Microsoft, and IBM - invest between 2 and 15%. In both kinds of industries, the concept of a research pipeline directly connected to future productivity is salient. By contrast, financial corporations and oil companies spend so little on research that research does not even make it as a line item onto consolidated financial statements.

By these comparisons, the federal budget for education is behaving more like the oil business than the discovery-oriented technology and pharmaceutical firms. But higher education is not in the extraction business, digging students out of high schools and efficiently refining them for the labor market. If education seeks to be more discovery- and risk-oriented, it needs to see its calling as the creation of greater inspiration and opportunity for research.

The public's confidence in education will be increased by valid and reliable research. Through research, we can become more efficient transmitters of the knowledge we create. Each of you can imagine just how efficient we might become if the budget for, say, just one five-billion-dollar aircraft carrier were spent on how to better educate the nation. So, in the end, what is the payoff for research? That is today's topic.

POWERPOINT PRESENTATION
www.sunysb.edu/Reinventioncenter/Conference_04/Kuncl/Powerpoint.pdf

Research on Learning as a Foundation for Curricular Reform and Pedagogy

Speaker: *Elizabeth Ligon Bjork, Professor of Psychology, University of California, Los Angeles*

Over the last few decades, we have been learning a great deal about how people learn and the types of conditions that optimize long-term retention and transfer, and numerous findings from this research have important implications for ways in which we can improve instructional practice. In this presentation, I focus on those results indicating that in order to maximize the effectiveness of instruction and training, we need to pay greater attention to an old distinction in psychology—namely, the distinction between performance and learning—but in a slightly different way than researchers thought about this distinction in the past.

Early investigators of learning were forced to make a distinction between performance and learning when several, now classic, studies showed that—despite the lack of any evidence in an animal's performance during training—learning had nonetheless occurred and could be revealed under the right circumstances, such as when a food reward was introduced into the situation. More recently, a variety of results suggest that what we might think of as a corollary to this earlier distinction needs to be made. Specifically, whereas learning can be occurring with no apparent change in performance during training, improvements in performance during training can occur with little or no durable learning being achieved. Or, put slightly differently, conditions of instruction that make performance improve rapidly often fail to support long-term retention and transfer, while conditions of instruction that appear to create difficulties for the learner, often slowing the rate of *apparent* learning, can actually optimize long-term retention and transfer.

As a consequence of this corollary, performance during training can be a poor and unreliable guide to whether the type of learning that is the goal of our instruction—that is, learning that will be both durable and support transfer—has actually occurred. But, of course, what is readily observable to us as instructors is the performance of our students during instruction and training. Consequently, as instructors, we can easily be misled into using manipulations of training and instruction having the property of enhancing performance during training and instruction, but failing to support learning as measured by long-term retention and the transfer of skills and knowledge. And, conversely, as instructors, we can easily be led away from using conditions that introduce difficulties for the learner and appear to slow the rate of learning, but that are actually enhancing post-training retention and transfer.

A discussion of these latter types of conditions—originally labeled as "desirable difficulties" by Robert A. Bork (1994) to indicate their property of being conditions of instruction that seem to present difficulties for the learner, that appear to slow down the rate of acquisition, but actually result in better long-term learning and transfer—constitutes the remainder of this presentation. In this discussion, I hope to accomplish two main goals. First, I hope to give you a feeling for a few types of desirable difficulties, one of which I will also illustrate with experimental findings. And, second, in this context, I want to point out the potential for teachers and trainers—as well as students and trainees—to be misled as to what are and are not good educational practices or good conditions of learning.

As instructors, we can often be misled in this determination because what is readily available to us is the performance of our students during instruction, which can be a poor indicator of whether durable learning is actually occurring. If, for example, all we consider is the rapidity and apparent ease of their learning during training and instruction, we can easily be led into preferring poorer conditions of learning to better conditions of learning. Additionally, as learners, it seems that we do not develop—through the trials and errors of everyday living—an accurate mental model, so to speak, of those operations that result in learning and those that do not. Furthermore, we are fooled by certain indices—such as how fluently we process information during the re-reading of to-be-learned material—into illusions of learning and/or competence that then lead us to prefer poorer conditions of learning to better conditions of learning.

