# Providing Learners with Feedback—Part 1:

## Research-based recommendations for training, education, and e-learning.

WORK-LEARNING RESEARCH, Inc.

A Work-Learning Research Publication. May 2008.
© Copyright 2008 by Will Thalheimer. All rights are reserved.

You may share this report as long as you share it intact and whole and don't charge a price for it or include it in any priced event or product.

This is a two-part report. Part 1 provides background and recommendations. Part 2 focuses on research support.

How to cite this report using APA style:

Thalheimer, W. (2008, May). *Providing Learners with Feedback—Part 1: Research-based recommendations for training, education, and e-learning.* Retrieved November 31, 2008, from http://www.work-learning.com/catalog/

Obviously, you should substitute the date on which the document was downloaded for the fictitious November 31 date.

Published May 2008

**The Author:** Will Thalheimer is a learning-and-performance consultant and research psychologist specializing in learning, cognition, memory, and performance. Dr. Thalheimer has worked in the workplace learning field, beginning in 1985, as an instructional designer, simulation architect, project manager, product leader, trainer, consultant, and researcher. He has a PhD from Columbia University and an MBA from Drexel University. He founded Work-Learning Research in 1998 to help client organizations create learning environments to maximize performance and help learning professionals—instructional designers, e-learning developers, trainers, performance consultants, talent managers, and chief learning officers—utilize research-based knowledge to build effective learning-and-performance solutions.

Dr. Thalheimer can be contacted at info@work-learning.com or at the Work-Learning Research phone number below for inquiries about learning audits, workshops, speaking engagements, high-level instructional design, and consulting on e-learning, classroom training, learning measurement, learning strategy, and business strategy for learning.

Work-Learning Research, Inc.
Somerville, Massachusetts USA
(617) 666-9637

www.work-learning.com

All information contained in this document is copyrighted.
© Copyright 2008 by Will Thalheimer. All rights are reserved.

***Sharing Rules:*** You may copy and distribute this document as long as you share it intact and whole and don't charge a price for it or include it in any priced event or product. (this includes generating good will that might be translated in the future into tangible financial benefits).

**We do not receive a salary for translating research into practical recommendations. If you abuse the spirit of the "Sharing Rules" describe above, you are not only stealing, you are limiting our ability to continue to support you and others in the learning-and-performance field!**

# Introduction

Hi. I'm Dr. Will Thalheimer, a consultant and researcher specializing in learning fundamentals, instructional design, performance improvement, learning measurement, and workplace learning. I help people create more effective learning interventions by building bridges between learning research and learning practice. There is wisdom in both camps, but only by integrating the two can we maximize our learning outcomes.

I've been researching the topic of feedback for almost 10 years (off and on), studying research articles from the world's preeminent refereed journals, translating that research to provide practical wisdom for learning professionals, and compiling what I've learned to make it available to others. Although I've shared my tentative findings with a few others in the past, this document represents my first published research-to-practice report dealing directly with the topic of feedback.

The feedback literature is probably the most difficult learning-research area I have had the pleasure to research and translate. In short, the research landscape on feedback is a war zone, complete with countless shattered structures, broken windows, and unexploded ordnance. Those poetically inclined can parse these symbols for their full meaning, but let me just say this: as Hemingway and others found some wisdom and solace in war, the feedback literature has its own bounty of truth to tell, even in the midst of its unfolding story (as new research continues to become available).

This report is designed to help you improve your feedback methods. I certainly won't claim to have all the answers, nor do I think simple recipes are available in dealing with feedback. I do believe strongly, however, that all of us can improve our feedback methods substantially, and by so doing improve the practice of education and training.

I would like to thank Questionmark for agreeing in advance of my final writing efforts to license this report for the benefit of their clients. Questionmark is available on the Web at www.questionmark.com and by phone at 800-863-3950 (North America), +44 (0)20 7263 7575 (United Kingdom) or +32 2 298 02 01 (Europe).

## This Report's Design

This is a research-to-practice report, written in two parts. Part 1 is written for a general audience. It provides perspective on feedback's place in learning and it makes specific recommendations for practice. It also includes a two-page summary of the practical recommendations. Part 2 contains a research review that supports the recommendations made in Part 1. Research references are included in Part 2.

Because some of the concepts and terminology in this report may introduce you to new paradigms, you can expect that they will seem foreign at first. As their power becomes evident—as you build new mental models of these concepts—the initial fog will clear.

## Contents

Introduction ..... 3
This Report's Design ..... 3
The Feedback Context ..... 6
The Three Phases of Learning Something ..... 8
The Power of Retrieval Practice ..... 11
Retrieval Example ..... 12
Retrieval Enables Future Retrieval ..... 13
Methods for Inducing Retrieval ..... 15
Open-ended vs. Forced Choice Retrieval ..... 16
Why is Feedback Valuable? ..... 17
The Value of Feedback ..... 17
Feedback Corrects Misconceptions ..... 18
Cognitive Mechanisms of Incorrect and Correct Feedback ..... 20
Feedback When Learners Know the Correct Answer ..... 21
Feedback When Learners Answer Incorrectly ..... 22
Summary (Before We Get to the Recommendations) ..... 22
Recommendations for Feedback ..... 23
Feedback Should Help Learners Develop Appropriate Conceptions ..... 24
More Extensive Feedback May be Needed to Build Understanding ..... 26
Providing Feedback Before Retrieval is Complete ..... 28
Feedback for Incorrect Answers ..... 30
After-Feedback Retrieval Practice ..... 30
Feedback for Correct Answers ..... 32
The Timing of Feedback ..... 35
Deliver Corrective Feedback before Additional Retrieval Opportunities ..... 35
Immediate versus Delayed Feedback ..... 38
Three Simple Heuristics for Feedback Timing ..... 40
Getting Learners to Pay Attention to the Feedback ..... 41
Revisiting the Distinction between Building Understanding and Supporting Retrieval and Fluency ..... 42
Other Considerations ..... 43
Using Confidence Ratings to Augment Feedback ..... 43

Using the Answer-Until-Correct Feedback Mechanism ..... 43
The Dangers of Evaluative Feedback ..... 44
Feedback May Also Have Subsequent Learning Effects ..... 45
On-the-job Follow-up Feedback ..... 47
Training Follow-Through as Feedback ..... 47
Summarizing Principles and Recommendations ..... 49

## The Feedback Context

To understand feedback—to have wisdom about how to use it—we need to understand it in its full context. Feedback is never an isolated event. It always comes after other learning events and before other learning and performance events. Here is the typical progression. Note that "Feedback" is the third of the four event boxes depicted.

First, we present our learners with learning materials and messages. Second, we give them opportunities to retrieve information from memory (sometimes involving simple retrieval and sometimes including decision-making or practice). Retrieval is basically a process in which learners encounter a cue (for example, a test question) and are prompted to attempt to retrieve information stored in long-term memory (for example, answering a test question). Only after our learners attempt to retrieve information in this way do we give them feedback. Finally, after learners receive feedback, they may encounter or be directed to additional learning events. They may also engage in real-world situations that require them to retrieve and utilize what they've learned. These situations may include such things as job tasks, certification examinations, and conversational dialogue. Please take a look at where feedback is placed in this process (see the above diagram).

The full context of feedback has some obvious implications. The better the original learning events are in creating understanding, the less the need for feedback (because learners will be more likely to retrieve correct information). The more the retrieval opportunity is aligned with what was learned, the less the need for feedback (because learners will be more likely to perform well in retrieving information from memory).

On the other hand, there are some subtleties that are easy to overlook. The most important has to do with retrieval practice. To put it bluntly, feedback is not as important as retrieval practice in supporting later learning. If our learners are able to confidently

retrieve information from memory, feedback provides little benefit. If our learners can't retrieve information, feedback is critical but it is not enough. The best preparation for later retrieval is a successful current retrieval, especially when the current retrieval situation mirrors the future retrieval situation. So, if we give feedback on an incorrect answer, we have to also give another retrieval opportunity. More about this later. For now, the point is that the whole context of feedback is important.

The only variant to the basic flow illustrated on the previous page is the flow indicated below. It shows the same sequence without the initial learning presentation.

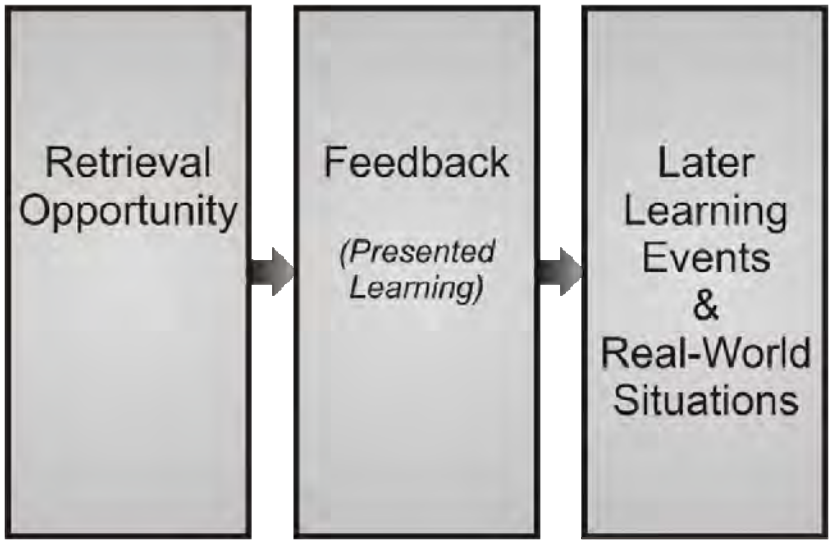


Such a feedback context is common when using a prequestion-first instructional flow. This type of learning design can be very effective in engaging learners and ensuring that they know what aspects of subsequent learning material to give their fullest attention.

As you might imagine, feedback here may require much more information and instructional support than feedback provided after learners have already been introduced to the learning concepts. Moreover, this sequence does not provide the type of retrieval that directly supports long-term retrievability of the information being learned. In a prequestion-first instructional flow, more retrieval practice will be needed after the material is fully understood.

## The Three Phases of Learning Something

Learning is not a one-dimensional process. Learners need different cognitive supports when they are new to material than when they are more experienced with material. If our learners don't know how to throw a baseball, we won't be able to help them throw a curve. If our learners can't quickly comprehend spoken words in a foreign language, we can't use a conversation to teach them phrasing.

Different instructional supports are needed in different phases of learning. I offer the following model to help make sense of this. The diagram below overlays three labeled arrows on a typical learning curve. As learners are introduced to a topic, the most important thing we can do is support our learners in building understanding of the topic. It is often critical in these early phases to break the learning down into constituent parts, give learners extra time to process the material, help the learners test their understanding, provide worked examples, challenge typical misconceptions, and be careful not to overload working-memory capacity, among other things.

After our learners "get" the basic concepts, we can provide them with retrieval practice to increase the likelihood that they will be able to retrieve the information in the future. Finally, if it is important that learners are fluent in being able to retrieve information from memory, we can give them additional practice to automate their retrieval.

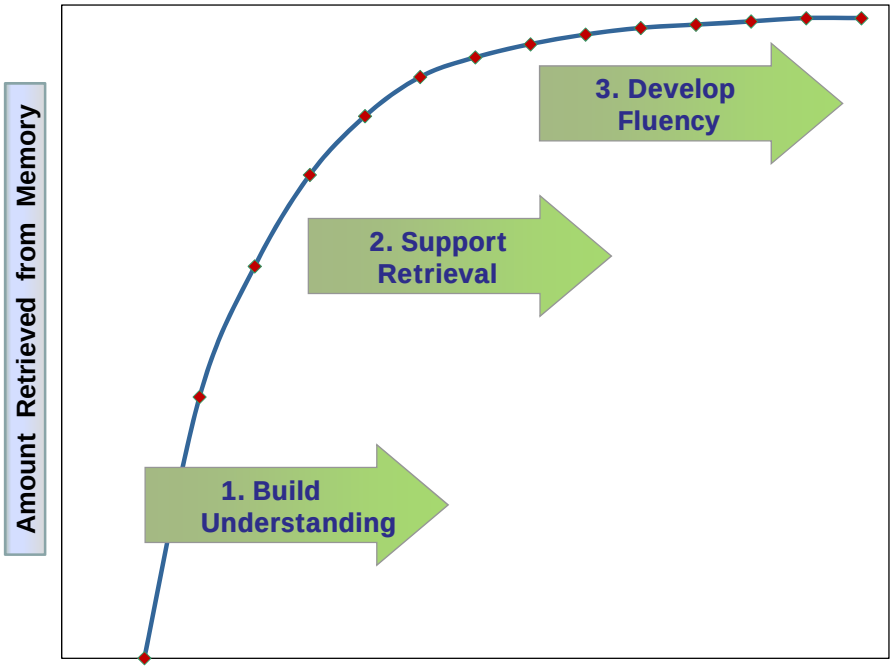

These phases don't have definitive dividing lines, but overlap in the real world of learning. For example, we may be able to support retrieval at the same time we build understanding of boundary conditions or special contingencies. Building fluency begins with the initial phases of retrieval practice.

