

5

Learning

Making Things Hard on Yourself, But in a Good Way: Creating Desirable Difficulties to Enhance Learning

Elizabeth L. Bjork and Robert Bjork

University of California, Los Angeles

- *Please describe your current position and research interests.*

Elizabeth Ligon Bjork: I am Professor and Senior Vice Chair of Psychology at the University of California, Los Angeles, Academic Senate. My research interests have included visual attention and developmental processes but now focus on practical and theoretical issues in human memory and learning, particularly the role that inhibitory processes play in an efficient memory system.

Robert A. Bjork: I am Distinguished Research Professor of Psychology at the University of California, Los Angeles. My research focuses on human learning and memory and on the implications of the science of learning for instruction and training.

- *How did you get interested in studying the facilitating effect of apparent impediments to learning?*

Elizabeth Bjork: My interests in optimizing learning were triggered by interactions with students lamenting during office hours how hard they had studied, only then to perform poorly on a just-given exam. This motivated me to examine why students' study activities were sometimes so ineffective.

Robert Bjork: My interests go back to my efforts—as a graduate student—to understand the relationship of forgetting and learning, especially why inducing forgetting often enhances subsequent learning. My interests in the application of “desirable difficulties” were fanned by my experiences teaching and coaching and from what I learned as Chair of the National Research Council Committee on Techniques for the Enhancement of Human Performance (1988–1994).

- *What has been the real-world impact of this work?*

Overall, the impact has been modest, but there are multiple indications that the impact of basic research findings on educational practices is increasing. It is a slow process, in part, because optimizing instruction requires some unintuitive innovations in how the conditions of instruction are structured.

*Making Things Hard on Yourself, But in a Good Way:
Creating Desirable Difficulties to Enhance Learning*

As teachers—and learners—the two of us have had both a professional and personal interest in identifying the activities that make learning most effective and efficient. What we have discovered, broadly, across our careers in research, is that optimizing learning and instruction often requires going against one's intuitions, deviating from standard instructional practices, and managing one's own learning activities in new ways. Somewhat surprisingly, the trials and errors of everyday living and learning do not seem to result in the development of an accurate mental model of the self as learner or an appreciation of the activities that do and do not foster learning.

The basic problem learners confront is that we can easily be misled as to whether we are learning effectively and have or have not achieved a level of learning and comprehension that will support our subsequent access to information or skills we are trying to learn. We can be misled by our subjective impressions. Rereading a chapter a second time, for example, can provide a sense of familiarity or perceptual fluency that we interpret as understanding or comprehension, but may actually be a product of low-level perceptual priming. Similarly, information coming readily to mind can be interpreted as evidence of learning, but could instead be a product of cues that are present in the study situation, but that are unlikely to be present at a later time. We can also be misled by our current performance. Conditions of learning that make performance improve rapidly often fail to support long-term retention and transfer, whereas conditions that create challenges and slow the rate of apparent learning often optimize long-term retention and transfer.

Learning versus Performance

This apparent paradox is a new twist on an old and time-honored distinction in psychology—namely, the distinction between learning and performance. Performance is what we can observe and measure during instruction or training. Learning—that is, the more or less permanent change in knowledge or understanding that is the target of instruction—is something we must try to infer, and current performance can be a highly unreliable index of whether learning has occurred. For a review of the history and current status of the learning versus performance distinction, see Soderstrom and Bjork (in press).

Learning Without Performance and Performance Without Learning

Decades ago, learning theorists were forced to distinguish between learning and performance because experiments revealed that considerable learning could happen across a period when no change was apparent in performance. In latent-learning experiments with animals, for example, periods of free exploration of a maze, during which the animal's behavior seemed aimless, were shown—once reward was introduced—to have produced considerable learning. Similarly, in research on motor skills, investigators found that learning continued across trials during which the build-up of fatigue suppressed performance.

More recently, a variety of experiments—some of which we summarize below—have demonstrated that the converse is true as well: Namely, substantial improvements in performance across practice or training sessions can occur without significant learning (as revealed after a delay or in another context). To the extent, therefore, that people interpret current performance as a valid measure of learning, they become susceptible to misassessing whether learning has or has not occurred.