So, what are some of these manipulations or conditions of instruction that introduce desirable difficulties for the learner? I briefly describe five of them. Then, I illustrate one—providing contextual interference for the learner—with some experimental findings. Finally, I present a number of points that, as instructors, we should keep in mind to try to introduce some of these desirable difficulties into the design of our undergraduate courses and curricula.

Manipulations that Introduce Desirable Difficulties for the Learner

1. *Varying the Conditions of Practice*. When instruction occurs under conditions that are constant and predictable, learning appears to become what might be called contextualized. That is, while it looks very good in that context, the learning acquired in that context does not support retention later when tested in other contexts, and it does not transfer well to different contexts. In contrast, varying conditions of practice—even just the place where you study (as illustrated by Smith, Glenberg, & Bjork, 1978, and by Smith & Rothkopf, 1984)—can enhance recall at a later time. With respect to these findings, it is interesting to note that a how-to-study hint frequently given to students is that they should find a quiet, convenient place to study and then do all their studying in that same place.
2. *Providing Contextual Interference during Learning*. If when trying to learn several different things, you intertwine the learning of those things in such a way as to cause interference among them during acquisition, long-term performance on them will be enhanced. This type of desirable difficulty, often accomplished by interleaving the practice of the various things to be learned, rather than blocking their practice, is the desirable difficulty that I will illustrate with

some relevant experimental findings.

3. *Distributing or Spacing Study and Practice*. The effects of distributed practice on learning are somewhat complex. Although massing practice (e.g., cramming for exams) supports short-term performance, spacing practice (e.g., distributing presentations, study attempts, or training trials) supports long-term retention. That the spacing of practice enhances long-term performance is among one of the more robust and general findings in learning research, holding across a variety of spacing intervals, types of materials, and types of learners. Unfortunately, however, because massed practice or study can support short-term performance, students can be rewarded by good test performance following an all-night cramming session. Little of what they were able to recall after such a short delay, however, will still be recallable after a more substantial delay; whereas, had they distributed their study, much more of the to-be-learned material would still be recallable after a long delay. If throughout the duration of a course, students simply cram for each exam and there is no cumulative final for which they must go back and re-study information already tested, it is little wonder that most students appear to retain very little of the content of a course they had presumably mastered within even a moderate delay from having completed it.
4. *Reducing Feedback to the Learner*. That reducing feedback to the learner during acquisition could be a desirable difficulty seems very strange. Indeed, for many years in the area of motor-skills learning, it was thought that the more feedback you give the learner, the faster and better the learning would be. More recent work, however, has shown that by reducing the feedback you actually enhance the long-term retention and generalizability of motor skills—that is, the ability to produce those skills accurately after a long delay and under different circumstances. (For reviews of the work supporting this new view of feedback and why reduced feedback leads to more durable and flexible learning, see Schmidt & Bjork, 1992, and Christina & Bjork, 1991.)
5. *Using Tests (rather than presentations) as Learning Events*. Much research in the laboratory (e,g., Landauer & Bjork, 1978; Carrier and Pashler, 1992) has demonstrated the power of tests as learning events and, indeed, in terms of long-term retention, such research has demonstrated that a test or retrieval attempt, even when no corrective feedback is given, can be far more effective than a second presentation or study opportunity. In addition, much current research is being addressed to questions concerning test effects, such as the optimal distribution of tests, the optimal form of tests for different types of delays and materials, and the optimal use of feedback with respect to testing outcomes. I do not have time to cover this work in today's talk, but before leaving this topic, I do want to make two points relevant to testing effects.

First, it seems clear that the value of tests as learning events is greatly underappreciated in most educational contexts, where, instead, tests are primarily viewed as assessment tools. Clearly, those of us who study learning in the laboratory must do a more effective job of communicating to teachers and instructors, in general, about the power of tests to promote learning, not just assess it. To address this need, Roediger and Karpicke (2005) at Washington University are currently looking at testing effects with educationally realistic materials and are obtaining dramatic and compelling evidence concerning the benefits of testing over representations of material. As more of these types of results, obtained with such materials, become available, our ability to communicate to teachers and instructors regarding the effectiveness of tests as learning events should be greatly improved. (For references demonstrating the effectiveness of tests as learning events and discussions of why tests are so effective, see Bjork, 1975; Bjork & Bjork, 1992; & Carrier & Pashler, 1992; and for a review of this literature, see Dempster, 1996.)