Note that several factors intersect to determine when learners engage in each of the three phases. On the graph below, I've illustrated how two factors interact: (a) The complexity of the concepts being learned (on the left axis) and (b) the timeframe of learning—whether learners are relatively early or late in their exposure to a particular concept—(on the bottom axis). While other factors may also contribute, concept complexity and length of learner exposure seem the most critical and leverageable from a learning design standpoint.

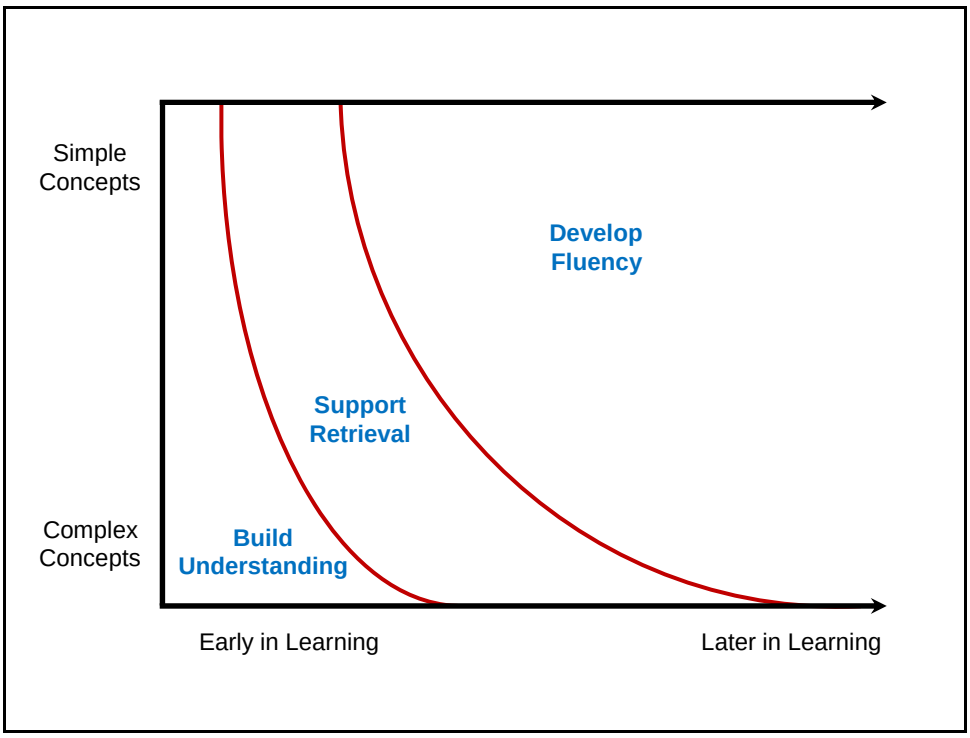


As you can see in the graph above, simple concepts require shorter timeframes for all phases. With complex concepts, it takes longer to build understanding, longer to support retrieval, and longer to begin to develop fluency[1]. Note that the rightmost part of the graph continues into the future, as additional repetitions may be required to develop fluency and then maintain it over time.

[1] This graphic is meant to show how (a) the complexity of concepts being learned and (b) the amount of time the learner has been exposed to the concepts being learned influence a learner's processing stage. The axes could have been switched, of course, but I kept the timeframe on the bottom axis to keep the visual perspective the same as the typical learning curve. Note what would happen if you overlaid the learning curve on the diagram. It would show how, from a learner's perspective, complex topics become simple topics with the passing of time.

Don't worry if the graph above seems overly complicated. The key thing to take away from this discussion is that there are three stages of learning. This three-phase conceptualization is particularly relevant to feedback as we'll see later when we discuss specifics of the feedback process[2].

The sneak preview is this. We may want to deliver different types of feedback to learners when we're helping them *build understanding* compared with times when they are moving toward *retrieval* and *fluency*. Specifically, when they are building understanding they may need more feedback, feedback on correct answers (where later they may need very little feedback on correct answers), more explanatory feedback, and more immediate feedback, among other things.

Dr. Will Thalheimer

*"Although I regularly produce these research reports, I earn my livelihood by consulting, speaking, and doing learning audits.*

*If I can help transform your organizations' learning practices, let me know."*

*info@work-learning.com*

[2] Unfortunately, much of the academic research on learning relies on situations where learners are learning very simple concepts. This is a reasonable approach up to a point, but it is not analogous to all learning situations or methods. When concepts are easy to learn, the process of "building understanding" is so trivial that the focus of the research—and the learning—shifts almost entirely to retrieval support. Experiments also rarely focus on building fluency because the time of experimental participants is relatively costly and getting participants to return over several sessions is difficult. The experimental flattening of learning into a focus on retrieval support is not especially troublesome in some research areas (for example, when we look at repetition, spacing, or retrieval practice), but it is problematic for feedback because learners are likely to benefit from different feedback at different points in their learning. For example, it's not hard to imagine that when learners are building understanding that they might need more extensive feedback than if they are simply practicing retrieval. Because the research on feedback has tended to focus on retrieval support, it has tended to find that less feedback information is better than more feedback. There are hints that more extensive feedback is beneficial when the material is difficult, but as far as I can tell, researchers haven't focused on the distinction between "building understanding" and "supporting retrieval." The bottom line here is that wisdom about how to interpret the research is needed. From my wrinkled gray matter, I think it important that we distinguish between feedback given when learners are building understanding and when they are supporting retrieval or developing fluency.

## The Power of Retrieval Practice

We've already hinted at the power of retrieval, but because the concept is so important—and because it challenges many long-held views of learning, I'm going to provide more detail here. The following diagram provides an outline of the process.

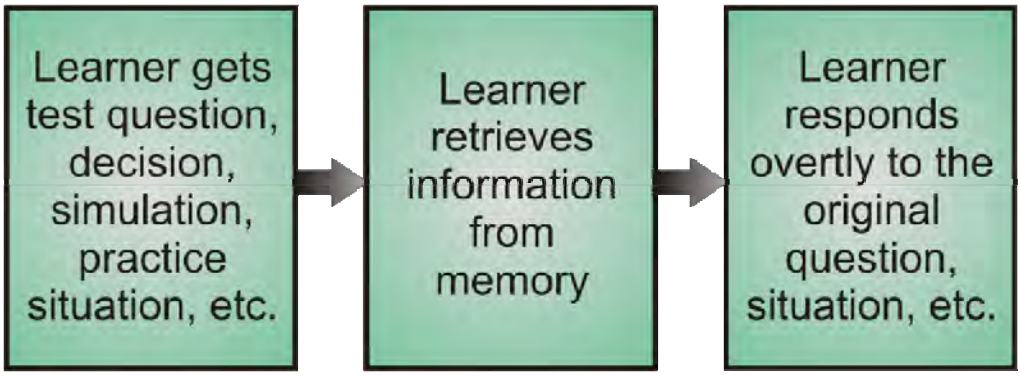


The first box provides a cue—an initiating event that spurs the retrieval process. In formal learning situations, retrieval is usually initiated through tests, quizzes, questions from an instructor, decision scenarios, case studies, problems to solve, simulation decisions, etc. But cues trigger memory retrieval in all sorts of activities, including hitting a baseball, driving a forklift, listening to music, smelling the fragrance of an old barn, watching a movie, and speaking with friends.

The second box represents the learners' memory-retrieval process. If we ask our learners to make a management decision—after having taught them some supervisory principles in the morning—they may retrieve what we taught them about managing others. More simply, if we ask our English-speaking learners what "muchas gracias" means, they may retrieve the English equivalent, "thank you very much."

The third box represents the overt response the learner makes to demonstrate their ability to retrieve. I distinguish between the cognitive process and the overt process to highlight where the power of retrieval comes from—it comes from the cognitive process. So, for example, in most circumstances learners can get benefits from retrieval without making overt responses.

Let me use a simple example to make this whole process clear. Suppose I taught my learners that "a fosse is a ditch." Later, I gave them a quiz, consisting of the question, "What is a fosse?" and they retrieved the information "a fosse is a ditch." This retrieval process could be represented in the diagram on the following page:

## Retrieval Example

1. After learners encode information and store it in long-term memory, the information remains there until it is triggered by external events. Here, the stored information is “a fosse is a ditch.”

2. When the learner encounters a cue, for example the word “fosse” or the question, “What is a fosse?,” the retrieval process is triggered.

3. The learner searches memory. Because the information was well learned, the learner is able to access the information.

4. The learner then retrieves the information (into working memory) and is able to make an overt response (for example, saying “a fosse is a ditch.”)

### *Retrieval Enables Future Retrieval*

While it's common to think that our ultimate goal as learning professionals is to enable our learners to learn, the fuzziness of the word "learn" often gives us the wrong impression. What we really want to do is help our learners retrieve what they've learned—at an appropriate time in the future, and in an appropriate situation.

The best way to enable future retrieval is to figure out the salient cues in the future situation, and then give learners practice retrieving information given those cues (or realistic facsimiles of those cues). In short, the best way to enable future retrieval is to provide retrieval practice.

The power of retrieval practice has been well documented in the learning research. For example, in a now classic study, Jones (1923-1924) did a series of classroom experiments and found that when the last five minutes of a one-hour session were used for retrieval practice, significant improvements were produced in later retrieval. See the graph below.

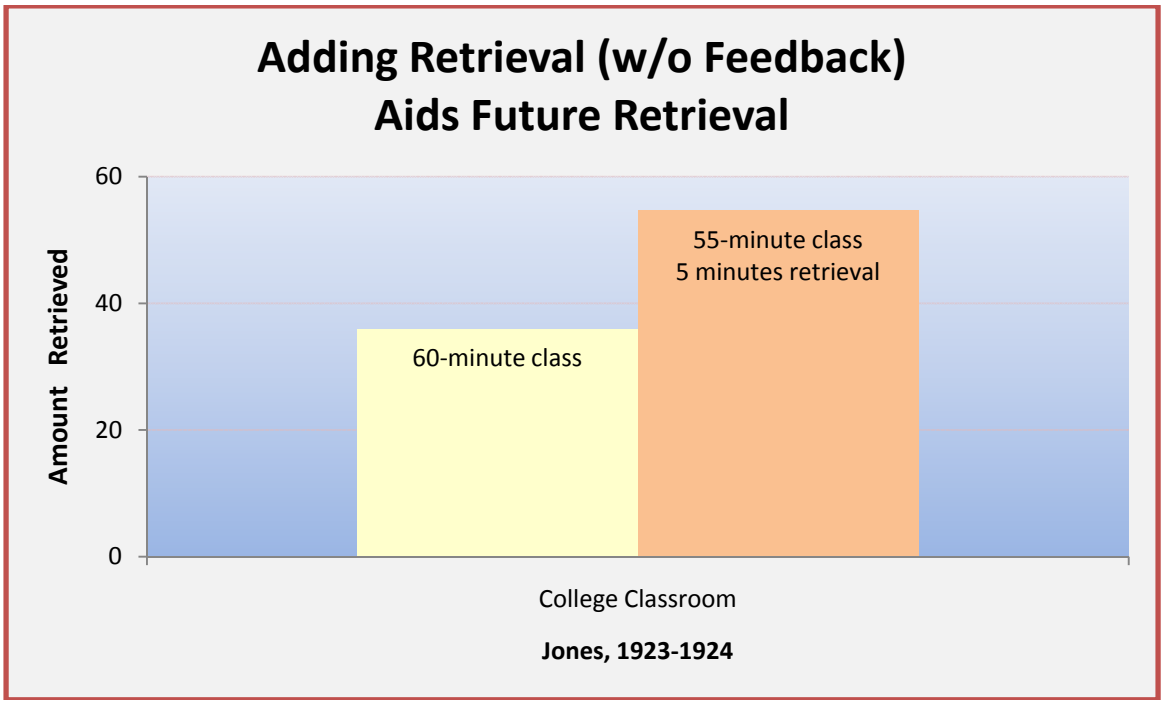


This pattern of results has been found consistently. While Jones' study is a classic, research for many years has found similar results. For example, see the graph for the Roediger and Karpicke (2006b) study on the following page. The power of retrieval practice is so evident that researchers have coined a term for it. They call it the "testing effect."