Storage Strength Versus Retrieval Strength

At a theoretical level, we (Bjork & Bjork, 1992) distinguish between the storage strength and the retrieval strength of information or skills stored in memory. Storage strength reflects how entrenched or interassociated a memory representation is with related knowledge and skills, whereas retrieval strength reflects the current activation or accessibility of that representation and is heavily influenced by factors such as situational cues and recency of study or exposure. Importantly, we assume that current performance is entirely a function of current retrieval strength, but that storage strength acts to retard the loss (forgetting) and enhance the gain (relearning) of retrieval strength. The key idea for present purposes is that conditions that most rapidly increase retrieval strength differ from the conditions that maximize the gain of storage strength. In other words, if learners interpret current retrieval strength as storage strength, they become susceptible to preferring *poorer* conditions of learning to *better* conditions of learning.

Introducing Desirable Difficulties to Enhance Learning and Instruction

So what are these better conditions of learning that, while apparently creating difficulty, actually lead to more durable and flexible learning? Such desirable difficulties (Bjork, 1994; 2013)

include varying the conditions of learning, rather than keeping them constant and predictable; interleaving instruction on separate topics, rather than grouping instruction by topic (called blocking); spacing, rather than massing, study sessions on a given topic; and using tests, rather than presentations, as study events.

Before proceeding further, we need to emphasize the importance of the word *desirable*. Many difficulties are undesirable during instruction and forever after. Desirable difficulties, versus the array of undesirable difficulties, are desirable because they trigger encoding and retrieval processes that support learning, comprehension, and remembering. If, however, the learner does not have the background knowledge or skills to respond to them successfully, they become undesirable difficulties.

Varying the Conditions of Practice

When instruction occurs under conditions that are constrained and predictable, learning tends to become contextualized. Material is easily retrieved in that context, but the learning does not support later performance if tested at a delay, in a different context, or both. In contrast, varying conditions of practice—even varying the environmental setting in which study sessions take place—can enhance recall on a later test. For example, studying the same material in two different rooms rather than twice in the same room leads to increased recall of that material (Smith, Glenberg, & Bjork, 1978)—an empirical result that flies in the face of the common how-to-study suggestion to find a quiet, convenient place and do all your studying there.

A study of children's learning provides a striking illustration of the benefits of varying conditions of practice. Eight-year-olds and 12-year-olds practiced throwing beanbags at a target on the floor with their vision occluded at the time of each throw. For each age group, half of the children did all their practicing throwing to a target at a fixed distance (for example, 3 feet for the 8-year-olds), while the other half threw to targets that were closer or farther away. After the learning sessions and a delay, all children were tested at the distance used in the fixed-practice condition for their age group (Kerr & Booth, 1978).

Common sense would suggest that the children who practiced at the tested distance would perform better than those who had never practiced at that distance, but the opposite was true for both age groups. The benefits of variation—perhaps learning something about adjusting the parameters of the motor program that corresponded to the throwing motion—outweighed any benefits of being tested at the practiced distance. Many other studies have shown that when

testing after training takes place under novel conditions, the benefits of variation during learning are even larger.

Spacing Study or Practice Sessions

The effects of distributed practice on learning are complex. Although massing practice (for example, cramming for exams) supports short-term performance, spacing practice (for example, distributing presentations, study attempts, or training trials) supports long-term retention. The benefits of spacing on long-term retention, called the spacing effect, have been demonstrated for all manner of materials and tasks, types of learners (human and animal), and time scales; it is one of the most general and robust effects from across the entire history of experimental research on learning and memory.

Rather than describing any of the myriad studies that have demonstrated the benefits of spacing, we will simply stress the importance of incorporating spacing and avoiding massing in managing learning. Massing repeated-study activities is often not only convenient, but it can also seem logical from the standpoint of organizing one's learning of different topics, and it frequently results in rapid gains in apparent learning. Good test performance following an all-night cramming session is certainly rewarding, but little of what was recallable on the test will remain recallable over time. In contrast, a study schedule that spaces study sessions on a particular topic can produce *both* good exam performance *and* good long-term retention. Furthermore, because new learning depends on prior learning, spacing study sessions optimally can also enhance transfer of knowledge and provide a foundation for subsequent new learning.

