GLOBAL EDITION

Applied Behavior Analysis

THIRD EDITION

John O. Cooper • Timothy E. Heron • William L. Heward

Applied Behavior Analysis

Third Edition

Global Edition

John O. Cooper

Timothy E. Heron

William L. Heward

All, The Ohio State University

Pearson

Harlow, England • London • New York • Boston • San Francisco • Toronto • Sydney • Dubai • Singapore • Hong Kong
Tokyo • Seoul • Taipei • New Delhi • Cape Town • São Paulo • Mexico City • Madrid • Amsterdam • Munich • Paris • Milan

PART 4

Reinforcement

The three chapters in Part 4 are devoted to reinforcement, the most important and widely applied principle of behavior analysis. Reinforcement, a deceptively simple behavior–consequence relation, is the fundamental building block of operant behavior. Chapter 11 presents the operation and defining effect of reinforcement, explains how reinforcers are classified and identified, describes briefly how antecedent stimulus conditions modulate the effects of reinforcement, discusses factors that influence the effectiveness of reinforcement, outlines experimental control techniques for verifying whether a positive reinforcement contingency is responsible for increased responding, and offers guidelines for using reinforcement effectively.

In Chapter 12, Richard Smith and Brian Iwata describe one of the most consistently misunderstood principles of behavior. Smith and Iwata define negative reinforcement, an operant contingency in which responding increases as a result of the termination, reduction, or postponement of a stimulus as a behavioral consequence; compare and contrast it with positive reinforcement and punishment; distinguish between escape and avoidance contingencies; describe events that may serve as negative reinforcers; illustrate how the application of negative reinforcement can be used to strengthen desired behavior; and discuss ethical issues that arise when using negative reinforcement.

One of Skinner's most important discoveries was that reinforcement does not have to follow every response. In fact, under many intermittent schedules of reinforcement—in which reinforcement follows some, but not all, occurrences of the target behavior—responding occurs at higher, more consistent, rates than under a continuous schedule of reinforcement, in which each response is reinforced. Chapter 13 describes how intermittent reinforcement can be scheduled based on various combinations of response and/or temporal requirements and identifies characteristic patterns of responding associated with each schedule. Practitioners who understand how different schedules of reinforcement influence behavior can program reinforcement for more effective and efficient acquisition of new skills and for improved performance and endurance of established skills.

CHAPTER 11

Positive Reinforcement

LEARNING OBJECTIVES

- Define and provide examples of positive reinforcement.
- Define and provide examples of conditioned and unconditioned reinforcement.
- Describe and provide examples of the operant conditioning paradigm (i.e., the three-term and four-term contingencies).
- Identify potential reinforcers.
- Use appropriate parameters and schedules of reinforcement to identify reinforcers.
- Use response-deprivation procedures (e.g., the Premack principle).
- Identify control procedures for positive reinforcement.
- Use positive reinforcement effectively.

In looking back, it seems to me that the most important thing I learned in graduate school was from another student, Burrhus Frederic Skinner (I called him Burrhus, others called him Fred). This man had a box, within which was a smaller box, within which he would place a hungry laboratory rat. When the animal, in its explorations, would depress a lever that projected from one wall, a pellet of food would be discharged into a tray beneath the lever. Under such conditions, the rat would learn, in a matter of minutes, sometimes seconds, how to get its meal by depression of the lever. It would even keep on pressing, sometimes at a rapid rate, when pellets were delivered only now and then; and if the food supply was cut off entirely, the animal would still keep working for awhile.

—Fred Keller (1982, p. 7)

Positive reinforcement is the most important and most widely applied principle of behavior. Although some people continue to believe that the results of laboratory research with animals is not applicable to human learning, by the mid-1960s researchers established the significant role of positive reinforcement in education and treatment. "It is safe to say that without Skinner's detailed laboratory analyses of reinforcement (Skinner, 1938), there would be no field of 'applied behavior analysis' today, least not as we know it" (Vollmer & Hackenberg, 2001, p. 241).

Fittingly, the lead article in the first issue of the *Journal of Applied Behavior Analysis* reported an experiment showing the effects of positive reinforcement on student behavior (Hall, Lund, & Jackson, 1968). Six elementary students who were disruptive or dawdled frequently participated in this classic study. The dependent variable, study behavior, was defined individually for each student, depending on the subject matter being taught, but it generally consisted of the student being seated and oriented toward the appropriate object or person (e.g., looking at course materials or the lecturing teacher) and class participation (e.g., writing the assignment, answering the teacher's question). The independent variable was teacher attention, cued by an observer who held up a small square of colored paper not likely to be noticed by the target student. On this signal, the teacher attended to the child by moving to his desk and making a verbal comment or giving him a pat on the shoulder.