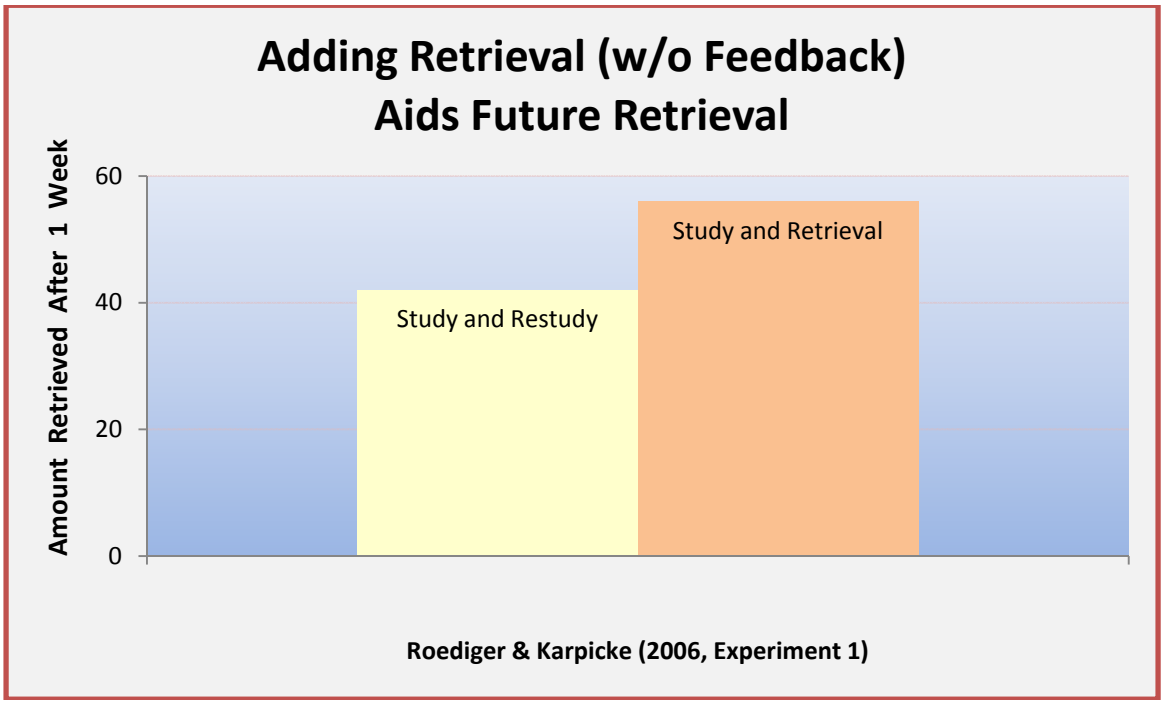


Roediger and Karpicke (2006b, Experiment 1) also looked at retrieval results after two days and after five minutes—in addition to the one-week retention interval graphed above. For the two-day retention interval the results were the same; retrieval produced better results than extra study time. However, for the relatively short (and generally unrealistic) five-minute retention interval, extra study time was better than retrieval practice. This finding shows that retrieval practice is especially potent in supporting long-term memory—the kind of result we learning professionals tend to target.

It should be noted that even when no feedback has been given to learners, retrieval practice is found to produce significant increases in long-term retrieval. None of the research cited above (on this page or the previous one) provided learners with feedback!

Of course, learners need to be successful in a large proportion of their retrieval-practice opportunities for the benefits of retrieval to accrue without feedback. For example, Kang, McDermott, & Roediger (2007) found that when retrieval practice was difficult (when learners had to produce a response as opposed to selecting from multiple-choice alternatives), feedback was critical in helping them benefit from retrieval practice.

### *Methods for Inducing Retrieval*

Retrieval happens all the time. When a friend asks us a question, retrieval is triggered as we think about how to answer the question. When we try to hit a softball, we retrieve information about how to perform that task. When we smell a unique fragrance, we may retrieve memories from childhood. When we read a newspaper, we retrieve all kinds of information to help make sense of the words we're reading. When we see a photograph of our first love, we may retrieve a barrage of emotional memories. When we decide how to drive to grandma's house, we retrieve information about the route, about grandma, and about the objects we pass along the way. When we make a decision, we retrieve information to guide our decision making.

Note that retrieval isn't always a conscious or intentional process. In fact, it is never a solely conscious process. Even when we intentionally search memory to answer a test question, much of our search processing goes on without our conscious awareness.

In some sense, simply presenting people with learning material induces retrieval. While such undirected retrieval can be beneficial in facilitating learning, it is not anywhere near as effective as retrieval processing that is intentionally well-designed to support future retrieval.

As learning professionals we have a number of directed retrieval methods we can employ. The following list includes some of the most commonly used methods:

- True-False questions
- Multiple-choice questions
- Matching questions
- Recall questions
- Essay questions
- Problems to solve
- Case studies
- Simulations
- Decision scenarios
- Skill demonstrations
- Hands-on practice
- Discussions
- Oral responding
- Action planning

The best retrieval-practice events are designed to mirror future retrieval situations. Real-world practice and high-fidelity simulations are powerful because they utilize realistic retrieval practice. Using multiple simulated contexts is beneficial as well, especially in supporting later retrieval in various situations. On the other hand, retrieval can actually impair later retrieval if it focuses on irrelevant concepts. For example, asking restaurant workers to define terminology is counter-productive in helping them to remember how to ensure food safety.

Inspired by the excellent work of Shrock and Coscarelli (2007)[3], I created (Thalheimer, 2007) the following taxonomy of retrieval authenticity, starting with the highest level of authenticity and ending with the lowest level of authenticity[4].

A. Real-World Performance
B. High-Fidelity Simulations
C. High-Fidelity Decision-Making Scenarios
D. Low-Fidelity Simulations
E. Low-Fidelity Decision-Making Scenarios
F. Memorization of Critical Information
G. Memorization of Perfunctory Information

The higher the level of authenticity in our retrieval-practice situations, the better the long-term retrieval. Again, this is true because when learners encounter real-world cues, those cues will be more likely to trigger appropriate memory retrieval if the learner has had previous successful experiences retrieving information when faced with analogous cues. As I have said for many years, one of our goals as learning designers is to promote this type of "spontaneous remembering."

## Open-ended vs. Forced Choice Retrieval

In assessing retrieval, we should also distinguish between forced-choice retrieval practice and open-ended retrieval practice. Forced-choice retrieval gives learners a retrieval opportunity, but prescribes the options learners may utilize. The most common forms of forced-choice retrieval are multiple-choice, true-false, and matching questions.

We probably rely too often on forced-choice questions, especially multiple-choice questions. The learning research is pretty clear that open-ended questions produce more powerful learning benefits than forced-choice questions (because they prompt deeper and more appropriate cognitive processing). Nevertheless, sometimes forced-choice questions are beneficial due to their logistical ease, scalability to large numbers of learners, and the general benefits of questioning. They also sometimes enable us to diagnose the reasons for learners' misconceptions. By carefully selecting the answer choices we provide, we may be able to provide corrective feedback tailored to the most common misconceptions.

[3] Shrock and Coscarelli (2007) use the following categories: Level A—Real World, Level B—High-Fidelity Simulation, Level C—Scenarios, Level D—Memorization, Level E—Attendance, Level F—Affiliation. While these categories are perfect for use in assessment certification decisions, I modified them because more gradations are helpful in talking about authenticity of assessment items.

[4] While the list puts high-fidelity decision-making scenarios above low-fidelity simulations, this ordering could be reversed depending on the level of fidelity of each in comparison with what is most important for the learner to be able to do.

# Why is Feedback Valuable?

We've now seen how feedback works in context. It typically comes after learners have been presented with learning events and after they have been given a retrieval practice opportunity. To reiterate, I present the diagram again:

We've seen the power of retrieval practice to support future retrieval and we've learned that feedback can support that retrieval. It's now time to turn our attention to feedback itself. We will return to the feedback context throughout the remainder of the discussion to highlight important principles.

### *The Value of Feedback*

Feedback is one of the most important learning factors because it helps learners overcome several difficulties. Feedback is beneficial to help learners overcome the following issues:

- Learners didn't learn the information in the first place.
- Learners learned the information incorrectly.
- Learners forgot the information.
- Learners remembered the information incorrectly.
- Learners incorrectly remembered co-presented information as correct; for example, they remember lures on multiple-choice tests.
- Learners remember a concept in the abstract, but can't retrieve the information when presented with certain (for example, realistic) cues.

When we say that feedback is needed, what we're really saying is that both a retrieval opportunity and feedback are needed. As the above list highlights, there are many fundamental problems that feedback can help alleviate.

### *Feedback Corrects Misconceptions*

If you look at the list immediately above—the one that lists the difficulties that feedback helps learners overcome—it will make perfect sense that feedback works by correcting misconceptions. This is one of the most important things to remember about feedback, it's all about supporting learners in making cognitive corrections.

The main implication is that feedback is much more important for incorrect answers than for correct answers. When learners retrieve incorrectly—when they fail to retrieve or when they retrieve the wrong information—feedback is critical because misconceptions must be corrected.

When learners retrieve correctly—when they follow the correct route through memory—they've already done the most important thing they can do to support later retrieval; they've correctly retrieved. Any additional feedback we might give them is almost redundant. The caveat, of course, is that a guessed correct answer is not really indicative of a correct retrieval route. Guessing is not the same as knowing. We'll discuss ways to handle correct-answer guessing when we get to the recommendations section.

Many research results point to the greater importance of feedback on incorrect answers than on correct answers. Kang, McDermott, and Roediger (2007)—detailed in the graph below—showed how feedback improves retrieval performance when learners got feedback on incorrect answers, but not when they got feedback on correct answers. The left side of the graph shows how feedback for initially-incorrect answers improved results by almost half (from 26% to 46% correct). The right side of the graph shows that for initially-correct answers, no improvement was evident (82% vs. 80% correct).

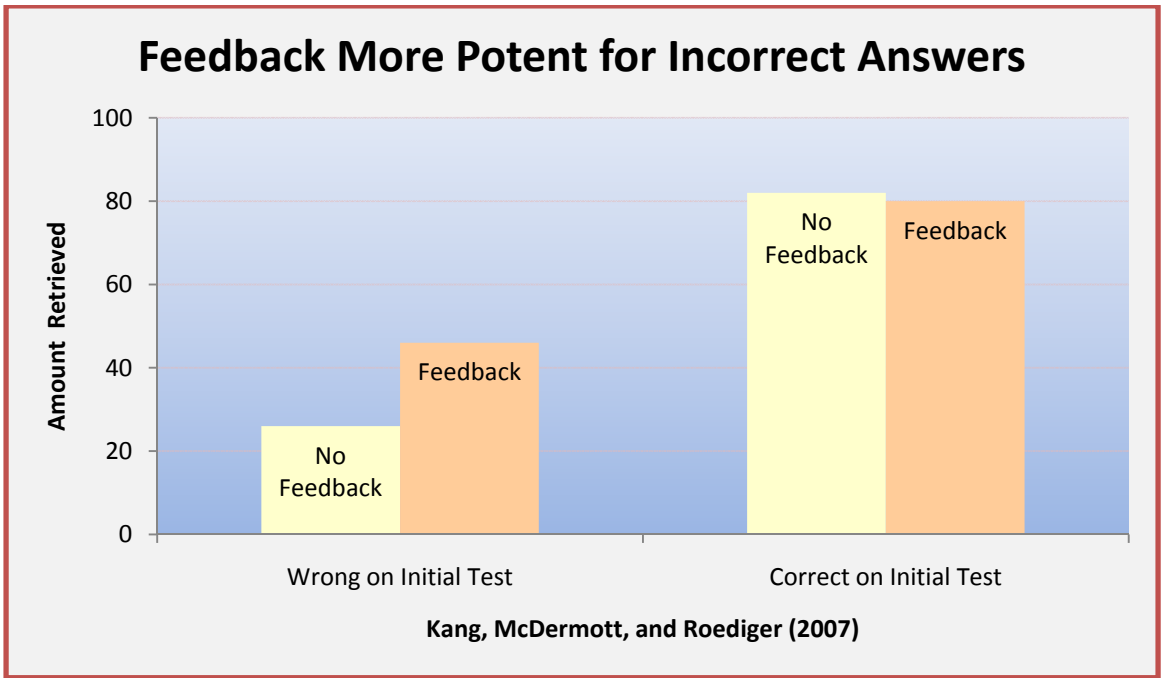

In similar research, Peeck, van den Bosch, and Kreupeling (1985) found even more profound improvements due to feedback on incorrect answers—and again, feedback on correct answers produced negligible improvements. The left side of the graph shows the power of feedback on incorrect answers. Feedback improved performance from 20% to 56% correct on initially-incorrect answers. The right side of the graph shows no improvement for initially-correct answers. Feedback barely made a difference, with retrieval performance at 88% with no feedback and 89% with feedback.

