principles and inferred new connections, helping them make sense of a challenge they had never faced before.

In Chapter 1, we looked at the concept of compression of a subject into a few critical ideas; then, in Chapters 2 to 5, we investigated how connections generate understanding, improve memory, and evoke interest in those ideas. This chapter completes the set of five critical skills a professor should have: enabling students to expand their understanding through reasoning. This ability is also the biggest differentiator between novices and experts.

REASONING UNLOCKS THE POWER OF UNDERSTANDING

Richard Feynman (2011, p.4) explained: "There is an enormous amount of information about the world if just a little imagination and thinking

are applied" to a few critical ideas. What he called 'imagination and thinking' I will simply refer to as 'reasoning'. Reasoning causes the seeds of a few critical ideas to bloom into a forest of smaller, connected ideas – it unlocks the full power of understanding.

Reasoning is commonly defined as the process of drawing new conclusions from a pre-existing set of information. It usually involves deductive logic, but in the cases of inductive and abductive reasoning, it may also require some imagination. Mercier *et al.* (2016) point out that this is more accurately described as 'inference'. They explain that reasoning is, in fact, a more specialized form of inference: it is inference guided by the process of finding, using, and evaluating reasons. It is not simply making claims (like "It is going to rain later."), but making arguments: making a claim, backed by evidence, warrants, counterarguments, and qualifications.

Conversations make reasoning particularly effective, because, when we make arguments, we reveal our reasoning, so that our reasons for claiming something can then be questioned and improved by others. As we are exposed to new arguments in this way, we start seeing weaknesses in our own thinking, and our understanding deepens. Mercier *et al.* (2016) call this public exchange of reasons 'argumentation'.

To be accepted by others, all ideas have to be argued. Without having any reasons, it would be impossible to respond to ideas, and it would also be difficult to know how to connect them to other ideas and build on them. As the connections are uncovered, it becomes possible to elaborate on the ideas. The imaginary conversation technique, encountered in Chapter 1, and which is itself a form of reasoning, showed how this happens.

Reasoning, but specifically argumentation, also overcomes a big problem that occurs when people feel they understand something completely – the tendency to absolutism. When compressing information to a few critical ideas, it is easy to get a sense that one now possesses the ultimate set of ideas that can explain anything, anywhere, at any time. Because reasoning exposes our thinking to others, it creates the opportunity for others to question us and for doubt to enter our minds. This keeps us open to the possibility that our understanding is imperfect and that there is still more knowledge to construct. Like the History professors, when we reason, we are made to realize that much of what we know is wrong, or at the very least, conditional, and that it often depends on context or perspective. It is this realization that makes us understand as experts do. Experts do not simply see the connections, they also recognize the conditions under which the connections may or may not be valid, and this provokes them to always search for ways to improve their understanding.

Finally, because argumentation forces us to explain ourselves publicly, it taps into the 'self-explanation effect' (see Box 2.3). It has now been

well established that, by making our thinking explicit and articulating ideas and reasons, we become aware of imperfections in our own understanding, and are thus encouraged to improve it and fill in the gaps.

Clearly then, active promotion of student reasoning is an important part of great professors' courses (Bain, 2004). The temptation is then to teach it. But that would be a mistake.

DO NOT TEACH STUDENTS HOW TO REASON

Human reasoning is an ability that evolved over millennia to solve the problem of creating and evaluating knowledge in groups. While not perfect, it has left almost every human with excellent reasoning skills, as long as these skills are exercised *under the right conditions*, specifically in settings where there is a public exchange of diverse ideas – that is, in argumentative settings.

Teaching reasoning skills through things like logical fallacies for example, not only distorts natural reasoning skills, it also wastes valuable time because one would be teaching something that is an innate ability. It is better to simply create the conditions that get students to actually use the natural abilities they already possess and guide them in applying these skills to the discipline. A detailed and well-researched defense of this view can be found in Mercier *et al.* (2016) and Mercier and Sperber's (2017) insightful book. For the purpose of this section, I review only some of the arguments with educational implications.

Kahneman (2011) and other researchers, especially in economics and psychology, create the impression that human reasoning is flawed and biased. They seem to ignore the vast body of literature that finds that people are actually naturally competent reasoners. Humans can easily spot fallacies and biases in the arguments of other people, are quite good at evaluating arguments that they care about, and become critical thinkers when confronting arguments with which they disagree. The reason for this is that our ability to evaluate arguments mainly depends on whether we are exposed to counterarguments and also how easy it is for us to produce or anticipate such counterarguments.

Indeed, as individuals, we are lazy reasoners and prone to reach biased conclusions. However, in a conversation, the arguments I generate for my case are the counterarguments against your case. So, when we get together, your (counter)arguments cause me to change the initial evaluation of my own arguments or cause me to produce better arguments, which, in turn, causes you to re-evaluate your arguments.

If we are bad at reasoning, it is not because of biases or a lack of knowledge, but because we have not been able to think of sufficient

counterarguments. This is something easily fixed in groups containing people who have different perspectives. In fact, Mercier *et al.* (2016) highlight that the most persuasive evidence that humans are good at reasoning comes from studies of group discussion. The research agrees that a group can out-reason individuals, *but only if* it contains people with *different* views who are willing to *voice* their opinions and criticize the opinions of others. Due to its ability to generate more counterarguments, group reasoning has been found to consistently outperform individual reasoning. Groups containing a variety of views are more likely to produce counterarguments, and, furthermore, in cases of inquiry-based reasoning, groups are up to five times better at reasoning than individuals.

Very similar findings exist with regard to the use of evidence. People fail to use evidence only when they find it difficult to think of any, which is often the case in informal discussions. However, when evidence is available, people are likely to use it and use it intelligently (Brem & Rips, 2000).

Studies on collaborative learning agree. Students' discussions are effective when they are exposed to alternative views and feel free to state their views and change their minds. Such discussions improve student performance and deepen conceptual understanding (Johnson & Johnson, 2009; Nussbaum, 2008). This assumes that groups are small and that certain rules are in place, such as avoiding personal attacks, staying on the topic and fair representation of arguments. It is here where templates – such as those of Graff and Birkenstein (2006), the sentence openers of AcademicTalk (McAlister, Ravenscroft & Scanlon, 2004) and the prompt and response frames of Zwiers & Crawford (2011) are useful.

When reasoning skills are taught as something that is separate from the discipline, students rarely transfer those skills to other ideas. However, when reasoning is integrated into the teaching of one topic, students not only apply it to that topic but also transfer it to new topics as they learn to anticipate counterarguments and improve their own arguments (Kuhn and Crowell, 2011). The more students talk to each other in argumentative settings, the more they learn (Resnick, Asterhan & Clarke, 2013, Henderson *et al.*, 2015).

UNDER THE RIGHT CONDITIONS, EVERYONE REASONS WELL

From the above, we can derive the conditions under which reasoning is naturally effective. If we can create these conditions, not only will students develop and extend their natural reasoning skills, but they will also remember and understand more.

Most importantly, reasoning is best in groups consisting of individuals who exhibit different perspectives. In contrast to the directives of brainstorming proponents, people in such groups should feel free to criticize each other, because the quality of our reasoning is significantly enhanced as we are exposed to more counterarguments. After a while, the realization that our arguments are easily refuted causes us to adopt a more critical attitude toward our own reasoning and anticipate counterarguments on our own.

This interplay of argument and counterargument produces optimal results in small groups of four to five people. If the group is too small, the variety of views may be too limited, and if it is too large, not everyone gets an opportunity to share their arguments.

Argumentation works best when there is an exploration of differences rather than an attempt to eliminate those differences. Individuals in a group should care less about being right and more about making sense of ideas. To achieve this, the reasoning process should ideally start with an idea about which, even among experts, there is a knowledge gap or conflict. If students know that even the experts are unsure, they are less likely to be dogmatic and more motivated to pursue new understanding.

Finally, reasoning is better practiced in the context of knowledge evaluation and production within the discipline, as opposed to teaching it as a separate topic or course. This applies even in cases where the disciplinary aspects of reasoning need to be taught.

IF YOU HAVE TO TEACH, THEN TEACH MOVES AND VARIATIONS

Reasoning is a skill that is generally expected in all disciplines, and fairly generic. In fact, when Arnold Aarons shares his list of reasoning abilities that comprise critical thinking, many professors think that because the list is so accurate, he must know their discipline (see Bain, 2004, pp. 85–87). But these generic skills usually need not be taught. What should be taught are the "disciplinary ways of thinking and knowing" (Middendorf & Pace, 2004, p.1) and the discipline-specific variations of reasoning.

Variations are found in the way in which different disciplines search for evidence and in the different sources used. Another variation lies in drawing inferences from data, because, not only is the quality of data judged differently, different systematic methods of data analysis are used in different disciplines. These, and other disciplinary differences, are not innate to human beings, and therefore have to be taught.

There is one general aspect of reasoning that needs explicit instruction, and that is the language used in talking about it. This includes the

vocabulary used in describing different elements and structure of arguments and the different argumentative moves. The Toulmin model is an excellent way to acquire this vocabulary, while argumentative moves can be learned from various sources, such as the templates mentioned earlier, or from theoretical categorizations, such as appraisal theory (Martin & White, 2005) or dialog theory (Walton, 2000).

BOX 6.1 ELEMENTS OF ARGUMENT AND ARGUMENTATIVE MOVES

The most commonly used model to learn about the elements and structure of individual arguments is that of Stephen Toulmin. To find information on that, search for **"Toulmin model"**. There are some good summaries of Graff and Birkenstein's (2006) book of argumentative moves, which you can find by searching for: **summary "moves that matter"**.

HERE IS ONE WAY TO PROMOTE LEARNING THROUGH ARGUMENTATION

While there are countless books and courses that aim at learning-to-reason, there are far fewer reasoning-to-learn approaches. Some of the best approaches I have come across are discussed by Andriessen and Baker (2014) and Scardamalia and Bereiter (2014). To this valuable literature, I will add one more approach that is suited to controversial topics without obvious answers. It draws quite heavily on the techniques described in Wentzel (2017) and Chapters 3 and 4 of Wentzel (2018).

Step 1: Select a Topic

Select a topic within the syllabus that contains ideas that may lead to very different, even conflicting, conclusions. This will not always be communicated in the textbook, so you often have to draw on your understanding of the controversies in the discipline.

For example, when teaching strategic management, one issue is which approach to developing strategy is superior. While most textbooks support a prescriptive approach, where strategy is developed and defined by a small group at the top and then communicated downwards for implementation, in practice, one finds that strategy is not pre-determined, but rather emerges as the organization continuously adapts to changes.

Step 2: Generate Uncertainty in Students' Minds

Avoid adversarial reasoning: where students become invested in a position and defend it at all costs, so stay clear of religious and politically sensitive topics. Instead, encourage inquiry-based reasoning: where students become less interested in being right and more interested in making sense of the information that generates controversy.

Make a convincing case for the different views (ideally limited to two) without hints as to which one you agree with. Make it as clear as possible why these views are in conflict or inconsistent with each other, why it is important to take a position and why it is difficult to agree with both. If there are gaps in experts' knowledge, strong disagreements among them or puzzles they have not yet solved, these will help to generate further uncertainty.

If the textbook takes one view, you need to draw on history, practice or alternative views to show that there is no clearly correct view and that even the experts are unsure or in disagreement. You don't want students to approach reasoning as a multiple-choice test or as a competition, but rather to see other arguments, even opposing ones, as useful in helping them to improve their understanding.

Uncertainty is also related to some of the triggers of interest from Chapter 4 and gets students to care about the arguments. This helps because humans become more competent reasoners when dealing with topics they care about.

Step 3: Derive an Incomplete Summary

Based on your explanations of the views, summarize the result (as shown in Figure 6.1). This summary is only a start and gets completed in step 5.

To create the diagram, write the opposing views in the boxes V_1 and V_2. Then ask yourself: If we take this view or put it in practice, how will it help those involved? For the system in question, what need does it satisfy? The need satisfied by V_1 is written in N_1, and the need satisfied by V_2 goes in N_2. For example, those who agree that strategy must be pre-determined at the top, believe this is necessary in order to execute and evaluate strategy as fast and as consistently as possible.

The diagram is most likely correct if: (1) you can read 'if-then' statements from both branches in a way that makes sense; (2) it is very difficult to agree with both views (V_1 and V_2); and (3) both sides are consistent with credible literature or practice.

Figure 6.1 seems to work because an 'if-then' statement that makes sense can be derived from both branches. The first statement is: "If we

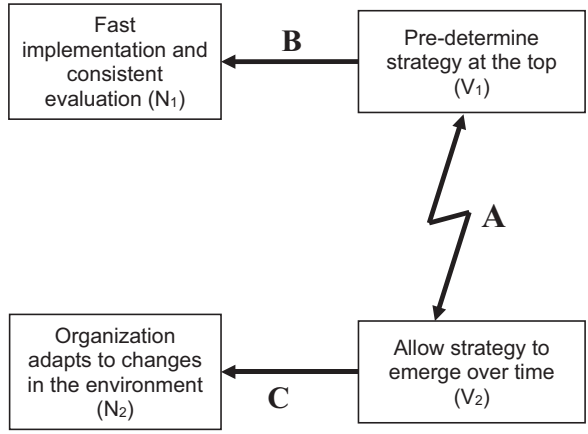

Figure 6.1 An Incomplete Summary (Completed in Figure 6.3)

pre-determine strategy at the top, then we will have fast implementation and consistent evaluation (of strategy)"; and the second: "If we allow strategy to emerge over time, then the organization will adapt to changes in the environment." Based on the literature, the two views are definitely in opposition. Once you have generated this summary from your explanation, keep it up where it is visible.

Steps 4–7 build on this using a technique called 'think-pair-square-share' (Bain, 2004, p.130). It works because, as mentioned earlier, humans seem to reason best in small groups (consisting of 4–5 people). In such groups, the counterarguments of everyone are heard, and this is important because it is the availability of counterarguments that makes reasoning effective

Step 4: Individual Thinking with Evidence

Ask individuals to take a *provisional* position on the different views and to find reasons for their position. If they agree with V_1, they will find the reasons in arrow B of the incomplete summary, but if they agree with V_2, they will find these in arrow C. To find the reasons, they can simply ask themselves questions such as: Why do I agree with this 'if-then' statement? What needs to be true for this statement to be true? What is the statement assuming?

For example, the reasons underlying $V_1 \rightarrow N_1$ may be that top managers have better knowledge of the company's challenges and should,

therefore, decide on strategy, or that the participation of other employees in the strategy-making process would slow it down. Reasons underlying $V_2 \rightarrow N_2$ may include: that strategy does not have to be a formal document; that strategy should be influenced by those who have to implement it; or that the pace of change is accelerating, so strategy should remain open to change.

These reasons should then be transformed into arguments, using something like the Toulmin model (see Box 6.1). This means that reasons have to be supported by evidence and explanations of how the evidence supports the reason (warrants) and qualifications. If there is time, also ask students to think of possible counterarguments to their position, and how they would rebut such arguments.

People will reason with evidence if the evidence is available. So, for this step, prepare a sheet containing relevant evidence in advance, and hand it out to students before continuing. The evidence should not be skewed in favor of one view. By anticipating as many of the prospective reasons as possible, you can gather, in advance, the evidence that confirms or contradicts each of these. The result should be a fact sheet that contains information that both confirms and questions the validity of the reasons on both sides of the diagram. Alternatively, you can allow students to search for evidence during the lecture.

Step 5: Pair for Maximal Difference

Have students position themselves physically on a line such as the one in Figure 6.2. Place them in pairs in such a way that an equal difference (more or less) exists within every pair. The dotted lines in the figure indicate how this would be done: match students F and L, students G and M, and so on. Matching students from the outside in (for example, F-S, G-P, H-N) would not work because some pairs would have very large differences, while others

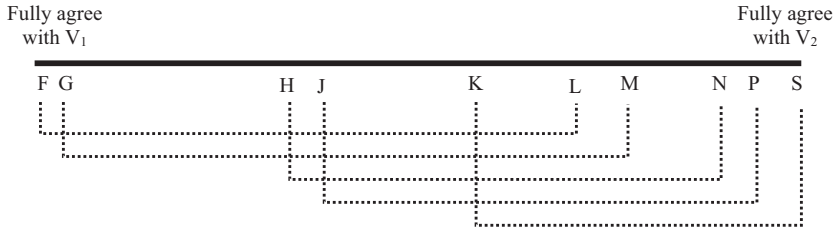

Figure 6.2 Students Take a Position

would have small differences; in such a case, not all pairs will exhibit the same quality of reasoning. Reasoning naturally works better when there are differences that are large enough to expose people to counterarguments they would not have thought of by themselves. Of course, it is ideal if one could get fewer students to initially take a middle position, but this is not always possible, and it is best not to artificially force students to be extreme.

The first thing that pairs should do is get into the right mindset to pursue inquiry-based reasoning (as opposed to adversarial). Ask students to find common ground by finding their common goal and extending Figure 6.1 into Figure 6.3.

Figure 6.3 shows that this is done by recognizing that both sides, even though they disagree, are trying to achieve the same goal. The common goal is usually quite broad: the broader and more general, the easier it is for both sides to agree with it. In the figure, for example, the common goal may be something like: "An organization that is profitable in the long-run". Test the common goal by checking if it generates if-then statements that make sense. In this case, both make sense. $N_1 \rightarrow G$ makes sense: "If there is fast implementation and consistent evaluation of strategy, then the organization will be profitable in the long-run"; and so does $N_2 \rightarrow G$: "If the organization adapts to changes in the environment, then it will be profitable in the long-run".

Recognizing that both sides have a common goal will go a long way to evoke a more collaborative attitude between both sides, as captured by the "principle of charity" (Wilson, 1959, p.532). The principle of charity encourages us not to assume that those who disagree with us are

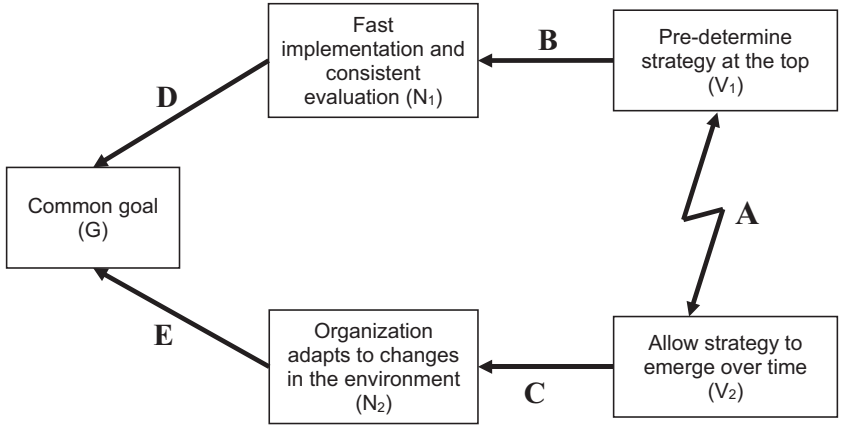

Figure 6.3 Find Common Ground

incompetent or ignorant, but rather to assume that they have good reasons for their view. Instead of nit-picking on tiny inconsistencies in an opponent's argument, we should try to understand its overall intent and construct the best possible version of such an argument with which to argue.

BOX 6.2 PRINCIPLE OF CHARITY

The Wikipedia entry for the 'principle of charity' is not so easy to understand, but you can find many clearer explanations when your search for "**principle of charity**" on the web. Several of them explain how this principle enables us to understand more, how it makes our arguments more likely to produce valuable knowledge, and how it stretches and improves our ability to formulate stronger arguments.