Second, because students, by and large, do not realize that tests—or attempts to retrieve information—are more effective in promoting learning than are repeated presentations of the material to be learned, they are led to adopt highly inefficient study activities. Were we, for example, to follow some typical students around campus and watch how they went about studying, we would find that they spend way too much time representing information to themselves—reading a chapter over and over again, highlighting passages in different colors, and so forth—and far too little time trying to retrieve information. Or, put slightly different, they would be spending far too much time on the input side of learning and far too little time on the output side of learning. That this mode of studying is so typical among students stems, at least in part, from a faulty mental model of how we learn and remember. They, as many of us do, tend to think of memory as being too much like a tape recorder. Thus they feel that if they just present materials over and over again to themselves, eventually it will write itself on their memories. As it turns out, however, nothing could be further from the way we actually learn and remember.

<u>Contextual Interference as a Desirable Difficulty</u>

I turn now to the desirable difficulty of contextual interference and to demonstrate it with some empirical studies. In the first study I discuss, by Shea and Morgan (1979), contextual inference during learning was provided by having some subjects learn three different movement patterns in an interleaved manner, while others learned them in a blocked manner. The apparatus used by Shea and Morgan looked somewhat like a pinball machine, having two vertical rows of hinged paddles on each side with a start button and a hole containing a tennis ball located between these two rows. In addition, located at the back of the apparatus were three differently patterned stimulus lights, each of which was associated with a different movement pattern that the participant was to learn. When one of the lights came on, the participant was to: 1) push the start button; 2) pick up the tennis ball; 3) while holding it, knock down the paddles in the manner associated with that particular light (e.g., knocking down the first paddle in the left row, then the middle paddle in the right row, and then the rear paddle in the left row); and, 4) when finished, return the ball to its initial location, which turned off a response timer.

In the blocked condition, participants learned the three movements by practicing only one pattern at a time in a blocked manner. For example, a given participant would practice the first pattern to be learned, say A, for many times in a row, then movement pattern B for the same number of trials, and then movement pattern C, also for the same number of trials. For participants learning in the interleaved (or random) condition, the light designating a given movement, say A, might come on for the first practice trial, then the light designating movement C, then A again, then B, then C, and so forth, in a random order, until the participant had practiced each movement pattern for the same number of trials as had the participants in the blocked condition.

As might be expected, during training, the performance of the participants given blocked practice improved much more rapidly than did that of the participants in the interleaved or random condition. Although performance in the interleaved condition eventually caught up to that in the blocked condition, it took quite a while for it to do so—essentially, twice as long to attain the same asymptotic level of performance. If Shea and Morgan had ended their study at this point, and, thus, all the results available to us would have been the participants' performance during acquisition or training, it would seem clear that blocking of practice trials was the superior leaning procedure. But, fortunately, Shea and Morgan did not stop their study at this point. Rather, they had participants return after 10 days at which time they were given a retention test on the movement patterns—a final exam, so to speak. What happened on this exam was quite dramatic!

Shea and Morgan tested their participants in two ways: either under conditions that matched those present during training or under conditions that did not. Thus, for participants trained initially in the blocked condition, half were tested under blocked conditions again and half were tested under interleaved or random conditions. Similarly, for participants trained under interleaved or random conditions, half were tested under the interleaved conditions again and half under blocked conditions. When testing was done under interleaved conditions, the participants who had been trained under those conditions performed essentially as well as they had on their last day of training—that is, they showed little or no forgetting of the three movement patterns. In dramatic contrast, those participants who had been trained under blocked conditions—the participants who had looked the best during training—performed exceptionally poorly on the test. Indeed, their performance was so poor as to look like they had never been trained in the first place. When participants were tested under blocked conditions, the performance of participants trained under blocked conditions was much better, showing only a small amount of forgetting, but—of greater importance—the performance of participants trained under interleaved conditions also showed little or no forgetting. Indeed, if anything, their performance was better—even when tested under blocked conditions—than that of the participants originally trained in that manner.

