CHAPTER 5

What Kinds of Practice and Feedback Enhance Learning?

When Practice Does Not Make Perfect ...

I teach a public policy course to juniors, and I believe strong communication skills are essential to moving up the ranks in the public sector. As a result, I require my students to write frequently. The three papers I assign focus on the different types of writing my students will potentially do: a policy briefing, a persuasive memo to their boss, and an editorial for a newspaper. I had expected the students' writing on these assignments to be at least decent because all of our students are required to take two writing courses in their first year. Then, when I saw the serious problems in their first papers, I thought at least I could help them improve. So I have been spending an enormous amount of time grading and writing margin comments throughout their papers, but it does not seem to be doing any good: the second and third assignments are just as bad as the first. As much as I think these assignments are useful because they prepare students for their future professional lives, I am ready to nix them because the students' writing is so poor and my efforts are bringing about little or no improvement.

Professor Norman Cox

They Just Do Not Listen!

Last semester, when I taught Medical Anthropology, the students' research presentations were all glitz and very little substance. So this time, because this project is worth 50 percent of their final grade, I tried to forewarn my students: "Do not be seduced by technology; focus on substantive anthropological arguments and create engaging presentations." And yet, it happened again. Last Tuesday, student after student got up in front of the class with what *they* believed to be engaging presentations—fancy fonts in their PowerPoint slides, lots of pictures swishing on and off the screen, embedded video clips, and so on. It was clear they had spent hours perfecting the visuals. Unfortunately, although their presentations were visually stunning, the content was very weak. Some of the students had not done thorough research, and those who did tended merely to describe their findings rather than craft an argument. In other cases, students' arguments were not supported by sufficient evidence, and most of the images they included were not even connected to the research findings. I thought I was clear in telling them what I wanted and did not want. What is it going to take to make them listen?

Professor Tanya Strait

WHAT IS GOING ON IN THESE STORIES?

In both stories, the professors and their students seem to be putting in time and effort without reaping much benefit. For example, Professor Cox makes lengthy comments on his students' writing but fails to see any improvement across assignments. Professor Strait's students spend an inordinate amount of time on aspects of the presentation that actually matter least to her,

despite the guidance she gave them. And both professors are understandably frustrated that students' learning and performance is not up to expectations. A theme running through both stories is that time is being misspent—just the kind of mistake that neither students nor instructors can afford to make.

In the first story, Professor Cox's students probably enter his course with only basic writing skills. Unfortunately, even though the students may begin to develop additional writing skills through the practice they get during the first writing assignment, these new skills are not built upon through the later assignments. Recall that Professor Cox's assignments involve different genres (policy briefing, memo, and editorial). This means they involve somewhat different writing skills to address the distinct goals, audiences, and writing styles specific to each (see Chapter Four). Moreover, even though Professor Cox gives plenty of comments on his students' papers, the students probably have little opportunity to incorporate this feedback into further practice because each subsequent assignment is so different from the previous ones.

In the second story, Professor Strait tells her students that their arguments should have substance and their presentations should be engaging. However, her students seem not to understand what constitutes a substantive anthropological argument based on thorough research or what characteristics she identifies with engaging presentations. Although it is true that Professor Strait's students have spent the bulk of the semester reading and analyzing anthropological arguments, they have had relatively little opportunity to conduct library research and construct arguments of their own. So this partly explains their disconnect. Similarly, although these students have accumulated a good deal of prior experience giving oral presentations, they have not done so earlier in her course, so they mistakenly equate putting glitz in their presentations with what Professor Strait wants. Thus, the students probably have only minimal skill at argument

construction and yet great familiarity with applying technical skills to prepare PowerPoint slides (for example, adding animations, pictures, and sound). Thus, it appears that these students are falling back on the more comfortable task of working on visuals at the expense of articulating an argument in their presentations. Professor Strait reasonably assumes that her warnings should be sufficient to guide students, but students often need significantly more guidance and structure than we would expect in order to direct their efforts productively. With only one chance to "get it right" with regard to this large-scale project, these students end up losing a key learning opportunity.

WHAT PRINCIPLE OF LEARNING IS AT WORK HERE?