Whether students so frequently mass their study activities because they believe that massing is good, or simply because they get behind and are forced into massed study sessions before exams, is not clear. There is recent evidence, though, that students believe that information that is more important or valuable should be studied again sooner, rather than later, perhaps owing to the intuition that restudying should happen before too much forgetting has occurred (Cohen, Yan, Halamish, & Bjork, in press). Such an intuition, however, leads students to avoid spacing repeated study sessions—and, of course, also to avoid the benefits of spacing.

Interleaving versus Blocking Instruction on Separate To-Be-Learned Tasks

Interleaving the practice of separate topics or tasks is an excellent way to introduce spacing and other learning dynamics. In a classic comparison of interleaving and blocking (Shea & Morgan,

1979), participants practiced three different movement patterns, each requiring the participants to knock down three of six hinged barriers rapidly on a pinball-like apparatus in a prescribed order. All participants received 18 trials on each pattern, but in the interleaved condition, practice on a given trial was randomly determined, whereas in the blocked condition, one pattern at a time was practiced.

As you probably suspect, participants given blocked practice improved more rapidly than those given interleaved/random practice. Thus, if the researchers had stopped their study at the end of training, blocking of practice would have seemed the superior learning procedure. But, instead, participants returned 10 days later and were retested under either blocked or interleaved/random conditions. Under interleaved/random testing conditions, participants who had practiced under interleaved conditions performed far better than did the blocked-practice participants, who appeared, when tested under a random schedule, to have learned virtually nothing. Under blocked testing conditions, performance was essentially the same for both groups, but the small difference still favored the interleaved group.

The skills literature includes many replications of the pattern that blocked practice *appears* optimal for learning, but interleaved practice actually results in superior long-term retention and transfer of skills (for a review, see Lee, 2012), and research illustrates that learners—as well as instructors—are at risk of being fooled by that pattern. For example, when participants who had learned three different keystroke patterns were asked to predict their performance on a test the next day, those given interleaved practice predicted their performance quite closely, whereas those given blocked practice were markedly overconfident (Simon & Bjork, 2001). In effect, the blocked-practice group misinterpreted their good performance during practice as evidence of long-term learning, rather than a product of the local (that is, blocked) conditions. Said differently, they misinterpreted the retrieval strength of a given keystroke pattern as an index of its storage strength.

Other results illustrate that the benefits of interleaved practice extend beyond the learning of motor skills. For example, when participants were asked to learn formulas for calculating the volumes of different solids, such as a truncated cone, in either a blocked or interleaved manner, interleaved instruction enhanced performance on a delayed test. The size of the long-term advantage of interleaved practice was striking: 63 percent versus 20 percent of new problems worked correctly a week later (Rohrer & Taylor, 2007).

More recently and surprisingly, we have found that interleaving even enhances inductive learning (Kornell & Bjork, 2008). When participants were asked to learn the styles of each of 12

artists based on a sample of 6 paintings by each artist, interleaving a given artist’s paintings among the paintings by other artists—versus presenting that artist’s paintings one after another (blocking)—enhanced participants’ later ability to identify the artist responsible for each of a series of new paintings. This result is surprising because blocking would seem to make it easier to note the commonalities that characterize a particular artist’s style. Indeed, as illustrated in Figure 1, the majority of participants—when asked after the test whether interleaving or blocking had helped them learn an artist’s style better—definitely had the impression that blocking had been more effective than interleaving, the opposite of their actual learning. Blocking may indeed have facilitated noticing commonalities, but the final test required distinguishing among the artists, and interleaving may have fostered learning the differences as well as similarities among the styles of different artists.




Figure 1. The left panel shows the proportion of participants who selected Blocked, Interleaved, or The Same in response to the question: Under which condition do you believe you learned better? The right panel indicates the proportion of participants who actually performed better in the blocked or interleaved conditions or performed the same in the two conditions.

Why might interleaving enhance long-term retention and transfer? One theory suggests that having to resolve the interference among the different things under study forces learners to notice similarities and differences among them, resulting in the encoding of higher-order representations, which then foster both retention and transfer. Another explanation suggests that

interleaving forces learners to reload memories: If required to do A, then B, then C, and then A again, for example, the memory for how to do A must be reloaded a second time, whereas doing A and then A again does not involve the same kind of reloading. Such repeated reloadings are presumed to foster learning and transfer to the reloading that will be required when that knowledge or skill is needed at a later time.