In other words, when participants trained under blocked conditions were later tested under conditions not identical to those present during their training, their performance was extremely poor, essentially looking like they had never been trained at all. In contrast, participants trained under interleaved conditions were not only able to perform with little or no forgetting when tested under the same conditions, they were also able to perform with little or no forgetting under changed conditions. This pattern of results thus provides a dramatic illustration of the benefits of introducing contextual interference into the learning process. Although slowing acquisition during training relative to blocked practice, the contextual interference introduced by the random practice procedure served to enhance performance at a delay and in a different context.

Several possibilities have been advanced in the literature to explain why interleaving might be so beneficial for long-term retention and transfer. One of these (e.g., Battig, 1966) is in terms of the learner having to resolve the interference among the different things that he or she is trying to learn. To accomplish this resolution, the learner has to notice similarities and differences among them and to schematize or develop a more abstract representation of each item or movement. This higher-order type of learning is what permits both long-term retention and transfer. Another explanation assumes that what is beneficial in the interleaving procedure is that it forces us, as learners, to reload our memories for the different things we are trying to learn over and over again. If required to do A, then B, then C, and then B again, the memory for how to do B is not just sitting there in short-term memory waiting for us to access with no effort. Instead, we have to retrieve it again from long-term memory. These successive attempts to retrieve things that have been forgotten from short-term memory are what lead to the enhanced long-term retention in the interleaved situation. (For a discussion of forgetting as a condition for learning, see Bjork, 1994; Estes, 1955; and Cuddy & Jacoby, 1982.)

While the results of the Shea and Morgan study illustrate how we, as instructors, could easily be misled by the performance of our students during instruction or training into preferring a condition of instruction that is actually not supportive of long-term retention and transfer over one that is, the next study I describe illustrates how we, as learners, can similarly be misled into preferring poorer conditions of learning to better conditions of learning. In this study, conducted by Simon and Bjork (2001), participants also learned three different movement patterns, and they also learned them in either a blocked or interleaved (random) order. Rather than knocking down paddles, however, participants in the Simon and Bjork study learned to execute three different movement patterns on a computer number pad in a specific amount of time (i.e., 900, 1200, and 1500 milliseconds), and they were given feedback on how close they had come to the required duration after each trial. Twenty-four hours after their training, participants returned to the lab and were tested on the three movements. Consistent with the results of Shea and Morgan, participants who learned under blocked training performed better during acquisition; but 24 hours later, they performed more poorly than the participants who had received the random or interleaved training.

The new wrinkle in the Simon and Bjork study was that participants were periodically stopped during training and asked to take a reading on how well they were learning the task. They were asked, if you were to stop training right now and come back in 24 hours, how well do you think you would do—that is, how close do you think you could come to the correct movement time. Participants in the blocked condition all predicted that they would do better than the participants in the interleaved condition predicted that they would do. In other words, their meta-cognitive assessment of how well they were going to do later was exactly wrong. Participants in the blocked condition most likely mistook the rapidity and apparent ease of their being able to perform the required movement patterns—made possible by the blocking of practice trials—as indicating that they were actually learning them well; whereas, the participants in the interleaved condition most likely mistook the slowness and apparent difficulty with which they were being able to perform the required motor pattern as indicating that they were not learning them well. (For a relevant discussion of such confusion between performance and learning in terms of the difference between the retrieval strength and storage strength of memories, as hypothesized in a new theory of disuse, see Bjork & Bjork, 1992.)