We all know that practice and feedback are essential for learning. Unfortunately, the biggest constraint in providing sufficient practice and feedback to students is the time it takes—both on the part of students and faculty. Although we cannot control the length of a semester or class period, we can be more *efficient* in designing practice opportunities and giving feedback. Thus, this chapter focuses on ways to "work smarter" by exploring what kinds of practice and feedback are most productive.

It is important to acknowledge that all practice is not equal. In particular, there are more and less effective ways students can practice. Consider two music students who spend the same amount of time practicing a piece after having made several errors in a difficult passage. If one of the students practices for an hour, spending the majority of that time working on the difficult passage and then playing that passage in the context of the whole piece, this student will be likely to show sizeable performance gains. However, if the other student spends the same hour but

uses that time to play through the whole piece a few times, much of that time will be spent suboptimally by practicing parts of the piece that were already mastered. This is reminiscent of Professor Strait's students, who seem to spend much of their time on what they already know—how to make fancy PowerPoint slides—only to miss their main chance at practicing less developed skills. In other words, *how* students spend their time on a learning activity (either in or out of class) determines the benefits they gain.

This problem of unproductive practice is even worse when students fail to receive sufficient feedback along the way. Think about the first music student who spent considerable time on the problematic passage rather than playing the whole piece multiple times. Even though this student's approach had greater potential to fix all the errors, this student could have introduced new errors without realizing it because no feedback was provided. In this way, lacking feedback, the first student's practice actually could have entrenched new, bad habits. This example highlights the critical role that feedback plays in keeping learners' practice moving toward improvement. In other words, students need both productive practice *and* effective feedback.

***Principle**:* Goal-directed *practice coupled with* targeted *feedback are critical to learning.*

At one level, this principle states the obvious: practice is important, and feedback is helpful to learning. To be clear about terminology, we define "practice" as any activity in which students engage their knowledge or skills (for example, creating an argument, solving a problem, or writing a paper). We define "feedback" as information given to students about their performance that guides future behavior. However, the full potential of practice and

feedback is not realized unless the two are effectively combined. For example, Professor Cox provides an enormous amount of feedback, but it is not coordinated with practice opportunities in which students could incorporate the feedback and refine a repeated set of skills. In contrast, when practice and feedback are focused on the same aspects of students' performance, students have the chance to practice and refine a consistent body of new knowledge and skill. Figure 5.1 depicts this interaction as a cycle: practice produces observed performance that, in turn, allows for

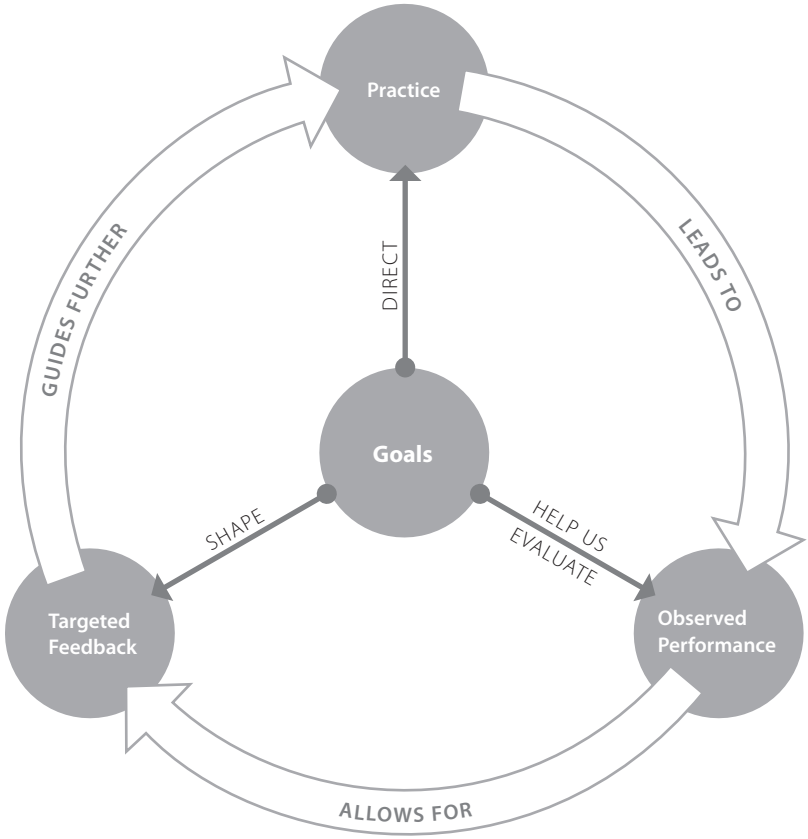

Figure 5.1. Cycle of Practice and Feedback

targeted feedback, and then the feedback guides further practice. This cycle is embedded within the context of learning goals that ideally influence each aspect of the cycle. For example, goals can direct the nature of focused practice, provide the basis for evaluating observed performance, and shape the targeted feedback that guides students' future efforts.