From the standpoint of our theoretical framework (Bjork & Bjork, 1992), learning from reloading is an instance of a broader dynamic in human memory: Namely, that forgetting (losing retrieval strength) creates the opportunity for increasing the storage strength of to-be-learned information or skills. Said differently, when some skill or knowledge is maximally accessible from memory, little or no learning results from additional instruction or practice.

Generation Effects and Using Tests (Rather Than Presentations) as Learning Events

An effect that rivals the spacing effect for its generality and its significance for instruction and learning is the generation effect, which refers to the long-term benefit of generating an answer, solution, or procedure versus being presented that answer, solution, or procedure. Basically, any time that you, as a learner, look up an answer or have somebody tell or show you something that you could, drawing on current cues and your past knowledge, generate instead, you rob yourself of a powerful learning opportunity. Retrieval, in effect, is a powerful “memory modifier” (Bjork, 1975).

Closely related to the generation effect are the benefits that accompany retrieving information studied earlier. Much laboratory research (for example, Landauer & Bjork, 1978; Carrier & Pashler, 1992) has demonstrated the power of tests as learning events, and, in fact, a test or retrieval attempt, even when no corrective feedback is given, can be considerably more effective in the long term than reading material over and over. The reason why rereading is such a typical mode of studying derives, we believe, from a faulty model of how we learn and remember: We tend to think of our memories as working much like an audio/video recorder, so if we read and reread or take verbatim notes, the information will eventually write itself on our memories. Nothing, however, could be further from the way we actually learn and remember.

Unfortunately, the effectiveness of tests as learning events remains largely underappreciated, in part because testing is typically viewed as a vehicle of assessment, not a vehicle of learning. As Henry L. Roediger, Kathleen B. -McDermott, and Mark A. McDaniel describe in their essay in this chapter, however, recent research using more educationally realistic materials and retention intervals has clearly demonstrated the pedagogical benefits of tests (for

example, Roediger & Karpicke, 2006). Similar to the pattern with variation, spacing, and interleaving, repeated study opportunities appear, in the short term, to be more effective than repeated testing, but testing produces better recall in the long term.

Two other pedagogical benefits of tests must be mentioned: First, tests have metacognitive benefits in terms of identifying whether information has or has not been understood and/or learned. A student's ability, for example, when going back over a chapter in a textbook, to judge whether information will be recallable on an upcoming examination is severely limited, whereas attempting to answer a fellow student's questions on the chapter can identify what has and has not been learned.

The second, related benefit is that tests can potentiate the effectiveness of subsequent study opportunities even under conditions that insure learners will be incorrect on the test (Kornell, Hays, & Bjork, 2009). Even taking a test before a first reading of to-be-learned material can potentiate one's learning of it, and not only for the specifically pretested information but also for related information (e.g., Little & Bjork, 2011). Thus, when needing to read a new chapter in your textbook that is followed by test questions, you may want to try answering those questions first and then read the chapter. Again, the basic message is that we need to spend less time restudying and more time testing ourselves.

Concluding Comments

For those of you who are students, we hope we have convinced you to take a more active role in your learning by introducing desirable difficulties into your own study activities. Above all, try to rid yourself of the idea that memory works like a tape or video recorder and that re-exposing yourself to the same material over and over again will somehow write it onto your memory. Rather, assume that learning requires an active process of interpretation—that is, mapping new things we are trying to learn onto what we already know. (There's a lesson here for those of you who are teachers—or parents—as well: Consider how you might introduce desirable difficulties into the teaching of your students or children.)

Be aware, too, when rereading a chapter or your notes, that prior exposures create a sense of familiarity that can easily be confused with understanding. And perhaps most importantly, keep in mind that retrieval—much more than restudying—acts to modify your memory by making the information you practice retrieving more likely to be recallable again in the future and in different contexts. In short, try to spend less time on the input side and more time on the output

side, such as summarizing what you have read from memory or getting together with friends and asking each other questions. Any activities that involve testing yourself—that is, activities that require you to retrieve or generate information, rather than just representing information to yourself—will make your learning both more durable and flexible.

Finally, we cannot overstate the importance of learning how to manage your own learning activities. In a world that is ever more complex and rapidly changing, and in which learning on one's own is becoming ever more important, learning how to learn is the ultimate survival tool.

Suggested Further Reading

Bjork, R. A. (2011). On the symbiosis of learning, remembering, and forgetting. In A. S. Benjamin (Ed.), *Successful remembering and successful forgetting: a Festschrift in honor of Robert A. Bjork* (pp. 1-22). London, UK: Psychology Press.