Give students some time to *individually* find a couple of reasons for why arrows D and E could be valid. A possible reason underlying D include may be that strategy has to be implemented fast, so the organization can see if it is doing the right things; and a reason for why arrow E could be valid may be that when the environment changes fast, organizations can only survive if they adapt quickly. While doing this, students will find that some of the reasons underlying D and E are sources of disagreement. Ask them to take the reasons for the side they are closest to, and transform them into arguments, using the fact sheet or other sources of evidence.

Once they have done this, ask them to exchange these arguments, in addition to the arguments for their side from step 4. So, the student closest to V_1 should share her arguments based on the reasons underlying arrows B and D; and the one closest to V_2 should share his arguments underlying arrows C and E.

Before they engage in argumentation, have the students simply ask each other the following questions on the arguments they wish to challenge:

- Is this reason true for all time?
- Is this reason true under all conditions?
- Is there really evidence for this reason? If so, is the evidence credible?

Then, let them respond to each other with additional counterarguments using the language, pre-defined rules, and templates of argumentation. If the two sides are inconsistent or in conflict, the counterarguments

should be easy to find. The reasons underlying arrow B usually evoke counterarguments to those underlying arrow C, and vice versa, and the same applies to arrows D and E.

Next, ask each side to refine their position based on the counterarguments of the other side, especially those counterarguments they did not consider in step 4. This may involve shifting their position, filling gaps in their reasoning, making concessions to the other side and, most importantly, more clearly specifying the context or conditions under which their argument would be true.

The diagram that results from steps 2–5 is based on the work of Goldratt (1994). It has many other uses, including finding original research contributions (Wentzel, 2018); designing productive questions; and systematically identifying assumptions (Wentzel, 2017).

Step 6: Square for Greater Understanding

Now ask any two pairs to come together. If the class number is not divisible by four, try to keep the groups between four and six students. In the case of Figure 6.2, where there are 10 students, I would probably ask students F, G, N, P and S to join together in one group; and H, J, K, L and M in another. In this case, the size of the differences between groups is less important because students will have already opened their minds to a wider variety of arguments as a result of step 5. The purpose of step 6 is to use the arguments to prepare students for the creative synthesis in step 7.

Let students briefly exchange their positions and main arguments so that everyone is familiar with each other's views. Now ask them to identify what both sides agree on. For example, if the argument on one side is that: "Strategy has to be implemented fast so the organization can see if it is doing the right things"; while the other side argues that "When the environment changes fast, organizations can only survive if they adapt quickly"; then they should be able to see that both sides acknowledge that strategy deals with change – even if they disagree on the kind of change.

If there is time, ask students to define what information would be necessary to decide between the most critical opposing arguments. Based on this, let them design potential research projects or experiments that would be able to generate this information.

Step 7: Share to Prepare for Synthesis

Bring the whole class together. Allow a little bit of time for people to share the most interesting arguments they encountered. But this can go even further.

Explain to them that, even when we disagree with a view, if that view exists in practice or in the minds of other experts, it contains valuable information. By synthesizing opposing views, we can take advantage of the information on both sides, instead of trying to destroy the opposing side and the information it holds. Many creative breakthroughs are often the result of new ideas that emerge from synthesis. Here, you can mention some examples from the discipline, for example, some best-selling books on strategy, like Ries (2011), have found ways to synthesize the opposing sides found in Figure 6.3.

Figure 6.4 shows the three questions that guide the class discussion towards synthesizing the two sides and generating new knowledge from their understanding at this point:

- How would it be possible to have fast implementation and consistent evaluation of strategy by allowing strategy to emerge over time? ($N_1 \rightarrow V_2$)
- How would it be possible for the organization to adapt to changes in the environment by pre-determining strategy at the top? ($N_2 \rightarrow V_1$)
- How would it be possible to achieve *both*: allow strategy to emerge over time and also pre-determine strategy at the top? ($V_1 \leftrightarrow V_2$)

Produce a record of the ideas that emerge. Guide the class to develop arguments for the most promising new ideas, and to establish where

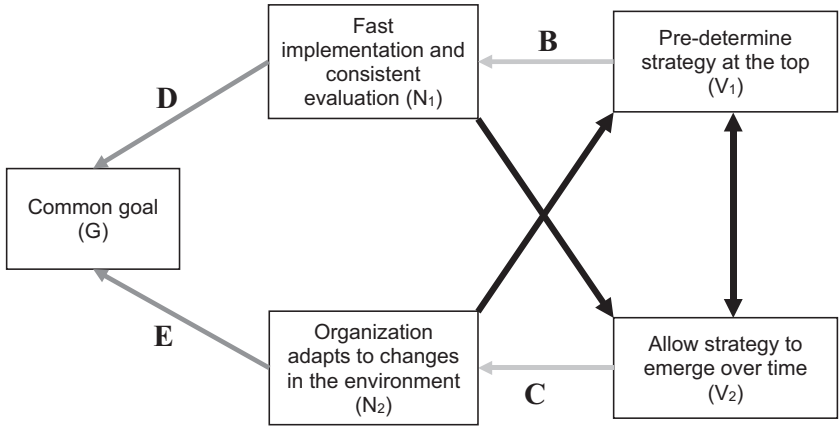

Figure 6.4 Creative Synthesis

these ideas fit within the conversation between the experts in the discipline. Gaipa (2004) suggests some ways to think about this. Also, when facilitating a class discussion, it is useful to mainly use Socratic questions such as those discussed by Paul and Elder (2016).

Step 8: Debriefing

Using questions, debrief the class. Review how they may have modified their original views, and what they learned. Get a sense of how many people shifted their position – maybe get them to take a position on the line again. Ask them to identify what surprised them during this process, and to draw some overall conclusions. Help them to see how their thinking is starting to approach the kind that is found among experts.

A Summary of the Procedure

This procedure is summarized in Table 6.1.

Table 6.1 A Procedure to Promote Learning through Reasoning

Step	Description of action	Minutes
1	Announce a controversial topic	<1
2	Explain the views, generate uncertainty, show the difficulty of a middle position and why taking a position is important	15–30
3	Generate an incomplete summary and keep it visible	5
	Hand out a fact sheet and let students review it	5–10
4	Students take a provisional position	5
	Students engage in individual reasoning in order to formulate arguments underlying the arrow closest to their position (and possible counterarguments*)	10–15
5	Students position themselves on a line	5
	Divide them into pairs that contain different views	5
	Pairs define their common ground	5–10
	Individually formulate arguments underlying arrows D and E	10–15
	Exchange arguments within each pair	5–10
	Students ask each other challenging questions and try to answer them	10–15
	Engage in further counterarguments	10–15
	Students individually refine their argument and change their positions	5–10
6	Pairs combine into groups of 4–6 people	5
	Groups familiarize themselves with each other's arguments (no debating)	5–10
	Groups identify points or principles that both sides agree on	10–15

Step	Description of action	Minutes
	Propose research projects or experiments that can settle some differences*	5–10
7	Bring the whole class together and explain the value of synthesis	5
	Facilitate a discussion to discover new ideas that synthesize opposing views	10–15
	Help students to understand where their new ideas fit into the discipline	5–10
8	Debrief to review and summarize what was learned	10–15

** optional*

The whole process takes two to three hours. If all the steps cannot be completed during one session, then it is best to complete steps 1–4 in one session, step 5 in one session and steps 6–8 in one session.

CONCLUSION

We become experts, not by accumulating facts, but by connecting them and then reasoning with them to make sense of situations and ideas that were previously unknown to us. In this way, reasoning enables understanding to expand to its full potential. To help students achieve this, do not try to teach reasoning. It is much better to place students in argumentative contexts that allow them to amplify their natural reasoning abilities. Doing this will not only improve memory but deepen understanding as students become increasingly aware of the imperfections and conditionality of human knowledge. In many cases, this will encourage them to contribute to the body of knowledge: their own, that of the classroom community, and ultimately, that of the world.

As professors, we play a critical role in making this possible, because as Schwarz (2009, p.102) argued, "Without meticulous planning concerning tools, initial knowledge of the discussants, their social arrangement etc., talk is generally non-argumentative; argumentative talk emerges generally *when structured by the teacher* and/or by representational tools" (emphasis added). For this reason, this chapter concluded with an approach to planning and facilitating learning through reasoning.

EXERCISES

1. Follow the procedure laid out in this chapter and prepare the materials for a class where students will learn by means of reasoning. At the very least, do the following:

a) Select a topic.
b) Prepare notes for your explanation that will reveal the opposing views; make both views convincing, generate uncertainty about which one is true, show why taking a position matters, and why it is difficult to take a middle position.
c) Work out the complete diagram (the incomplete summary and the common goal).
d) From the complete diagram, identify reasons and assumptions that underlie each arrow and find evidence that exists to support both sides. Mix it up and put it on one page if possible.
e) Consider possible designs of research projects or experiments that could help to decide between the opposing views.
f) Prepare notes for your explanation of why synthesis is important and examples from the discipline.
g) Anticipate possible synthesizing ideas that students might generate and prepare notes on where each idea might fit into the disciplinary conversation.
h) Prepare questions for the debriefing.

2. To prepare for the class in (1) above, what do you think you should teach students about reasoning? Mention what you think is necessary and how you would integrate it in your teaching (not as an isolated topic). This may, or may not, include the following:
   a) The language and structure of argument, as well as argumentative moves.
   b) Ways of knowing that are unique to the discipline – e.g. what the discipline regards as debatable and non-debatable, what counts as evidence in the discipline or acceptable methods to generate new evidence etc.
   c) Rules of verbal argumentation, useful templates or particular conventions in the discipline relating to verbal argumentation.

REFERENCES

Andriessen, J. & Baker M. 2014. Arguing to Learn. In: Sawyer, R.K. (ed.), *The Cambridge Handbook of the Learning Sciences*, 2nd edition, New York: Cambridge University Press, 439–460.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Brem, S.K. & Rips, L.J. 2000. Explanation and evidence in informal argument. *Cognitive Science*, 24:573–604.

Feynman, R.P. 2011. *Six Easy Pieces: Essentials of Physics Explained by Its Most Brilliant Teacher*. New York: Basic Books.

Gaipa, M. 2004. Breaking into the conversation: How students can acquire authority for their writing. *Pedagogy: Critical Approaches to Teaching Literature, Language, Composition, and Culture*, 4(3):419–437.

Goldratt, E.M. 1994. *It's Not Luck*. Great Barrington, MA: North River Press.

Graff, G. & Birkenstein, C. 2006. *They Say/I Say: The Moves that Matter in Academic Writing*. New York: Norton.

Henderson, J.B., MacPherson, A., Osborne, J. & Wild, A. 2015. Beyond construction: Five arguments for the role and value of critique in learning science. *International Journal of Science Education*, 37(10):1668–1697.

Johnson, D.W. & Johnson, R.T. 2009. Energizing learning: The instructional power of conflict. *Educational Researcher*, 38(1):37–51.

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar, Straus & Giroux.

Kuhn, D. & Crowell, A. 2011. Dialogic argumentation as a vehicle for developing young adolescents' thinking. *Psychological Science*, 22:545–552.

Martin, J.R. and White, P.R., 2005. *The Language of Evaluation: Appraisal in English*. New York: Palgrave Macmillan.

McAlister, S., Ravenscroft, A. & Scanlon, E. 2004. Combining interaction and context design to support collaborative argumentation using a tool for synchronous CMC. *Journal of Computer Assisted Learning*, 20(3):194–204.

Mercier, H., Boudry, M., Paglieri, F. & Trouche, E. 2016. Natural-born arguers: Teaching how to make the best of our reasoning abilities. *Educational Psychologist*, 1–16.

Mercier, H. & Sperber, D. 2017. *The Enigma of Reason: A New Theory of Human Understanding*. London: Allen Lane.

Middendorf, J. & Pace, D. 2004. Decoding the disciplines: A model for helping students learn disciplinary ways of thinking. *New Directions for Teaching and Learning*, 98 (Summer):1–12.

Nussbaum, E.M. 2008. Collaborative discourse, argumentation, and learning: Preface and literature review. *Contemporary Educational Psychology*, 33:345–359.

Paul, R. & Elder, L. 2016. *The Art of Socratic Questioning*. Dillon Beach: Foundation for Critical Thinking.

Resnick, L.B., Asterhan, C.S. & Clarke, S.N. (editors). 2013. *Socializing Intelligence through Academic Talk and Dialogue*. Washington, DC: American Educational Research Association.

Ries, E. 2011. *The Lean Startup*. New York: Crown Publishing.

Scardamalia, M. & Bereiter, C. 2014. Knowledge building and knowledge creation: Theory, pedagogy, and technology. In: Sawyer, R.K. (ed.), *The Cambridge Handbook of the Learning Sciences*, 2nd edition, New York: Cambridge University Press, 397–417.

Schwarz, B.B. 2009. Argumentation and Learning. In: Muller Mirza, N. & Perret-Clermont A-N. (eds), *Argumentation and Education: Theoretical Foundations and Practices*, London: Springer, 91–126.

Walton, D. 2000. The place of dialogue theory in logic, computer science and communication studies. *Synthese*, 123:327–346.

Wentzel, A. 2017. *Creative Research in Economics*. London: Routledge.

Wentzel, A. 2018. *A Guide to Argumentative Research Writing and Thinking*. New York: Routledge.

Wilson, N.L. 1959. Substances without substrata. *The Review of Metaphysics*, 12(4):521–539.

Zwiers, J. & Crawford, M. 2011. *Academic Conversations: Classroom Talk That Fosters Critical Thinking and Content Understandings*. Portland: Stenhouse Publishers.

Chapter Seven

Transform Assessments into Learning Experiences and Eliminate Cheating

Paul Samuelson (1990, p.ix), the writer of an influential economics textbook, once said: "I don't care who writes a nation's laws … if I can write its economics textbooks." I will do him better, and add: "I don't care who writes a nation's textbooks, if I can set its assessments." Assessments encompass things like tests, exams, and assignments, and are much more than simply a way to evaluate students or gather information on how much they learned. In fact, this evaluation function is the least important job of assessment.

Once we start to recognize the other important functions of assessment, we shall see more learning in our courses. Indeed, studies on the best professors (such as Ambrose, Bridges & DiPietro, 2010 and Bain, 2004) agree that their students learn more because these professors place much less emphasis on assessments as a way to judge and coerce students.

This chapter shows that assessments themselves can be learning experiences, perhaps even better than lectures. It does so with the aid of ideas that have been around for over a century, but have not found their way into popular theories of teaching, unlike erroneous ideas like learning styles which have led to misguided approaches to assessment (see Box 7.1).

BOX 7.1 THE MYTH OF LEARNING STYLES

For some time now, evidence has shown that the idea that students have learning styles is a myth. Evidence can be found in studies such as Riener and Willingham (2010), Pashler *et al*. (2008) and Rohrer and Pashler (2012). Simply searching for **"myth of learning styles"** will yield useful information. As suggested in Chapter 4, there is some evidence that there are memory styles, but this is not the same as learning styles.

When done correctly, assessments provide accurate information about student learning and this can be used to promote future learning. But, seen like this, assessment simply *reacts* to learning; and this perspective does not support the claim above that assessments can be more powerful than textbooks.

What is often not recognized is that assessment *determines* learning, and this happens regardless of whether it is used appropriately or not. Our learning is determined by what we do and what we think about, not by what we listen to. What is remembered longer is rarely the content, but the memory of the ways of being that follow from what we are asked to do and think about.

Since it is during assessments that students engage in the most thought and action, it is here where students learn most. Students who simply have to memorize content for assessments learn that life is a trivia quiz and that success depends on finding out what authorities think and repeating their thoughts back to them. In contrast, assessments that call for reasoning and require students to engage with the contingency of knowledge, encourage them to participate in knowledge creation. This understanding of learning is often discussed under the banner of 'the hidden curriculum' and suggests that some of the most damaging and life-transforming ideas that students learn are the ones that are not explicitly taught (see Box 7.2).

BOX 7.2 THE HIDDEN CURRICULUM

The hidden curriculum consists of the unstated and hidden messages carried by the practices, behaviors, and structures in an educational setting. Simply search for **"hidden curriculum"** to obtain information. The communication theorist, Marshall McLuhan, summarized the idea succinctly in his aphorism "the medium is the message". Basically, everything said, not said, done, and not done, carries a message. Since these messages are unspoken, they are rarely refuted and hence quite powerful.

This is quite different from that of the view of those who say: "If I don't put it in the exam, students won't learn it." The fact is that when something appears in an assessment designed only to *evaluate* students, it is too late to expect learning. Very little learning happens under the pressure of evaluative assessments. Such assessments only tell you what students have *already* learned, and are better known as summative assessments (see Chapter 5). If we want to direct students'

learning through assessments, we need to carry out more than just summative assessments.

BOX 7.3 SUMMATIVE AND FORMATIVE ASSESSMENTS

For more information on the differences between the two kinds of assessments, simply type **summative versus formative assessment** in the web search box.

Summative assessments only determine learning to the extent that they influence students' expectations. If students expect memorization in a summative assessment, they will learn to pack as much information into their short-term memory as possible. Whereas, if they expect summative assessments to focus on the application of critical ideas, they will instead aim to *understand* what is important.

While summative assessment may determine learning, formative assessment goes even further. Formative assessment *is* learning – formative assessment when executed correctly, is almost indistinguishable from learning. As I explained in the Instructor Manual to Wentzel (2018), students actually learn comparatively little during my lectures, instead, the vast majority of their learning occurs during formative assessments. This only happens if the formative assessment forms part of an iterative cycle that involves timely, targeted feedback and further improvement.

If done in the wrong way, even formative assessment will misdirect or even sabotage learning. So, before turning to the practicalities of assessment in the next chapter, it is important to first understand the importance and power of assessment. This will not only ensure that you use assessment correctly, but tap into its potential – not simply as a way to evaluate learning, but as an *act of learning in and of itself*.

ASSESSMENT DETERMINES LEARNING

As professors, we sometimes tend to exaggerate the importance of our teaching. No doubt the quality of our teaching can sometimes make or break students, but as John Biggs explains, it is what our students do (not what we do) that has the biggest effect on learning.

However, we have to be careful here. Daniel Willingham (2009) refines this point by adding that it is what we expect students to think about that determines learning. If we ask students to do something, like making a video about the Cambrian period, we have to anticipate what

they will think about. We want them to learn about this geological period, but if they spend more time thinking about the technology of producing a video, they will learn very little about this period. The best way to get them to think is through assessments that direct their thoughts mainly to the ideas we want them to learn.

The reality is that students often just want to pass our courses, and for them, the most important question is: "Will this be in the test?" What students will learn is determined by what they expect to see in an assessment, and what they expect to see is determined by what professors *actually* require from them in tests. So, if you state that you expect students to apply their knowledge, but in assessments you mainly test memorization, students will study for memorization purposes regardless of what you say.

If we want students to study for understanding, we need to teach for understanding *and also assess* understanding. John Biggs calls this 'constructive alignment', which occurs if an assessment makes students think about what the teacher intends them to learn.

BOX 7.4 CONSTRUCTIVE ALIGNMENT

There are three excellent short videos about constructive alignment on YouTube. To find them, go to YouTube and enter: **teaching teaching & understanding understanding**. For further information search for **"constructive alignment" AND Biggs**.

Without constructive alignment, the intended learning will not happen. If the message of a professor in class is different from the message he sends through his assessments, students will listen to the message of the assessment, because it is the assessment that determines whether they pass or not. So, assessments, to a large degree, determine learning. We should, therefore, align our assessments with how we teach and with the kind of learning we want from students.

This alignment applies not only to the kind of thinking we expect but also to the content of the course. Understanding is more likely not only if we teach the critical ideas, but when we assess them too (Ambrose, Bridges & DiPietro, 2010). To ensure that understanding of the critical ideas occurs, it is better to assess less memorization and more reasoning, especially in the form of authentic assessments (as suggested in Chapter 5).