Although practice and feedback ideally go hand in hand—as this chapter's principle and Figure 5.1 indicate—each has a sizeable body of literature. So we discuss the research in two major sections below—one on practice and the other on feedback—and highlight the importance of their coordination.

WHAT DOES THE RESEARCH TELL US ABOUT PRACTICE?

Research has shown that learning and performance are best fostered when students engage in practice that (a) focuses on a specific goal or criterion for performance, (b) targets an appropriate level of challenge relative to students' current performance, and (c) is of sufficient quantity and frequency to meet the performance criteria. The following sections focus on these three characteristics of practice.

Focusing Practice on a Specific Goal or Criterion

Research shows that the amount of time someone spends in *deliberate practice* is what predicts continued learning in a given field, rather than time spent in more generic practice (Ericsson, Krampe, & Tescher-Romer, 2003). One of the key features of deliberate practice is that it involves working toward specific goals. As an illustration of the power of such goal-oriented practice, research shows that world-class musicians spend much of their time

engaging in rather demanding practice activities, continually monitoring their performance toward a particular goal, and then, once it is achieved, pushing themselves to strive for a new goal (Ericsson & Lehmann, 1996; Ericsson & Charness, 1994). In contrast, we all know of people who have studied a musical instrument—even spending considerable time practicing it—but who do not achieve a very high level of performance. Ericsson's explanation of these contrasting paths is that those who spend their considerable practice time working deliberately toward a specific goal tend to go on to be expert musicians, whereas those who do not engage in such deliberate practice do not.

Intuitively, it makes sense that having specific goals for practice would be helpful to learning. Goals provide students with a focus for their learning, which leads to more time and energy going to that area of focus. Consistent with this, Rothkopf and Billington (1979) found that students who had specific goals when they were learning from a text paid more attention to passages that were relevant to their goals and hence learned those passages better. Another advantage of having a goal to direct one's learning is that one can monitor (and hence adjust) one's progress toward that goal along the way (see Chapter Seven).

A key challenge in providing goal-directed practice is that instructors often think they are conveying specific goals to students when, in fact, they are not. This is natural because, as experts, we often see things very differently from our students (see Chapter Four), and so we tend not to recognize when our stated goals are unclear to students or when students are likely to misinterpret our criteria. A case in point is Professor Strait, who thought she was being clear by advising her students to focus on "substantive anthropological arguments" and "engaging presentations"—two ideas that carried specific meaning in her field of expertise. However, her students did not share that expertise, so they did not

share her sense of the specific goals for their work. Without a clear idea of what Professor Strait wanted, the students "filled in the blanks" based on their prior experience (see Chapter One). Unfortunately, in this case, students' interpretations of the goals led them to spend their time in a way that gave more practice to skills they already had developed (such as creating glitzy PowerPoint presentations) and less practice to skills they needed to develop (such as creating anthropological arguments).

When instructors do not clearly articulate their goals, it is difficult for students to know what (or how) to practice. For example, giving students the goal of "understanding a key concept" tells rather little about the nature or level of understanding students should be trying to attain. In contrast, the goals of "recognizing when a key concept is at issue" or "explaining the key concept to a particular audience" or "applying the key concept to solve problems" are more concrete and directive. Note that these more specifically stated goals share several key features. First, they all are stated in terms of something students *do*, which automatically leads to more concrete specifications that students can more easily interpret correctly. Second, all of these goals are stated in such a way that students' performance can be monitored and measured (by instructors as well as students themselves), which enables the provision of feedback to help students refine their performance or learning. For more information on articulating effective learning goals (also called learning outcomes or objectives), see Appendix D.