Bjork, R. A., & Bjork, E. L. (2006). Optimizing treatment and instruction: Implications of a new theory of disuse. In L-G. Nilsson and N. Ohta (Eds.), *Memory and society: Psychological perspectives* (pp. 109–133). Hove, East Sussex, England, and New York: Psychology Press.

Bjork, R. A., Dunlosky, J., & Kornell, N. (2013). Self-regulated learning: Beliefs, techniques, and illusions. *Annual Review of Psychology*, *64*, 417-444.

References

Bjork, R. A. (1975). Retrieval as a memory modifier. In R. Solso (Ed.), *Information processing and cognition: The Loyola Symposium*, pp. 123–144. Hillsdale, NJ: Erlbaum.

Bjork, R. A. (1994). Memory and metamemory considerations in the training of human beings. In J. Metcalfe and A. Shimamura (Eds.), *Metacognition: Knowing about knowing* (pp. 185–205). Cambridge, MA: MIT Press.

Bjork, R. A. (2013). Desirable difficulties perspective on learning. In H. Pashler (Ed.), *Encyclopedia of the mind*. Thousand Oaks: Sage Reference.

Bjork, R. A., & Bjork, E. L. (1992). A new theory of disuse and an old theory of stimulus fluctuation. In A. Healy, S. Kosslyn, & R. Shiffrin (Eds.), *From learning processes to cognitive processes: Essays in honor of William K. Estes* (Vol. 2, pp. 35–67). Hillsdale, NJ: Erlbaum.

- Carrier, M., & Pashler, H. (1992). The influence of retrieval on retention. *Memory & Cognition*, 20, 633–642.
- Kerr, R., & Booth, B. (1978). Specific and varied practice of a motor skill. *Perceptual and Motor Skills*, 46, 395–401.
- Kornell, N., & Bjork, R. A. (2008). Learning concepts and categories: Is spacing the “enemy of induction”? *Psychological Science*, 19, 585–592.
- Kornell, N., Hays, M. J., & Bjork, R. A. (2009). Unsuccessful retrieval attempts enhance subsequent learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 35, 989–998.
- Landauer, T. K., & Bjork, R. A. (1978). Optimum rehearsal patterns and name learning. In M. M. Gruneberg, P. E. Morris, & R. N. Sykes (Eds.), *Practical aspects of memory* (pp. 625–632). London: Academic Press.
- Lee, T.D. (2012). Contextual interference: Generalizability and limitations. In N.J. Hodges & A.M. Williams (Eds.), *Skill acquisition in sport: Research, theory, and practice* (2nd ed) (pp. 79-93). London, UK: Routledge.
- Little, J. L., & Bjork, E. L. (2011). Pretesting with multiple-choice questions facilitates learning. In L. Carlson, C. Hölscher, & T. Shipley, (Eds.), *Proceedings of the 33rd Annual Conference of the Cognitive Science Society* (pp. 294-299).
- Roediger, H.L., & Karpicke, J.D. (2006). Test-enhanced learning: Taking memory tests improves long-term retention. *Psychological Science*, 17, 249–255.
- Rohrer, D., & Taylor, K. (2007). The shuffling of mathematics practice problems improves learning. *Instructional Science*, 35, 481–498.
- Shea, J.B., & Morgan, R.L. (1979). Contextual interference effects on the acquisition, retention, and transfer of a motor skill. *Journal of Experimental Psychology: Human Learning and Memory*, 5, 179–187.
- Simon, D., & Bjork, R. A. (2001). Metacognition in motor learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 907–912.
- Smith, S. M., Glenberg, A. M., & Bjork, R. A. (1978). Environmental context and human memory. *Memory & Cognition*, 6, 342–353.

Soderstrom, N. C., & Bjork, R. A. (in press). Learning versus performance. In D. S. Dunn (Ed.), *Oxford bibliographies online: Psychology*. New York: Oxford University Press.

Tauber, S. K., Dunlosky, J., Rawson, K. A., Wahlheim, C. N., & Jacoby, L. L. (2013). Self-regulated learning of a natural category: Do people interleave or block exemplars during study? *Psychonomic Bulletin & Review*, *20*, 356-363.