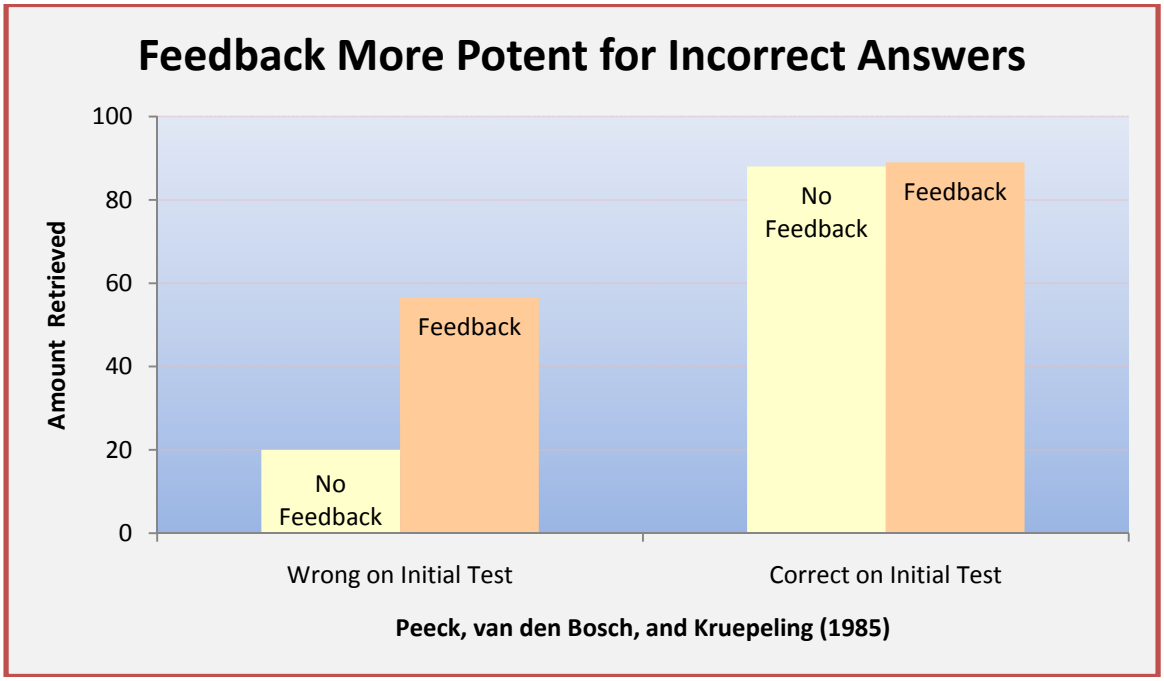


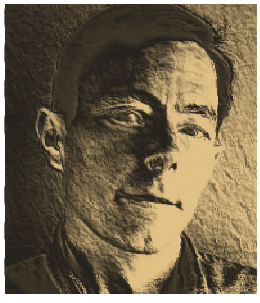
Dr. Will Thalheimer

*"One of the biggest problems I hear about in learning-design shops, is that nobody's on the same page, everybody's speaking a different language.*

*With my research-based approach, I can help get your team headed in the same direction. Effectiveness!"*

*info@work-learning.com*

Feedback isn't always ineffective for correct answers. Often people guess correctly, and so need feedback to reinforce their correct responding. For example, in a recent study by Butler, Karpicke, and Roediger (submitted for 2008b), while feedback on incorrect answers was more important, feedback on correct answers was important too.

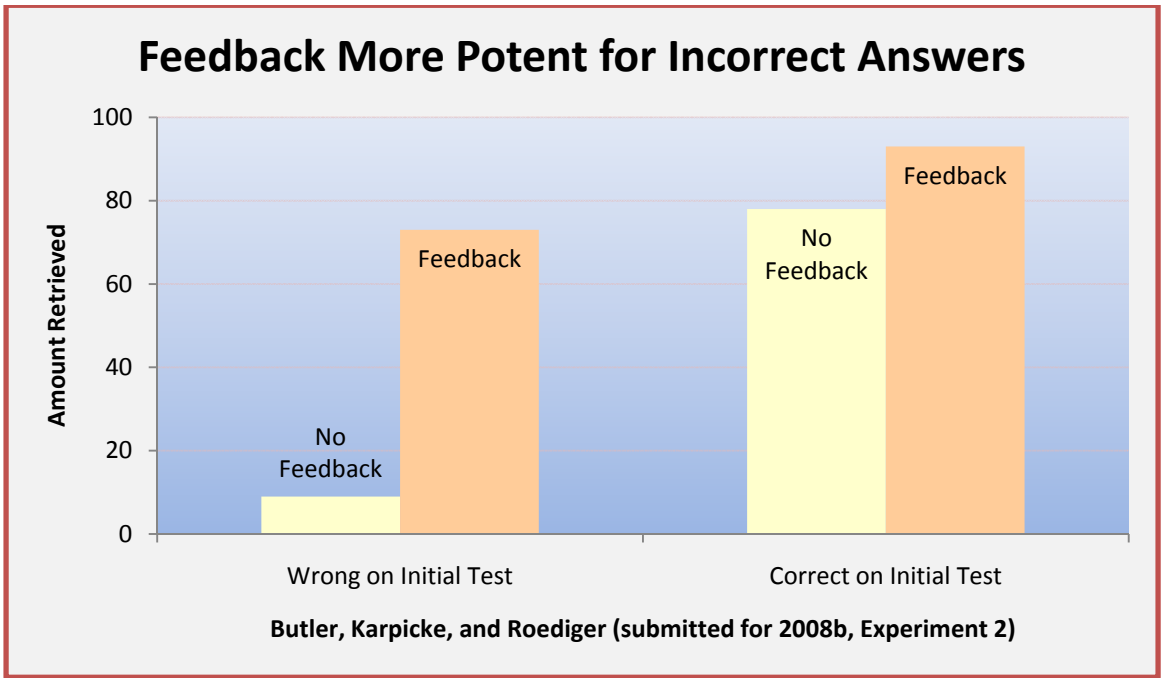


As you can see, feedback on correct answers provides some benefits, but still, feedback on incorrect answers provides significantly more.

## Cognitive Mechanisms of Incorrect and Correct Feedback

So that we can make the most informed decisions about how to deliver feedback, it can be very helpful to examine the cognitive differences between feedback on correct answers and feedback on incorrect answers. As you will see, there are at least one or two insights in this analysis that will lead us to unique and powerful instructional-design recommendations.

To make these distinctions as clear as possible, I offer the diagrams on the following pages. The first set of diagrams details what happens when learners know the correct answer. The second shows what happens when learners get an answer wrong.

## Feedback When Learners Know the Correct Answer

1. The learner successfully retrieves the information and recalls the correct answer. In other words, the learner follows the correct retrieval route through memory.

2. We provide feedback. The diagram barely registers the cognitive effects of the feedback. In some sense, feedback is almost redundant for learners who know the answer and correctly retrieve.

Dr. Will Thalheimer

*"I used to sell these research reports for money, but even though they provided research-validated information from the world's best refereed journals, not that many people actually benefitted from my efforts. Now, legions are benefitting."*

*Want to advertise?*

*info@work-learning.com*

## Feedback When Learners Answer Incorrectly

1. The learner retrieves the wrong information from memory. The correct information, while it may be in memory (upper left), is not accessible. Feedback is needed to correct this.

2. We provide feedback and the learner uses that feedback to encode (or re-encode) the information in memory.

### *Summary (Before We Get to the Recommendations)*

Human learning is seriously complex. Feedback is similarly complicated. To make optimal learning-design decisions, we need deep mental models of how learning works. As professionals, we have to make tradeoffs. Without understanding, we can't make intelligent tradeoffs. I have provided the pages before this one so that you can make more informed feedback-design decisions.

We've learned the following fundamentals:

1. Retrieval is one of the most important learning factors.
2. Feedback is delivered after a retrieval attempt.
3. Retrieval is generally more important than feedback.
4. Some retrieval opportunities produce better learning than others.
5. Learning can involve the phases: (a) Building understanding, (b) Supporting retrieval, and (c) Developing fluency.
6. Feedback is valuable because it corrects misconceptions.
7. Correct and incorrect responses produce different cognitive effects.
8. Feedback is more important for incorrect responses.

## Recommendations for Feedback

It's time to get to the recommendations. If you skimmed through the introductory material to get to these recommendations, you're making a mistake. To make good learning-design decisions, we must understand feedback in its full perplexity. While many of us in the learning profession have come to believe that we can rely on simple rules and general principles, the complexity of the human learning system suggests otherwise. Instead of simplistic bromides, we need a deep understanding of human learning. This is especially true with feedback, as you will see.

In this section, I will attempt to answer the following questions:

1. Is feedback always helpful?
2. How much information should be provided as feedback?
3. What kind of information should be provided as feedback?
4. Is it okay if feedback simply points learners back to the learning material?
5. How should feedback be provided on correct answers?
6. How should feedback be provided on incorrect answers?
7. When should feedback be provided—immediately or after a delay?
8. Is feedback sufficient to correct misconceptions?
9. Can feedback short-circuit a learner's metacognitive processing?
10. Can learners use "feedback" inappropriately to hurt their own learning?

Now that you have a good idea of the nature of feedback, it's time to get to the recommendations. These recommendations are based on research and practical wisdom. The research on feedback is not always definitive, so we have to be somewhat circumspect in our practices. You will produce the best results by pilot testing your feedback methods in your real learning situations—with your real learners and real learning materials, and then testing the learning results with authentic measurement instruments.

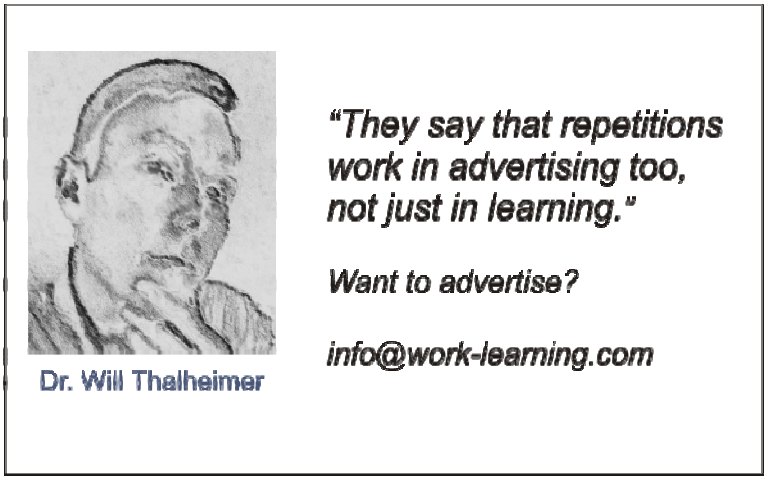

### *Feedback Should Help Learners Develop Appropriate Conceptions*

Earlier in the paper, we discussed the truism that feedback works because it corrects misconceptions. Researchers have found that to correct misconceptions, it's important that learners are given specific corrective information. For example, it's not as helpful to tell learners that they are right or wrong, as to tell them what the right answer is. It's not as helpful to direct them back to the original learning material, as it is to provide feedback targeted to their particular misconception.

The overarching principle is to give learners the feedback they need. Give them the feedback that will help them build appropriate mental models of the concepts to be learned. Don't waste their time with extra tangential information. Don't overload their limited working-memory capacity. Don't give them partial information. Don't provide them with verbatim repetitions of the original learning material. Don't confuse them with diversions. Don't exhaust them with extra obligations to attend to unnecessary feedback. Instead, give them exactly what they need to develop appropriate conceptions.