The notion of articulating goals in a measureable way still leaves open the question (to students and instructors) of *how much* of a particular measureable quality is enough for the goal to be achieved. Research has shown that clearly specified performance criteria can help direct students' practice and ultimately their learning. For example, Goodrich Andrade (2001) found that creating a rubric (a clear description of the characteristics associated

with different levels of performance; see Appendix C) and sharing it with students when an assignment is distributed leads to better outcomes—both in terms of the quality of work produced and students' knowledge of the qualities associated with good work.

An important caveat here, however, is that the goals one specifies must be in accord with what one really wants students to learn. For example, Nelson (1990) studied a case in which students were given detailed specifications for a research paper, such as the requirement to include at least three pieces of evidence supporting their argument. In writing their papers, students took this and other similar prescriptions to heart and included the required pieces of evidence in their writing. An important missing piece, however, was that the paper assignment did not specify higher-level goals such as having a well-organized paper or making a coherent argument. Thus, although these students included the required pieces of evidence in their papers, they tended to fall short on other important criteria. A key implication of this work is that explicitly communicating goals for students' performance can indeed guide their work, but one must be sure that those goals are ones that will support students in what they need to do and learn.

Identifying the Appropriate Level of Challenge for Practice

Specifying goals and criteria is not enough. To ensure that students' practice has a significant effect on learning, the practice they do should be at an appropriate level of challenge and, as necessary, accompanied by the appropriate amount and type of support. An appropriate level of challenge is neither too hard (the student struggles, makes many errors, and possibly gives up) nor too easy (the student completes the goal without much effort and

is not pushed to improve). This relates to the notion of deliberate practice mentioned earlier. As it turns out, deliberate practice more specifically is defined as working toward a *reasonable yet challenging goal* (Ericsson, Krampe, & Tesch-Romer, 2003).

Identifying the appropriate level of challenge seems possible, albeit potentially time consuming, to accomplish in one-on-one teaching and learning situations. Indeed, research has shown that the success of one-on-one tutoring is in large part driven by this capacity to tailor instruction to an individual student's needs (Anderson, Corbett, Koedinger, & Pelletier, 1995; Bloom, 1984; Merrill, Reiser, Ranney, & Trafton, 1992). Instructors who, given practical constraints, cannot provide different levels of challenge for individual students will be glad to know that research has also shown benefits from adjusting the difficulty of a practice task to fit students' needs at the group level. In one study, Clarke, Ayres, and Sweller (2005) designed an instructional unit to teach students mathematical concepts and procedures through the use of a spreadsheet application. Instruction was either sequential (focused on learning spreadsheet skills first and then using those skills to learn the mathematics) or concurrent (learning and using these skills simultaneously). They found that, for students with little prior knowledge of spreadsheets, the concurrent learning condition was too demanding; these students showed better mathematics learning and performance in the sequential condition, where the tasks were presented in isolation, making the challenge level more reasonable. Correspondingly, the opposite pattern held for more knowledgeable students. These results reinforce the idea that when novices are given too great a challenge, learning is hampered. This was probably part of the problem faced by Professor Strait's students, who were asked to take on challenges they had not practiced before (doing research in medical anthropology, constructing an argument of their own, and creating an engaging presentation).

Given a particular instructional activity, then, how can one effectively adjust it to target the appropriate level of challenge for different students, particularly those students who might not be quite ready to take on the activity in its full form? Research has shown that adding structure and support—also called *instructional scaffolding*—to a practice activity in or out of class promotes learning when it helps students practice the target skills at an appropriate level of challenge. This relates to Vygotsky's Zone of Proximal Development, which defines the optimal level of challenge for a student's learning in terms of a task that the student cannot perform successfully on his or her own but could perform successfully with some help from another person or group. A research study by Palincsar and Brown (1984) shows the success of this approach in helping students who were learning to read texts actively rather than passively. In particular, the researchers developed a protocol for pairs of students to follow in which students switched back and forth between the role of teacher and student, with the "teacher" asking the "student" a set of questions designed to exercise four strategic subskills of active reading—questioning, clarifying, summarizing, and predicting. These researchers found that when active reading skills were explicitly supported in this way, students' overall comprehension and retention improved markedly.