When constructive alignment exists, students no longer need to guess what will be in assessments. This means that assessments will no longer test how well students can read examiners' minds, instead, they will test what they are intended to test: understanding of the ideas.

ASSESSMENT IS LEARNING

We can go one step further and state that assessment is, in itself, a way to learn. To understand this and put it in practice, requires an awareness of a relatively unknown idea that has been around for decades.

Testing Effect

Roediger and Karpicke (2006), in a well-cited paper, are known for spreading the news about one of the most powerful teaching strategies that exist: the testing effect. In a recent reflection article (Roediger and Karpicke, 2018), they wonder why their article could have had such a big impact, given that the strategy had been known for about a century. It remains a puzzle, but an even greater one is why so few professors use it in their courses.

The testing effect shows that formative assessments have a bigger and longer-term effect on learning than teaching, revision or cramming – and it allows people to remember more and with less effort. In experiment after experiment, the testing effect has been confirmed. Here are some typical results (for more information, consult the sources in Box 7.5).

In a typical experiment, students are divided into two groups. Both are given the same science text. Group A is allowed to read it, take a break, and then read it again. Group B is allowed to read it once, then is given a short quiz on it. One week later, both groups are given the same test on the work. Which group did best? Group B always outperforms Group A. By how much? They usually obtain a mark that is 50% higher. The conclusion is that repetitive reading (what we mistakenly call 'studying') is not as good as being tested on what you just read when it comes to remembering something.

What if we allow students to cram, that is read a text as many times as they wish within a given time? In another experiment, Group A was allowed to cram, but Group B was allowed to read the text only once and then told to quiz themselves. In the test that followed immediately after the students finished, Group A did better (as expected). But when they were tested two days later, Group A only scored an average of 50%, while Group B scored more than 80%. The same conclusion follows – you can cram as much as you want, but it will not be nearly as effective as a test (even a self-quiz) in helping you remember. With cramming, people forget more and forget faster, so if you cram for the test and you are tested on it again sometime later, you will need to cram everything all over again. Cramming is painful, boring and inefficient; while the simple act of testing, or self-quizzing, connects the ideas to your memory with less suffering and keeps it there longer.

BOX 7.5 THE TESTING EFFECT

There is much research on the testing effect that explodes many myths about how we learn and the role of assessment in learning. For general summaries, search for **"testing effect"** on Google. Also consult the research by Roediger and Karpicke (2006) and that of McDaniel and co-authors (McDaniel *et al.*, 2011) as provided in the list of references. Authors like Lang and Willingham also provide good summaries (see the list of references). For more information about how to use the testing effect, search for **"using spaced repetition"**.

Figure 7.1 is derived from the results of Roediger and Karpicke (2006) and shows that even though there is always some forgetting, testing slows this down.

Figure 7.1 shows that the groups who read repeatedly (study-study) without being tested will, after five minutes, still remember around 80% of what they read, but after one year, will barely remember 20% of it (perhaps less). In contrast, groups that have been tested even only once on the work (study-test) will remember it for longer.

Testing only has a positive effect if it takes place while the students can still remember what they have read or heard. Testing should ideally take place after a lecture or after students tried to read with understanding. If students are tested only after lots of cramming, testing will have

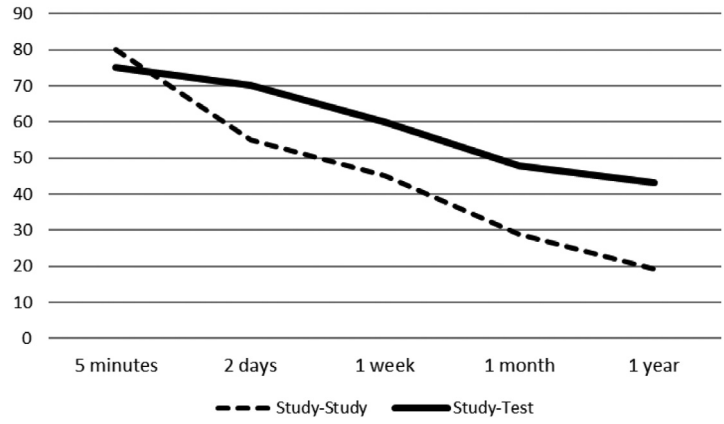

Figure 7.1 The Effect of Testing on Recall

Source: Derived and extrapolated from Karpicke and Roediger (2006)

little effect, especially if students know that those tests will only assess memorization. In such cases, they do not think about what they read, and after a year, they will remember so little that it would be as if they never studied that work before. Professors often complain about this, but it is easily solved: give students short quizzes on the work straight after a lecture.

Many believe that revision is important – either after a lecture or just before a test, but research suggests that testing produces better results than revision too. If you divided a class into two groups, Group A (who are taught the content followed by revision) and Group B (who are taught the content and then given a short quiz on it), then Group B will obtain an average mark that is more than 10% higher than Group A in the test a few days later. Experiments show that if you tested both groups eight months later, Group B's average would still be higher. In other words, if you have time after a lecture, don't use it to revise the work, rather give a short test. Perhaps, even better, *make* time to do this, because it will generate disproportional improvements in students' results.

The research by Karpicke and Roediger (2006) also demonstrates that the number of tests matter. Even one quiz on the work you just taught students will raise their future learning performance. If you test them repeatedly, you can strengthen the effect. They found that: students who read the same content four times in a row forgot 50% after a week; those who read it three times and then were tested once afterwards forgot 30%; but those who could only read it once, but were then given three tests in a row after reading, forgot only 15% after a week.

Most teachers don't have the time to ask students to write three tests in a row, and this is not the best way to use multiple testing anyway. As Figure 7.1 shows, you will still forget work you were tested on, even if it happens more slowly. But, if you space tests out over gradually longer intervals, you can almost reverse the forgetting, as shown in Table 7.1. Every self-test or other kind of repetition reverses the forgetting curve until you forget almost nothing. The table is a stylized version of the pattern.

Table 7.1 The Effect of Spaced Repetition

	Day 1 Test 1	Day 2 Test 2	Day 4 Test 3	Day 8 Test 4	Day 16 Test 5	Day 32 Test 6
Forgotten before the test	–	50%	35%	25%	15%	5%
Remembered after the test	100%	100%	100%	100%	100%	100%

Spaced repetition is why language learning programs like Pimsleur and the free Anki software are so effective, as Wyner (2014) explains. I can attest to it. I used to think that I had little talent for languages until I used spaced repetition to memorize the top 5,000 words in Spanish. This happened in less than a year, with surprisingly little effort and without the need for complicated mnemonic systems.

BOX 7.6 SOFTWARE FOR SPACED REPETITION

The simplest, most flexible and free spaced repetition software is Anki which you can download free of charge at https://apps.ankiweb.net/. Many web sources explain how to use Anki for language learning, but it works for anything where you need to remember facts. When you search for it, try to direct the search results toward your area of interest, for example by searching **using Anki for [insert your area]**.

The repetitions should happen at intervals, and as Table 7.1 shows, they should occur at increasingly longer intervals. Repetition should take place just as you are about to forget something (i.e. while it is still on the tip of your tongue). A repetition at this point is most powerful at cementing it in your memory. Difficult material should initially be tested at shorter intervals. If you wait until you've completely forgotten a fact, repetition is not so effective, and this is what usually happens with cramming. Students tend to cram for a test, and then forget, only to cram the same work again before the next test. The result is that almost nothing is remembered despite many hours of effort. Fortunately, spaced repetition systems are pre-programmed to help get your timing right.

Assessment is a kind of spaced repetition, but reading is not. Simply reading the content does not count as a repetition because it does not usually involve thinking. But assessments, if designed correctly, require thinking, and this is what activates the benefits of the testing effect.

If memory were a container or storage device, the testing effect would not work. But the storage view is wrong. Neuroscience has shown that memories are not found in specific places in the brain, but are distributed throughout the brain, and that remembering occurs by changing the connections between brain cells.

From this perspective, the testing effect makes a lot of sense. For a memory to be formed, a connection has to be created. At first, the connection is like a path in an overgrown jungle. Every time you think about the idea, you travel the path and make it less overgrown, and a

bit wider. After traveling the path a few times, the ground becomes more worn, the connection between points is secured and the memory remains. By forcing you to retrieve ideas, assessments make you travel along the paths (connections) that (re)create those ideas, and soon those connections are too strong to be lost again.

Since the testing effect improves one's memory of facts, some people may think that it would be of little use to promote understanding of ideas. However, Brame and Biel's (2015) review of the main studies on the testing effect reveals that it promotes understanding too. In fact, with questions that require inference or reasoning, the testing effect improves students' performance more than many other highly effective techniques, such as concept mapping.

How can something that improves the ability to remember facts also improve the understanding of ideas? Given that memory is necessary for understanding to occur (see Chapter 4), this is not so surprising. Memory of old ideas facilitates the forming of new connections, because a connection can only be formed if the old ideas are held in one's mind long enough to become aware of the connections between them.

While writing this, I remember an experience I had during my undergraduate studies. There was a subject, Monetary Economics, about which I thought I understood virtually nothing due to the impenetrable textbook and confusing lectures. Fortunately, I was able to obtain the class notes of a more diligent student, and using the memory techniques that I described before, I memorized these notes. I not only ended up obtaining a distinction for this subject but obtained the highest mark in the course. The reason for this was that the lecturer used to ask questions that required us to make connections between different parts of the work, and because I was able to remember so much of the work with little effort, I could hold it all in my mind while looking for connections.

Consequently, the other finding that Brame and Biel highlight is not so surprising: that those who benefit from the testing effect also learn new material faster. This is to be expected because memory makes it possible to recognize more connections, and so increases the pace of learning. All of this is confirmed by my own experience in terms of learning Spanish. Not only did spaced repetition help me to learn vocabulary with little effort, I also found that, after a while, I learned new words faster. After a certain point, I started to recognize the patterns in Spanish words and phrases, so new words were less strange to me than before.

A further advantage of the testing effect is that students get a more realistic view of their progress in the course much faster, instead of having to wait for the first high-stakes summative assessment, at which

point it may already be too late for them. Students who only read the material and are not tested on the material can easily fall victim to the 'illusion of knowledge' or the 'feeling of knowing' – something that will not happen if they are forced to retrieve and apply new material early on. Professors who let their students experience the testing effect also spot problems faster and are rated more positively by their students, partly because such students perform better in their subjects.

Before turning to how we put the testing effect to use, I want to emphasize that it is no use to have excellent explanations if students never think about the ideas explained. Without assessment, all the effort that a teacher puts into explanations is lost because students do not form strong connections. By putting the testing effect to use, not only is the effect of good explanations maintained, it is, in fact, amplified.

Using the Testing Effect

It is best to give short tests or tasks on new material directly after you taught it. These tasks should be completed while the ideas are still in the students' minds. Tasks or questions that involve writing and thinking are better than short-answer or recognition items, although any kind of assessment can activate the benefits of the testing effect. The tasks should be formative and low-stakes (i.e. they should count very little). This is because the purpose is not to see if students paid attention or how much they know, it is to make them retrieve the material that was taught and start the process of forming connections. The task can, therefore, be as simple as asking the students to reflect on what was taught.

The testing effect is even more powerful if students receive slightly delayed feedback. Ideally, students should not get feedback immediately, because the testing effect is more powerful while the ideas are still on the tip of your tongue. When students receive immediate feedback, they don't learn as much. Even so, feedback should be given while students can still remember what they were thinking when they made mistakes, and this will do a better job of embedding the ideas in their memory.

Initially, students will not like writing many short assessments, because assessments are more difficult than listening to an explanation. The difficulty experienced can actually be a good thing, and cognitive scientists call this 'desirable difficulty'. Robert Bjork's research (see Box 7.7) suggests that anything that slows down desirable learning, and makes it a little more difficult, results in students forming stronger connections. In fact, the experience of difficulty is often a sign of learning, while

something that is easy (like reading a text or just listening to a clear explanation) creates the 'illusion of knowing', and is almost always a sign that no long-term learning is happening. Assessments that are challenging will enhance learning *provided they are constructively aligned* and provided that students believe that the difficulties can be overcome. If a somewhat difficult assessment takes place after a clear explanation, it reinforces the benefits of the explanation, which would otherwise be lost.

BOX 7.7 DESIRABLE DIFFICULTY

For a comprehensive review of the idea of 'desirable difficulty' from Robert Bjork who coined the concept (with several videos), see: http://bjorklab.psych.ucla.edu/research.html. For a less technical overview, simply search for **"desirable difficulty" AND (classroom OR learning)**. Based on this principle, scientists at RMIT University in Australia recently developed a new typeface called 'Sans Forgetica' that help people to remember what they read for longer.

Given that students may resist having to write tests, we need to explain the testing effect to students and teach them how to tap into the power of the testing effect through self-quizzing. This is because the testing effect works regardless of who administers it and whether it is formal or informal.

Finally, since there is insufficient time in real classrooms to assess everything or have long assessments, it is important to apply the testing effect to those facts or ideas around which strong connections need to form. Tests should, therefore, be focused on getting students to think about the critical ideas and getting them to remember only those facts that are most important to memorize. Testing other ideas and facts would be a waste of the power of the testing effect.

Iterative Formative Assessment and Targeted Feedback

Any graphical representation of learning (like the Kolb learning cycle) shows it as being a circular and iterative process, not a linear one. It is safe to say then that very little learning takes place in courses with few assessments that are also summative. In such assessments, students often receive feedback that contains little information, such as a standard mark or grade. Usually, there is also no opportunity to submit improvements that incorporate the feedback. As a result, there is no incentive for them to think about the ideas again, and the learning ceases with the assessment.

In contrast, the testing effect is amplified in courses where learning is allowed to take its natural iterative form as far as the constraints of the institution allow. The key to transforming a linear form into a cycle is giving informative feedback with sufficient time for students to respond to such feedback.

Feedback is informative when it highlights not only required actions and strategies for improvement but also what the student did well. This kind of positive feedback should not be given merely to motivate students; it plays a much more important role. For students confronting new ideas, feedback acts like a GPS that not only tells them what not to do, but what to do more of. All of us will start getting nervous, confused and perhaps overwhelmed if a GPS only gave us information when we needed to make a change. We also need to know that we are still on track and if we should continue what we are doing (Ambrose, Bridges & DiPietro, 2010).

Paradoxically, feedback loses its informative value when there is too much of it. Feedback should focus on the critical aspects or ideas. It is better to give feedback in layers, almost like the gradual complexification method mentioned in Chapter 2. If you are able to distinguish between the critical and non-critical, you will find it easier to offer such feedback.

Eventually, one will want students to be able to generate their own feedback. To make this possible, it is useful to be very explicit about the goal of an assessment, why it is important and how it fits into the key objectives of the course. The criteria by which the assessment will be judged should also be as explicit as possible. Ambrose, Bridges & DiPietro (2010) suggest that you should be more explicit than you think is necessary. They also suggest giving students models: annotated models of good performance that show what you expect, and annotated models of what students should avoid.

Since learning is very much learning from mistakes, it is not a good idea to punish mistakes. After feedback, students should have at least one more opportunity to use it to improve their performance. This may take the form of having students re-submit the same assessment, or to create a series of assessments, where each assessment builds on the prior assessments.

To make learning iterative implies using as few summative assessments as possible, and relying mainly on formative assessments. Such assessments should not happen under exam-like conditions, but should be tasks that allow for thinking and can be completed under the least pressure possible. If not, one may create incentives for cheating, as explained next.

HOW ASSESSMENT ENCOURAGES CHEATING

Assessment is powerful in promoting learning, but at the same time, even more powerful in impairing or even blocking learning. This usually happens when it encourages students to cheat. When there is cheating, the intended learning does not happen.

I cheated several times in school and university, and every time I learned none of the things the teacher wanted me to learn. Unless you are close to being a saint, you have probably also cheated or plagiarized as a student. Maybe you had secret notes in a test, copied homework or plagiarized. Think about those times: What did they all have in common? What motivated you to plagiarize or cheat?

The reason for cheating, as Lang (2013) argues in his insightful book on the topic, is not so much that we spot opportunities to cheat, but rather that the learning environment is designed to encourage it. This means deterrence is an ineffective way to combat cheating in assessments and that the way professors approach assessment plays the biggest role in encouraging cheating.

No matter what happens, there is always learning in classrooms, but not always the learning that we intend (Box 7.2). When a professor teaches by simply reading from the textbook, the learners learn something very different from what was intended. The same thing happens when professors expect students to simply copy notes from dictation, from the board or from slides. And the same thing happens when assessments only require memorization. Every time, the same harmful message is conveyed, and students learn it very well.

What is the message? It is: "Please cheat and plagiarize!" This is because, in each case, the students are really just copying and pasting. When a professor just reads from lecture slides or the textbook, he is essentially saying it is acceptable to just to repeat other sources and not think for yourself. When a professor expects students to just copy notes, he is again saying it is not just acceptable to copy and paste, but actually required. And just asking for memorization tells students that copying and pasting are necessary for success – except that students should first 'copy' the information into their brains before 'pasting' it on the test script. In fact, with such assessments, the student who copies and pastes the best is the one who gets rewarded with the highest mark. With this message being reinforced by teachers, year after year, it is not surprising that we find so many students cheating and plagiarizing.

In the rest of this section, I draw extensively on the book by Lang (2013) about how we can change the way we assess students to make cheating less likely. Though I will discuss the main messages of his book (see Table 7.2), I cannot do it justice here, and therefore recommend that you read the book yourself.

Table 7.2 The Conditions That Promote Cheating and Learning

When cheating happens	When learning happens
Emphasis on performance (often a single one)	Emphasis on the process of learning
High stakes (make-or-break performance)	A progression of low-stakes learning opportunities
Extrinsic motivation (perform for others)	Intrinsic motivation (because I want to learn)
Low expectation of success	A realistic view of self and belief in one's ability to learn

Source: Derived from Lang (2013)

Emphasis

If you have only one or very few chances to prove that you have learned something, and all that matters is the outcome that you deliver at the end, then you are more likely to cheat. Cheating is common in most major sports events because all that matters is the final outcome (whether you win or lose) and there is often only one race or match that determines it. Make a single mistake, and you may lose. But learning is not a product, it is a process – a process that needs us to make mistakes and learn from them. Treating it like a sports event encourages the same widespread cheating we see in professional sports.

Professors who focus on the process of learning tend to have less cheating. They pay less attention to how well a student does in any single performance and place more emphasis on how well the student is progressing towards mastery. The goal is not a specific mark but a well-specified goal of mastery. Those who emphasize mastery will give students many assessment opportunities and also give them feedback on their progress after each opportunity. Such opportunities will be mostly formative assessments that aim to move the student closer to mastery and exploit the testing effect.

Emphasizing only a single performance, like a big exam, sends the message to students that only the final product matters – no matter how you get there, even if it is through cheating. However, emphasizing the process of learning sends the message that learning itself matters. In fact, the emphasis on mastery makes cheating quite useless because it is not any single performance that determines whether you pass or fail, but rather your proven progress towards mastery.

Stakes

Closely related to the issue of emphasis is what is at stake. When a single performance can make or break you, it is a high-stakes assessment. It is then a matter of a rational risk-versus-reward calculation that brings you to the point of cheating. If you cannot afford to fail, and if all that matters is whether you pass or not, then most of us will gamble to avoid failure. And the fewer assessment opportunities there are, the smaller the perceived risk because you only have to cheat once or twice to get away with it.

One can reduce the likelihood of cheating simply by lowering the stakes. If students have multiple opportunities, and none of them alone can cause failure, then cheating will decline for several reasons: (1) the reward for cheating is smaller because no assessment can break you; (2) the risk of cheating is greater, because you have to cheat several times in a row to get away with it; and (3) if the opportunities gradually build your skills toward a clear learning goal, you receive the message that learning is more important than performance.