Thus, taken together, these two studies illustrate both how we, as instructors, can be misled if we only attend to or only have available to us the performance of our students during acquisition, and how we, as learners, can be misled into thinking that we are learning better under one condition than another when, in fact, the opposite is true. Unfortunately, as learners, we do not seem to be very good at assessing our actual state of competence or knowledge during training and seem easily misled concerning the conditions of training and instruction that are optimal. We seem, for example, to intuit that we are learning better under massed as opposed to spaced conditions of practice, or when the conditions of learning are kept constant as opposed to varied, or when we are given more rather than less feedback. Apparently, these conditions—because they support our performance during training—give us a sense of ease and a sense of learning that turns out to be misleading as far as the actual long-term learning that we are achieving. Whether or not, we, as learners, could be made to be more meta-cognitively sophisticated with respect to when we are or are not learning well is a topic of considerable research interest right now. (For a more thorough discussion of factors that can lead to such "illusions" of knowledge and/or competency, see Bjork, 1999, and Jacoby, Bjork, & Kelley, 1994).

Now, in case by the studies I have used so far to illustrate the benefits of contextual interference, I have created the impression that this desirable difficulty only works with motor learning or simple materials, I end by describing two studies using more educationally relevant materials. In the first study, Mannes and Kintsch (1987) examined the effects of contextual interference on learning from the reading of text. Participants were given a certain period of time to study a technical, but somewhat interesting article on the industrial use of microbes and bacteria with the clever title, “Industry in Ferment.” Prior to studying this article, however, participants had either been given a consistent or an inconsistent outline to read. The consistent outline had the same structure as the article and 25% of the information in the article was presented in the outline; thus, it was very much like the type of advanced organizer frequently used in educational settings. The inconsistent outline had all the same factual information—thus it too had 25% overlap with the “Industry in Ferment” article—but it was actually the outline of an *Encyclopedia Britannica* article on microbes and, thus, it mismatched the article in a number of ways. After participants had studied their assigned outline and then the article, different types of tests were administered. When given a straightforward, verbatim recall kind of test, participants who had received the consistent outline performed better. When given a test that involved problem solving and a deeper understanding of the article, however, the participants who had received the inconsistent outline performed better.

How can we explain this pattern of results? Mannes and Kintsch argued that the inconsistent outline created contextual interference for the readers, forcing them to engage in more active processing of the material in order to resolve this interference. To make peace, so to speak, between the two sources of information, these readers were forced to notice similarities and differences between them and to make inferences in order to bridge gaps between them. Consequently, the readers in the inconsistent-outline condition achieved a deeper understanding of the material than did those in the inconsistent-outline condition.

Although Mannes and Kintsch did not do so in this study, it is interesting to speculate what they would have discovered had they asked their participants how helpful they had found their outlines to be. Participants receiving the consistent outline would probably have given the outline high marks. But what about the participants in the inconsistent condition? Most likely, they would not have given their outline high marks. In fact, they would probably have complained about its inconsistency with the article, even though it was probably in the resolution of these inconsistencies between the outline and the article that learning of a deeper kind was taking place. Almost certainly, however, like the participants in the interleaved versus the blocked conditions of the Simon and Bjork study, these participants too would not have been able to appreciate the better learning being produced by the inconsistent versus the consistent condition.

Finally, in the last study that I want to share with you; McNamara, Kintsch, Songer, and Kintsch (1996) introduced desirable difficulties into their participants processing of text by creating two different levels of coherency in a text about heart disease. Additionally and interestingly, they also had participants with different levels of background knowledge in the domain of biology read the two different levels of text. They then tested their participants regarding the text in a variety of ways by asking them different types of questions - some text-based and some requiring the making of inferences or the solving of problems. Although it was more complicated study than I am describing now, the two hypotheses of relevance to the present discussion were that (a) for both types of students, the consistent outline should be better for the straight recall of text information, but (b) for students with the requisite background knowledge, the text with low coherence could be more beneficial than the test with high coherence. Similar to the reasoning as to why the inconsistent outline was beneficial for deeper learning, the idea behind the second hypothesis was that such students may learn better when they have to provide the coherence themselves (e.g., make the inferences and provide the explanatory connections that are not explicitly provided in the text, thus integrating the information in the text with the information they already have stored in long-term memory.) In contrast, students without the requisite background knowledge would not be able to make the necessary inferences nor fill in the gaps. For them, then, the low coherence in the text would not be a desirable difficulty, as it would present them with difficulties that they would not be able to overcome.

