

Chapter 7

Surprises from cognitive psychology and neuroscience

Why making things more difficult and less enjoyable for students in the short term can enhance long-term learning

By the end of this chapter you will:

- understand what is meant by the phrase “desirable difficulties” and where strategies that appear to make learning more difficult in the short term may in fact result in better long-term retention and recall
- recognise the important distinction between performance and learning and how having this in mind can help you to avoid deceiving yourself about whether long-term learning is actually taking place
- know about five key approaches and have reflected on how you might go about applying them in your own classroom: varying the way you teach, allowing and deliberately building in forgetting (through “spaced learning”), testing to produce better long-term learning, mixing things up by interleaving content and the “generation effect”.



Why would you want to know about this?

So far, we have talked about areas of neuroscience and cognitive psychology that have many parallels in the existing teacher effectiveness research. In this chapter, we will explore a number of areas of evidence that contradict some of the assumptions that many teachers make every day.

Knowing about these things has the potential to help you go beyond your current practice and enable you to enhance long-term learning and recall in your students. Furthermore, some of the evidence (particularly that related to “spaced learning”) suggests that if neuroscientists and teachers worked more closely together, novel, sophisticated and highly effective pedagogies could begin to develop within laboratory-style classroom studies. The topics would also make excellent areas

for a teacher-led randomised controlled trial (see Chapter 8 and Richard and Eleanor's book *Teacher-Led Research* (2016)).

“Desirable difficulties” – what does this mean?

This chapter summarises and provides practical examples of what Elizabeth and Robert Bjork (two psychology professors from the University of California) describe as “**desirable difficulties**” in their book chapter, ‘Making things hard on yourself, but in a good way: creating desirable difficulties to enhance learning’ (Bjork and Bjork, 2011). Most of the research in this area has come from the field of cognitive psychology, but because the neuroscience evidence about learning underpins and supports these approaches we have chosen to include these findings here.

Robert Bjork first introduced the idea of desirable difficulties in 1994 as an umbrella term covering a range of effects that had been found in a number of different studies. In essence, the evidence showed that teaching and training approaches that require more effort on the part of learners in the short term can improve long-term learning. These “difficulties” can therefore be considered “desirable”. The phrase “desirable difficulties” applies to a number of teaching strategies:

- **Varying the conditions of practice** instead of keeping the teaching approach constant (such as changing the way you teach when you repeat information and processes to learners so the children practise the same information in different ways) (Smith et al., 1978; Katak et al., 2010).
- **Spaced learning** (or “distributed practice”) rather than the massed delivery of information (Cepeda et al., 2006; Kramár et al., 2012; Kelley and Watson, 2013; Gerbier and Toppino, 2015).
- Seeing **testing as a learning experience** in its own right (e.g. using testing, rather than the re-studying of material, as a revision approach) (Roediger and Karpicke, 2006a; Halamish and Bjork, 2011; Vestergren and Nyberg, 2014).
- **Interleaving** different content rather than block delivery of instruction (Shea and Morgan, 1979; Kornell and Bjork, 2008; Rau et al., 2013).
- Making use of the **generation effect** (getting learners to generate content and explanations themselves) (Slamecka and Graf, 1978; DeWinstanley and Bjork, 2004).

It is easy to recognise why there are positive effects for some of these things if you think back to what we learned in Chapter 2 about learning and remembering. For a start, something which requires a degree of concentration by definition enhances attentional processing, while breaking things up into smaller episodes (with gaps between them) is likely to generate more primacy and recency effects at the beginning and end of these episodes. Other things, like the positive effects of testing, require a bit more explanation which will be provided later.

The difference between learning and performance, and why this distinction is so important

In order to understand the importance of the desirable difficulties research you first need to understand the distinction between **learning** and **performance** and the research that supports the existence of such a distinction. Evidence from a wide range of studies within psychology and related learning sciences has recognised for some time the need to make a clear distinction between learning and performance (for a review see Soderstrom and Bjork, 2015).