One way to provide learners with exactly the information they need is to categorize their responses and provide specific feedback to help them overcome those specific misconceptions. This is done most easily using multiple-choice questions[5]. One of the insights that cognitive science brought to the field of learning is that learners usually get answers wrong for a reason. In other words, mistakes are typically not random. If we can figure out the misconceptions learners have, we can provide them with better feedback. The following question is an example of this:

> **Question: Why is the word "it's" not grammatically correct in the sentence below? Select as many answers as are correct.**
>
> ***The puppy chewed it's food before taking a drink of water.***
>
> A. The word "it's" should be replaced with the word "the."
> B. The word "it's" is a possessive pronoun, and food cannot be possessed or owned.
> C. The word "it's" is a contraction for "it is" and "it is" does not make sense in the sentence.
> D. The sentence ***is*** grammatically correct. The word "it's" is correct because the apostrophe indicates that "it's" is a possessive pronoun.

[5] While multiple-choice questions may provide the easiest way to provide directed feedback, other methods may be equally effective or even more effective. Open-ended questions, for example, can be parsed with a natural-language processor, with appropriate feedback delivered depending on the response. Tutors can provide learners with individualized feedback. Expert systems can diagnose learner understanding and provide appropriate feedback.

Examples of feedback for each answer choice are provided below. The answer choices are shown on the left and feedback options are presented on the right, tailored to each answer choice. Note how the feedback text attempts to anticipate the reasons the learners misunderstand the concept and provide specific corrective feedback[6].

| Answer Choice | Feedback |
|---|---|
| A. The word "it's" should be replaced with the word "the." | • Correct.<br>• It appears that you understood that "it's" and "its" are confusing to use in practice and might easily be used incorrectly. Replacing "it's" with "the" is a strategic move that solves the problem, though you may not be able to use such a clever maneuver in all sentences.<br>• Here's an easy way to remember which to use. The word "it's" is a contraction for "it is." Use "it's" when you mean "it is" and "its" when you don't mean "it is." |
| B. The word "it's" is a possessive pronoun, and food cannot be possessed or owned. | • Incorrect.<br>• Apostrophes can indicate possession, for example when we say, "The puppy's food spilled all over the floor," the apostrophe in the word "puppy's" indicates that the puppy possesses the food.<br>• The word "it's" might look like it signals possessive traits because it has an apostrophe, but "it's" really is a contraction for "it is."<br>• Possessive pronouns (like "its") don't have apostrophes. Some examples: "its," "yours," and "hers."<br>• Here's an easy way to remember which to use. Because the word "it's" is short for "it is," you should use "it's" when you mean "it is" and "its" when you don't mean "it is." |
| C. The word "it's" is a contraction for "it is" and "it is" does not make sense in the sentence. | • Correct.<br>• You seem to understand the concept.<br>• You might also be interested in knowing that the word "its" is a possessive pronoun. Possessive pronouns don't need apostrophes. Some examples: "its," "yours," and "hers." |
| D. The sentence ***is*** grammatically correct. The word "it's" is correct because the apostrophe indicates that "it's" is a possessive pronoun. | • Incorrect.<br>• The word "it's" is NOT a possessive pronoun.<br>• Possessive pronouns don't need apostrophes. Some examples: "its," "yours," and "hers."<br>• Sometimes an apostrophe does indicate possession. For example, we might say, "The puppy's food spilled all over the floor." Because the food is possessed or owned by the puppy, we use an apostrophe. Note that "puppy" is not a pronoun; it's a regular noun.<br>• Here's an easy way to remember which to use. Because the word "it's" is short for "it is," you should use "it's" when you mean "it is" and "its" when you don't mean "it is." |

[6] I created the feedback with only a rudimentary knowledge of the reasons that learners might misunderstand this concept. One advantage I had in writing this, however, is that I misunderstood this when I was young, and then finally I learned a method to distinguish between the two options. While sometimes we can write good feedback based on our knowledge of the subject matter, a more effective way is to analyze actual learners and the misunderstandings that they most commonly exhibit. Such extra effort may or may not be worth the investment, depending on many factors.

## More Extensive Feedback May be Needed to Build Understanding

Learners may need more extensive feedback when they are building understanding than when they are supporting retrieval or developing fluency. Imagine giving your learners a prequestion before you've even introduced the concept. You ask them a question, and then you give them feedback by introducing the topic. Look at the following example.

> **Question: You've recently been hired as a manager for a new Asian-fusion restaurant in town. During the first week you are on the job, a customer says to a waitperson, "I'd like the veggie stir-fry, but I'm allergic to gluten. Can you make sure there's no wheat or gluten in my food?" The waitperson said he would talk to the chef, but because the chef seemed fairly casual about it, the waitperson came to you (as the person in charge), worried that the customer might have an allergic reaction. You wish you knew more about gluten. Are any of the following ingredients in the veggie stir-fry likely to cause the customer a problem?**
>
> 1. Steamed Rice
> 2. Extra-Firm Tofu
> 3. Soy Sauce
> 4. Vegetables
> 5. All are safe

In using this question, you are teaching the topic of Celiac Disease (gluten sensitivity) to restaurant managers in a course on food allergies and sensitivities. You've never introduced the topic before and most of your class is unlikely to know much about Celiac Disease or gluten sensitivity. To introduce the topic, you give them the question above. They answer the question.

Isn't it obvious that you'll have to give them extensive feedback to help them understand all they need to understand to answer the question above? Here's a short list of the kinds of information they may need to learn:

- Gluten is a substance found in wheat, rye, and barley.
- Gluten can cause intense and/or chronic problems for people with Celiac Disease.
- About 1 in 150 people have Celiac Disease. In the typical restaurant, every day will see at least several people who have Celiac Disease.
- Soy sauce usually contains wheat. Most cooks have no idea that it contains wheat.
- Given this, all wait staff and cooks should be trained to know that when a customer says they can't have wheat or gluten, no

soy sauce should be used in his or her food (or wheat-free soy sauce should be used).
- Wait staff have to be trained on how to respond to customers who express a food allergy—to make customers feel comfortable that their food will be safe.
- Soy sauce isn't the only hidden source of gluten, so much more has to be learned as well. For example, many restaurant workers are not aware that regular flour contains wheat.

Giving feedback on prequestions may be the extreme that proves the point. The general principle is the following. When learners are building understanding, they are likely to need more extensive feedback. This won't be true if the learners are memorizing definitions or foreign-language vocabulary, but it will be true if they are learning concepts that are more complicated than simple associations.

Providing learners with more extensive feedback when they are building understanding is a tricky business. Clearly you don't want to overload them with information. Too much information will dilute their attention away from the critical information. To help the learners build understanding, it is often helpful to provide the correct response with a simple explanation of why the response is correct. On the other hand, your learners may require worked examples, graphical animations, or simply more time to fully process the learning material.

Unfortunately, simple recommendations are probably more dangerous than helpful in this instance. Your ability to empathize with your learners and understand what information they need will be particularly helpful in supporting their understanding. You can gain this intelligence through extensive experience with your learners and your subject matter, through pilot testing, or through the use of insightful and experienced subject-matter experts[7].

[7] Many subject matter experts are surprisingly poor judges of what learners need to know. Beware of subject matter experts who organize their expertise by topic rather than with a situation-based appreciation of the real-world cue-action contingencies that decision-making requires.

### *Providing Feedback Before Retrieval is Complete*

Usually when we provide learners with questions or tasks, we wait for them to answer the questions or complete the tasks. We enable them to complete their retrieval attempts. On the other hand, sometimes we enable learners to get feedback without first requiring retrieval; for example, short-circuiting retrieval in the following ways:

1. We put a question on one page and the answer on the following page, allowing the learner to peek at the answer without actually answering the question.
2. In the classroom, we ask a question and quickly call on the first person who raises a hand, short-circuiting retrieval for everyone else.
3. On the computer, we give learners a question and allow them to hit the "continue" button to move forward without answering.
4. On the computer, we give learners feedback if they move their mouse cursor over an answer choice.
5. When coaching some sort of task performance, we give learners hints before letting them go off in the wrong direction.

Researchers discovered several decades ago[8] that when learners are given a choice to short-circuit retrieval to get to feedback, they'll do it. If the answer is on the next page, they'll peek. If it's a click away, they'll click. Learners don't seem to know that retrieval practice is critical to support long-term retrieval. As they try to retrieve information, they may tell themselves, "I know that; let me just see the answer." Perhaps they think, "I can wait here all day, but I'm not going to remember the information." For whatever reason, learners may avoid the benefits of retrieval if given the opportunity to look at the feedback.

The question then becomes, is this short-circuiting helpful or harmful? Or, is it helpful in some situations and harmful in others? Here's the bottom line: Short-circuiting retrieval can sometimes be beneficial early in learning when learners are building understanding, but it's counterproductive when they are practicing retrieval and building fluency.

When learners are building understanding, they may benefit from having extra instructional support and coaching. Similarly, if they attempt to retrieve and have difficulty, this in itself may spur appropriate attention as they subsequently process feedback information. When learning is focused on complex tasks or problem solving, the learning process may be more efficient when we provide learners with feedback and coaching before they stray too far down the wrong path.

On the other hand, after learners understand the concepts being learned, retrieval practice provides substantial benefits, and so anything that short-circuits retrieval is generally counterproductive. Researchers, for example, have found that when learners have a

[8] Kulhavy (1977).

chance to peek at answers, their ability to retrieve at a later time suffers. So, when learners are engaged in retrieval practice, it's important that we prompt them to complete the retrieval opportunity before we provide them with feedback. This means that we don't allow them to peek ahead to answers, we wait after asking a question in the classroom before calling on someone (or we use audience-response technology to ensure that everyone answers the questions), we design our online questions to force a response before delivering feedback, we avoid interrupting learners practicing complex skills.

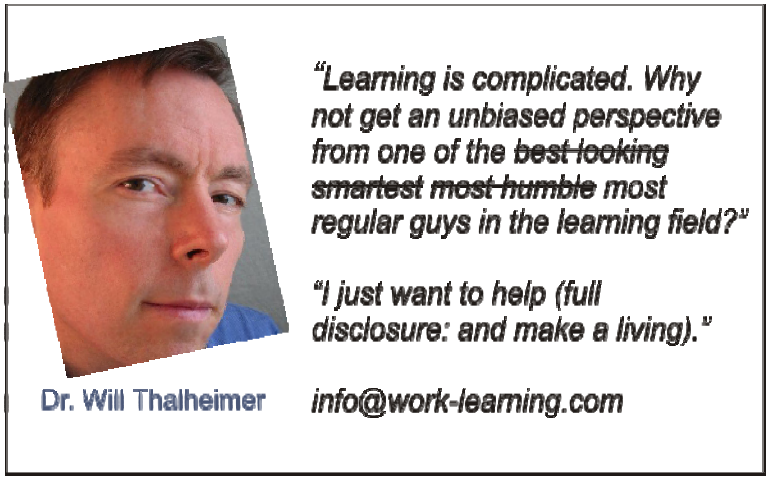

### *Feedback for Incorrect Answers*

We've already discussed how feedback is most important for incorrect answers (and inadequately demonstrated skills). A recommendation might seem fairly straightforward—provide learners with feedback when they respond inadequately. In general this is true, but there is more to it. We've already learned about some of the complications, but we're going to learn about more as we go forward.

We've also learned that the bottom line is that feedback should provide the learner with specific corrective feedback. Much of the time this means providing the learner with the correct answer. However, we've also seen how concept complexity and newness may mediate this recommendation.

Later in this report we'll cover an additional complicating factor, the timing of feedback. For now, we will turn our attention to one of the most overlooked phenomenon regarding feedback for incorrect answers.

## After-Feedback Retrieval Practice

As you should recall by now—I hope this notion is now deep in the marrow of your memory—retrieval practice is one of the most important learning factors there is. If we want our learners to be able to retrieve information later, we have to give them retrieval practice during our learning interventions. Unfortunately, when a learner gets an answer wrong, even if we give them feedback, they don't get a successful retrieval practice. Look again at the following diagram. It represents the learner's retrieval processing for an incorrect answer. Do you see in the diagram any indication that the learner got a cue and used it to successfully search memory?

The learner unsuccessfully searched memory. Notice how the search path never goes from the cue to the correct information. Even if we give our learners feedback, they still haven't traveled down the correct path through memory.

When I first noticed this, I was stunned at the ramifications. Almost nobody provides follow-up retrieval practice to learners after giving them feedback. We've all assumed that feedback was enough. This is another example of how important it is to understand feedback in its context and the human learning system in all its depth and wonder. Only when we know how potent retrieval practice is, can we fully appreciate the importance of providing our learners with an after-feedback retrieval opportunity.