I realized the importance of having multiple assessment opportunities when I taught a course in research writing skills. I used to expect students to produce a research proposal at the end of the course, and they had only one opportunity to submit it. It was a typical single-performance high-stakes assessment, so it was not surprising then, that despite many warnings of the dire consequences of plagiarism, many students plagiarized year after year, even when I used plagiarism detection software.

But then, one year, I changed my approach. I gave students 17 short assessments in the course of the semester. Except for the final assessment (the complete research proposal), none of the assessments counted for much. Their purpose was to give students the opportunity to build the necessary skills and gradually develop the content they needed to produce the final proposal. After each assessment, they received feedback that enabled them to improve in the next one. By just making these changes, the quality of their work not only improved dramatically, I saw significantly less plagiarism. And because of the many assessments, I was able to detect cheating much faster, before it was too late.

As in the case of my research writing course, sometimes you cannot avoid having at least one high-stakes assessment (like a final exam). In such cases, you should still have multiple low-stakes formative assessments that build the skills students will need in the final assessment. This gradual progression of skills will result in less cheating because students will have greater (and justified) confidence in their ability to pass without needing to resort to cheating. Whether such an expectation of success exists is the next determinant of cheating.

Expectation

A person who believes that their chance of success is small, and that they cannot improve their chance, will find it more rational to cheat, especially in the case of high-stakes assessments. The choice in such a case is either: (1) don't cheat and fail, or (2) cheat and have a chance to pass. Under conditions when a loss seems almost guaranteed, humans become *risk-seeking*, even if there is only a small chance of avoiding the loss (see Box 7.8). So, if you make assessments as difficult as possible and talk about them in terms of loss (failure) avoidance, you are inviting students to cheat.

However, if a student believes that they have the ability to overcome the challenges presented by the assessment, it will not seem necessary to take the risk of cheating. Through a series of low-stakes assessments that gradually build the skills of the students through informative and targeted feedback, students will start to develop an expectation that they can succeed in assessments, and, more importantly, develop the confidence in their own learning ability.

BOX 7.8 BEHAVIORAL BIASES

Though I was critical of the literature on biases in Chapter 6, they contain valuable ideas to make sense of behavior under conditions of uncertainty with rewards and costs. Search for **"loss aversion"** and **"risk aversion"**, and, for more technical detail, search for **"prospect theory"**. Kahneman (2011) is a good introduction, and an interesting application to cheating is Grolleau, Kocher & Sutan (2016).

It is important that students' confidence is justified. Having an overly low level of confidence may be more dangerous than over-confidence. Some students have an unrealistically low opinion of their ability and may be tempted to cheat even though they do not need to do so. If assessments are formative, multiple, and progressive, they help students to develop the right opinion of themselves, and this enables them to regulate their own learning. It helps to recognize early if your high confidence is unjustified, or if your low confidence is misplaced. And realizing that you have the ability to overcome meaningful challenges can be highly motivating.

Motivation

Not all kinds of motivation are equal. When we are motivated by things outside us (rewards or pleasing other people), we have extrinsic motivation. This kind of motivation is very fragile. With extrinsic motivation,

all that matters is the outcome – the final mark or others' approval – and again, if the outcome is all that matters, then the process of learning is perceived as less important. Extrinsically motivated students are therefore more likely to cheat.

Assessments play an important role in the kind of motivation students possess. It is possible to design assessments that promote intrinsic motivation, which is motivation that comes from within and connects to what matters to students. It is easier to be intrinsically motivated when you can use what you learn. In the first chapter, I mentioned that we can find 'critical ideas' by asking *why* someone would want to learn about the topic and how it would be useful to them. Teaching and assessing such critical ideas should create more intrinsic motivation than simply expecting students to repeat content from a textbook. Making ideas interesting, as discussed in Chapter 5, also promotes intrinsic motivation; especially the triggers of fascination (which connect to what matters to the student) and authentic assessments (which connect to what is happening in the world).

As Figure 7.2 suggests, our assessments will promote intrinsic motivation if they offer meaningful and authentic challenges for students. Such assessments make students think about the questions of the subject by connecting this to the questions in their own lives and questions in the real world.

CONCLUSION: THE TEACHER AS DESIGNER

Assessment is more than simple evaluation and measurement of student learning, in fact, that is the least important role of assessment. Used well, it can be the most powerful tool for determining learning, because assessment is, in itself, an act of learning.

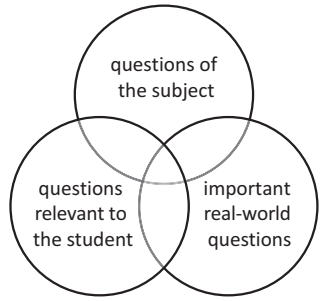

Figure 7.2 Assessments That Promote Intrinsic Motivation

It is professors, as the designers of assessments, who create an environment conducive to learning. Instead of fighting against the behavior of students, we should recognize, like Buckminster Fuller did, that, "You never change things by fighting the existing reality. To change something, build a new model that makes the existing model obsolete" (Ray, 2018, p.9). By changing our thinking about assessments, we can design a new model of learning in our courses and start thinking less like enforcers and more like designers (IDEO, 2012).

EXERCISES

1. Consider the assessments you have set for your students in the past.
   a) Would you say they were constructively aligned? Why or why not?
   b) What can you do to improve them so they become more aligned in the future?
2. Design three assessments that will take less than five minutes to complete, which you can use after teaching a topic in your next few classes to activate the benefits of the testing effect for your students.
3. Given the constraints of your courses, where are there opportunities to make learning in your courses more iterative?
4. Consider the assessments culture in your courses according to conditions that encourage cheating.
   a) Where do you think you may be encouraging cheating? What can you do about it?
   b) Where do you think you are encouraging honest learning, and what can you do to strengthen that?

REFERENCES

Ambrose, S.A, Bridges, M.W. & DiPietro, M.2010. *How Learning Works: Seven Research-Based Principles for Smart Teaching*. San Francisco: Jossey Bass.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Brame, C.J. and Biel, R. 2015. Test-enhanced learning: Using retrieval practice to promote learning. *CBE – LifeSciences Education*, 14:1–12.

Grolleau, G., Kocher, M.G. & Sutan, A. 2016. Cheating and loss aversion: Do people cheat more to avoid a loss? *Management Science*, 62(12):3428–3438.

IDEO. 2012. *Design thinking for educators*, 2nd edition [available at www.designthinkingforeducators.com, accessed 16 January, 2019].

Kahneman, D. 2011. *Thinking, Fast and Slow*. New York: Farrar Straus & Giroux.

Lang, J.M. 2013. *Cheating Lessons*. Cambridge: Harvard University Press.

McDaniel, M.A., Agarwal, P.K., Huelser, B.J., McDermott, K.B., & Roediger, H. L. 2011. Test-enhanced learning in a middle school science classroom: The effects of quiz frequency and placement. *Journal of Educational Psychology*. 103(2):399–414.

Pashler, H., McDaniel, M., Rohrer, D. & Bjork, R. 2008. Learning styles: Concepts and evidence. *Psychological Science in the Public Interest*, 9(3):105–119.

Ray, A.L. 2018. *Fulfilling the Needs of Teachers: Five Stepping Stones to Professional Learning*. Lanham: Rowman & Littlefield.

Riener, C.R. & Willingham, D. 2010. The myth of learning styles. *Change: The Magazine of Higher Learning*, August, 33–35.

Roediger, H.L. & Karpicke, J.D. 2006. Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*. 17(3):249–255.

Roediger, H.L. & Karpicke, J.D. 2018. Reflections on the resurgence of interest in the testing effect. *Perspectives on Psychological Science*, 13(2):236–241.

Rohrer, D. & Pashler. H. 2012. Learning styles: Where's the evidence? *Medical Education*, 46:634–635.

Samuelson, P.A. 1990. Foreword. In: Saunders, P. & Walstad, W. (eds), *The Principles of Economics Course: A Handbook for Instructors*. New York: McGraw-Hill Publishing, pp. ix–x.

Wentzel, A. 2018. *A Guide to Argumentative Research Writing and Thinking*. New York: Routledge.

Willingham, D. 2009. *Why Don't Students Like School?* San Francisco: Jossey-Bass.

Wyner, G. 2014. *Fluent Forever: How to Learn Any Language Fast and Never Forget It*. New York: Harmony.

Chapter Eight

Create Valid Conventional Assessments but also Consider Alternatives

While the previous chapter emphasized assessment as learning, this chapter focuses on the less important dimension of assessment as the measurement of learning, which is mostly the focus in practice. It approaches assessment from the perspective of validity, especially considering how validity may be harmed. Finally, it suggests that the biggest threat to the validity of assessments is the traditional approach to assessment itself, and that we should seriously consider alternatives.

Throughout the chapter, I draw extensively from Miller, Linn & Gronlund's (2009) superb book on assessment. Since I can only touch on very few of the issues in this very comprehensive book, I highly recommend buying this book, which is already in its 12th edition.

WHAT TO ASSESS (OR NOT)

What you intend for students to learn is also what you should teach, and this, in turn, should determine what is assessed. If an assessment accurately measures the intended learning, and only the intended learning, the assessment is said to be 'valid'.

Unfortunately, many things can cause an assessment to either be an inaccurate measure of the intended learning or to measure something other than the intended learning. For example, if the assessment is not constructively aligned, the result will reflect something else besides the intended learning. If the assessment focuses exclusively on memorization it will be invalid because the results will not reflect the degree of understanding. If the assessment focuses only on one or two ideas, or on the non-critical ideas, it will not reflect students' understanding of all the critical ideas.

One of the simplest ways to enhance validity is to create a 'table of specifications' to guide you when designing a complete assessment. Such a table ensures that you spread the required tasks across all important

Table 8.1 The Table of Intended Learning

Topic	Understanding	Memorization
Topic 1	List the 2–3 main critical ideas of Topic 1 that need to be understood here	List the ideas worth memorizing (foundational ideas, key details and important over-complex ideas as explained in Chapter 4) from Topic 1
Topic 2	2–3 critical ideas of Topic 2	The ideas that are worth memorizing from Topic 2
...	...	...
Topic *n*	2–3 critical ideas of Topic *n*	The ideas that are worth memorizing from Topic *n*

topics and cognitive skills, and give each topic and skill the appropriate weight. This way, the result of the assessment is more likely to measure whether students learned what the teacher intended for them to learn.

A table of specifications is done in two steps. First, construct a table of intended learning as shown in Table 8.1.

From the table of intended learning, one derives the table of specifications with its various totals. A hypothetical example is shown in Table 8.2.

To produce Table 8.2, decide how much the assessment will count (in this case, 50 marks) and how much weight you want to give each topic. In this example, Topic 1 may have been a very short topic, so it only gets a weight of 20% (10/50), and Topic 2, as the most important topic, is assigned a weight of 50% (25/50). It should add up to 100% of the total mark.

Next, decide how much weight to give each cognitive skill. How you break up memorization and understanding into different cognitive skills is up to you. In Table 8.2, I keep it simple and just have two skills: memorization and understanding. Memorization gets a weight of 40%

Table 8.2 Incomplete Table of Specifications

	Memorization	Understanding	Total
Topic 1			10
Topic 2			25
Topic 3			15
Total	20	30	50

and understanding a weight of 60%, so it again adds up to 100% (or 50 marks).

Others might prefer breaking the cognitive skills into the levels of Bloom's taxonomy, Wiggins and McTighe's six facets of understanding, or Biggs's Solo taxonomy. If you prefer one of these, you can extend the table accordingly (see Box 8.1).

BOX 8.1 COGNITIVE SKILL LEVELS

There are countless sources on the web that explain the levels of cognitive skills in Bloom's taxonomy and the SOLO taxonomy. Just search for: **Bloom's taxonomy** or for **SOLO taxonomy** in Google, or even better, do it in Google Images. For the six facets of understanding, simply search for: **"six facets of understanding"**.

Next, think about how you want to distribute the kinds of questions within each topic and so complete the table. Ensure that all the marks add up to the totals. If the numbers do not add up, change the weights slightly until it works out. Table 8.3 shows the final product.

The table of specifications guides you when setting the complete assessment. For example, I will look at the critical ideas of Topic 2 and create tasks that test the understanding of one or more of those ideas and that count for 14 marks. I will similarly set tasks for the other cells. It may not always be possible to stick to the exact marks as planned, but that is fine as long as you remain consistent to your intention.

To generate ideas of the kind of questions to ask under each cognitive skill, you can look up the 'action verbs' for whichever categorization of cognitive skills you choose to use. Box 8.2 provides more information.

Table 8.3 Complete Table of Specifications

	Memorization	Understanding	Total
Topic 1	7	3	10
Topic 2	11	14	25
Topic 3	2	13	15
Total	20	30	50

BOX 8.2 ACTION VERBS USED IN TASKS

There are many sources on the web that give the action verbs for the different levels. Search for: **"Bloom's taxonomy" AND "action verbs"**, and you can do the same with the six facets of understanding and the SOLO taxonomy.

To keep it simple, I use action verbs that are derived from the chunking technique in Chapter 3 (see Figure 3.2).

From this technique, one can derive many appropriate action verbs such as those given in Table 8.4.

These are only some of the verbs that can be used to generate ideas for tasks that assess understanding. Box 8.2 will lead you to more.

All this ensures validity by keeping an assessment focused on the intended learning and avoiding an overemphasis on certain topics or memorization. There are many other threats to validity, such as giving students too little time to complete the assessment, using language that is too difficult or asking unclear questions. If the result is significantly influenced by factors that do not reflect understanding – like spelling, grammar, and neatness – the assessment will also not be valid. There is insufficient space to address all of these threats, but they receive extensive treatment in Miller, Linn and Gronlund (2009). The next sections will deal only with two major threats: matching the wrong items to different cognitive levels and designing assessment items incorrectly.

MATCHING WHAT YOU ASSESS TO HOW YOU ASSESS

There are two main kinds of test items: objective test items and complex achievement items. Objective test items allow for no freedom of

Table 8.4 Action Verbs to Assess Understanding

Part-whole	Dis/similarity	Reason-action	Cause-effect
Classify	Compare	Demonstrate	Interpret
Identify	Contrast	Explain	Solve
Categorize	Illustrate	Apply	Test
Distinguish	Give examples	Justify	Discover
Divide	Relate	Motivate	Dis/prove
Analyze		Theorize	Conclude
Elaborate		Recommend	Adapt

response, while complex achievement items require interpretation and, sometimes, creativity from the student.

Objective test items have a clear and unambiguous correct answer. They include: alternate choice items – where the answer is one of two options such as true/false; completion or short-answer items – where a blank space is filled in with an exact answer; matching items – where a word or statement in one column is matched with only one other word or statement in another column; and multiple-choice items – where a range of answers is given but only one is correct.

Objective test items work best when assessing the lower level cognitive skills of Bloom's Taxonomy (remembering and comprehension) and are unsuitable to assess higher levels (evaluating and creating). This is because objective test items restrict freedom: the range of answers is predetermined and there is usually only one correct answer per item. Trying to test higher cognitive levels with objective test items will produce invalid results. So, when planning an assessment with a table of specifications, use objective test items for the marks allocated to memorization and the lowest level of understanding.

Complex achievement items have no single correct answer and the full range of possible answers cannot be determined in advance. In a written assessment, these would include data/constructed response items (where students are asked to interpret information) and essay items. Data response items would contain mainly restricted response questions, where very clear guidance is given about the kind of answer that is expected, while essays are extended response questions that are open-ended and allow for the greatest freedom of response.

Data response items are suitable for assessing the middle levels of Bloom's taxonomy (applying and analyzing), and essay questions are best suited to higher levels such as evaluating and creating. Using essay questions to assess low cognitive levels is a waste of their potential. If you are going to test memorization, it is much more efficient and accurate to use objective test items.

Complex achievement items need not be restricted to conventional assessment situations. They can also include learning portfolios, research projects, posters, other kinds of performances in the class or productions outside the class that require the use of higher cognitive levels.

Validity can be harmed by overusing one kind of item. Because objective test items are easy to score, it is tempting to overuse these and so limit the range of cognitive skills that can be assessed. It is tempting to overuse essay items because they are faster to design, but this results in using essay questions to assess lower-level cognitive skills that are more accurately tested by objective test items.

DESIGNING OBJECTIVE TEST ITEMS

While objective test items are easy to score, it is difficult to design good ones. Badly designed objective test items will confuse students and not accurately reflect learning. This section explains the use and design of the three most common objective test items. I cannot do justice to them in the space I have here, so please refer to Box 8.3 for more information.

BOX 8.3 DESIGNING TEST ITEMS

There are many good textbooks on designing assessment items such as Miller, Linn and Gronlund (2009) and later editions, as well as several good web sources, which can be found by searching for: **"improving your test questions"**.

Short Answer Items

Short-answer items are questions that can be answered with a word, number, symbol or short phrase. According to Miller, Linn and Gronlund (2009), they are best used to test knowledge of the following: terminology ("What do you call a …?"); facts ("Which principle is used to …?"); principles ("What will happen if …?"); or procedures ("What do you use for …"?).

Besides knowledge, they can be used to test interpretation of data by providing students with information and following it up with a series of short-answer items. For example, one can provide a table or case study with data, and then ask students to calculate the values of certain variables. I might give students the annual values of a price index and then ask them: "Calculate the inflation rate for 2018." In the case of calculation questions, it is usually best to indicate the unit ($, kilogram etc.) in the blank space where students have to write the answer. For example, after asking them to calculate the inflation rate, I add a blank space like this: "_____%" to indicate the unit of measurement is a percentage. When you have several short answer questions, ensure that the blank spaces where students fill in the answers are of equal length.

Avoid questions where there may be degrees of correctness – there should be a single, unambiguously correct answer. For this reason, Miller, Linn and Gronlund (2009) advise that one avoids items that involve completing the sentence, such as:

A document showing a family's income minus expenses is a _____

Complete-the-sentence questions can have many right answers. The expected answer is 'budget', but the student could also write other answers that are not as correct, but still make sense (such as 'good thing' or 'spreadsheet').

Questions that involve supplying the answer are better and do not allow for a range of correct answers, so the question is better stated as:

What does one call a document which shows a family's income minus expenses? _____

While short-answer items are easy to construct, they have limitations. They are suitable for assessing memorization and not effective for questions with many degrees of correctness. They can be used to assess basic application and interpretation, but not to assess more complex learning.

Alternate-Choice Items

Alternate-choice questions are those with two possible answers, like true-false, yes-no, fact-opinion, and many more. Below are three such items.

True	False	Someone's budget tells you the net worth of that person.
Fact	Opinion	It is more important for a family to try to control its expenses than its income.
Yes	No	Will a person who keeps to a strict budget incur rising debt?

There are variations that involve three options, such as True-False-Opinion or Correct-Incorrect-Not Known. One may combine alternate-choice with short answer questions, and a common variation is true-false with a reason. An example is shown below:

Indicate if the following statement is true or false by circling the correct option.
If the statement is false, explain why you think it is false.

True False Income minus liabilities gives you net wealth

Explanation:

Another less common variation is:

Indicate if the following statement is true or false by circling the correct option. Only if the statement is false, replace the underlined word with a word that would make the statement true, and write this word in the empty box:

True	False	Income minus liabilities gives you net wealth
True	False	Rent is an expense to a business

The result, if correctly answered would look like this:

True	(False)	Assets	Income minus liabilities gives you net wealth
(True)	False		Rent is an expense to a business

Alternate-choice questions are easy to construct but are most appropriate for assessing recognition of correct ideas, not understanding, except perhaps for the variations that are combined with short answers. One can cover much content with alternate-choice questions, but they sometimes enable students to pass by guessing, and that can harm the validity of an assessment.