Research also indicates that instructional support does not need to come directly from another person to be helpful. For instance, Bereiter and Scardamalia developed a set of written prompts to help writing students target their efforts on two oft-neglected stages of the writing process: planning and revision. Because students did not naturally engage in these two stages on their own, following the prompts shifted their attention and effort toward (a) generating, refining, and elaborating their ideas and (b) evaluating their own writing, diagnosing problems, and deciding on revisions. As a result, students' writing process and

product showed significant improvements, including a ten-fold increase in the frequency of idea-level revisions (Bereiter & Scardamalia, 1987). This set of research results suggests that if Professor Strait had employed various kinds of instructional scaffolds to support her students in completing their final project presentations, they probably would have spent their practice time more effectively, learned more from it, and lived up to her expectations on the final project presentation.

Another advantage of finding an appropriate challenge level for students' practice is that it can help students remain motivated to sustain their efforts (see Chapter Three). For example, if a challenge is too great, learners may have a negative expectation for success and hence become disengaged and apathetic. In contrast, if students feel that the challenge is reasonable, they will likely hold a positive expectation for success that will increase their tendency to persevere and work hard for the goal. Finally, engaging in a task that is at the right level of challenge for a person's knowledge and skills is one of the key predictors of *flow*—the state of consciousness in which a person is totally engaged in and experiencing deep enjoyment of a particular task (Csikszentmihalyi, 1991).

Accumulating Practice

In addition to identifying the two features that make practice most productive—goal-directed and appropriately challenging—research in this area also reiterates the importance of *time on task*. In other words, even if students have engaged in high-quality practice, they still need a sufficient *quantity* of practice for the benefits to accumulate (Healy, Clawson, & McNamara, 1993; Martin, Klein, & Sullivan, 2007). The idea that the benefits of practice accumulate only gradually may seem obvious, but the practical constraints of time and resources often lead faculty

to move from concept to concept or skill to skill rather quickly, giving students no more than a single opportunity to practice each. For example, Professor Cox is giving his students exposure to multiple genres, but this comes at the expense of giving students only a single opportunity to develop their skills at writing in each of the genres he has assigned. If his goal is to simply expose students to the three different genres, without expecting them to gain proficiency in any of them, then the design of his activities is appropriate. But if his goal is for students—by the end of the course—to be able to write in each of the three genres at a professional level, then they would need more time on task.

Generally speaking, both professors and students underestimate the need for practice. Students often assume that when they can perform a task on one occasion in one context, their knowledge is secure when, in fact, it is much more difficult than that (see Chapter Four). It takes much more than one trial to learn something new, especially if the goal is for that new knowledge to be retained across time and transferred to new contexts.

Although it is true that the benefits of practice accrue gradually, it is important to note that the knowledge or skill gained by a given amount of additional practice often depends on where the student is in his or her learning process. As Figure 5.2 indicates, the early and late phases of learning tend to show relatively little effect of practice relative to the middle phase. These flatter portions at both ends of the curve tend to occur for two reasons.

The first reason is that the measures students often use to monitor their learning, such as accuracy, tend to be less sensitive at the extremes. So even though learning may be occurring, students do not see evidence of the change and hence feel like they are at a plateau. For example, consider a student who has just started learning to play the violin. Even though this student may be improving in several ways (better recall of the finger positions for different notes, increased accuracy in placement of the bow),