When learners are supporting retrieval or developing fluency, the benefits of after-feedback retrieval will increase as we lengthen the time after feedback. An immediate retrieval opportunity—for example, right after a concept is taught—may produce slight benefits. Because the information to be retrieved will be so easy to access from memory, any practice effects of retrieval will be diminished. It would be analogous to having marathon runners practice by walking 10 meters—such easy practice would not prepare the runners for their difficult marathon journey. Giving learners retrieval practice immediately after they get feedback is similarly useful in helping them support retrieval and fluency—in other words, not very useful.

On the other hand, when learners are building understanding, immediate follow-up retrieval can help the learners clarify their conceptions. As they iterate through learning materials, retrieval practice, and feedback, they build correct mental models of the concepts being learned.

To reiterate, correcting incorrect answers to an optimal level requires that we provide learners with feedback and then give them a later opportunity to retrieve the information from memory. Only in this way have we given our learners a successful practice opportunity going from cue to retrieval.

Dr. Will Thalheimer

*"But seriously, one of the things that brings me the most pleasure is to help others see things they hadn't seen before, to utilize information they've learned from my research efforts, to transform their learning environment to one that really works."*

*info@work-learning.com*

### *Feedback for Correct Answers*

Earlier I hinted that there are two types of correct answers—those for which a learner knows the answer and those for which the learner guessed the answer. Of course, I oversimplified by dichotomizing a continuum into two extremes. There are many shades of gray in between knowing and guessing. Learners know some answers with great confidence, some with a medium amount of confidence, some with a little confidence, and some with no confidence.

Correct answers are not all the same. Some indicate that learners are able to follow correct routes through memory. Some correct answers are more indicative of guessing. If we knew that learners didn't guess at all, we wouldn't have to give feedback on correct answers. We could assume that the learners could traverse the memory route. On the other hand, when learners guess correctly, their lack of confidence indicates that the retrieval route is too tentative to be trusted.

Guessing is less of a problem when open-ended responding is required. With forced-choice questions (like multiple-choice questions) learners are more likely to guess when they don't know an answer. We can't assume, though, that feedback is unimportant for open-ended questions. Even with open-ended correct answers, some answers are probably given with medium levels of confidence, and these answers are likely to benefit from feedback.

For forced-choice questions (multiple-choice, true-false, and matching) the guessing problem is magnified because learners are almost certain to guess at low to mid levels of confidence. When a learner guesses correctly, we have no way of knowing how to tailor the feedback. We don't know how much feedback they need, nor do we know the nature of their misconceptions.

We could, of course, give the learners feedback on all correct answers. This is what we normally do. While this strategy is acceptable, the downside of this approach is that learners may get frustrated in getting feedback on correct answers, especially if that feedback is perceived as too extensive, unnecessary, or time-consuming to read through. Remember that one of the key feedback principles is to give learners the feedback they need—not too much, not too little. Remember also that learners who follow the correct retrieval route through memory—while doing it with confidence—don't benefit much from feedback. Feedback is almost redundant if learners know an answer and retrieve it correctly. Furthermore, if learners find correct-answer feedback frustrating, they may be more likely to skim all the feedback they receive, whether it's feedback they need or not. Also, time spent reading feedback comes with an opportunity cost. It is time that the learners could allocate to other learning activities. So, while feedback on correct answers will produce benefits, such benefits should be weighed against other learning opportunities.

Fortunately, feedback that simply describes the correct answer is less likely to frustrate learners and use valuable time than more extensive feedback. Unfortunately, when learners guess with low levels of confidence, they are likely to need more extensive feedback than simply knowing the correct answer.

Taking the opposite approach—not giving any feedback on correct answers—is imperfect as well. Those who guess will not get the feedback they need. The guessing issue requires us to provide some sort of feedback on correct answers.

In providing feedback for correct answers we need to find a balance between giving too much feedback and not giving enough. The importance of correct-answer feedback depends on how many questions (or tasks) learners are having difficulty with. If learners are answering 99% of the questions correctly, feedback is not that important (in the whole). With such a high success rate, learners are unlikely to be guessing at many answers. Of course, feedback will be critical for the few answers that are incorrect and for those that are guessed correctly.

On the other hand, if learners are answering only 40% of questions correctly for a particular topic, then it's likely that a significant percentage of those correct answers are being guessed with a low level of confidence. This may seem strange, but think about it this way. It is unlikely that learners know 40% of a topic really well and 60% of the same topic very poorly. In a multiple-choice test, when someone gets 40% correct, for example, they may know some of that information with great confidence, but some they are likely to know with medium confidence and some with low confidence.

Note also that when the percentage of correct answers is high, overdoing feedback on correct answers is more problematic because there are so many more correct-answer feedback opportunities to frustrate learners. When the percentage correct is low, there are relatively few such opportunities. On a 100-item test, a 99% score would require 99 correct-answer feedbacks while a 40% score would require only 40 correct-answer feedbacks. What kind of learner would want to slog through 99 extensive feedback presentations when they know the material really well?

Another way to balance these competing demands is to provide short feedback on correct answers and give learners the option of getting more detailed feedback if they feel they need it.

Alternatively, we can measure learner confidence in their answers and provide feedback accordingly. Low-confidence correct answers could get more extensive feedback than high-confidence correct answers. The downside of measuring confidence for each question is that it can involve an extra interaction for the learners. Learners first have to answer a question on the content. They then have to answer another question about their confidence in their answer to the first question. Good interface design can lessen this problem somewhat, but it still is an issue that should be considered. Schemes that

combine question-answering and confidence-rating may be workable for some types of learning materials.

Similarly, if we add an answer choice such as, "I'm not sure of the answer," we may be able to limit guessing somewhat[9]. With this "not sure" answer choice, we might decide to forgo feedback for correct answers, giving feedback only on wrong answers and the "not sure" answers. Of course, this won't completely solve the problem. While low-confidence guessing may be eliminated, medium-confidence guessing is still likely, and such answers will still not get the feedback support they may need.

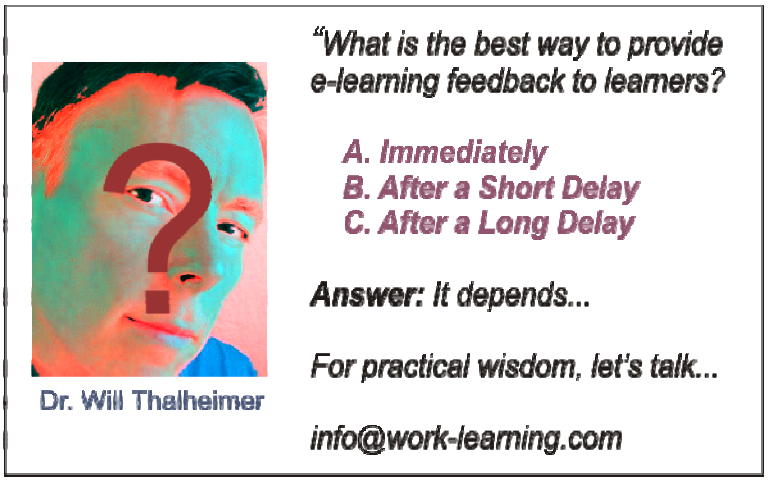


[9] Adding such an answer choice may be untenable on graded tests unless learners are somehow discouraged or disincentivized for guessing.

### *The Timing of Feedback*

When should feedback be delivered? Should we give our learners feedback immediately after they answer a question? After they finish a whole set of questions? Or should we wait for an hour, a day, or a week or more?

The timing of feedback is one of the most vexing questions in the learning field, and yet, unfortunately, the issue remains unsettled. Research support can be found for both immediate feedback and delayed feedback, depending upon the circumstances in which that feedback is employed. Although I have painstakingly combed the research in an attempt to conjure clear recommendations from the fog of conflicting experimental results, I have been unable to elicit general heuristics for all feedback situations.

Nevertheless, there is wisdom in the research. First, it tells us that we should be skeptical of absolutism. In particular, it would be perilous for us to say, "Immediate feedback is always better," or, "Delayed feedback is always better." Before more research is done, all we can say is that neither case is proven.

Later, I will provide my recommendations to help you decide between immediate and delayed feedback. But before I do that, let me turn to one recommendation about the timing of feedback that is clearly consistent with what we know about human memory.

## Deliver Corrective Feedback before Additional Retrieval Opportunities

In general, corrective feedback should be delivered before learners get additional retrieval opportunities. Because of the extraordinary power of retrieval to strengthen memory, it's important that we don't enable our learners to reinforce inappropriate memory traces. Because learners don't usually change their answers or approaches spontaneously without a reason, they're likely to continue answering or practicing incorrectly if not given corrective feedback. When learners get an answer wrong or practice a skill inappropriately, we ought to give them feedback before they attempt to re-answer the question or re-attempt the skill. This doesn't necessarily mean that we should give them immediate feedback, but it does mean that we don't want to delay feedback until after they are faced with additional retrieval opportunities.

There are exceptions to this rule—almost all rules of learning design have exceptions. Sometimes it may be beneficial to let learners hold their misconceptions for a period of time—specifically to drive home learning points regarding meta-issues (issues that go beyond the specific learning points being taught). I once took a class in college where the professor had us read and discuss a whole book over several weeks before he had us read another book which completely debunked the first book. It was a powerful experience because while reading the first book we reinforced our inappropriate conceptions of the topic. Almost all of us believed fully in the conclusions of the first book before we

received corrective feedback from the second book. The second book not only cured our misconceptions, but because of the feedback delay—the second book could be considered feedback—the experience delivered a powerful lesson. We learned that even though a well-argued point of view may seem irrefutably persuasive, we should always remain open to the possibility that better data or information may disprove the whole conception.

It may also be beneficial to delay feedback after retrieval opportunities in simulated or real-world environments to encourage learners to notice natural feedback inherent in those situations. Similarly, it may be helpful sometimes to delay corrective feedback to maintain the learners' feeling of being engaged in an immersive experience. When such engagement prompts learners to exert greater attention to important information, learning can be strengthened[10].

While there may be other exceptions as well, in general corrective feedback should come before subsequent retrieval opportunities. This applies not only to inappropriate responses, but also to appropriate responses that are guessed at, or made with little confidence. Thus, while corrective feedback before subsequent retrievals is more important for incorrect answers than for correct answers, some correct answers may require corrective feedback as well.

Immediate feedback prevents or lessens the problems discussed above because, by definition, it provides feedback before learners face additional opportunities for review or practice. Note that—because retrieval is stronger than feedback—immediate feedback won't always overcome this problem. Sometimes learners get feedback but then subsequently forget the feedback and retrieve their previously incorrect response. Despite this possibility, immediate feedback limits the probability that learners will reinforce inappropriate conceptions. Delayed feedback can do the same thing, but only when it is delivered before learners encounter additional review and practice opportunities. While it thus may seem advisable to avoid delayed feedback altogether, delayed feedback has several advantages over immediate feedback—as we'll discuss later—making it critical that we fully understand how delayed feedback works in practice.

So, to reiterate, when we delay feedback, we may enable learners to reinforce incorrect retrieval routes through memory. In other words, when we delay feedback, learners may continue to reinforce the wrong conceptions—either by incorrectly practicing retrieval or simply by thinking about the information incorrectly. In deciding between immediate and delayed feedback, we have to weigh this possibility. Specifically, we might ask ourselves, what is the likelihood that our learners will spend time thinking about the learning material or practicing retrieval between the time we give them a retrieval opportunity and the time we provide them with feedback?

[10] It's important to note that immersive experiences can easily go astray by focusing learners on inappropriate or less important aspects of the simulated environment. Please let me emphasize that in the text, I said, "*greater* attention to important information."

Most college students are very likely to study material between quizzes and the final exam. Employees taking a four-Mondays-in-a-row supervisory training class may practice inappropriate techniques before getting feedback from their instructor. In these types of cases, delaying feedback may cause difficulties because learners may reinforce inappropriate retrieval routes.

On the other hand, third-graders studying history, and employees taking sexual harassment training are much less likely to be subsequently engaged with the material—third-graders because they aren't yet doing homework and employees taking a course on sexual harassment because there are relatively few opportunities on the job that prompt thinking about this issue (in comparison to employees' many other task demands). In general, the more likely our learners are to re-engage learning concepts after they have been involved in retrieval, the more helpful it may be to provide immediate feedback to prevent reinforcement of misconceptions. If re-engagement is unlikely, delayed feedback may be preferred.

Note how some real-world learning situations can provide their own feedback, and so may not be subject to the same influences. For example, employees who take a course over several weeks to learn a software program may make many attempts to use the program during the weeks they are learning about it. As they do this, they'll receive feedback from trying things out. Because they are essentially creating their own feedback—or their interactions with the software is producing feedback—delaying the formal-learning feedback is not nearly as likely to cause problems due to the reinforcement of incorrect retrievals.