There are some dangers to validity to avoid. Do not follow a pattern in the answers to make it easy to score. In such a case, the result would not reflect the students' learning in the subject, but their ability to recognize patterns. To make guessing harder, true statements should not be longer than false statements. Sometimes true statements contain more words and this immediately gives away the answer. Also, do not simply copy statements directly from the textbook and turn them into alternate choice items. This promotes mere memorization of the textbook and reduces the thinking that is required.

To ensure clarity, keep the statements as short as possible. Long statements can be confusing and contain more than one thing that could be true or false. In such a case, students will not know what they need to respond to, and the validity of the item then declines because students have to guess which idea you are testing. For example, a poor item is the following:

True	False	Wants are greater than needs, because wants can be created through advertising.

In this item, it is not clear if the teacher is testing whether students know the truth of the statement "wants are greater than needs" or of the statement "wants can be created through advertising".

Questions can be very confusing if the statements are negative. For example, which statement do you find easier to understand: (1) is it true or false that it is not unusual for firms not to raise their prices when consumers are not insensitive to the prices? Or (2), is it true or false that firms usually drop their prices when consumers are sensitive to the price? The negatives in the first statement made it difficult to understand, while the second statement is easier. So, avoid this:

True	False	Wants are not greater than needs

Rather restate it in positive terms like this:

True	False	There are fewer wants than needs

If a student is confused by a question, it does not test their actual learning, but rather their ability to read the examiner's mind, and this makes the result invalid. As is evident then, alternate-choice questions may seem easy to design, but they are full of pitfalls.

Multiple-Choice Items

A multiple-choice question consists of a stem, followed by a response. The response is a list of alternatives to select from, with only one alternative being correct and the remaining ones being distracters. This is shown below:

1. This is the stem followed by the response which consists of a:
   a) distracter
   b) distracter
   c) distracter
   d) correct alternative

Multiple-choice items are similar to alternate-choice items which also have a stem, but only two alternatives. To eliminate the benefit from guessing, multiple-choice questions should ideally have four alternatives, or perhaps five if you like to make extra work for yourself.

When designing stems, it is best if the stem makes sense by itself and is not simply a word or a short phrase. In fact, multiple choice items are

clearer when the stems are longer and the responses shorter. The longer the responses, the easier it is to become confused while answering the question.

Again, avoid negative stems, since they can be confusing. There should be no pattern in the correct responses. Examiners often favor the third alternative, which means that if a student does not know the answer, it is a good strategy to always guess the third alternative, and this harms validity.

Ensure that responses are more or less equal in length and that the correct response is not obviously longer. In the question below, the examiner puts so much into the correct response, that (c) stands out as being correct:

---

2. The experimental studies of strobocophenamine found that:
   a) Everyone exposed to it becomes anxious
   b) No one treated with it becomes anxious
   c) The chance of menopausal women suffering from long-lasting anxiety after a low-dosage treatment is high compared to alternatives
   d) It is a beam of light in a club

This item also suffers from a further problem – alternative (d) is absurd. Maybe the examiners wanted to be funny, but the result is that students can eliminate (d), and with the relevant alternatives reduced, guessing becomes more profitable.

Stems should be grammatically consistent with the response. If not, students can see the correct answer without needing to know the material. For example, if you use the verb 'are' in the stem to refer to something in the response, then all the alternatives should be plural. Here is another way in which the stem gives away the answer:

---

3. The percentage change in CPI from year to year is called an:
   a) Price growth rate
   b) Inflation rate
   c) Gini coefficient
   d) Balance of payments deficit

The stem, in this case, ends with 'an' as the article, which tells you the correct response starts with a vowel, that is, alternative (b). This item tests students' knowledge of grammar, not their knowledge of economics, as intended.

Furthermore, check that all responses are similar. One mistake, for example, is to state all the distracters in absolute terms but to qualify

the correct alternative. This happens in the next question where the correct response is obviously (d).

4. In his research on nonhuman primates, Jones discovered that the outcome of attachment therapy was that:
   a) A change in attitude was always seen
   b) All of them reacted with anger
   c) They were never violent
   d) They tended to isolate themselves

Just one more: look at how the responses follow from the stem. Sometimes the distracters do not follow so well from the stem, because examiners usually write the correct alternative first and do not pay as much attention to the distracters. Take the next question as an example – alternative (c) reads better grammatically than any of the other alternatives, as it follows best from the stem, and thus stands out as the correct one.

5. The effect of strobocophenamine on studying is:
   a) Eye movements are associated with reading about 80% of the time
   b) Reduces median performance
   c) That strobocophenamine treatment helps to increase concentration
   d) Strobocophenamine results in a lower impact of future treatments

There are so many ways to make mistakes that give away the answer that it would take several more pages for me to expound on these. Box 8.4 contains more information on the kinds of mistakes that allow students to pass multiple choice questions through intelligent guessing rather than learning.

BOX 8.4 PASSING MULTIPLE-CHOICE TESTS WITH INTELLIGENT GUESSING

Bad multiple-choice tests can be passed with good guessing when teachers make certain mistakes. Search for: **intelligent guessing on multiple-choice exams** for more strategies used to out-guess examiners who set poor questions.

Multiple-choice items are used to assess recognition of the correct ideas and are best used to test knowledge, terminology, facts, principles, and procedures. With a little extra thought, you can use them to test the application of knowledge, specifically to:

- Identify applications of ideas ("Which of the following is an example of…?")
- Interpret cause-and-effect relationships ("If expenses exceed income it will lead to…")
- Justify actions or methods ("Why do governments create…?")

Multiple-choice items have more pitfalls that can harm their validity than other kinds of objective test items. But, if well-designed, they can be highly effective. For example, because of the wider range of responses, a teacher can better pinpoint what exactly students struggle with, provided that the distracters capture common mistakes.

BOX 8.5 MORE INFORMATION ON OBJECTIVE TEST ITEMS

There is more to the design of objective test items than discussed here. The web contains many good sources with guidelines, which you can find by searching for: **manual for constructing objective test items**. There are also some interesting sources that are critical of objective test items, which you can find by searching for anything like this: why multiple choice questions suck.

DESIGNING COMPLEX ACHIEVEMENT ITEMS

Complex achievement items assess higher cognitive levels and they range from interpretive exercises (a stimulus followed up with objective response items) to data response and essay questions. Authentic assessments usually consist of complex achievement items. They do not cover as much content as objective test items, and marking them involves subjective judgment. In this section, I will only focus on data response items and essay questions.

Data Response Items

Data response (also known as 'constructed response') items are open-ended questions in response to an authentic stimulus, such as a diagram, cartoon or case study (like a newspaper extract). The questions are usually restricted response ones, so they should give clear guidance on what

is required. They should also assess cognitive skills from the middle levels of Bloom's taxonomy, namely, application, interpretation and analysis. Do not ask questions that test memorization (use objective test items for that) or ask for wide-ranging opinions (use essay questions for that, if you have to).

Data response questions should be short – they could ask for a calculated value or a sentence, but not more than a paragraph. Questions should focus on connecting the knowledge in the textbook to the given stimulus.

It should not be possible to answer the questions by simply repeating what is already in the stimulus or by repeating textbook content from memory. A data response item is not an English comprehension test, so if your stimulus is a text, do not ask questions that relate only to the stimulus – students should be required to *apply* their subject knowledge to the stimulus.

It should also not be possible to answer the question without the stimulus, because that means the stimulus is redundant. If a question can be answered from the textbook alone, it should not be part of a data response item.

Take the next example of a short stimulus text in economics.

Lower Interest Rate Unlikely This Year

Inflation in South Africa is likely to rise in the next six months, while spending will continue to rise as well. This will make almost impossible for the central bank to cut the interest rate this year.

So far this year, the central bank kept the interest rate constant, despite the fact that many labor unions negotiated above-inflation wage increases. Economists predict that the interest rate will remain unchanged, while inflation will increase to 5.9%. If this forecast is correct, the inflation rate will come dangerously close to the upper range of the central bank's target of 6%.

Economist Kganya Kgare stated that: "The recent depreciation of the South African currency will put upward pressure on inflation." He believes that the prospects for economic growth will be limited this year. "The economy will remain stable for now… but growth will likely pick up by next year, at which point interest rates will start to rise", he said.

The following three are poor data response questions:

1.1 Define inflation
1.2 List five factors that influence inflation
1.3 What does 'upper range' mean?

Questions 1.1 and 1.2 can be answered without referring to the stimulus: they can be answered from memory (or the textbook) alone, so there is no reason for the stimulus to be there. Question 1.3 does not assess understanding of Economics, but rather English vocabulary, so it will not produce valid results. If a question does not require application of subject knowledge, there is no reason for the question to be there.

The following three questions are preferable. All three require that the student read the extract, understand it and then apply their knowledge to it when answering the questions:

1.1 What will the central bank do if the inflation rate is expected to exceed the "upper range" of its target?
1.2 How would a "depreciation of the South African currency" help to keep inflation high?
1.3 What is the most important reason why the central bank cares about wage increases?

Since data response items assess complex achievement, there will be several correct answers, so you need to know in advance what you will look for in the answers you receive. Rubrics are useful for this purpose (see Box 8.6). If you have several complex questions, design a different rubric for each one.

BOX 8.6 RUBRICS

The Cult of Pedagogy (www.cultofpedagogy.com) is a wonderful resource and has some good explanations about rubrics. On the site, you can find articles on different aspects: holistic, analytic versus single-point rubrics, rubric codes and constructing rubrics.

Essay Questions

Essay questions are suited to assess the highest levels of cognitive skills because they are open-ended questions that often allow for original thinking and reasoning. Their power should not be wasted by asking students to simply repeat things from memory alone. In essay questions, you want to see if students can make the connections between different ideas, if they can think about and organize information, and whether they can communicate the ideas. There is never only one correct answer for a good essay question and many degrees of correctness. In fact, it

should not be possible to know in advance exactly what answers you could receive.

An essay question could be a series of restricted response questions that build on each other, or just a single extended response question. What you choose depends on your purpose. You would, for instance, use a series of restricted response questions if you wanted to cover more content and give students clearer guidance on what you expect. Or, you would use a single extended response question to give students the maximum freedom to select the information they think is relevant, organize it in the way they think is best and make their own connections. Obviously, the single extended response question allows for more creativity and freedom.

Here is an example of a poor essay question that is a single extended response question:

List and discuss the advantages and disadvantages of inflation targeting.

This question suffers from at least two weaknesses. It can be answered by just repeating ideas from memory as they appear in the textbook. Also, the word 'discuss' is too vague and does not indicate what is expected. Without clear guidance, students will assume 'discuss' means 'just write down everything I can think of and can remember from the textbook'.

Here is an example of an improved version of the same essay question:

You are the CEO of a reputable economic consulting firm in South Africa. You have noticed that the Minister of Finance is becoming increasingly unsure about whether inflation targeting is appropriate for an emerging market like South Africa. He has called a press conference to announce whether South Africa will continue with inflation targeting or not. Last week, he called you for advice.

Write a memo in which you advise the Minister on whether the country should continue with inflation targeting. In your memo, summarize the arguments on both sides (for and against) so that the Minister can see the different views. You need to help the Minister choose a side, so ensure that you refute the arguments of the side you disagree with. At the end of your memo, you need to reach a clear conclusion and recommendation based on your discussion. Since you know the Minister quite well, start your memo with 'Dear Tito' and end it with 'Your trusted advisor' and sign your name.

(30 marks, 45 minutes)

Extended response questions need not be as long as this one, but notice that students receive some explicit guidance. This ensures that we assess students' understanding, not their ability to guess what the examiner had in mind. An added benefit is that it is a more authentic task. Students can decide which information they will use and are not told what conclusion to reach.

The same question could have been asked as a series of restricted response questions, as seen below. This gives even more explicit guidance to students.

---

1. Recently, prominent South African politicians have expressed doubt about whether inflation targeting is appropriate for an emerging market like South Africa. Based on this, answer the questions below:
   a) Why is South Africa classified as an emerging market? (4)
   b) What is, in your opinion, the strongest argument for an emerging market to use inflation targeting? Explain the argument and also explain why you think it is the strongest argument. (10)
   c) What is, in your opinion, the strongest argument against an emerging market to use inflation targeting? Explain the argument and also explain why you think it is the strongest counter-argument. (10)
   d) Should South Africa continue using inflation targeting? Reason toward a conclusion by weighing your arguments in (b) and (c) against each other. (6)

(30 marks, 45 minutes)

It is usually a good idea to indicate the marks and the time students should spend on a question, as done above. This tells students about how much you expect of them.

Do not finalize an essay question until you have constructed the marking rubric. After constructing the rubric, revise your question, align it with the rubric by removing any problems you noticed. Then use the rubric to communicate your criteria – the better students know what is expected, the better the validity of the question. Another way to enhance clarity is to use very specific and active verbs (see Box 8.2). If there is any lack of clarity or confusion, once again, students' results will not reflect their learning, but how well they were able to guess what you were thinking. Verbs like 'discuss' or 'examine' tend to be too vague unless you provide very specific follow-up instructions.

While essay questions assess higher cognitive skills very well, marking them can be quite challenging and time-consuming. Rubric codes (mentioned in Box 8.6) can speed up the scoring and grading, but the

potential for subjectivity remains a problem. How to deal with this is one of the things I address in the next section.

BOX 8.7 DESIGNING ESSAY QUESTIONS

On the web, there is a detailed guide on designing essay questions, written by Christian Reiner, Timothy Bothell, Richard Sudweeks and Bud Wood. You can find it by searching for: **preparing effective essay questions Reiner**.

BEFORE, DURING, AND AFTER THE ASSESSMENT

Once the individual assessment items are designed, the rest of the process needs to be executed, starting with putting the items together in one assessment. A simple rule when doing this is to start with the easiest and simplest items, and end with the more complex ones. If the first item is difficult, students may get confused, panicked and discouraged at the start, and their mark will not reflect what they actually learned.

Before printing or distributing the assessment, it is best to leave it for at least one or two days and then answer the assessment as if you were a student. This serves two purposes: (1) you work out what you expect from students; and (2) you will pick up mistakes and places where the items are unclear or not focused on what you expect from students. Correct these items, and, if you can, give it to someone else to check.

During the assessment, interrupt students as little as possible. An interruption may break students' concentration at critical times and can harm the validity of the assessment. Avoid giving hints to students who ask about items, as this introduces unfairness. If one student gets a hint during an assessment, all students should receive the same hint, but doing so will mean interrupting all students, so it is best to avoid it unless absolutely necessary.

Once the students complete the assessment, it is time to mark or score them. It is easy to do so with objective test items but complex achievement items are more difficult and subjective. Marking can become very unreliable if we are influenced by the identity of the students, so it is best to conceal student names when marking. You may also be influenced by your mood or the time of day, so try to mark each question in one sitting to minimize the influence of external changes. Even when I do this, I find my marking changes over time, so I try to overcome this in two ways: (1) after marking a question, I go back to the first few

papers and re-mark them because I tend to be either too strict or too lenient with the first few papers before determining a baseline; and (2) I shuffle student papers in a way that I don't mark them in the same order with every question. The influence of other things – like handwriting, spelling or neatness - is more difficult to eliminate. One way is to deal with this is to discuss the rubric or marking memorandum with students and have them check their marks against it.

Most of the time, marks are expressed as numerical values converted into percentages or letter grades. There are alternatives to such grading, and some of them are mentioned later (also see Boxes 8.9–8.11).

After marking, you need to reflect on the results to firstly get a sense of where students are struggling, and secondly, where your items may have been confusing. For this purpose, an item analysis is most useful. Item analysis can be as simple as checking which items most students got wrong and which ones most got right (see Box 8.8 for more information). If the results were different than expected, then look at how you designed the items and determine if you made mistakes that harmed the validity of those items.

BOX 8.8 ITEM ANALYSIS

For more information on item analysis, search for classroom assessment item analysis and you can also find a good article on the Cult of Pedagogy website if you search for: **how to make a quiz work harder for you**.

ASSESSMENT ALTERNATIVES

Traditional written tests and exams are not the only kinds of assessments and they rarely reflect the way learning is assessed in the real world. Any difference between the conventional assessment of learning and the way it is done in the real world may threaten validity, regardless of how well the assessment is designed. Fortunately, there are many alternatives that are more authentic. This section mentions only some.

Portfolios

In the real world, we do not sit down and write down everything we know. Rather, we build targeted skills and knowledge through a variety of authentic tasks. A learning portfolio tries to simulate this.

When students create a portfolio, they produce a collection of artifacts and reflections that are organized around a learning purpose. The artifacts need not all be written/typed; they can also be electronically-based, objects made or found, photos or recordings of events, or their own performances. The portfolio is not the learning itself, but the physical remains of the learning experience and this is what the teacher assesses.

To guide the students, a professor assigns tasks that prompt students to produce the artifacts over time. It is not necessary for students to complete all tasks, but whichever tasks students are allowed to select should show development and progress towards the learning purpose based on the critical ideas. Every artifact should, therefore, be dated and students should be required to periodically present their progress for assessment. This is done to avoid a situation where the student tries to complete all tasks the night before the deadline.

As mentioned before, we only learn that which we think about. With portfolios, you do not want students to think about how they make the portfolio. Rather, you want them to think about the ideas in the course, and also about their own learning. Students should be asked to reflect on each task and again at the end of the process. Their reflection can be guided by questions such as:

- What have I learned in completing this task? When did I learn most often?
- When, how, and why has my learning surprised me?
- What have been the highlights of my learning? The disappointments?

The process of creating a learning portfolio project could be as follows:

1. Identify the topic on which you want students to produce a learning portfolio
2. Draw up a planning table as follows:
   - What are the main understandings and/or skills you want learners to acquire? List the main critical ideas or skills in the first column of a table such as Table 8.5.
   - For each skill or idea, what kinds of evidence would prove that the learning took place? What possible tasks could students do to produce that evidence? It is not necessary that students complete all these tasks because you might want to give them the option to choose one per category.
   - What questions for reflection will you ask them after completing the task/s for each idea or skill?

Table 8.5 Learning Portfolio Planning Table

Understanding or skills	Evidence	Reflection
Critical idea$_1$ or skill$_1$	List what might count as evidence of learning of the critical idea or skill here and what possible tasks would produce the learning	List the possible reflection questions here
Critical idea$_2$ or skill$_2$	List what might count as evidence of learning of the critical idea or skill here and what possible tasks would produce the learning	List the possible reflection questions here
...	...	...
Critical idea$_x$ or skill$_x$	List what might count as evidence of learning of the critical idea or skill here and what possible tasks would produce the learning	List the possible reflection questions here

3. Communicate the planning table to students. Clearly explain the topic, the purpose of the portfolio, the tasks, the choices they have, the deadlines, the reflection questions and other instructions that would help them complete the portfolio.

When the final portfolio is presented, it should contain at least: a table of contents according to the instructions; an overall reflection; the individual artifacts and their corresponding reflections; and any feedback received from the professor during the process.

BOX 8.9 LEARNING PORTFOLIOS

Some useful sources on the use of learning portfolios are videos on YouTube such as one 6-minute one called **"Incorporating student portfolios into your project-based classroom"** and various articles on the web, such as **"The learning portfolio: a powerful idea for significant learning"**.

Self-Assessment

The most successful learners are those that can assess their own learning. They are the ones who can recognize their own needs for improvement and plan their own learning. Most experts and some of the most exceptional people on this planet are auto-didacts, and it is only possible to

become one if you have the ability to assess yourself. Self-assessment helps students to develop this meta-cognitive skill.

There are many misconceptions about what self-assessment is. Table 8.6 shows possible versions of what it means.

Is [A] sufficient, where students only check their own understanding or do they need to go as far as [B] and also mark their own assessment? Actually, most people associate [B] with self-assessment, but this does not go far enough. This is because as long as someone else decides what counts as learning, there is no real self-assessment.