Animal experiments in the early 20th century showed that rats exposed to a maze (without any other attempts to induce learning with reward) appeared to have undergone considerable learning by the time a reward was introduced (Tolman and Honzik, 1930). But other rats exposed to the same maze with a rewards system in place from the start did no better at finding their way through the maze successfully than the rats who had merely wandered aimlessly through the maze. Indeed, over time (after about a couple of weeks) the delayed reward rats caught up with the rewarded group and in fact did even better at finding their way through the maze. Although the use of a reward had created greater short-term performance in the rats who first received it, it had not yielded any particular long-term learning compared to having no reward. In the classroom, short-term performance can be created through giving the children engaging activities. But just because the children look busy and can produce instant answers related to a topic does not mean that they are going to have gained any long-term learning benefits from the process.

“... it is perfectly possible to have a situation in which learning takes place without performance, and indeed conversely where performance takes place without learning.”



Since Tolman and Honzik's studies, numerous pieces of evidence have shown that this distinction between learning and performance is also present in humans. Thus, it is perfectly possible to have a situation in which learning takes place without performance, and indeed conversely where performance takes place without learning. Again, as a classroom example, the latter might take place if children were engaged in what appears to be a meaningful activity using the key words needed to understand the curriculum and, although they appeared to be performing well in the actual lesson, you found that the following week they were still unable to remember the terminology and its meaning. Performance in this context is defined as the student knowledge, skills and understanding that can be observed and measured during the teacher's lesson; whereas learning is seen as outcomes of the teaching that can be observed after the instruction or training has been completed and/or at a later date (for a discussion see Soderstrom and Bjork, 2015).

Changing the way you teach so you vary the conditions to which your learners' brains are exposed

Varying the conditions for practice instead of keeping the teaching approach constant (i.e. changing the way you teach when you repeat information and processes to learners so the children rehearse the same information in different ways) can have a positive effect on learning. As one piece of research has shown (Smith et al., 1978), simply teaching the same material in two different rooms compared to teaching the same material in the same room twice enhances learning. With regard to this particular finding, it appears that some of the old wisdom about study skills (such as deciding on a single place to do your revision and always doing it there) may not be so useful after all. If you think about what we said regarding attention and learning in Chapter 2, you perhaps will not be too surprised about this finding. After all, your brain can get very bored very easily, and without grabbing its attention you are not going to trigger working memory as effectively as you need to. Varying the conditions of practice probably works either because it is restoring some of the salience to the stimulus that was there when the brain first learned the content or because it is increasing the range of associations.

The phenomenon of enhanced learning resulting from varying the conditions of practice similarly occurs when children are learning physical skills, not just academic content and processes. In another study (Kerr and Booth, 1977), 64 children were required to practise beanbag throwing at a target placed on the floor. The researchers used a screen to stop the children seeing the target just before they made their throw. Half of the children did their throws from a fixed distance, whereas the other half had the distance varied – some closer to, some further from the target. Following this, they tested all the children in both groups at the distance used for the fixed practice group. Contrary to what you might predict, the children who had never practised at the fixed distance outperformed the children who had practised at that distance. It also appears to be the case that the more unusual and novel the varied conditions are, the better the enhanced outcomes, with this effect present irrespective of gender (Kerr and Booth, 1977).

“The phenomenon of enhanced learning resulting from varying the conditions of practice similarly occurs when children are learning physical skills, not just academic content and processes.”



Reflection 7.1. Varying the conditions within which your learners practise their knowledge or the application of knowledge

Think about a class and an area of the curriculum that you are planning to teach soon, particularly one that your students often find difficult to grasp without repetition. Think about the different ways in which you could vary the way in which the children will re-experience the content and ideas that you need them to understand.

Think about things like:

- shifting the location you do the practice in
- varying room layout
- changing the type of pedagogical approach you use (whole class, group work, individual working)
- varying the way materials are presented (visual, auditory, physical activity, spatial activity)
- finding a way to approach the material from a completely different angle.