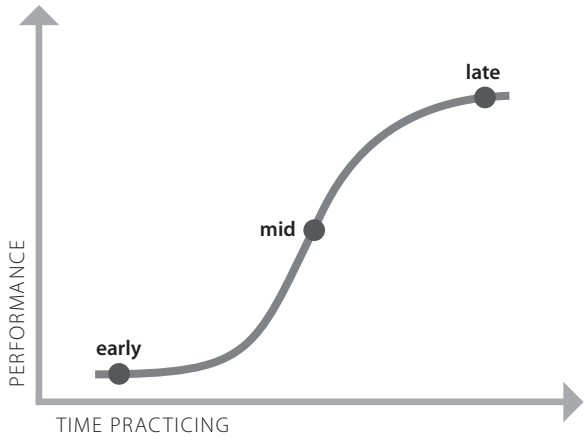

Figure 5.2. Unequal Effects of Practice on Performance

the sound produced may be so poor that improvements are hard to detect. Or imagine a student learning to program in a new computer language. Early on, the student may be making so many errors in programming syntax that it is hard to discern that he or she is formulating increasingly better algorithms. A similar lack of sensitivity to changing performance tends to occur on the upper end of learning because in this later phase students have managed to refine their performance to such a degree that they do not perceive changes, or the changes may occur in aspects of performance to which they are not attending. For example, advanced students may not recognize that they have actually improved in their ability to complete tasks more quickly and with less effort than they could before, or they may not realize that they are now able to reflect on their own processes *while* they complete complex tasks. Thus, because of this phenomenon at the early and late phases of learning, it is all the more important for instructors to highlight for students how their performance is changing or to provide more refined goals and criteria so that students can discern that they are improving.

The second reason that the learning curve in Figure 5.2 tends to be flatter at both ends is that the tasks we naturally assign for practice tend to pose too great a challenge for beginning students and too little challenge for accomplished students. As discussed earlier, when students engage in practice that is either too challenging or not challenging enough, their learning is hampered. This reason offers additional support for the notion of setting an appropriate level of challenge for students.

In contrast to the early and late phases of learning, the middle part of the curve in Figure 5.2 is steep, which indicates students are able to see large improvements in performance with additional practice. This is because students in this phase have a foundation of knowledge and skills upon which to build and because they are more likely to be able to detect improvements in their performance. This may also explain why students sometimes appear to "take off" in their development of knowledge and skill only after they have achieved a certain amount of learning.

Implications of This Research

Overall, the implications of the body of research on practice are that to achieve the most effective learning, students need *sufficient* practice that is *focused* on a specific goal or set of goals and is at an *appropriate level of challenge.* Given the constraints of time and resources that we must face, however, it is often difficult or impossible to increase students' practice time (either in or out of the class). Instead, the results in this chapter highlight the benefits of using a given amount of practice time more efficiently by focusing students' efforts on what they need to learn (rather than what they already know or may be more comfortable doing) and setting their goals for performance at a reasonable and productive level of challenge.

WHAT DOES THE RESEARCH TELL US ABOUT FEEDBACK?

Goal-directed practice alone is insufficient to foster students' learning. Goal-directed practice must be coordinated with targeted feedback in order to promote the greatest learning gains. The purpose of feedback is to help learners achieve a desired level of performance. Just as a map provides key information about a traveler's current position to help him or her find an efficient route to a destination, effective feedback provides information about a learner's current state of knowledge and performance that can guide him or her in working toward the learning goal. In other words, effective feedback can tell students *what* they are or are not understanding, *where* their performance is going well or poorly, and *how* they should direct their subsequent efforts.

Taking this map analogy a step further, imagine trying to find your way through a maze without any guiding information as to where you are relative to the entrance or exit; you could wander in circles without even realizing it, waste time, and become confused—even if you ultimately do find your way out of the maze. This situation is akin to the position that students are in without effective feedback. It is not surprising, then, that effective feedback can greatly facilitate students' learning. For example, consider two students who have the same misconception that leads them to solve several problems incorrectly. Suppose, however, that these two students receive feedback on their work at different times and with different content. One student solves all of these problems in a single, large homework assignment and, after submitting the assignment, gets it back a week later with the letter grade "C." He notices from the points marked off that he failed to get full credit for even a single problem, so he infers that he is totally lost on this topic. Suppose the other student is in a course

where the instructor includes a bit of problem-solving practice in each class session and then highlights some natural mistakes and how to remediate them after students have a chance to try a couple of problems. This student rather quickly gets some input from the instructor, indicating that in two of the practice problems he was making the same error. Once this is identified, the student is able to correct his understanding and then go on to solve that week's homework problems with this in mind.

Note that based on the different timing and content of the feedback, these two students may take very different paths from this point onward in the course. The first student, not realizing that it was only a single misconception that led to his level of performance, may believe he is unable to learn the current topic and hence skip any opportunities for further practice (for example, not bothering to study for the upcoming exam). The second student, armed with information about where he went wrong, can work on additional problems to strengthen his new understanding of this tricky issue. In other words, feedback at the right time and of the right nature can promote students' learning not only in the present but also in the future.