One leading expert on self-assessment, David Boud (2003, p.12) argues that to "…eliminate completely students' consideration of criteria is to remove them from participation in the core processes of learning…". Having students decide how they will assess themselves is critical to learning. Self-assessment goes at least as far as [C], where they also decide what counts as learning by creating their own marking memorandum or rubric and marking themselves accordingly.

In Box 8.10 below, I refer to my experience with self-assessment, and it gives a sense of how to do it and overcome some of the challenges. In short, designing a self-assessment involves the following:

1. Design an authentic task.
2. Someone facilitates a *participative* process in which the students generate their own rubric or criteria for judging their work. Their professor may facilitate but is not allowed to dictate or suggest anything.
3. Students complete the task, and each one scores or marks him/herself according to the rubric or criteria. They have to provide reasons for their marks.
4. The professor marks the same tasks and provides reasons.
5. If there is a large discrepancy, a negotiation takes place, but, if not, the highest mark counts. There should be no hint of coercion during this process, and, if necessary, an arbitrator can be involved.

Table 8.6 What Counts as Self-Assessment?

	A	B	C	D
Check own understanding	X	X	X	X
Grade own assessment		X	X	X
Decide how to mark own assessment			X	X
Design own assessment				X

Because of steps [4] and [5], self-assessment actually involves the same or more work for a professor than traditional assessment. In fact, it creates more work for students as well and this is what my students in the past identified as one of the things they disliked about self-assessment. Table 8.7 summarizes what they liked and disliked.

When I interviewed these students, all of them told me that they thought self-assessment was a valuable skill, and even though it was more demanding, it was worth the trouble. Most of them indicated that self-assessment gave them a more realistic awareness of their own abilities and motivated them to produce good work.

BOX 8.10 SELF-ASSESSMENT

My experience with using self-assessment in a course is described here: https://www.academia.edu/25112851/Teaching_Econometrics_through_self-assessment. Many sources can be found on the web by searching for "student self-assessment".

Alternatives to Grading

In the real world, we do not have people giving us grades; assessment still takes place, but it is very different from grading. Grading is the practice of assigning letters or numbers to students' work. Despite its widespread use, we have known for some time that grading damages learning (Kohn, 1999).

In summary, the research suggests that grading distracts students from the learning process and makes them focus on the outcome (the mark or symbol) alone. This was identified in the previous chapter as something that encourages cheating. As a result of this distraction, students tend to: pay less attention to the feedback they receive from teachers, especially when they get good marks; think less deeply, which makes them remember less; and feel encouraged to avoid challenging tasks that may result in lower grades, even though these are the very tasks that promote

Table 8.7 Students' Views of Self-Assessment

Why they liked self-assessment	Why they disliked self-assessment
We have a clearer idea of what to do and know better how to improve. It is fair and transparent. The process generates information that the professor wouldn't have known.	It takes more time. We have to work harder and think more.

learning. Also, grading sends the hidden message that ability is more important than effort and improvement, and so they never fully grasp the importance of learning from mistakes.

BOX 8.11 AGAINST GRADING

You should start by reading at least one of the articles by the most outspoken critic of grading on his website: http://www.alfiekohn.org. Then, review the wide range of alternatives to grading by searching for alternatives to letter grades in education. If you want to explore some more, you can search for other alternatives like: **"contract grading"** and also **"standard-based grading"**.

Instead of being 'institutions of learning', grading turns schools and universities into 'institutions of achievement'. No doubt, achievement is important in the real world, but institutions of learning should aim to *prepare* students to achieve in the future by *learning* today. By focusing on achievement now, students never really learn how to learn.

CONCLUSION

The biggest threat to effective assessment is lack of validity, which is what happens when you do not accurately assess the intended learning. There are many things that can endanger validity: assessing the wrong ideas and skills; misaligned assessments; incorrect design of assessment items; inappropriate actions during and around the assessment process; and even sticking too closely to traditional assessment and not considering alternatives. This chapter explained how to overcome all these threats.

EXERCISES

1. For a topic of your choice: design two alternate choice items, two short answer items, and two multiple choice questions.
   a) In the first one of each kind, deliberately make a mistake.
   b) In the second one of each kind, eliminate the mistake/s and do it correctly.
2. For a topic of your choice: create a data response item starting with an appropriate stimulus.
   a) Design three poor questions and explain what is wrong with each question.
   b) Then design three good questions without any mistakes.

3. For a topic of your choice: create four essay questions as follows:
   a) Design two poor questions and explain what is wrong with each question.
   b) Then eliminate the mistakes and rewrite the poor questions as two good questions.
4. For a topic of your choice: which of the alternatives to conventional assessment and grading can you introduce in your course? How would you do it?

REFERENCES

Boud, D. 2003. *Enhancing Learning through Self-Assessment*. London: Routledge Falmer.

Kohn, A. 1999. From degrading to de-grading. *High School Magazine*, 6(5):38–43, Mar 1999.

Miller, M.D., Linn, R.L. & Gronlund, N.E. 2009. *Measurement and Assessment in Teaching*, 10th edition. Upper Saddle River.

Chapter Nine

Rethink Educational Technology by Noticing What Everyone Else Misses

Chay taught English in a suffocating, bare-walled room, where the footprints of one class remained in the dust until the next group replaced them with their own. The worn-out desks were neatly organized, but there was no electricity. Actually, she was in a country that did not even have a television service. Yet, she understood more about educational technology than anyone I have ever met. Without computers or much printed material, she found some of the most effective ways to teach her students – ways that would never have entered the minds of the pampered professors in the best-equipped venues at elite universities. In fact, for them, such conditions would have been a nightmare.

Like most young people throughout the world, students love staring at screens. One day, Chay thought that a television would be the ideal way to not only engage her students but also increase their reading speed. But how does one do this when a television would not work, even if you had one? This is the kind of problem that resourceful teachers solve every day, using only what they find around them.

Before reading further, think for a moment what you would have done. How would you have gotten a 'television' to keep the students interested and motivate them to read faster and with more comprehension?

I base this illustration on one of many true stories of resourceful teachers from Brian Tomlinson (2005 p.6), and I'll let him reveal what the teacher did:

> She wrote out in large print on a scroll of cardboard the English version of a local folktale. She then inserted a rolling pin into the scroll and then placed it into a cardboard box that she had made to look like a television by cutting out a screen and painting knobs below the screen

She walked into the class with this contraption and the students sat down in front of it. As she turned the rolling pin, the story scrolled up the screen. The class was so captivated that they had regular 'television-watching' sessions thereafter. It worked even better than a television: she found that her rolling pin allowed her to control the scrolling speed, and as a result, her students' reading speed gradually increased over the course of the year.

Stories like this make one rethink the meaning of 'technology', especially educational technology. The term is broader than the constricted modern definition of it. Words carry implicit definitions and associations that limit what we perceive and how we act in the world. This is also true of the word 'technology'. When most people think of technology in an educational setting, what comes to mind is information technology and mobile devices. I will go further and define technology not merely as something physical, but as the *art* of finding and using things, ideas, people, and methods that help us to achieve objectives.

With this in mind, I will argue that many of us do not use educational technology to its full potential because, unlike Chay, we do not perceive most of the technology already present around us. Often, we notice only the electricity-hungry devices. We allow this kind technology to control and constrain our teaching, and become immobilized when it does not work, or malfunctions.

This chapter presents a different perspective on the use of educational technology, be it a lecture hall or under a tree. It tries to explain what Brian Tomlinson (2005, p.4) meant when he said that teachers and professors in "well-resourced classrooms are not necessarily at an advantage over those in under-resourced classrooms… in fact, the latter might even be at an advantage". Indeed, if we were able to transport Chay from Vanuatu to the best-equipped lecture hall at Harvard, and gave her time to catch up, she would take far better advantage of the available technology than the most technology-savvy professor there. After this chapter, you should be more aware of the vast amount of hidden educational technology at your disposal, and as a result, this should enable you to better exploit technology of any kind.

THERE IS MORE TO EDUCATIONAL TECHNOLOGY THAN YOU REALIZE

If we regard technology as the art of finding and using things, ideas, people and methods that help us to achieve objectives, then to understand technology we must first know our objectives. In a teaching situation, these objectives are derived from the learning a professor hopes to bring about during the lecture.

Learning only happens when we think, and it is this thinking that generates the new connections necessary to remember, understand and reason. From the research, we now realize that thinking is not limited to what happens inside our heads. While I will refer to this idea as 'distributed cognition', it comes in many versions (see Box 9.1). It shows that we spread (distribute) the hard work of thinking (cognition) over many things outside of our heads. This is not limited to humans alone; for instance, even spiders are known to use their webs to think (Japyassú & Hilton, 2017).

This distribution allows us to offload cognition and think faster, more accurately and more efficiently as we are able to focus our limited attention on what is critical. We offload cognition onto body parts (e.g. when we count on our fingers), objects (e.g. when we use a book or a dashboard), other people (e.g. when we ask for directions), or even nature (e.g. using the stars to navigate). The objects that carry our cognitive load need not be tangible; they can even be mental objects, such as analogies, diagrams or mnemonics.

BOX 9.1 DISTRIBUTED COGNITION

Distributed cognition is now well accepted in cognitive science and education. For more information consult the Wikipedia entry on 'distributed cognition'. Classic sources on the topic are Hutchins (1995) and the talk by David Chalmers on YouTube, entitled "Is your phone part of your mind?" You can also search for related concepts like: **enacted cognition, embodied cognition, extended cognition, social cognition**, and **situated cognition**.

Educational technology then, as seen through the lenses of the broader definition above, is the art of finding and using anything (tangible or intangible) that can carry cognitive load, and so enhance cognition and accelerate our learning.

To better understand the arguments that follow, I will divide technology into conventional technology (such as devices) and various kinds of 'resources' that remain unused. Using the concept of 'affordances', I will relate those resources to learning objectives, and so ensure that you become aware of the under-utilized technological potential present in any location.

FREE RESOURCES CREATE THE IDEAL CLASSROOM

Awareness of the potential of unused resources, personified by teachers like Chay, leads to very resourceful teaching. The invention methodology TRIZ is a useful tool to structure one's thinking about such unused

resources. From the perspective of TRIZ, conventional technology is not the only tool in a classroom – there are many other resources. In fact, a resource includes any *unexploited or under-utilized* object, energy, ability or state that is available to the professor and students in, or around, the teaching venue.

TRIZ is interested in resources because this is how one creates an ideal system. The ideal system is one that generates only the desired benefits, with no harmful effects and without any costs whatsoever. Costs can only be eliminated by using *free* resources, that is resources that can be used at no (or negligible) cost. To help people find free resources, TRIZ categorizes these resources, as seen below, with modifications for how they may be seen in a teaching venue:

- All inputs to the teaching process: materials brought by the professor, students or others (in them, on them, held by them), mental content brought by them, movements, and interactions between various people and objects before teaching starts.
- Anything present or perceived in the classroom, such as systems, sun, gravity, pressure, noise, temperature, reactions, locations, gaps, volumes, light, distance, smells, and even harmful effects.
- Features: shapes of objects, potential or actual behavior of people or objects, colors, surfaces, textures, stickiness, stability, density, strength, states, and more.
- Time: sequences, triggers, endpoints, synchronization, dependencies, preparation, time spent, time remaining, changes, and the effects of time.
- System interactions: anything that creates or enables interaction – e.g. information, heat, pressure, energy carriers, fear, irritation, joy, humor, and so on.
- Scale: super-systems up to the highest accessible level and subsystems down to the smallest accessible level.

Scale, the last resource category, is especially fruitful when used together with the 'nine boxes technique' shown in Figure 9.1. This technique is really about raising awareness of resources.

To use it, place the present system in the center of the nine boxes, which is the venue during the lecture. Then consider all the resources in this system: inputs; anything present; features; time; and interactions. This may include: notes; papers; pens; phones; clothes; shoes; bags; desks; chairs; tiles; conversations; gaps between desks; the distance between you and the students; stability of chairs; location of the doors; irritation with the professor; prior interactions with students; students'

	Past	Present	Future
Super-system	Resources outside the venue or in the larger world outside the venue before the lecture	Resources available outside the venue or in the larger world outside the venue during the lecture	Resources outside the venue or in the larger world outside the venue after the lecture
System	Resources in the venue before the lecture	**Resources in the venue during the lecture**	Resources in the venue after the lecture
Sub-system	Resources found in the smaller parts of the resources before the lecture	Resources found in the smaller parts of the resources in the venue during the lecture	Resources found in the smaller parts of the resources after the lecture

Figure 9.1 Nine Boxes in Which to Search for Resources

experiences; current news… this can go on and on. One then searches in the same way for resources that exist prior to the lecture (system-past box) and will exist after it (system-future box).

Then consider the supersystem, which is comprised of the bigger systems of which the venue is part: immediate surroundings, the institution, the city, and so on. Within these supersystems, we may find further accessible resources: during the lecture (supersystem-present), before it (supersystem-past) and after it (supersystem-future). Finally, search for resources in the subsystem, which comprise the parts of the lecture hall, and also the parts of the people, objects and processes, during, before and after the lecture.

As you identify the resources in each of the nine boxes, you will realize the vast amount of free resources hidden in plain sight. If one was able to produce all the benefits with such free resources only, the system would no longer be needed, which is why they say in TRIZ that the ideal system is no system at all.

Let's return to Chay to illustrate this point. She needed a television to engage the students and encourage them to read at a constant and increasing pace. But the ideal television is no television at all, that is to say, it is something that delivers all the benefits of a television, without a

television, by using free resources only. What Chay wanted was not the television, but its benefits: learner engagement and reading speed consistency and control. The lack of conventional technology pushed her to be resourceful and consider how she could develop a system that could generate all the benefits with only free resources. The result was something quite close to the ideal, one which was more robust and worked even better than an actual television. Return to the story at the start and see how many resources you notice.

Many free resources must have been in Chay's house, which was part of the super-system, for example writing material, cardboard, scissors, glue, paint and a rolling pin in her kitchen. Then there were the folktales in the community and Chay's own knowledge of English, television, the basic mechanics of objects and her understanding of her students' behavior and skills. Within the classroom, she exploited: gravity; the shape of the rolling pin and cardboard; the flexibility of the paper; the empty space; the arrangement of the objects and students; as well as the students' reactions.

It is incredible how much better one can become at achieving learning objectives by exploiting free resources. Resourceful thinking makes one think more deeply about what is critical to learning, instead of being distracted by technological features.

The way John Lello (2001, p.vi) taught a group of children how to be historians exemplifies how productive this focus on resources can be. First, he asked them what they thought history was, and:

> One said it was Granny; another said it was the church; a small girl at the front said it was Snow White. So I asked what they would think if Snow White met Granny in the church, and the little girl at the front said that she thought it would make a good story.

Using this, he found it easy to help them understand the nature of history. Then, to help them understand how historians explore evidence to tell stories, he continued by using another thing close at hand:

> I threw my wallet on the classroom floor and invited [them] to find out things about me from what they could extract.... After half an hour they knew a lot about me and they had handled and touched a considerable amount of material and had interpreted much of it fairly easily.

They quickly understood how evidence could be interpreted in a variety of ways and how different stories need to be evaluated. By only using

free resources – such as students' own experience, well-known stories, their innate curiosity about strangers, gravity, noise, the floor, his wallet, the contents of his wallet, and many more – these children learned more about how to be a historian in one hour than many university students learn in a semester from doing quizzes, frantically typing notes on mobile devices, mindlessly searching the web and passively staring at PowerPoint slides.

If Lello had the resources, he might have taken the students to a library that contained historical documents, but what he created was much better. He created the ideal library: a system that generated the educational benefits he wanted, but without the library.

BOX 9.2 THE IDEAL SYSTEM AND FREE RESOURCES

TRIZ is one of the most powerful invention and problem-solving methodologies available. It is based on a systematic study of millions of patents. The application of its tools, especially those related to ideality and resources, has led to resourceful solutions to difficult problems. You can find much information on the web, by searching: **TRIZ AND "ideal final result" AND "free resources"**. The material can be technical, but some accessible books are Straker and Rawlinson (2011), Hipple (2012) and Boyd and Goldenburg (2014). Another useful approach, but one that focuses mainly on intangible free resources, is the 'funds of knowledge' approach (Moll *et al.*, 1992, González, Moll & Amanti, 2005).

THE AFFORDANCES OF RESOURCES

One cannot fully exploit resources or any kind of technology, or even appreciate their potential, without understanding the concept of affordances. Affordances are opportunities for action generated by an object, which a person can recognize and exploit. For example, a shoe can hold something and can protect what it holds. So, a shoe affords holding and protecting feet while walking. The concept of affordance connects an object to a person's needs or objectives and capabilities. In other words, you would only exploit the opportunities for action that the shoe affords if you *wanted* protection while walking and if you were *able* to walk.

The problem is that culture and conventions condition us to notice only a small number of obvious affordances. We notice that the shoes afford walking with comfort, because they allow us to protect our feet

and keep them comfortable, or that they allow us to look more attractive or meet certain cultural expectations. However, many other affordances remain unnoticed.

The biggest difference between teachers like Chay and John Lello and the rest of us is that they see many more affordances. They notice the incredibly vast number of affordances hidden not only in the objects we regard as technology, but also in a much bigger collection of unused resources. Chay saw affordances in a rolling pin and local folktales; and Lello saw affordances in his wallet and in himself as a stranger, which, in turn, enabled students to think more efficiently about history.

Getting back to the shoe, what are its hidden affordances? It affords holding objects other than a foot, but also measuring, hitting, drinking, floating, throwing, stacking, marking, making a noise, and many more. Its different parts – like the sole, the leather, the lining, and the laces – also have affordances, and even more when these are combined with other objects.

We innovate, not only as we create new technologies, but as we become aware of the hidden affordances in existing technology. One could describe the evolution of any technology as the process of how humans gradually, and usually by accident, discovered affordances that they did not perceive before.

However, the process need not be accidental. There are ways to find hidden affordances in any object (see Box 9.3). This, combined with the approach explained in this chapter, shows that educational innovation can be deliberate and accelerated, with or without conventional technology, and allows a professor to take advantage of distributed cognition.

BOX 9.3 DISCOVERING AFFORDANCES

The concept of affordances was coined by James Gibson (1966) and made famous by Don Norman (1988). Soler and Santacana (2013) make an interesting distinction between finding 'chair-burning' versus 'secret door affordance', but there is little literature on how to systematically find or discover affordances. The little I could find included: Maier and Fadel (2007), Srivastava and Shu (2012) and Shu *et al.* (2015). For more entertaining approaches to resources and affordances, I can recommend the television series MacGyver (preferably the original) and any book in Cy Tymony's series of 'sneaky' books, such as Tymony (2007).

LINKING RESOURCES AND THEIR AFFORDANCES TO LEARNING IN THE LECTURE HALL

Let's bring it all together to see how you can tap into the full potential of technology during a class. Firstly, you need to determine what you want students to think about or do during the class, and also what you need to do to encourage it. This obviously includes explaining ideas *a, b* and *c*, but also performing actions – like getting students to articulate or apply these ideas; facilitating their participation in a discussion; teaching them to collaborate with other students, and more. Insert them as 'desired actions' in the first row, as shown in Table 9.1.

Secondly, list all the available resources in the venue in the first column of the table. There are obviously too many to list, but the mere act of listing some of them serves to open your mind to new possibilities. You can always add more as necessary. Do not forget that even harmful effects are resources that have affordances. I have seen master teachers incorporate student restlessness, insults, accidents, and ringing mobile phones into the

Table 9.1 Exploiting the Full Technological Potential in Any Venue

Desired actions / Resources	Explaining ideas *a, b* and *c*	Student application of ideas *m, n* and *p* to achieve *x*	Student participation in the class discussion to achieve *y* and *z*
Desks and chairs			
Light			
Wind			
Arrangement			
Clothes			
Door			
Dust			
Bags and contents			
Body parts			
News events			
Space			
Textbooks			
Pens and paper			
Time before class			
Hairstyles			
Tiredness			
Noise			

flow of a lecture – as Paglia (1992) also explains in her memoir of Milton Kessler, as well as in the supplementary chapter on this book's website.