In the past we would probably have avoided doing some of these sort of things in the classroom because we might have thought them distracting. But from the perspective of the desirable difficulties research, making things less easy in the short term is more likely to enhance longer term learning.

Building in time for forgetting by using spaced learning

Spaced learning involves having breaks between repeated content rather than massing those repetitions into a single blocked learning episode (Bjork and Allen, 1970; Cepeda et al., 2006; Kelley and Watson, 2013). As we discussed in detail in Chapter 2, the brain's attentional systems select stimuli from the environment for encoding into memory. One frequent way that long-term memory encoding of information takes place occurs when memory systems are stimulated through repetition of the input separated by spaces of time without that stimulus. Research has demonstrated this in a variety of species including insects (see Menzel et al., 2001 for

“Spaced learning involves having breaks between repeated content rather than massing those repetitions into a single blocked learning episode.”



Neuroscience for teachers

a study of spaced learning with honey bees). The evidence suggests that fundamental mechanisms may be influencing spaced learning, mechanisms which may have been important for survival going back to the early history of life.

Like many areas of neuroscience and learning, translating the neuroscience findings that support the use of spaced learning by teachers into actual teacher/classroom practice is challenging. Kelley and Watson (2013) conducted one such successful study (see Research Zone 7.1).

Research Zone 7.1. Making long-term memories in a matter of minutes

Kelley and Watson (2013) conducted a study involving 440 urban secondary school students in England between the ages of 13 and 15. Their research demonstrated that two approaches to spaced learning had a significantly greater impact on test performance than a considerable amount of traditional teaching time. Specifically, the spaced learning groups had only 60 minutes of instruction compared to the control group who received the same content embedded over four months (a total of 23 hours of direct teacher instruction).

The type of spaced learning applied in the study involved three intensive instructional episodes repeating the same content with some minor variations. Each of these episodes lasted for a maximum of 20 minutes. Spaced in-between the instruction were two “distractor” activities of 10 minutes. The teachers developed a range of distractor options such as juggling, basketball practice and modelling with clay-type materials (Kelley and Watson, 2013: 4). Distractor activities are used to create a clear break between the learning episodes, so the learners’ brains are likely to perceive the next episode as a second encounter with the learning (see our discussions in Chapter 2 regarding the benefits of breaking learning up into different episodes). The important thing is that they contrast to the actual learning rather than being of any particular type. One interesting aspect of the approach emerged from the challenge of delivering the volume of content (at speed) within the instructional sessions. This needed careful planning and in some cases the delivery of the content by pairs of teachers.

Remarkably, for the children in the experimental group, the effect of one hour of spaced learning was equivalent to four months’ teaching in the control group. This suggests that spaced learning approaches could provide more efficient ways of delivering subject areas where the acquisition of facts is required, thus releasing teaching time to allow for the coverage of more collaborative and enquiry-based learning.

Teachers and educators have begun to develop several approaches to spaced learning in response to studies like Kelley and Watson’s. One approach is the 20/10 method (see Reflection 7.2).

Reflection 7.2. Applying the 20/10 method

One approach to spaced learning is the 20/10 method. The name 20/10 is derived from the simple fact that the approach involves the alternation of 20 minutes of intensive rapid teacher input with 10 minutes of “distractor” activity.

In this form of 20/10 teaching the following sequence has been applied:

- 1 Intensive and condensed study of the target information for a maximum of 20 minutes.
- 2 Immediately following step 1, there is a 10 minute physical or distractor activity to induce forgetting.
- 3 Repeated study of the target content with focus on recall (maximum 20 minutes).
- 4 A 10 minute physical or distractor activity to induce forgetting.
- 5 Repeated study of the target content with an application focus (maximum 20 minutes).

Think about an area of the curriculum that you are planning to teach next week. Use the planning grid below to plan the lesson so that it follows a 20/10 structure.