Consistent with this example, research points to two features of feedback that make students' learning more effective and efficient: content and timing. First, feedback should communicate to students where they are relative to the stated goals and what they need to do to improve. Second, feedback should provide this information when students can make the most use of it, based on the learning goals and structure of activities you have set for them. Like so many aspects of teaching and learning, there is no single approach to feedback that will work across the variety of situations students and instructors encounter. Rather, the content and timing of feedback need to be considered in terms of the learning goals we have for our students, students' incoming level of knowledge and proficiency, and the practical constraints of the course.

Research on what tends to make the content and timing of feedback most effective is discussed in the following two sections.

Communicating Progress and Directing Subsequent Effort

Feedback is most effective when it explicitly communicates to students about some specific aspects of their performance relative to specific target criteria, and when it provides information that helps students progress toward meeting those criteria. This kind of feedback, which informs students' subsequent learning, is often called *formative* feedback. In contrast, *summative* feedback is that which gives a final judgment or evaluation of proficiency, such as grades or scores.

Extending our earlier analogy between using a map to navigate and receiving feedback to learn, consider a more sophisticated navigational aid such as a global positioning system (GPS). A GPS has the capability of tracking a traveler's current position relative to a destination. To be helpful, a GPS needs to communicate more than the fact that the traveler is far away from the destination; ideally, it needs to identify how far the traveler is from the destination and provide directions to help the traveler reach it. Similarly, effective feedback needs to do more than simply tell a student that he or she is wrong; effective feedback involves giving students a clear picture of how their current knowledge or performance differs from the goal and providing information on adjustments that can help students adjust to reach the goal.

Research has long shown that feedback is more effective when it identifies particular aspects of students' performance they need to improve rather than providing a generic evaluation of performance, such as a grade or abstract praise or discouragement (Black & William, 1998; Cardelle & Corno, 1981). As illustrated by the example of the student who received a "C" with no

comments on his homework, giving only a letter grade or numerical score tends not to be effective feedback. Although grades and scores provide some information on the *degree to which* students' performance has met the criteria, they do not explain *which aspects* did or did not meet the criteria and *how*. Moreover, feedback that is specific to the processes students are engaging in (for example, helping students to properly approach a problem or to detect their own errors; see Chapter Seven) has been associated with deeper learning (Balzer et al., 1989). In one study, students were learning to solve geometry problems on the computer, with feedback automatically provided whenever the computer detected an error in students' solutions. One group of students received generic messages indicating that they had made an error, and another group received specific information about their errors and how to remediate them. The group with the more targeted feedback significantly outperformed the generic feedback group on a post-test assessing problem-solving skills (McKendree, 1990).

At the other extreme, simply giving students lots of feedback about their performance is also not necessarily an example of effective feedback. This is because too much feedback tends to overwhelm students and fails to communicate which aspects of their performance deviate most from the goal and where they should focus their future efforts. For example, research has shown that too many comments in the form of margin notes on student writing are often counterproductive because students are either overwhelmed by the number of items to consider or because they focus their revision on a subset of the comments that involve detailed, easy-to-fix elements rather than more important conceptual or structure changes (Lamburg, 1980; Shuman, 1979).

Remember Professor Cox's lament of spending so much time making comments on his students' papers but seeing no improve-

ment in later assignments? Providing too much information in his comments may have been part of the reason. In his case, giving fewer comments that addressed one or two top-priority issues probably would have provided his students with more targeted feedback. However, it is important to note that even if Professor Cox had given this kind of targeted feedback, it might not have been fully effective unless his students also had an opportunity to use the feedback in a rewrite or related assignment. The key idea here is that *targeted* feedback gives students prioritized information about how their performance does or does not meet the criteria so they can understand how to improve their future performance.

Indeed, the full benefits of feedback can only be realized when the feedback adequately directs students' subsequent practice *and* when students have the capacity to incorporate that feedback into further practice. Recall that in Professor Cox's course, students had only one opportunity to practice writing in each of the three genres he assigned. Although he may have conceived of this as repeated practice at the general skill of writing, these three assignments probably required rather different subsets of skills (see Chapter Four). So, even if Professor Cox had provided targeted feedback on the first assignment, students might not have benefited much from it unless they had an opportunity to carry it into the next assignment.