As predicted, for text-based recall of information, the high-coherence text was found to be better for both high and low knowledge students. And, also as hypothesized, for questions requiring problem solving or the making of bridging inferences, the high-knowledge students did profit from having to deal with the low coherent text. In contrast, but as predicted, for the low-knowledge students, the low-coherence text created difficulties that they could not overcome. Thus, for them, the low coherency of the text was not a *desirable* difficulty.

Concluding Comments

I hope in this discussion, I have been able to convince you of the need for us to take a new look at our own methods of instruction and how we design and organize our courses with an eye for introducing desirable difficulties for our students. In doing so, however, we need to keep a few points in mind. First, we need to be mindful of how easy it is for us, as instructors, to be misled regarding the optimal conditions of instruction. In particular, we need to be wary of preferring conditions that speed acquisition and seem to make the learning process too easy, as these conditions may simply be propping up the temporary performance of our students and not creating the type of learning that can lead to long-term retention and transfer. Furthermore, in making decisions regarding how to optimize the learning of our students, we must keep in mind that we cannot rely on the meta-cognitive reports of our students, who themselves—as learners—are often misled into preferring non-optimal to optimal conditions of learning. We want to introduce procedures that present difficulties for the learner—in general, difficulties that force the learner to be a more active participant in the acquisition process. At the same time, however, we need to insure that the difficulties we introduce are, in fact, desirable difficulties, that is, ones that the learner is capable of overcoming.

Resources/References

Publications

1. Battig, W. F. (1972). Intratask Interference as a Source of Facilitation in Transfer and Retention. In R. F. Thompson and J. F. Voss (Eds.), *Topics in Learning and Performance*, pp. 131-159. New York: Academic Press.
2. 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.
3. 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.
4. Bjork, R. A. (1999). Assessing Our Own Competence: Heuristics and Illusions. In D. Gopher and A. Koriat (Eds.), *Attention and Performance XVII. Cognitive Regulation of Performance: Interaction of Theory and Application* (pp. 435-459). Cambridge, MA: MIT Press.
5. Bjork, R. A., and Bjork, E. L. (1992). A New Theory of Disuse and an Old Theory of Stimulus Fluctuation. In A. Healy, S. Kosslyn, and R. Shiffrin (Eds.), *From Learning Processes to Cognitive Processes: Essays in Honor of William K. Estes*, vol. 2, pp. 35-67. Hillsdale, NJ: Erlbaum.

6. Carrier, M., and Pashler, H. (1992). "The Influence of Retrieval on Retention" *Memory & Cognition*, 20, 633-642.
7. Christina, R. A., and Bjork, R. A. (1991). Optimizing Long-Term Retention and Transfer. In D. Druckman and R. A. Bjork (Eds.), *In the Mind's Eye: Enhancing Human Performance*, pp. 23-56. Washington, DC: National Academy Press.
8. Cuddy, L. J., and Jacoby, L. L. (1982). When Forgetting Helps Memory: Analysis of Repetition Effects. *Journal of Verbal Memory and Verbal Behavior,* 21, 451-456.
9. Dempster, F. N. (1996). Distributing and Managing the Conditions of Encoding and Practice. In E. L. Bjork and R. A. Bjork (Eds.), *Handbook of Perception and Cognition.* Vol. 10, Memory, pp. 317-344. New York: Academic Press.
10. Estes, W. K. (1955). "Statistical Theory of Distributional Phenomena in Learning." *Psychological Review*, 62, 369-377.
11. Jacoby, L. L., Bjork, R. A., and Kelley, C. M. (1994). Illusions of Comprehension and Competence. In D. Druckman and R. A. Bjork (Eds.), *Learning, Remembering, Believing: Enhancing Individual and Team Performance*. (pp. 57-80). Washington, DC: National Academy Press.
12. Mannes, S. M., and Kintsch, W. (1987). "Knowledge Organization and Text Organization" *Cognition and Instruction*, 4, 91-115.
13. McNamara, D. S., Kintsch, E., Songer, N. B., and Kintsch, W. (1996). "Are Good Texts Always Better? Text Coherence, Background Knowledge, and Levels of Understanding in Learning from Text" *Cognition and Instruction*, 14, 1-43.
14. Roediger, H. L., and Karpicke, J. D. (2005). Test-Enhanced Learning: Taking Memory Tests Improves Long-Term Retention. Manuscript Submitted for Publication.
15. Schmidt, R. A., and Bjork, R. A. (1992). "New Conceptualization of Practice: Common Principles in Three Paradigms Suggest New Concepts for Training" *Psychological Science*, 3.
16. Shea, J. B., and 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.
17. Simon, D., and Bjork, R. A. (2001). "Metacognition in Motor Learning" *Journal of Experimental Psychology: Learning, Memory, and Cognition*, 27, 907-912.
18. Smith, S. M., and Rothkopf, E. Z. (1984). "Contextual Enrichment and Distribution of Practice in the Classroom" *Cognition and Instruction*, 1, 341-358.
19. Smith, S. M., Glenberg, A. M., and Bjork, B. A. (1978). "Environmental Context and Human Memory" *Memory and Cognition*, 6, 342-353.