To summarize the past few paragraphs, I've added the discussion above to provide a real-world perspective on what might happen when we consider delaying feedback beyond additional retrieval opportunities. The discussion is important because there is evidence that delaying feedback can be beneficial. Traditionally, we haven't had to worry about the intricacies I've just described because in most academic and employee-training situations our learners have been unlikely to study or engage material before we give them feedback. Partly this is due to our over-reliance on immediate feedback. Partly this is due to the tendency of learners to withhold exertion unless it seems necessary[11].

We can now turn to the more traditional situation in which learners get feedback before they engage in subsequent retrieval efforts or additional study and review.

[11] While you were waiting a week for your college exam results, how likely were you to study the material you were just tested on?

## Immediate versus Delayed Feedback

Feedback can be given immediately, after a short wait (for example, after 10 seconds), after a group of questions (for example, at the end of a quiz), or after a longer period of time (for example, after a day or a week).

If we don't have to worry about the issues described in the section immediately above, which is better, immediate or delayed feedback? To put it simply, the research is not clear on this, but it favors delayed feedback.

Delayed feedback produces the same advantages as spaced learning. The spacing effect is one of the most researched findings in all of learning and memory research. When learning points are repeated after a delay, learners improve in their ability to remember the information that was repeated. There are several reasons that spacing produces its benefits. First, spaced learning requires more effortful cognitive processing. The additional effort leads to better learning. Second, spaced learning provides learners with more varied repetition contexts. Variations in the background context of learning helps learners develop additional retrieval routes to the information stored in memory. To learn more about the spacing effect, check out my research-to-practice report *Spacing Learning Over Time: What the Research Says*, available at www.work-learning.com/catalog/. The report provides practical recommendations regarding the spacing effect, but also describes the research in depth, providing over one hundred research citations.

Delayed feedback essentially provides a spaced learning opportunity. Where retrieval is the initial learning event, subsequent feedback acts as a repetition. Where retrieval follows an initial learning presentation, retrieval acts as a first repetition and feedback provides a second repetition. In either case, when feedback is delayed, it acts just like a spaced repetition, creating the benefits of the spacing effect.

Delayed feedback—when it is processed at all—is more likely to be processed at a deep level than immediate feedback[12]. Delayed feedback also provides learners with at least two background contexts. When feedback is delayed, the retrieval context and the feedback context differ from one another, whereas with immediate feedback they are practically the same.

When it is important that learners be alert to real-world cues related to the retrieval context, delayed feedback is better than immediate feedback because it enables learners to process those cues themselves. For example, a learner practicing for a drivers' test may benefit if the driving instructor waits fifteen seconds after an error before providing feedback. Such a delay enables the learner to take responsibility for fully processing the contextual cues, whereas immediate feedback may short-circuit that processing by placing the responsibility for noticing contextual cues in the hands of the instructor.

[12] Note that delayed feedback will only produce its effects when learners process the feedback. If your learners are unlikely to access feedback after a delay, then they won't receive the benefits suggested above.

Immediate feedback provides benefits—as already mentioned—to prevent learners from reinforcing inappropriate conceptions after retrieval practice. Immediate feedback may similarly help learners who will face a lot of interfering learning materials between retrieval practice opportunities. Interference can occur when learners are faced with additional information related to the same topic.

Good examples of high potential for interference include college students taking an introductory survey course on psychology, employees learning a new software program, and teenagers taking a geometry course. In each of these cases, the material may appear to the learners as a broad and extensive undifferentiated chaos. The learners simply don't have the experience with the material to avoid confusions. Note how interference is caused by the intersection of learners and materials. The less experienced the learners are with the discipline they are learning about, the more the likelihood of interference. The broader the coverage of the material and the more it lends itself to confusion, the greater the likelihood of interference. Learners who have more experience in a discipline are better able to perceive learning material from that discipline as differentiated than those with little experience. In other words, they have the mental models to help them make sense of it, whereas novices in that discipline are more likely to see the same material as a murky ocean of data.

Immediate feedback may help learners overcome interference by creating relatively stronger bonds between retrieval situations and their appropriate responses than delayed feedback can produce. So, if your learners are likely to involve themselves in extensive review of interference-laden materials before they have to produce responses, immediate feedback may prepare them better than delayed feedback.

Note how these interference-laden situations mirror the kinds of situations that I've described as "Building Understanding." On the surface of it, it just doesn't make sense that when a learner is piecing together arrays of building blocks into a fully-formed complex concept, they wouldn't need some sort of feedback as they build up from prerequisite concepts. If the conceptual foundation they build for themselves is wrong, adding to that faulty foundation is problematic. Feedback provided before these prerequisite mental modelettes are built should keep learners from flailing around too much. For this reason, I will tentatively recommend immediate feedback as learners build understanding.

An alternative argument could be made that learners who flail around may build richer mental models through trial and error. I'm sympathetic to this argument, but worry that such an approach is a time-waster. Besides, if concepts are complex, learners are likely to flail around when implementing them in the complexities of their real-world performance situations anyway, so the instructional time should be focused more efficiently.

## Three Simple Heuristics for Feedback Timing

Let me empathize with you, dear reader. Although I've tried to write clearly about the timing of feedback, I know that even my most elucidating prose is probably not adequate to the task of clarifying this difficult material[13]. For those of you who want a quick rule of thumb, instead of the intricacies detailed above, I provide the following. I offer it with the caveat that it may not be as effective as the more nuanced approaches detailed above. Of course, the best approach to ensure effectiveness might be for you to experiment with several methods and see which produces the best results with your learners and your materials in your specific situation. You should also feel free to engage me for a few hours of consulting to help you design your feedback strategy.

The first heuristic is based on two key ideas. The first is that retrieval practice is more important than feedback. The second is that the benefits of delayed feedback generally derive from the characteristics of spaced learning.

My simple suggestion is to provide immediate feedback after the first retrieval-practice opportunity and then later, to provide a second spaced retrieval-practice opportunity. This strategy provides the spaced-learning benefits of delayed feedback, it provides multiple retrieval opportunities, and it makes it likely that learners will receive at least one correct retrieval practice[14].

A second heuristic simply involves giving learners feedback after a short delay—for example, giving learners feedback at the end of a quiz rather than after each quiz question. This is a compromise solution. It provides almost-immediate feedback and utilizes a slight spacing effect.

Finally, let me reiterate that feedback—whether it be immediate or delayed—is significantly better than providing no feedback. So the third heuristic is to provide feedback—whether immediate or delayed—in one way or another. Do whatever the logistics of your situation allow to add feedback after retrieval practice.

[13] In fact, even when I reread my own paragraphs here, I'm inclined to scream.

[14] To maximize the benefits of this strategy, you may need to monitor the delay between the first and second retrieval-practice opportunities. If the delay is too short, you won't get the full benefits of spacing. If the delay is too long, the second retrieval practice may produce too many retrieval failures. You'll need to monitor how well your learners are performing on the second retrieval opportunity and how well they are able to retrieve information at an even later time. Remember, our ultimate goal of learning design is long-term retrieval. If, at a later time, your learners are getting, for example, 98% correct, things are good. If they're only getting 40% correct, you'll need to modify your tactics. You may need better upfront instruction, better feedback, or more retrieval-practice opportunities, for example.

### *Getting Learners to Pay Attention to the Feedback*

Feedback isn't worth a damn unless learners pay attention to it. While many researchers forget this subtlety because their learners are specially primed to pay attention to the learning task, real-world learners are not always so tuned into the feedback they are offered. For practical purposes, we as learning professionals have to figure out ways to get our learners to pay attention to our learning material, including the feedback we provide. This is one of the most important points in this document. In fact, if the most important point is that providing feedback is almost always better than not providing feedback, and the second most important point is that good feedback must be corrective, the third most important point is that the feedback must be presented so that learners will pay attention to it in a manner that is conducive to learning.

As I always suggest, you'll need to monitor your learners' behaviors to see what mechanisms are particularly effective. And, while you may conclude that your learners always respond well to a particular method of feedback, it is more likely that (a) you'll need to vary your feedback mechanisms to maintain learner attention at high levels, and (b) different feedback mechanisms will be given more or less attention depending on a whole host of factors—including such things as prerequisite knowledge, personal relevance, novelty of design, grading procedures, burnout, caffeine blood levels, and so forth, ad infinitum.

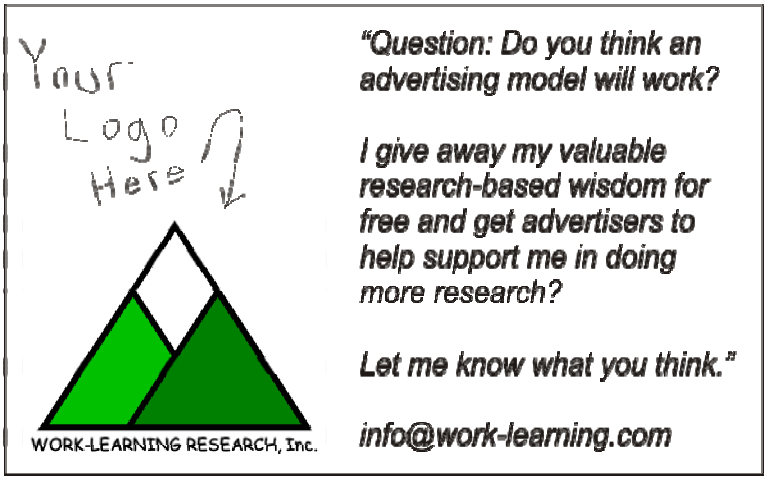

### *Revisiting the Distinction between Building Understanding and Supporting Retrieval and Fluency*

While my distinction between building understanding and supporting retrieval and fluency is inspired by the research, it is still a tentative formulation until additional research confirms its boundary conditions.

The following table summarizes my recommendations along several factors.

| | **To Build Understanding** | **To Support Retrieval and Fluency** |
|---|---|---|
| | Earlier in Learning or for Complex Concepts | Later in Learning or for Simple Concepts |
| **Amount of Feedback** | More Elaborate | More Brief and Succinct |
| **Timing of Feedback** | More Immediate | More Delayed |
| **Feedback on Correct Answers** | More Feedback | Less Feedback |
| **Empathy Given with Feedback** | More Important | Less Important |
| **Ensure Retrieval Before Giving Feedback** | Not Required | Required |
| **When to Give After-Feedback Retrieval Practice** | Soon After Feedback | With Delay After Feedback |

# Other Considerations

So far in this research-to-practice report, I've covered primary issues in feedback, especially those that have been extensively researched. As in almost all areas of human learning, there are other important factors that have yet to be given the research focus they deserve. If you only follow the recommendations above, you'll do great with your feedback methods.

As you can probably tell by now, I don't think it advisable to swallow the research whole. For each topic I research, I try to locate its strengths and limitations and develop some wisdom about boundary conditions, contingencies, and remaining questions. The following sections are my attempt to convey such wisdom to you.

### *Using Confidence Ratings to Augment Feedback*

Learners' confidence in an answer would seem, on the surface, to be related to feedback. Learners who are correct and highly confident would seem to need less feedback than people who are correct with low confidence. Don't guessers need feedback even if they are correct? Similarly, learners who answer incorrectly and confidently might need a different type of feedback than those who answer incorrectly with low confidence.

While research has been done that touches on confidence, there is less than adequate clarity to make many definitive recommendations. In fact, I am only confident in recommending that low-confident correct answers probably need more feedback than high-confident correct answers, which may not need any feedback at all.

### *Using the Answer-Until-Correct Feedback Mechanism*

While most forced-choice testing (for example multiple-choice testing) requires test takers to select one answer and move on to the next question, the answer-until-correct method forces learners to select answer choices until the correct answer is chosen. In short, not much good research has been done on this question, though some results are suggestive.

From a practical standpoint, one downside of this answer-until-correct feedback mechanism is that learners who aren't given incentives for taking the questions seriously may be likely to click through the answers without thinking about them, lessening the benefits of retrieval. Moreover, such designs can also create disdain for the learning programs in which the questions are embedded.

More research is required before recommendations can be made. In the meantime, it may be worth experimenting with an answer-until-correct methodology, especially if your learners can be prompted to take each answer choice seriously.