Thirdly, run through the resources quickly. Think about how each could be used, or transformed, to enable the first desired action, then do the same for the second desired action, and so on.

Suppose I wanted to explain an unusual concept like Edward de Bono's 'sur/petition' or Kim and Mauborgne's 'blue ocean strategy'. According to these concepts, successful companies do not try to compete with others, instead, they try to make the competition irrelevant. Looking through the resources, I may come up with the simple idea of a race between two students, while a third student reads a book as fast as possible. With this, I can show how the first two students exhaust themselves by trying to compete, while the third will not wear herself out like this because she is only competing with herself. All I need for this is space, a book, student movement, and maybe a desk or two. Once students understand this demonstration, I will refer to it repeatedly during the lecture, so it will carry a lot of the cognitive load until students fully grasp the concept. This also relates to the idea of teaching as improvisation as explained in the supplementary chapter on this book's website.

One organization that does an excellent job at helping teachers understand how to tap into the potential of available resources is the Council for Economic Education. I have seen lessons created that employ: the heights of students, their movement, rulers, water, paper, desks, paperclips, gravity, greed, common games on phones, old newspapers, clothing items, the floor, and much more; in helping students to understand some complex concepts (see www.econedlink.org).

CONCLUSION: THE ANTIFRAGILE PROFESSOR

Thinking in terms of resources and affordances often suggests simple teaching improvements. With practice, this makes one resourceful and able to see more possibilities in a lecture hall than those who are wedded to their laptops and PowerPoint slides. In fact, the ubiquitous PowerPoint presentation has probably done more harm to teaching than any other technology. However, using it correctly can enhance learning immensely, as the final chapter explains.

Even if you do use PowerPoint slides or Wi-Fi, resourcefulness will enable you to teach well when the unexpected happens, and even better when everything goes to plan. After all, surely a laptop or the light from a projector affords much more than merely showing lecture slides? Resourceful professors teach lectures that are more memorable and authentic, and develop better explanations. Since resourceful teaching is not dependent

only on electricity-hungry or screen-based technologies, it becomes 'antifragile', that is, performance improves when things go wrong (Hamer, 2017).

Too much of teaching today is fragile, that is, it gets worse when things go wrong. A power cut, a video that does not play or the unavailability of Wi-Fi is enough to send many of us into a panicked paralysis. Such teaching is often based on wrong assumptions about students being 'digital natives', and leads to flooding courses with ineffective technologies or more technology than students themselves would like (Kirschner & De Bruyckere, 2017; Escueta *et al*., 2017; Haselby, 2018). A professor who can recognize the affordances in a poorly equipped classroom will teach better in a well-equipped classroom than one who uses conventional technology without understanding its affordances.

EXERCISES

1. In the photograph below, of a teaching situation without conventional technology, identify as many resources as you can.

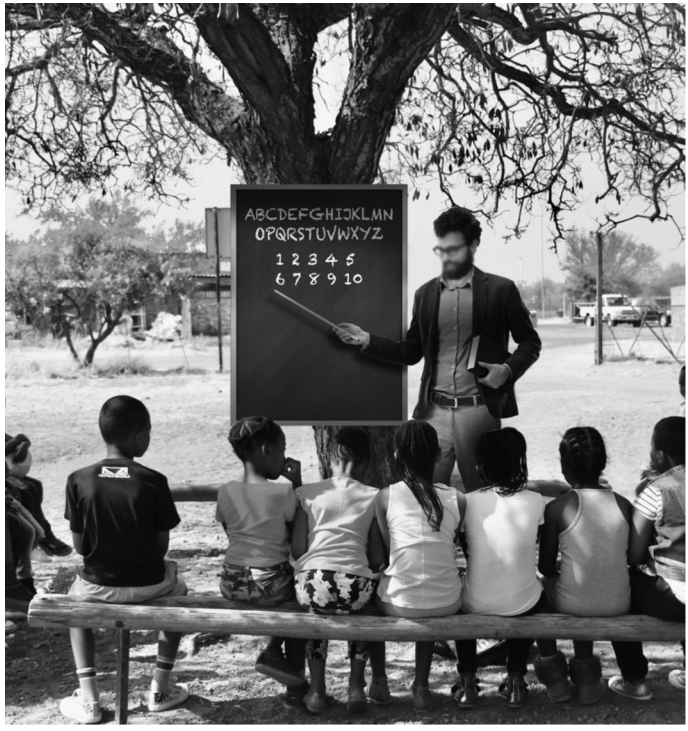

Figure 9.2 Classroom with Little Apparent Educational Technology

Photo credit: Ruan Klopper Photography

2. Consider a teaching situation with much more conventional technology, such as your own, and identify as many resources as you can.
3. Create a table of affordances (like Table 9.1) for an upcoming lecture. What new possibilities do you see for improving how you teach the topic?
4. Suppose in your next class there is a power cut and no internet access. Which resources will you use to ensure that you still achieve your learning objectives?

REFERENCES

Boyd, D. & Goldenberg. J. 2014. *Inside the Box: A Proven System of Creativity for Breakthrough Results*. New York: Simon & Schuster.

Escueta, M., Quan, V., Nickow, A.J. & Oreopoulos, P. 2017. Education technology: An evidence-based review. National Bureau for Economic Research, Working Paper no.23744.

Gibson, J.J. 1966. *The Senses Considered as Perceptual Systems*. London: Allen and Unwin.

González, N., Moll, L. & Amanti, C. (eds). 2005. *Funds of Knowledge: Theorizing Practices in Households, Communities and Classrooms*. Mahwah: L. Erlbaum Associates.

Haselby, S. 2018. Look up from your screen. *Aeon Essays* [available at aeon.co/essays/children-learn-best-when-engaged-in-the-living-world-not-on-screens, accessed 22 September, 2018].

Hipple, J. 2012. *The Ideal Result: What It Is and How to Achieve It*. New York: Springer.

Hutchins, E. 1995. *Cognition in the Wild*. Cambridge: MIT Press.

Hamer, J.F. 2017. *Parenting in the Age of Studies Have Shown*. Quebec: Likeville Press.

Japyassú, H.F. & Hilton, F. 2017. Extended spider cognition. *Animal Cognition*, 20:375-395.

Kirschner, P.A. & De Bruyckere, P. 2017. The myths of the digital native and the multitasker. *Teaching and Teacher Education*, 67:135–142.

Lello, J. 2001. *The Resourceful History Teacher*. London: Continuum.

Maier, J.R.A. & Fadel, G.M. 2007. Identifying affordances. International Conference on Engineering Design, 28–31 August, Paris, France.

Moll, L.C., Amanti, C., Neff, D. & Gonzalez, N. 1992. Funds of knowledge for teaching: Using a qualitative approach to connect homes and classrooms. *Theory into Practice*, 31(2):132–141.

Norman, D. 1988. *The Design of Everyday Things*. New York: Basic Books.

Paglia, C. 1992. *Sex, Art, and American Culture: Essays*. New York: Basic Books.

Shu, L.H., Srivastava, J., Chou, A. & Lai, S. 2015. Three methods for identifying novel affordances. *Artificial Intelligence for Engineering Design, Analysis and Manufacturing*, 29, 269–281.

Soler, E.A. & Santacana, A.B. 2013. Innovative scaffolding: Understanding innovation as the disclosure of hidden affordances. *Revista Iberoamericana de Argumentación*, 7:1–11.

Srivastava, J. & Shu, L.H. 2012. Affordances and environmentally significant behavior. *Proceedings of the ASME 2012 International Design Engineering Technical Conferences & Computers and Information in Engineering Conference*, 12–15 August, Chicago, USA.

Straker, D. & Rawlinson, G. 2011. *How to Invent (Almost) Anything: Serious Innovation Using Science and Psychology*. Self-published on Amazon.

Tomlinson, B. 2005. The Resourceful Language Teacher. TESOL Symposium on English Language Teaching (ELT) in Resource-Challenged Contexts, 16–17 December, Dakar, Senegal.

Tymony, C. 2007. *Sneakiest Uses for Everyday Things*. Kansas City: Andrews McMeel Publishing.

Chapter Ten

Design Presentations That Make Your Lectures Much More Useful

Despite being heavily criticized, the most common model of teaching remains a variant of direct instruction: the lecture. The best professors do not rely on lectures alone (Bain, 2004), but even if you do, you can still teach well if you improve where you are below average (as explained in the Preface). The source of a great number of problems is found in the inappropriate use of technology, especially presentations designed with software. The inappropriate use of presentations has turned countless potentially useful lectures into useless and boring ones.

While much of the criticism of lectures is justified, this chapter will not enter into that debate. This chapter will rather aim to explain how it is possible to make lectures more useful by changing how we design and deliver presentations, and combining it with the precepts of the first three chapters. While I will focus on the use of PowerPoint, the ideas can be applied to other presentation methods and software (like Prezi or Keynote) too. Building on the view of technology outlined in Chapter 9, it will suggest a rethinking of the role of lectures. This chapter is less concerned with the design of individual slides or the features of presentation software. Instead, it explains how to develop presentations that add value to the lecture as a *unique* learning experience.

PRESENTATIONS WORK IF LECTURES ARE ABLE TO BE USEFUL

Presentations are the most common method used to deliver lectures. If lectures are not useful, then by implication, the effectiveness of presentations also becomes questionable. First then, let us establish if lectures are able to be useful and in what ways. Based on this, we can grasp how presentations can enhance this usefulness.

On the one hand, there is increasing criticism of lectures, many even arguing that lectures are outdated and of little use (see Box 10.1). On the other hand, students still come to lectures – or, at least, some of them.

Why do students even attend lectures? Some come out of a sense of duty, but most arrive in the hope that lectures will improve their chances of passing the course (not necessarily to learn something). Often, it is a waste of their time, especially in courses that require only memorization to pass or in courses where professors add little value over and above the text or available web sources.

BOX 10.1 ARE LECTURES USEFUL? SOME OPPOSING VIEWS

An increasing number of people are arguing that lectures are no longer useful. Search for articles like: **"Ten reasons we should ditch university lectures"** and **"Lectures don't work, but we keep using them"**. While I disagree that lectures are always useless, the views of such critics are valuable in helping me think about what to avoid in my lectures. Lectures really are useless if they add nothing unique to the learning experience, but there are a handful of lectures that students do find useful. It is these kinds of lectures that the proponents of the lecture format have in mind, and you can find them by searching for articles such as: **"Lectures are an effective teaching method because they exploit human evolved 'human nature' to improve learning"**; **"In Defense of Lecturing"**; **"Lecture Me. Really"**; and **"Don't Give Up on the Lecture"**.

The lectures that are truly useful are those that enhance understanding. This means that assessments should be constructively aligned with understanding, because only then will students realize they need understanding to pass. However, even this does not guarantee that lectures will be useful. Professors also need to do things that develop understanding during a lecture. If lectures simply reproduce the text, they offer no additional value to students.

Lectures that develop understanding, firstly, reduce the cost of understanding by containing effective explanations; and, secondly, make students think. Understanding of ideas may be possible, albeit more difficult, without effective explanations. However, understanding is only possible if the student thinks about the ideas, and a lecture is an ideal place to make that happen.

Of course, it is possible that textbooks can make students think, so lectures also have to provoke and guide thinking in a way that a text

cannot. The question is whether lectures have something unique to offer – something unique that enhances understanding differently to texts or other impersonal sources.

Let us consider then what makes lectures different from such sources. Lectures are *verbal* presentations of content to an *audience*. From this, we identify two defining qualities of lectures that make them different from a text: (1) lectures are spoken; and (2) they are social events.

Considering the first quality, we can conclude that anything that harms the value of the verbal dimension of lectures should not be part of a presentation. Poor presentations often give primacy to the text by using slides that are so text-heavy that the text becomes the focal point. At best, such presentations detract from the uniqueness of lectures and make lectures redundant. At worst, text-heavy lectures cause listeners to remember less from *both* the text and the spoken word in lectures, in other words, they do more harm than good.

Jane Bozarth (2008) explains that different parts of the brain process textual information on a screen versus the auditory information coming from the presenter. The brain also processes textual and auditory information at different speeds and this splits attention. As a result, the two interfere and compete with each other. This creates cognitive overload and reduces the listener's ability to think about and understand the ideas being presented.

Humans have learned from other humans through speech for ages. Speech is more flexible than text: speakers can respond to situations and people, and also improvise interactions as needed. This differs from the written word where content is generally fixed. Unlike text, which is predetermined and linear, in a lecture, the speaker can return to previous ideas and repeat critical ideas when necessary, depending on the situation. Over-reliance on text-based slides eliminates much of this flexibility.

Learning is often a social activity. There is great value in seeing that another person has mastered the content, observing what they accentuate through volume, gestures, and expressions, and sharing the experience of understanding (or not understanding) with others. I have often found that students understand the content faster and better if I explain it to them in person, even when I stick closely to the text. This may be because social signals play an important role in human learning. Again, over-reliance on text-heavy slides obscures and hampers much of the social dimension of lectures.

While all of this is true, if we look more deeply, we can see that the real value of lectures, and the presentations that accompany them, lies in the fact that they enable distributed cognition. As explained in the previous chapter, when distributed cognition is activated, we become more efficient thinkers and problem-solvers, because it allows us to offload cognition onto our bodies, other people and real or abstract objects.

The literature (Karasavvidis, 2002; Scaife & Rogers, 2005; Heersminka & Knight, 2018) recognizes that technology enables distributed cognition in learning – through things like spreadsheets, phones, virtual reality, Google, all kinds of apps and so on. What is not recognized is that *lectures themselves* can activate the power of distributed cognition, but only if presentations are designed appropriately.

Many of the techniques from the first four chapters activate distributed cognition. For example, a well-chosen analogy used to explain a difficult concept allows a student to use less cognitive energy when thinking about a new concept. This is because the analogy not only highlights what is important, but also enables a student to think by interacting with the analogy instead of having to figure out everything himself from previous concepts. In a way, the analogy stores knowledge and releases it only when a person uses it to think about the new concept. The same can be said of demonstrations, applications and examples.

Other ideas from previous chapters – like organizing lectures around critical ideas, chunking complex ideas and finding an appropriate sequence – serve to manage the cognitive load that a student takes on while learning. Here, the cognition is initially distributed to the professor and the student's thinking is made more efficient by interacting with him or her. This interaction can be as simple as attentive listening but also includes more participative interaction.

Lectures that enable students to benefit from distributed cognition will firstly avoid anything that causes cognitive overload. Secondly, such lectures will use the resources of the lecture – the professor, fellow students, the actual presentation and the abstract objects that arise from the presentation – in a way that enables the student to spread their cognition, and so think more efficiently.

In summary, the best presentations during lectures are those that encourage and enable thinking, which is essential for understanding. But not only that – the best ones make you think *more efficiently* and so learn faster. They manage to do this by enhancing the qualities that make lectures different from the text, and so enhance understanding in ways that reading a text cannot.

PRINCIPLES FOR DESIGNING PRESENTATIONS FOR LECTURES

The best approach I have come across that applies these ideas is an adapted version of the 'assertion-evidence' approach developed by Michael Alley (2013). In this section, I explain how to use it.

The Assertion-Evidence Approach

In this approach, a presentation is designed around messages, not content. Every slide conveys just one message, and these messages flow directly from the critical ideas. Every slide contains just one sentence at the top, which conveys the message of the slide: this is the 'assertion'. Below the assertion, there should usually be a visual that: illustrates the assertion, captures the details of the message, and creates a focal point for the elaboration of the assertion. It is important that this visual *not* simply be decorative, because decorations add redundant information, which, in turn, adds to the cognitive load of the listener. Instead, it should be a visual that the presenter actually uses in the explanation of the assertion. This visual is the 'evidence' part in the assertion-evidence approach.

Adding a visual part to the slides is based on research on multimedia learning. Jane Bozarth (2008, p.31) explains how this is captured in Richard Mayer's multimedia principle: the idea that, "Learning is enhanced by the presentation of words and pictures rather than words alone". But this principle only applies if the visual is coherent with the assertion and useful in elaborating on it. Merely decorative visuals harm learning because they add irrelevant information. Richard Mayer (2007, p.135) refers to this as "seductive details"; and argues that, "When instructional messages contain extraneous material, learners may waste precious cognitive capacity on processing that material and therefore not have adequate capacity to engage in essential and generative processing" (p.136).

Figure 10.1 is an example of a slide based on the assertion-evidence approach. Note the assertion on top is written as a full sentence in no more than two lines. Also, observe that, while the visual contains some text, it is limited. The visual is used to explain the message that promot-

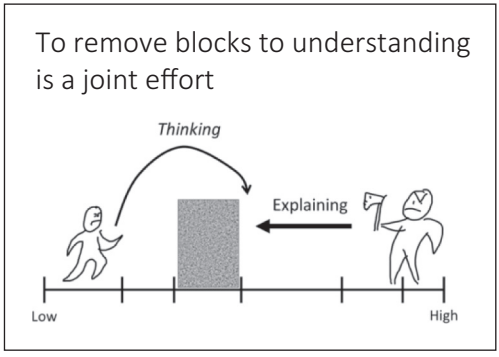

Figure 10.1 Example of a Slide in an Assertion-Evidence Presentation

ing understanding is a joint effort. In the visual, we see that the teacher, who has a high level of understanding, reduces the cost of understanding through explanation – imagine the axe in his hand is his explanation, and his facial expression is one of determination. The student, who starts from a low level of understanding, then overcomes the block to understanding by thinking about the ideas being explained. While talking from this slide, I would first highlight the message and then constantly refer to the visual as I elaborate on the message.

It is not always possible for the 'evidence' part of the assertion-evidence part to be visual. The rule is to avoid bullet point lists, and when you do use text, not to expose the audience to more than around 20 words a minute. If for example, you have to use a table with text, it would be best to minimize the text by gradually revealing it as your explanation progresses.

BOX 10.2 MICHAEL ALLEY'S 'ASSERTION-EVIDENCE APPROACH' TO PRESENTATIONS

For more information on the assertion-evidence approach, see Michael Alley's insightful book, *The Craft of Scientific Presentations*, or see the guidelines and the videos at www.craftofscientificpresentations.com. The approach also works well in disciplines other than science and engineering – for example, search for the talk by Robert Yale on the assertion-evidence approach on YouTube.

In Box 10.2, there are resources that explain the assertion-evidence approach in detail, and I would recommend at least watching the videos that explain it. Let me repeat, with this approach, you don't use bullet points and you minimize text. Every slide has a message, summarized by an assertion, and evidence (preferably visual) that guides both the speaker and the listener.

Adapting the Assertion-Evidence Approach

When I teach, I use an adapted version of the assertion-evidence approach that is based on the principles derived from the first three chapters (principles 1 and 2) and from the discussion above about the uniqueness of lectures (principles 3 and 4).

Principle 1: Presentations Should Be Organized around Critical Ideas

In the first chapters, you saw how critical ideas and their connections help to organize teaching. Naturally then, critical ideas should be the origin of the messages in lecture slides.

Some critical ideas may require more than one slide, but the principle is that only slides that convey messages derived from those critical ideas should be in a presentation. If you feel the need to include slides not directly derived from the critical ideas, then, either you are bringing in too much detail and should delete those slides, or your critical ideas are not correct, and you should revisit them.

Those critical ideas and their connections should be clear to students throughout the lecture. For this reason, I use a home slide that shows the critical ideas and their connections, and periodically return to this home slide as I transition from one idea to the next. This is consistent with Richard Mayer's (2007, p.138) signaling principle, which states that: "People learn more deeply from a multimedia message when cues are added that highlight the organization of the essential material."