Improving Student Learning: Moving from the Memory Laboratory to the Classroom

Speaker: Mark A. McDaniel, Professor of Psychology, Washington University in St. Louis

My primary focus this afternoon reflects the observation that at the college level, especially in foundation courses, much learning consists of the acquisition of factual information. For example, in introductory geology, students must master knowledge of multiple characteristics of types of rocks; in political science, students need to thoroughly learn a number of characteristics associated with each of a few types of political systems; and in developmental psychology, students must learn the attributes of a variety of theories of development (Pressley, Symons, McDaniel, Snyder, & Turnure, 1988).

From the perspective of a standard information processing model of memory, the challenge for the student is to transfer facts from the short-term memory (STM) store, where the facts reside in awareness when immediately attended, to a more permanent long-term memory (LTM) (see Figure 1 below). Extensive research in the memory laboratory has embellished this simple model. In this paper I will examine a number of implications from the basic memory model that potentially translate into improvements in student learning and classroom practice. To meet this objective, I will identify several key components of the memory model and briefly summarize the lessons learned in the memory laboratory. For each, using educational materials, I will then present translational research that informs techniques and approaches to improve student learning in fact-laden courses.

Figure 1

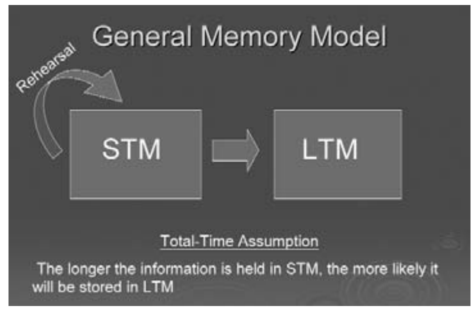

The likelihood of transfer from STM to LTM was originally assumed to be a positive function of the amount of time information resided in STM (the Total-Time Assumption). Because rehearsal—recycling information in STM—is the control mechanism by which the learner maintains information in STM, the total-time hypothesis implies that the more the learner rehearses target information, the more likely that the information will be stored in LTM. Basic memory research, however, has not supported the total-time hypothesis. Yet, for many students their typical study activities such as rereading text and lecture notes seem to heavily engage repetitive recycling of the information. Accordingly, the first implication for undergraduate education is that typical undergraduate study activities like rereading the text and notes may not be overly effective for learning and retention. The assumption here is that rereading for undergraduate students often involves repetitive recycling, and memory theory suggests this is not overly effective for increasing learning. Let's examine relevant research with educational type materials.

Roediger and Karpicke (in preparation; Figure 2 below) found that extensive rereading of a text in three study sessions produced only a modest gain in recall relative to several rereadings in one study session. Perhaps, most of the learning gain occurred in the first several rereadings and then reached a ceiling. Callender and McDaniel (Figures 3-4 below) showed that there was no apparent gain in learning from one to two readings, regardless of test type. Performance after one reading was relatively high, however. What about a text for which performance after one reading is at lower levels? Even in this situation, Amlund, Kardash, and Kulhavy (1986; Figure 5 below) reported only slight gains in learning of main ideas from one to three readings.