Stage in the process	Things to consider
Intensive and condensed study session 1	What content will you cover? How will you cover it?
Distractor activity 1	What will you get the children to do?
Repeated intensive and condensed study session 2	How will you repeat the content this time with an emphasis on recall?
Distractor activity 2	What will you get the children to do?
Repeated intensive and condensed study session 3	How will you repeat the content this time with an emphasis on application?

The sort of distractor activities that some teachers have used include breathing exercises, clay modelling, playing with Lego, basketball practice and playing musical chairs. As long as your distractor activity is sufficiently engaging to grab your learners’ full attention and thus wipe their working memory of the prior condensed content, you should be inducing the sort of forgetting needed to space the learning out.

Testing more often and seeing testing as a learning event, not just a way of measuring outcomes

There have been a number of debates about testing, with some teachers taking a negative view; although it is, of course, entirely true that children will not get better just by testing. In the same way that a farmer who just weighs the chickens, and does not feed them, is going to find that the chickens are not getting any fatter, so it is also the case with the measurement of outcomes in teaching. If testing is to be most effective, it clearly needs to take place with an overall diagnostic purpose in mind, so the results inform the teacher's next choices, practice and feedback in the classroom.

This said, the evidence from neuroscience and cognitive psychology does not support the assertion that there is no value in testing just as a classroom activity. In fact, the opposite is the case. The prevailing theory for this is that the very act of attempting to retrieve some knowledge, skill or understanding improves the likelihood of successful retrieval later (see Bjork and Bjork, 2011). Indeed, the positive effects of testing can be greater than the effects achieved from revision study sessions (Roediger and Karpicke, 2006a, 2006b). Furthermore, studies have shown that giving a test improves the effectiveness of later study sessions even when the children have been put in a position where they are guaranteed to provide incorrect answers (Kornell et al., 2009). The use of testing as a learning event in its own right to enhance long-term retention and recall is known as “retrieval practice” (Roediger and Butler, 2011) (see Research Zone 7.2).¹

Research Zone 7.2. Even when you cannot remember something, the very act of trying to recall the information may help that information be recalled in the future

As we have seen, testing enhances learning, but what happens if you cannot remember the answer? Does unsuccessful remembering have a negative effect on future recall or does it improve it? This is a question Nate Kornell and colleagues (2009) set out to investigate. As they noted in the conclusions to their paper, people often raise the fear that setting tests that are too hard could harm children. For instance, the United States Department of Education went as far as to raise the question, “Is it harmful for a learner to produce an answer that has a high likelihood of being an error?” (Pashler et al., 2007: 22).

¹ Effects related to this area have now been shown not only to exist at the level studied by the cognitive psychologist but also to be present at a neuronal level (Vestergren and Nyberg, 2014; van den Broek et al., 2016).

Kornell et al. investigated this area in a series of experimental studies. Contrary to some teacher and government beliefs, they found that, irrespective of whether a person is successful or not at retrieving information from their memory, the very act of attempting to retrieve the information enhances their subsequent learning of that information. This has important implications for the classroom and suggests that setting challenging tests may be better than trying to avoid error making with the testing that you use.

The ideas presented in Research Zone 7.2 are not new. William James proposed as early as the end of the 19th century that the active retrieval of information from long-term memory should improve memory (James, 1890). In the last 10 years, evidence has grown substantially across studies that have looked specifically at the value of testing in classroom contexts and with school-age children beyond the laboratory context or in a simulated classroom environment (Larsen et al., 2009; Butler, 2010; Little and Bjork, 2010; Orr and Foster, 2013; Benassi et al., 2014).

Multiple-choice tests and how they can help

You may similarly be surprised to know that exposure to a multiple-choice test that children do not know the answers to can improve the learning of that information on a later occasion and can even enhance the learning of information that is not part of the test (Little and Bjork, 2010). Little and Bjork interpret these findings by suggesting that when children do not know the answer to a multiple-choice question they may attempt to retrieve other relevant information in order to assess why the other answers are not correct in order to ultimately deselect the most unlikely answers. Engaging in this form of mental processing leads to the test becoming a learning event in itself – one that engages not only the tested information but also untested information that has been recalled in order to reject the least likely answers.