The next two figures show two kinds of home slides in a presentation based on the content of Chapter 2. Figure 10.2 shows the critical ideas in linear order with the connections between each shown in italics.

In contrast, Figure 10.3 is the home slide created to look more like a simplified concept map with the connections showing on the lines between the critical ideas. The map version is preferable for complex presentations that are not easily reduced to a linear sequence. You can either show the concept map every time you are about to move to the next idea, or add hyperlinks so that you can simply click on a critical idea or connection to go to its slide.

A home slide makes the knowledge structure very explicit and allows the student to offload some cognition onto it. Instead of having to hold the complete knowledge structure and all its connections in mind, students get exposed to this slide enough times so that they can spend their cognitive energies on understanding the details – and gradually assimilate the structure. Also, this home slide allows for some semi-linearity in

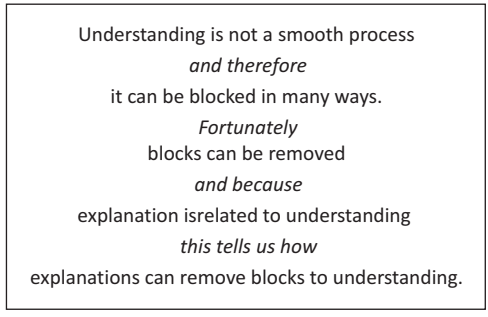

Figure 10.2 The Linear Version of a Home Slide

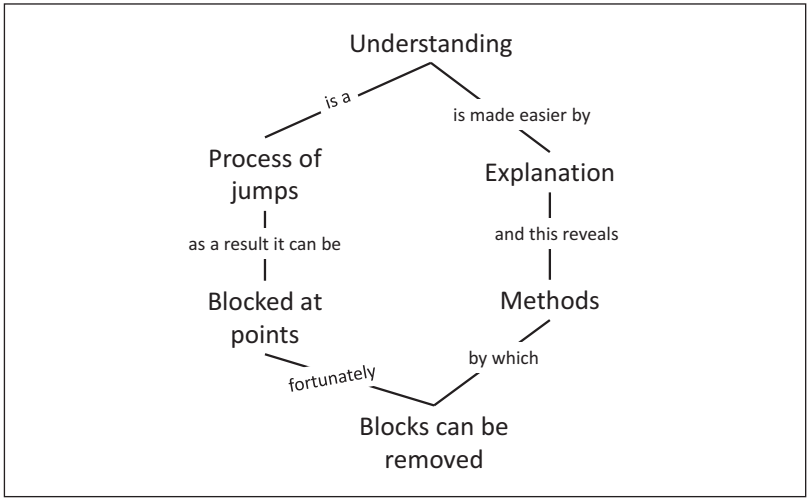

Figure 10.3 The Map Version of a Home Slide

the presentation, because it allows the teacher to return to previous ideas without students losing track of the flow between the ideas.

Principle 2: Presentations Should Follow the Sequence Dictated by the Hub-and-Spoke Knowledge Structure

With a home slide in place and slides that only elaborate on the critical ideas on the home slide, it becomes easier to follow the correct sequence. Following the hub-and-spoke structure explained in Chapter 2, presentations should start at the hub, move outwards and return to the hub. At the start of a presentation, you should show the home slide, highlight the first critical idea, go to the slides that elaborate on this critical idea and explain them, and then return to the home slide. Before moving to the next critical idea, you should explain its connection to the previous idea, and then proceed with the slides that elaborate on it.

When you transition to the next critical idea, ensure that you highlight the critical idea that you are about to elaborate on. An example is shown in Figure 10.4.

Note that the critical idea that is going to be explained next ('Blocks can be removed') is highlighted, as well as its connection ('fortunately') to the previous idea. Remember to explain the connection before explaining the next idea.

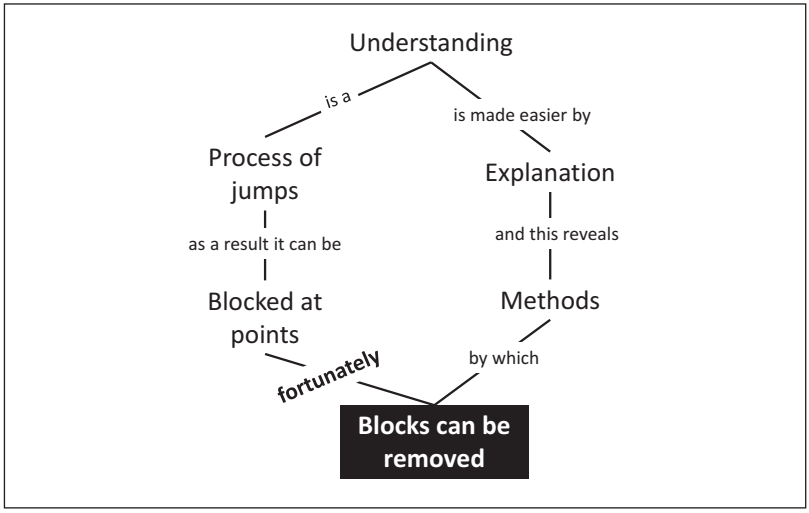

Figure 10.4 Home Slide Showing the Sequence and Current Location of the Lecture

Principle 3: Presentations Should Offload Cognition onto the Slides and the Professor

Too much text on a slide interferes with the decoding of the message and assumes that listeners know what the message is. As mentioned, when listeners receive the exact same message verbally and textually, it interferes with their memory and thought processes. If understanding is the purpose of a lecture, then anything that interferes with thinking will hamper understanding. And if there is interference with cognition, then there is no way that cognition can be distributed.

Like the home slide, the other individual slides should enable the student to offload cognition onto them. The assertion makes the message clear so that students need to spend less cognitive energy on trying to figure out what a professor is trying to say, and more on trying to understand the message itself. The visual evidence supports only the assertion, and so helps the student to remain focused on the message. When the student's attention wanders for short periods, both the assertion and the visual evidence help them to not to lose the essential message.

The visual evidence, in particular, is very effective at offloading cognition. The visual – if well-chosen, and not merely decorative – will store the content because it hints at the details of the assertion – very much like a good analogy. By listening to a professor interacting with the visual

evidence, the content is gradually released as necessary. When students want to recall the details, they simply call up the visual evidence used in the slide. This is much easier than trying to remember a bullet point list because, as explained in the fourth chapter, humans' visual memory is close to perfect.

What the visual evidence also does is represent the assertion in a way that highlights the critical details, instead of trying to convey all the details, as one finds in text-heavy presentations. This manages students' cognitive load: it helps them to recognize what is critical and frees up cognitive energy that can be spent on understanding them.

Principle 4: Presentations Should Enhance the Verbal Nature of Lectures

Too often a person talks about something in a presentation simply because it is on the next slide or bullet point. This leads to incoherent presentations that seem to jump to a new idea with no clear reasoning behind it. Anyone who has ever tried to turn a set of slides into a written article will know how many disconnected ideas there can be in a presentation.

This usually happens because the default bullet point format of popular presentation software does not encourage presenters to think deeply through the connections or flow between points, but simply to put down ideas as they occur. Even when there is a connection between bullet points, the connections are sometimes lost because one's mindset during preparation is most likely different from the mindset during the presentation. As a result, many often start talking about the next bullet point without establishing a connection between that and the prior point. In this way, the text also interferes with the presenter's thinking process and makes the verbal dimension of a lecture subservient to the textual.

However, in the assertion-evidence approach, the text serves the spoken word. Rather than reading the text, the audience is free to actively listen to, and think about, what is being said. This is because, instead of having speech and text duplicating and interfering with each other, one uses speech and visuals that complement and reinforce each other.

Even though visuals are anchored to the message, they are less constraining than text. This allows for more freedom to follow a coherent line of thought and to talk in a way that is closer to narration and responds to the *audience* rather than to the slides. This often leads to new insights. Many speakers, myself included, report seeing new connections and having novel insights while talking about ideas in this way. One is, therefore, more likely to say things that do not appear in the text that make the audience think. This cannot happen when the text strictly dictates what to talk and think about and when to do so.

This approach ensures that one follow Mayer's (2007, p.142) modality principle of effective presentations, which states that people learn more deeply from verbal narration than from printed words, and also more from visuals than from text. Speech not only is what makes lectures unique, but it is also the most effective way to help students learn during lectures.

Why Other Approaches Don't Work as Well

There is much good advice out there on designing presentations, but much of it is not developed with the uniqueness of lectures in mind or with an understanding of how lectures have the ability to exploit the power of distributed cognition. Any approach that encourages you to use even moderate amounts of text implicitly conveys the idea that presentations are there to deliver content. While, in some contexts, presentations are meant to deliver content, in teaching, this is inappropriate. Lecture presentations exist only to enhance understanding.

The only time when it might be acceptable to use even moderate amounts of text is when all the following conditions are met: (1) students do not have a textbook; (2) you cannot give them detailed notes; and (3) your assessments require them to mainly memorize content. If all three of these conditions are met, you are in an absurd teaching situation or you wrongly equate memorization with learning.

Some approaches propose that you use very little text – perhaps no more than three or four bullet points per slide or that you only use keywords or very short phrases to prompt you to speak about the right thing at the right time. I used this 'prompting' approach for many years, but I now realize that it too was wrong because such slides are designed to serve my needs rather than that of the students.

The purpose of presentations is to promote understanding, not to prompt the content delivery. I have three concerns with the 'prompting' approach: (1) it tends to contain too much text so there is still interference between the textual and verbal parts of the presentation; (2) the messages of the slides are still unclear or not derived from the critical ideas; and (3) the presentation does not tap into the power of distributed cognition. The prompting approach is only acceptable if the ideas are very complex and you did not prepare well. Even then, if your purpose is to enhance the understanding of students, it would be better to use notes that only you can see or to use something like the 'presenter view' function in PowerPoint.

The only time I would advocate relying on text when you are using examples that can only be shown as text – such as extracts from a document that you need to analyze - or when you are using new words and you need students to see their spelling. Even then, it is still better to

minimize the text, or in the case of text extracts, not to talk while the audience is reading it, and reveal the text as gradually as possible.

What the adapted assertion-evidence approach does better than any other approach is to use the uniqueness of lectures to enhance understanding. Much of the advice about designing slides becomes redundant with this approach. But there is still some useful advice for the design of individual slides, which we turn to next.

PRINCIPLES FOR DESIGNING INDIVIDUAL SLIDES IN A PRESENTATION

There are many good books on designing slides – such as Kosslyn (2007), Bozarth (2008) and Reynolds (2012). I am not going to try to compete with these excellent books, but will rather emphasize some important advice.

My advice is based on only one principle: maximize the signal-to-noise ratio, or translated into everyday language: if something does not convey important information, remove it! Basic communication theory tells us that when a sender sends a signal along a channel to a receiver, the signal is not always understood as intended because of things that distort the message (also known as noise). During lectures, the sender is the professor, the receiver is the student, the signal is the critical ideas and the channel is the lecture. The way the lecture is presented can create large amounts of noise, which can then either distort understanding or add much activity without adding new information. Let's look at each one.

Things That Distort the Signal

Many things can distort your message. Here, I just point out the ones that are most common, and easiest to avoid.

- Font size: When the font is small, audience members at the back or with bad eyesight will not be able to read it or will misread the text, so your message will not reach them. For this reason, many books advise you use a 28-point font size or larger. Also, font size reflects importance, so use changes in font sizes *only* if you want to communicate relative importance, not merely for variety or decoration. The same applies to the use of different font types.
- Color: When colors clash, are too bright or not sufficiently contrasted; people will struggle to make out the text or visual elements. Unless you know how to work with colors, it is best to stick with plain black text with a white background. Colors should only be used if they are the most efficient way to distinguish the important

parts on a visual. Color should be used only to convey information because a presentation is not there to look pretty – it is there only to promote understanding.

Things That Add Activity Without Adding Information

Many authors have written about guarding against things that add cognitive load without adding information. Kosslyn (2007, p.10) calls it the 'principle of informative changes' – that people expect any change to carry information. If a change (like an animation, a transition, different font or new color etc.) does not convey information about the critical ideas, the audience has to spend their limited attention on something that contributes nothing to their understanding. Hence, changes that are not informative make thinking less efficient.

Again, I am just going to highlight some of the most common mistakes.

- Animations: Only, and I mean only, use animations if they are necessary to convey information. The audience will assume that everything you put in a presentation is there for a reason. If an animation that is just there to make the presentation appear impressive, it adds to the audience's cognitive load and, instead of thinking about the critical ideas, they will be distracted by the redundant information that an animation adds.
- Transitions: In short, do not use the feature for transitions between slides. I have never come across a transition that contributes useful information.

One can mention many other things, such as avoiding all sorts of fancy-looking entrance and exit effects. However, if you bear the principle of maximizing the signal-to-noise ratio in mind, you will avoid the 20% of the mistakes that generate 80% of the distraction.

PROCEDURE FOR DEVELOPING GOOD PRESENTATIONS

The first three steps of the procedure for designing good teaching presentations follow from the first three chapters and are not part of the assertion-evidence approach. None of the steps require knowledge of the advanced features of presentation software.

Step 1: Find the critical ideas, write them down and connect them as explained before.

Step 2: Break the critical ideas up into smaller ideas using the chunking method (see Chapter 3) or some other method. Put these ideas in the appropriate sequence and ensure that there are logical connections between them.

Step 3: Create the home slide with the important ideas from steps one and two. The home slide should contain all the critical ideas, and possibly some of the important ideas derived from them. The connections between the ideas should also be shown. This will be the slide with the most text, so ensure that the text is at least a readable 28-point size. Show this slide at the beginning of the presentation, every time you transition to a new critical idea, and again at the end of the presentation.

Step 4: Using the assertion-evidence approach, develop slides that expand *only* on the ideas in the home slide. Remember the visuals do not have to look pretty. They only have to be presentable enough for you to use it to support and elaborate on your assertion, so even crude drawings will work. Consider my visual in Figure 10.1 – it is not aesthetically appealing, but it does the job.

Step 5: Rehearse your presentation and make notes for yourself that you can refer to during the lecture if you are afraid that you might forget important things. Don't do this for every slide – if you do this, it may be a sign that you are not yet comfortable with the topic, or more likely, that you are trying to say too much. Remember that 20% of the content produces 80% of the understanding, and it is this 20% that you want students to understand. They will probably forget the other details and can get these from the textbook anyway. Your presentation should draw on the strengths of the lecture: semi-linearity, flexibility, responsiveness to the audience and the use of social and verbal signals to emphasize what is important. Lectures are inferior substitutes for textbooks, so never try to imitate the textbook.

Step 6: Make the presentation (see the resources outlined in Box 10.2 for guidance). At the end of your presentation, reflect on it and make improvements.

CONCLUSION

Poor presentations make lectures useless, while good presentations make lectures useful. Good lectures – and, by extension, good presentations – are the ones that make students think and do so more efficiently than they could with a text (or any impersonal source) alone. Students need to think

because, without doing so, they will not understand. However, it is difficult to think about new ideas, and good presentations tap into the power of distributed cognition to manage one's cognitive load, make thinking more efficient, and enable one to think only about that which is critical to achieving understanding. This requires a different approach: one that focuses on the critical ideas and exploits the uniqueness of lectures.

EXERCISES

1. For any one of the topics you teach, design a set of slides (using your preferred presentation software or even plain transparencies) that apply the principles of this workshop.
2. Teach the class using these slides and reflect. What differences did you notice compared to your normal way of presentation in the way:
   a. you prepared;
   b. you taught;
   c. the students responded?

REFERENCES

Alley, M. 2013. *The Craft of Scientific Presentations*, 2nd edition. New York: Spring-Verlag.

Bain, K. 2004. *What the Best College Teachers Do*. Cambridge: Harvard University Press.

Bozarth, J. 2008. *Better Than Bullet Points*. San Francisco: Pfeiffer.

Heersminka, R. & Knight, S. 2018. Distributed learning: Educating and assessing extended cognitive systems. *Philosophical Psychology*, 1–22.

Karasavvidis, I. 2002. Distributed cognition and educational practice. *Journal of Interactive Learning Research*, 13(1/2):11–29.

Kosslyn, S.M. 2007. *Clear and To the Point*. New York: Oxford University Press.

Mayer, R.E. 2007. Research-based guidelines for multimedia instruction. *Reviews of Human Factors and Ergonomics*, 3(1):127–147.

Reynolds, G. 2012. *Presentation Zen: Simple Ideas on Presentation Design and Delivery*, 2nd edition. Berkeley: New Riders.

Scaife, M. & Rogers, Y. 2005. External cognition, innovative technologies, and effective learning. In: P. Gardenfors & P. Johansson (eds), *Cognition, Education and Communication Technology*. Mahwah: L. Erlbaum Associates, 181–202.

Chapter Eleven

The Future of Professor-Experts

While reading a draft of this book, a colleague asked me if there is a future for professors, especially given recent technological changes, the need for lifelong learning and the perpetual complaints that universities do not prepare students to think critically and be productive in the real world. As I was responding, I realized that this brought together many of the ideas in this book.

The future tends to move toward the ideal. The ideal university is one that delivers all the benefits with none of the costs, that is: the ideal university is no university at all. In this scenario, it is not the professors who will ultimately disappear, but the university itself. Universities will ultimately become processes of validation and assessment of learning, with professors remaining an important part, but performing a very different, probably less formal, role than they assume today. Even as universities will gradually seem to disappear, professors will not, especially not the good ones who are also experts.

Being an expert is not a state, it is a process. It is intertwined with the process of becoming a good professor – the two go hand in hand. Richard Feynman showed us that the greatest experts are those who perceive and understand the essence of their subjects. He believed that many of us underestimate the value of teaching, and discount its role in making us think deeply about the essence of our subjects from different perspectives. Experts who teach well are scholars, but not ones who limit their scholarly activity to writing numerous, unread journal articles that, in the majority of cases, simply add refinements and fill ever-shrinking gaps in our knowledge. Instead, they are the ones who: keep up with the general trends; search out the surprises; expose their cherished ideas in conversations with practitioners; continuously deepen their understanding of the ongoing controversies in their discipline; and, maybe most importantly, use teaching as a way to gain and share insights.

The final two chapters were perhaps the most direct response to my colleague's question. There I argued that lectures are unique learning experiences that brim with potential. This potential arises from the conversational and socially interactive dimensions of learning. If we do not appreciate this uniqueness and potential, our value will rightly be questioned. But there is a bright future for those who cannot be replaced with a textbook or a video recording.

Chapter 6 responded to the issue of thinking skills and demonstrated that that critical thinking, and its teaching, is misunderstood. If we grasp the real purpose of critical thinking, instead of over-simplifying it to instruction in syllogisms and fallacies, we are more likely to transform our students into critical thinkers.

Assessment, as suggested in Chapters 7 and 8, is highly underestimated compared to lectures, but will become increasingly important as we move toward the ideal university. When assessments are done correctly, students should learn much more from the process than from lectures, because well-designed assessments and feedback come closer to how learning happens in the real world.

Finally, why do the complaints persist that universities do not prepare students for the real world? It is not as if universities are passive: courses are updated and expanded, driven by various accreditation processes. As I argued in Chapter 1, this is part of the problem. Instead of teaching more and more, universities should rather teach less. Students will learn more outside the university, so it is pointless for universities to try to cover all possible content and skills. It is much more important that students learn a few critical ideas more deeply, which will, in turn, *enable them to take advantage of experience*.

To summarize this book: professor-experts are those that help students grasp the critical ideas so that they continue to expand that understanding for as long as they live. They teach less, explain better, encourage forgetting, upset preconceptions, release natural reasoning, assess authentically, notice more and exploit the uniqueness of learning events.

Learning does not end with professor-experts, it begins there.

Index

Note page references in *italics* indicate figures; **bold** indicates tables.