### *The Dangers of Evaluative Feedback*

Feedback can consist purely of information or it can contain value statements as well. For example, we could give the feedback, "Correct, antibacterial soap is not recommended for general use—simple soap and water is recommended." Alternatively, we could add an evaluative element, "Yes, excellent answer, only 25% of people get this correct, antibacterial soap is not recommended for general use—simple soap and water is recommended." Research on question feedback has found that positive reinforcement, like saying, "Good Job!" creates no learning benefits above and beyond the information that is provided. In short, it's the information that provides benefits.

Similarly, handholding is not generally beneficial either. For example, saying, "Sorry, that's incorrect, you'll do better next time" is unlikely to produce benefits over and above telling the learners that their answer is wrong and explaining the misconception.

In fact, there are suggestions from the research that focusing learners' attention on how well they're doing, especially in comparison with others—as opposed to focusing on what they need to learn—can promote future metacognitive thinking that hurts learning. Carol Dweck (Dweck, 1986, 2006; Elliott & Dweck, 1988; Dweck & Leggett, 1988) and others have found that learners who focus on what they need to learn utilize more effective learning strategies than learners who focus on how well they're doing. Learners who just want to look good may take shortcuts that hurt their learning. For instance, these ego-focused learners may focus their study time on being able to regurgitate information to do well on a test, whereas learning-focused learners may let their curiosity prompt them to build rich mental models of the topic area.

As with all learning advice, we need to understand the general recommendations with some wisdom. Some learners may thrive, for example, on competition. It may be the only thing that motivates them to pay attention. On the other hand, we ought to realize that if we use competition we may create immediate benefits due to heightened attention, but we may be short-circuiting the opportunity to help our learners create richer networks of understanding about a topic.

The gap between learning-to-learn and learning-to-look-good can be lessened somewhat if we create performance mechanisms (tests, quizzes, exercises) that require deeper learning and significantly less regurgitation of information.

Is it okay to use the exhortations, "Good job!" or "Sorry, wrong answer," or similar short evaluative phrases? For most people, these interjections are okay. The majority of learners will read those statements without triggering their evaluative antennae. If you're worried that certain statements will prompt your learners to focus on how they are being perceived, you can simply avoid using them. It's simple enough to say, "Correct, the right answer is…" or "Incorrect, the correct answer is…" or simply, "The correct answer is…" For variety, you can begin your feedback with other relatively non-evaluative responses,

such as good, right, yes, okay, no, wrong, not correct, oops, etc. Of course, as always, I suggest you experiment with different methods to see what works best with your learners.

### *Feedback May Also Have Subsequent Learning Effects*

In designing learning interventions, one thing that is critical to understand—but which is usually overlooked—is that learning methods not only impact our learners' current learning practices, but can have profound effects on their future learning as well. For example, inserting questions into a history text that ask learners to recall which year historic events took place (e.g., "In what year did the Battle of Hastings take place?") will prompt learners to focus on historic dates in future history reading assignments as well, even if no questions are provided in those texts. Such questions not only influence learners in their current learning efforts, but can prompt learners to turn away from more critical information and themes in future information as well.

Feedback is no different. As a learning factor, it too can influence future learning. While the research shows that short and very specific feedback about the correctness of responses is most effective in helping learners learn those responses, the research is largely silent on how such short feedback might affect future learning[15]. What I worry about is that such specificity may not support learners in being able to generalize to broader circumstances. It could be, for example, that slightly longer feedback that describes the general application of the specific learning point might enable learners to transfer their learning to broader circumstances.

Another potential issue with feedback is that it may diminish the incentive for learners to produce their own feedback. This concern isn't really relevant for simple questions or tasks. Your learners aren't going to gain very much if you force them to look up the capital of Massachusetts, the date of the Battle of Hastings, or whether Abraham Lincoln was the 15th or 16th or 17th or 18th President of the United States. On the other hand, learners who are using a hands-on simulation to learn how to diagnose and fix electrical failures may benefit from testing other hypotheses after their first attempt fails. By withholding feedback, we provide more authentic practice, and thus, we better prepare them for their future roles in unsupported situations.

This concern is relevant, I think, mostly to complex skills—for example, solving an algebra problem, learning a foreign language, or playing shortstop—because it's in these types of activities that learners have to respond to complex and subtle situational cues. By withholding feedback and enabling learners to explore complex concepts, systems, or procedures, we may also be giving them an opportunity to build richer mental models. Of course, we have to balance these benefits with the benefits that feedback can provide in preventing learners from being too long lost in inappropriate situations. Perhaps

[15] The research is largely silent on feedback's effects on future learning because virtually no research has been done on this. My analysis is based on other learning research that shows how learners' future attention is influenced by previous learning methodologies.

withholding feedback should only be done after learners have developed enough expertise so that they can work their way through unclear situations, learning as they do.

I have written this section to acknowledge that feedback may have future effects beyond the current learning situation. Be careful not to take these recommendations too far. There's not enough research on these questions to be definitive—though research on feedback in learning physical movement skills is suggestive—but we might want to be a little bit circumspect in providing feedback, avoiding a willy-nilly overuse of feedback when it might prompt our learners to ignore critical situational cues. Second, the future effects of feedback probably play a secondary role to feedback's effect on what is currently being learned. You may not want to minimize feedback's benefits in your current learning intervention just to support the prospects of future learning.

This section may seem to conflict with the general principle that direct corrective feedback is best. The subtlety is that corrective feedback is best to support the learning of the concepts or skills directly relevant to the feedback, but that such feedback might produce unintended effects on future learning events, for example, making it less likely that learners will subsequently be able to notice relevant situational cues. Thus, feedback has direct backward effects and indirect future effects.

Consider what people have to do in their real-world occupations. They are constantly faced with Situations. They have to Evaluate those situations—in other words, they have to make sense of those situations. Once they understand the situation, they have to make a Decision about it. Then they have to take an Action. This can be diagrammed as follows:

Situation → Evaluation → Decision → Action

I call this the SEDA model, after the acronym formed by the first letters. This model nicely simplifies what people do in the real world, especially when we consider that each Action creates a new Situation that people must deal with. This new Situation provides feedback on a person's previous efforts at Evaluation, Decision, and Action.

In the real world, then, learning at optimal levels requires people to become experts in Evaluating particular Situations, making particular Decisions, and taking particular Actions. To prepare our learners to handle these activities when they enter their real-world performance situations, we need to give them practice in each one of these activities. This is where our feedback comes into play. If our feedback undermines such practice, it will hurt our learners in their ability to Evaluate, Decide, and take Action. For example, if we tell an emergency room nurse that she'd better refill the IV bag, are we leaving her unprepared to monitor the IV bag on her own?

### *On-the-job Follow-up Feedback*

Although this report focuses on feedback for formal learning situations—feedback on test questions, quizzes, and practice opportunities—feedback can also be delivered to people when they're in a doing mode, as opposed to a learning mode. The most obvious example of this is when managers give their direct reports feedback on their job performance. If the new salesperson doesn't book enough sales visits, the district sales manager may work with her on how to make more high-quality phone calls to sales prospects. When a project-team member hurts team morale by continually blaming problems on others, the project manager can intervene with feedback.

For this type of feedback, it's generally best to provide people feedback that is (a) specific, (b) focused on the behavior to be corrected (as opposed to focusing on the person's general ability or tendencies), (c) non-evaluative, and (d) delivered soon after the inappropriate behavior. You may note the similarities between this type of feedback and feedback for formal learning situations.

One subtle difference is that whereas on-the-job feedback should generally be delivered soon after the targeted incident, feedback on questions and for other learning events can often provide more benefits if delivered after a delay. Behavioral feedback requires some immediacy because it is difficult for those getting the feedback to remember all the contextual cues that guided their behavior at the time of the incident. In most formal learning situations, the contextual cues can be reinstated in working memory after a delay. For example, a test question can be shown again. The point here is that, while there are similarities, we may be in danger of generalizing too widely if we assume that learning feedback and behavioral feedback always require the same recommendations for action.

### *Training Follow-Through as Feedback*

Learning events don't have to be isolated from the situations they are intended to improve. Often, however, K-12 and higher education courses are devoid of real-world relevance. Similarly, employee training sessions are all too often seen as one-time events that learners can forget about when they return to their jobs.

Recently, a movement in the training field has begun to bridge the gap between training and workplace implementation of what was learned. Software has been developed to track training-implementation goals well after the training sessions have concluded[16]. Learners are given goals—or develop them on their own—regarding how they'll take their learning back to the job. Later, the software reminds them of their goals and tracks their progress toward their goals. The learners' managers and the administrators of the

[16] See, for example, the Fort Hill Company's Friday5s software, and ZengerFolkman's Action Plan Mapper.

learning can be kept apprised of learners' progress toward their implementation goals as well. In addition, some of these packages can provide learners with follow-up learning events and retrieval practice.

In a very real sense, these training follow-through systems provide learners and learning designers with a variety of feedback mechanisms related to learning results.

Of course, software isn't necessary to provide learners with training follow-through even if it does formalize the process. Managers can take responsibility for ensuring that what was learned is properly implemented on the job. Either way, feedback is possible not just for traditional learning events like tests and quizzes, but also can be utilized as learners try to implement what they've learned.

### *End of Practical Overview*

My summary of recommendations for feedback can be found on the following two pages.

Part 2 of the report is devoted to the detailed research results on which Part 1 of this report was based. It includes the research references cited in both parts of the report.

While reading Part 2 may be more difficult, it can reinforce what you learned in Part 1.

## Summarizing Principles and Recommendations

1. The most important thing to remember about feedback is that it is generally beneficial for learners.

2. The second most important thing to remember about feedback is that it should be corrective. Typically, this means that feedback ought to specify what the correct answer is. When learners are still building understanding, however, this could also mean that learners might benefit from additional statements describing the "whys" and "wherefores."

3. The third most important thing to remember about feedback is that it must be paid attention to in a manner that is conducive to learning.

4. Feedback works by correcting errors, whether those errors are detected or hidden.

5. Feedback works through two separate mechanisms: (a) supporting learners in correctly understanding concepts, and (b) supporting learners in retrieval.

6. Early in learning a topic, learners need to focus on building understanding. Later in learning a topic, learners need to focus on supporting long-term retrieval and fluency.

7. Complex concepts will usually require more time for learners to build understanding.

8. To help learners build understanding, feedback should diagnose learners' incorrect mental models and specifically correct those misconceptions, thereby enabling additional correct retrieval practice opportunities.

9. To prepare learners for future long-term retrieval and fluency, learners need practice in retrieving. For this purpose, retrieval practice is generally more important than feedback.

10. Elaborative feedback may be more beneficial as learners build understanding, whereas brief feedback may be more beneficial as learners practice retrieval.

11. Immediate feedback prevents subsequent confusion and limits the likelihood for continued inappropriate retrieval practice.

12. Delayed feedback creates a beneficial spacing effect.

13. When in doubt about the timing of feedback, you can (a) give immediate feedback and then a subsequent delayed retrieval opportunity, (b) delay feedback slightly, and/or (c) just be sure to give some kind of feedback.

14. Feedback should usually be provided before learners get another chance to retrieve incorrectly again.

15. Provide feedback on correct responses when:

    a. Learners experience difficulty in responding to questions or decisions.
    b. Learners respond correctly with less-than-high confidence.
    c. All the information learned is of critical importance.
    d. Learners are relatively new to the subject material.
    e. The concepts are very complex.

16. Provide feedback on incorrect responses:
    a. Almost always.
    b. Except:
        i. When feedback would disrupt the learning event.
        ii. When it would be better to wait to provide feedback.

17. When learners seek out and/or encounter relevant learning material either before or after feedback, this can modify the benefits of the feedback itself.

18. When learners are working to support retrieval or fluency, short-circuiting their retrieval practice attempts by enabling them to access feedback in advance of retrieval can seriously hurt their learning results.

19. When learners retrieve incorrectly and get subsequent well-designed feedback, they still have not retrieved successfully; so they need at least one additional opportunity to retrieve—preferably after a delay.

20. On-the-job support from managers, mentors, coaches, learning administrators, or performance-support tools can be considered a potentially powerful form of feedback.

21. Training follow-through software—that keeps track of learners' implementation goals—provides another opportunity for feedback.

22. Feedback can affect future learning by focusing learners on certain aspects of learning material at the expense of other aspects of learning material. Learners may take the hint from the feedback to guide their attention in subsequent learning efforts.

23. Extra acknowledgements (when learners are correct) and extra handholding (when learners are wrong) are generally not effective (depending on the learners). In fact, when feedback encourages learners to think about how well they appear to be doing, future learning can suffer as learners aim to look good instead of working to build rich mental models of the learning concepts.