“... exposure to a multiple-choice test that children do not know the answers to can improve the learning of that information on a later occasion.”



In addition, research has shown other forms of testing to be more effective than the classroom activities often used by teachers as revision tools. For example, retrieval practice using test questions assessing comprehension and requiring students to make inferences is more effective than concept-mapping approaches (Karpicke and Blunt, 2011). This said, and as we discussed above, the diagnostic use of testing is also important. Furthermore, it is vital with regard to multiple-choice questions to give effective feedback to mitigate potential negative effects (such as feelings of failure) (Butler and Roediger, 2008).

Connecting the testing effect to learning and remembering

On the surface, the testing effect (discussed above) might appear somewhat surprising. But if you take a moment to remember how learning and remembering take place in the brain, it is clear that testing is going to help with memory (particularly long-term memory). You will recall from Chapter 2 that rehearsal (salient real-world forms of repetition and application) is a key way of getting knowledge, skills and understanding into long-term memory. Likewise, when you recall something from long-term memory, you place it back into working memory where you can think about that knowledge and process it again. In essence, when you retrieve some knowledge, skill or understanding from long-term memory, its representation in working memory is modified in such a way that it becomes easier to recall in the future (Bjork, 1975). Research Zone 7.3 gives an example of how more frequent testing could be beneficial.

Research Zone 7.3. An exam a day may keep the failure away (with older students)

Leeming applied an “exam-a-day” approach to the teaching of an introductory psychology course (Leeming, 2002). His study showed that students who completed a test every day, compared to those taking a single large exam, achieved results that were significantly greater when both groups took their final end-of-semester test.

In his research, he gave students in four classes a short exam at the beginning of every class they attended. The end-of-semester test results were significantly better for these students than the results that students had achieved in previous classes with the same teaching materials. Indeed, the positive effect was present after only four tests. Furthermore, Leeming found fewer withdrawals from the classes that had an “exam a day”. Although Leeming does not appear to have controlled for individual difference between the classes, the students were all of a similar academic level and he was careful to ensure that teaching style and the format of the lessons were the same.

In addition to carrying out an outcome test, Leeming surveyed the students who had experienced the “exam-a-day” procedure. Most of the students thought that having a test every day led them to study more and to learn more effectively. They also said that they liked the approach.

There is also good evidence from education studies to support the use of practice testing as a learning experience in its own right. But one thing to bear in mind is that the testing effect probably decreases – and may even disappear – as the complexity of learning materials increases (van Gog and Sweller, 2015), suggesting the effect may be useful for the acquisition of basic knowledge but less useful for content that requires high-order thinking and processing. Research Zone 7.4 summarises the evidence that is discussed in a recent meta-analysis.

Research Zone 7.4. Practice testing – the education evidence

The evidence associated with the learning benefits of taking practice tests, compared to other forms of non-test learning conditions (like filler activities, general practice, re-studying or control conditions such as no re-presentation of material), is becoming increasingly strong. Adesope et al. (2017) completed a meta-analysis of 272 independent effect sizes across 188 different experiments. The use of practice tests was associated with a significant moderate weighted mean effect size compared to re-studying (0.51). A much larger effect size was detected when practice tests were compared to filler activities or no re-presentation of material (0.93). It also appears that the amount, frequency and formatting of practice tests is important. Practice tests with multiple-choice options yield a larger effect size (0.70) than short-answer tests (0.48). Furthermore, it appears that the administration of a single test before a final test point is more effective than the taking of several tests prior to that point, while leaving a gap of less than a single day between the practice test and final test produced a smaller effect size (0.56) than leaving a gap of between one and six days (0.82).