How could Professor Cox use feedback in a way that ties in with students' opportunities for further practice? One option is that he could have included more repetition of assignments within the same genre and then asked students to incorporate his feedback into subsequent assignments. Alternatively, he could have asked students to submit a rough draft of each assignment, made targeted comments on those drafts, and then explicitly articulated that the final draft's goal was to address his comments in the

revision. This scenario highlights the interaction between feedback and practice. Indeed, one can conceive of the practice that follows targeted feedback as a particularly tailored form of goal-directed practice.

Timing Feedback Appropriately

Whereas the research just discussed involves the content of feedback, it is also important to consider the appropriate timing of feedback. This involves both *how soon* feedback is given (typically, earlier is better) as well as *how often* (typically, more frequently is better). The ideal timing of feedback, however, cannot be determined by any general rule. Rather, it is best decided in terms of what would best support the goals you have set for students' learning. For example, going back to our GPS analogy, it is clear that one of the key features of these devices is that they give feedback *when the driver needs it* to support the goal of reaching a particular destination as quickly as possible.

Generally, more frequent feedback leads to more efficient learning because it helps students stay on track and address their errors before they become entrenched. Ample research supports this conclusion (see Hattie & Timperley, 2007, for a review). However, given practical constraints, this is often difficult. Fortunately, research shows that even minimal feedback on students' writing can lead to better second drafts because the feedback gives students a better sense of what their readers do and do not understand (Traxler & Gernsbacher, 1992). This result highlights that giving even a modest amount of feedback, especially when it is given early, can be helpful. This result also suggests that if Professor Strait had established milestones early in her students' project work, it could have enabled her to offer feedback earlier on in the process, before her students went off track.

This research does not mean, however, that greater frequency of feedback is always better. Again, *timeliness* of the feedback is a significant factor. For example, consider a study in which college students were learning to write mathematical functions in a spreadsheet application (Mathan & Koedinger, 2005). The particular goal for students' learning in this situation was not only that they be able to write these functions accurately but also that they be able to recognize and fix their own errors. Students who received feedback immediately after they made a mistake scored lower on final assessments compared to students who received "delayed" feedback. Although surprising at first, this result makes sense when one realizes that the immediate feedback group was missing the opportunity to practice recognizing and repairing their own errors. In contrast, the students receiving delayed feedback had a chance to fix their own errors so they had more practice at the corresponding skills. That is, when the delayed feedback group made errors, feedback was given only when they (a) showed sufficient signs of not having recognized their error or (b) made multiple failed attempts at fixing their error. In this way, one could argue that even though it was not immediate, their feedback was given in a more timely manner relative to the learning goals at hand.

Implications of This Research

There are three key implications of this research on what makes feedback more effective. The feedback must (1) focus students on the key knowledge and skills you want them to learn, (2) be provided at a time and frequency when students will be most likely to use it, and (3) be linked to additional practice opportunities for students. As we saw in some of the sections above, each of these aspects of feedback must align with the goals you have set for students' learning. It is best to find a type and frequency of

feedback that allows students to reap the benefits of feedback while staying actively engaged in monitoring their own learning—in other words, feedback that does not undermine students' progress in becoming independent, self-regulated learners. Giving too little detail in feedback can leave students unclear on what they need to do to improve, whereas giving too much detail can overwhelm them or mislead them as to what aspects are higher priority. Similarly, giving feedback too infrequently can leave students floundering without enough information to direct their learning, whereas giving feedback too frequently can potentially irritate students or lead them to depend on the feedback rather than on themselves.

In addition to balancing the amount and timing of feedback to make it most effective, it is often necessary to pay attention to the practical aspects of giving feedback. For example, the instructor's time in composing or tailoring feedback and the students' time in processing and responding to feedback is a key consideration in guiding how and when to give feedback. We must always consider both the pedagogical and practical consequences of feedback. Also, all feedback need not be tailored to individual students, and it need not all come from the instructor. We discuss a variety of strategies for feedback that are effective and feasible, including peer response, group feedback, and more.

WHAT STRATEGIES DOES THE RESEARCH SUGGEST?

Here we present strategies that can help you provide students with (1) goal-directed practice and (2) targeted feedback. In both cases, the focus is on how to do so in effective and efficient ways.