The benefits of more frequent testing: some emerging conclusions

Drawing all the recent evidence together, Roediger and colleagues (2011) suggest ten ways teachers should see testing as a beneficial additional classroom strategy. In summary, and as well as supporting the reasons why many teachers test already, testing:

- can be used to identify gaps in knowledge
- can produce better organisation of knowledge in students' minds
- improves the transfer of knowledge to new contexts
- provides teachers with feedback on the effects of their delivery and the approaches they choose
- if used frequently, encourages students to undertake independent study.

The new evidence from cognitive psychology and neuroscience suggests that effective use of testing can also:

- act as a form of retrieval practice in itself and so aid later retention
- lead to students learning more from the next episode of study
- help to establish prior learning and prevent it from interfering with new learning
- facilitate the retrieval of knowledge that was not tested
- improve metacognitive monitoring (students' awareness of what they are learning).

Reflection 7.3. Thinking about ways in which you could use test-enhanced learning in your classroom

Obviously, you are not going to replace your usual practice completely and just do test-enhanced learning. As we have seen, it is important to have a variety of teaching strategies and to vary the way you approach a topic. This said, the evidence is overwhelming that there is advantage in building the use of testing into your practice more frequently and seeing this as another form of learning event.

There are a number of ways that you could do this – you can find some suggestions below. This is not an exclusive list and you will be able to think of more ideas the more you use this.

- Explain the testing effect to your learners so they too see testing as a learning event in its own right and understand how it can help them to learn more effectively.
- Build quizzes into your lesson plans more frequently – these could be short-answer quizzes or some form of multiple-choice questionnaire. Remember that the evidence suggests that providing feedback on the answers students give is also important and increases the effect.
- Give the children a pre-test that exposes them to the learning they are going to have on a topic, explaining to them that you do not expect them to get all the answers right. Remember that even getting an answer wrong on a pre-test can enhance future learning by requiring the learners to search for possible answers from their existing knowledge.
- Get your learners to sum up the learning and provide summary bullet points that they share with each other as a form of self-induced testing; remembering, of course, to do this in a way that does not make the children feel self-conscious or embarrassed. You could do this by putting aside part of the end of the lesson to get the children to organise their learning and describe it in their own way.

Looking at the four approaches above, think about a series of lessons (or scheme of work) that might benefit from you using this approach for the first time.

Interleaving the content you teach

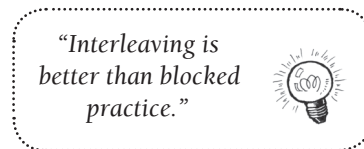
Several recent studies have shown that interleaving content, rather than the massed blocking of instruction, is a way of enhancing the learning and retention of skills, knowledge and understanding (Figure 7.1) (see e.g. Taylor and Rohrer, 2010; Rau et al., 2013), although there is not yet as much evidence to support this as there is for spaced learning.



Figure 7.1. Interleaving content.

Interleaving content is probably effective for similar reasons to spaced learning, but instead of the alternative episode of activity being a “distractor” activity, different actual subject content or tasks are used between the episodes rather than introducing a break from the learning.

Studies on interleaving, and its effect on short-term performance compared to long-term learning, have given us some of the best evidence for the potential of including desirable difficulties in learning situations. Shea and Morgan (1979), in a motor skills study, initially showed that participants who undertook block practice during the learning episode improved more rapidly than participants who had experienced interleaving. But when these participants returned after 10 days and were re-tested, the participants who had experienced the interleaved conditions performed significantly better than those who had had the blocked practice. There is also evidence for the positive effects of interleaving in maths lessons (Rohrer and Taylor, 2007) and with inductive reasoning contexts (such as learning to recognise an artist’s style) (Kornell and Bjork, 2008).



What is particularly interesting about interleaving is the way in which it affects the perception of learners with regard to their actual learning outcomes, compared with blocking. For example, in one study (Kornell and Bjork, 2008), when learners were asked if they had performed better or the same in blocked compared with interleaved sessions, it was clear that although participants thought they had learned more in blocked sessions (compared with interleaved sessions), in fact they performed far higher in the interleaved learning conditions. This finding further illustrates the way in which people tend to assume that their performance has been lower in more challenging situations (such as interleaved content) when in fact the opposite may be the case.

Getting the learners to make things up themselves – the generation effect

Researchers use the term “generation effect” to describe the advantage of using learner-created materials over materials created by the teacher or lecturer (Slamecka and Graf, 1978). Looking at outcomes for students, these effects rival spaced learning in terms of their potential to enhance existing effective teacher strategies (Bertsch et al., 2007).

From a classroom perspective, the evidence above suggests that any time you can get learners to come up with the answer themselves is going to be better than just telling them the answer. This maps strongly across to the evidence supporting the use of effective questioning and leaving waiting time (rather than jumping in) to encourage thinking (Creemers and Kyriakides, 2006; Coe et al., 2014). Other ways to create a generation effect in the classroom could include:

- answering review questions at the end of a chapter of a textbook without looking back at the pages
- getting learners to read a section of a book, close the book and then generate and answer questions about what they have just read
- creating original flashcards and using them.

Drawing on the cognitive psychology research from the wider field of the science of learning, you could in addition:

- Make more frequent use of **elaborate interrogation** (Pressley et al., 1989; Smith and Kimball, 2010) – getting children to explain why a particular fact or concept is true.
- Include **self-explanation** in learning activities – requiring the children to explain how new information relates to already learned information or getting them to explain or generate the steps that have to be taken during any problem-solving (Rittle-Johnson, 2006; de Koning et al., 2011).²

“... any time you can get learners to come up with the answer themselves is going to be better than just telling them the answer.”



With regard to the “desirable difficulties” research, Bloom and Lamkin (2006) showed positive effects on long-term learning with students taught to use the acrostic first-letter mnemonic method. This method involves the teacher asking the students to take the first letters of words in a list that needs to be learned and allowing the students to create their own unusual (or meaningful) sentence – or use one provided by the instructor. One well-known example of this approach is the way that music teachers frequently teach the order of sharps in a key signature (F, C, G, D, A, E, B) – “Father Charles Goes Down And Ends Battle”. In one study, Bloom and Lamkin (2006) found that those

² For a wider review of cognitive psychology evidence helpful to teachers, see Dunlosky et al. (2013) and Benassi et al. (2014).


learners who self-generated acrostics attained more on post-intervention tests than a control group and showed no decline in learning between two and ten weeks (in contrast with the control which showed a clear decline in learning).

Other studies appear to confirm that self-generated mnemonics may be even more effective than instructor-generated mnemonics (Slamecka and Graf, 1978; Hirshman and Bjork, 1988), and that teachers can similarly achieve this effect by getting learners to generate target materials through the use of a puzzle rather than simply reading (McDaniel et al., 1994). In McDaniel and colleagues' research this involved reordering randomly ordered sentences into a coherent structure. In the classroom this might involve taking a set text (such as a poem or description of something from a textbook), cutting up the text into individual sentences and getting the children to see if they can sort the sentences into their previous order and paragraphs.

You can tap into the same phenomenon that underpins the generation effect by making the materials you share with your learners less clearly organised and/or by providing only some of the background knowledge (McNamara et al., 1996). There is evidence that this effect extends to things such as using fonts that are slightly harder to read (Diemand-Yauman et al., 2011). In their paper, entitled 'Fortune favors the bold (and the italicized): effects of disfluency on educational outcomes', Diemand-Yauman and colleagues report on two studies. These build on prior research which shows that disfluency (the subjective experience of things being difficult that is associated with cognitive operations) results in deeper processing. In the first study, the researchers found that when they presented people with information in harder-to-read fonts they were better able to remember the information compared to easier-to-read fonts. In a second study, they extended these findings to high school students in real-life classrooms and found similar results.

The worked example effect versus the generation effect

An interesting conundrum exists with regard to the generation effect, namely that it seems to sit in opposition to evidence that the provision of full guidance on how to solve a problem is more likely to result in higher test performance compared to giving no guidance (known as the **worked example effect** – see Retnowati et al., 2010; van Gog et al., 2010). In a recent study, Chen and colleagues (Chen et al., 2015) tested the hypothesis that worked example effects are more likely to take place with complex materials that place a greater load on working memory, in contrast to the generation effect that is present with simpler content. They ran two experiments which suggested that in situations where students are required to learn material that places lower demands on working memory, learning appears to be enhanced by the generation effect, but where the material is more complex and more demanding, a worked example study approach is likely to be more effective. The worked example effect could also be seen as a form of in-depth modelling (see the discussions in Chapter 3). You can find more examples of research into the generation effect in Research Zone 7.5.



Research Zone 7.5. Jumbling up and leaving a gap – evidence supporting the use of the generation effect in the classroom

The term “generation effect” refers to the increased ability of people to learn and remember things if they have had to “generate” some of the answer for themselves when they are learning the content for the first time. One way that this has been shown to be possible with individual word learning is by giving children a letter-stem cue (like ar__ for “argon”) or by getting the children to unscramble an anagram (agrno) when they are learning the words. This technique has been shown to produce better recall when the children are asked to recall the words or answers at a later date.

In a study that examined the generation effect using typical classroom text materials, DeWinstanley and Bjork (2004) found that when children were told to generate certain “target words” within a paragraph, they had better recall of those words than if they were just read. Furthermore, when the children were then, on a later occasion, asked to read similar paragraphs containing the target words (with these either read or generated), both forms of word resulted in the same higher levels of recall.

DeWinstanley and Bjork argue that these findings show that having an earlier experience of word generation resulted in the children becoming more effective at remembering the paragraphs as a whole. Further research into this phenomenon has implied that this enhanced effect is caused (at least to some extent) by enhanced memory of the context that the target words were placed in (Little et al., 2011).

Next steps

The desirable difficulties research suggests a number of possible areas for inclusion in classroom practice that could benefit student outcomes. It also makes us question some of the preconceptions and ideas about what effective pedagogy is that have developed over the last 40 years or so.

In particular, the evidence we have presented in this chapter challenges teachers to think about the fact that by making short-term performance easy, they may be making long-term learning harder. From this perspective, the wealth of popular teaching books illustrating 1001 activities to use with children may have overemphasised short-term performance and the appearance of learning rather than the achievement of long-term learning. Not that there is anything wrong with making teaching engaging, interesting and fun, of course – as we have explored in earlier chapters; nor is there anything wrong with having many different approaches (something that the varied practice aspect of the desirable difficulties research supports). This said, if teachers overemphasise ease of performance then they may not be doing the children any favours in terms of helping them to achieve their long-term potential.

By extension, the research reminds us of the importance of effective classroom dialogue and feedback. Specifically, it reminds us of the need to make that dialogue challenging in order to get children to generate their own answers and make them think about the content they are learning – for instance, by using more open questioning and avoiding just giving the children the answers. Overall, the findings challenge us to devise classroom strategies that encourage deeper processing and avoid surface level activity. Thinking about this, you may want to take a moment to consider those areas of your practice that are currently focused on creating a comfortable environment, at the expense of generating the sort of benefits that can be achieved by varying your approaches, interleaving, spacing content and the accessing of retrieval effects through testing and getting learners to generate content.

Neuroscience for teachers

Now that you have read about the five main strategies that form part of this strand of research, think about each one and how you might be able to apply them or the principles that underpin them:

- Varying the conditions of practice instead of keeping the teaching approach constant.
- Spacing learning rather than the massed delivery of information – perhaps by including some short distractions such as “mind breaks” in-between the episodes.
- Using testing as a learning experience in its own right.
- Interleaving different content rather than block delivery of instruction in order to draw on the “power of forgetting” (see Chapter 2 and our discussions of the primacy and recency effect).
- Making use of the generation effect by getting learners to generate content and explanations for themselves more often.