The effects of contingent teacher attention on the behavior of all six students were striking. Figure 11.1 shows the results for Robbie, a third-grader who was "a particularly disruptive student who studied very little" (p. 3). During baseline, Robbie engaged in study behavior for an average of 25% of the observed intervals. The remainder of the time he snapped rubber bands, played with objects in his pocket, talked and laughed with classmates, and played with an empty milk carton from his earlier-served drink. The majority of attention Robbie received followed nonstudy, disruptive behaviors.

Following baseline, the experimenters showed the teacher a graph of Robbie's study behavior, presented the results of previous studies in which contingent adult attention had improved child behavior, and discussed the fundamentals of providing contingent social reinforcement.

The results of contingent reinforcement were as follows: During Reinforcement 1, Robbie's study behavior increased to a mean of 71%. When a reversal to baseline conditions was introduced, his study behavior decreased to a mean of 50%; but when Robbie's teacher again provided contingent attention for study behavior (Reinforcement 2), his study behavior recovered and stabilized at a level ranging between 70% and 80% of the observed intervals. Followup observations over a 14-week period showed that Robbie's study behavior had maintained at 79%. The teacher reported positive behavior changes associated with Robbie's increased study behavior. By the final week of

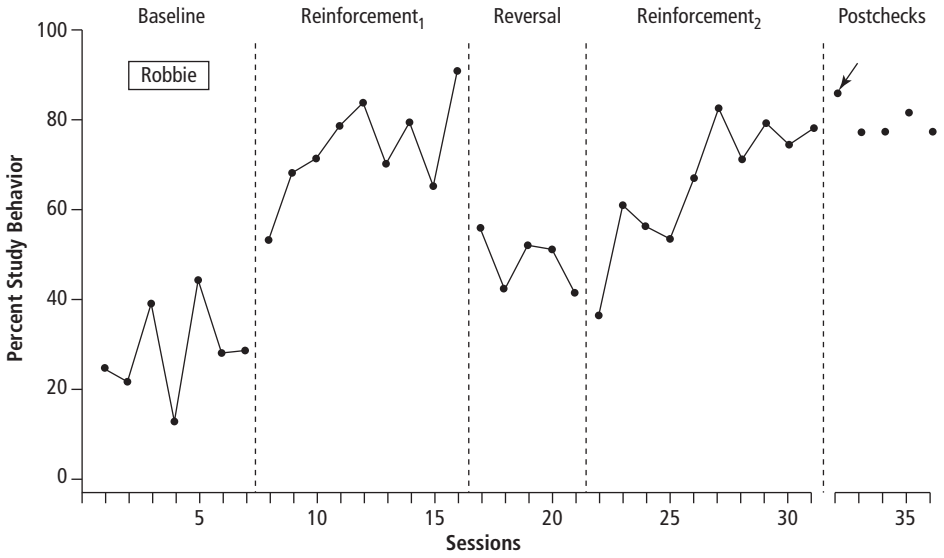

Figure 11.1 Percentage of intervals of study behavior by a third-grade student during baseline and reinforcement conditions. The arrow to the first postcheck data point shows when cueing the teacher to provide attention was discontinued.

Based on "Effects of Teacher Attention on Study Behavior" by R. V. Hall, D. Lund, and D. Jackson, 1968, *Journal of Applied Behavior Analysis*, *1*, p. 3. Copyright 1968 by the Society for the Experimental Analysis of Behavior, Inc.

Reinforcement 2, Robbie completed his spelling assignments more consistently, his disruptive behavior had diminished, and he continued to study while drinking his milk and did not play with the carton afterwards. Robbie's intervention was based on the principle of positive reinforcement.

This chapter examines the definition and nature of positive reinforcement, describes methods for identifying potential reinforcers and assessing their effects, outlines experimental control techniques for verifying whether a positive reinforcement contingency is responsible for increased responding, and offers guidelines for using positive reinforcement effectively.

POSITIVE REINFORCEMENT DEFINED

The principle of reinforcement is deceptively simple. "The basic operant functional relation for reinforcement is the following: When a type of behavior (R) is followed by reinforcement (S^R) there will be an increased future frequency of that type of behavior" (Michael, 2004, p. 30).[1] However, as Michael pointed out, three qualifications must be considered regarding the conditions under which the effects of reinforcement will occur: (a) the timing between the end of a given response and the onset of the stimulus change (i.e., the presentation of the reinforcer), (b) the relationship between stimulus conditions present when the response was emitted, and (c) the role of motivation. In this section we examine these qualifications and several other concepts requisite to acquiring a full understanding of reinforcement.

Operation and Defining Effect of Positive Reinforcement

Positive reinforcement occurs when a response is followed immediately by the presentation of a stimulus change that increases the future occurrence of similar responses. Figure 11.2 illustrates the two-term contingency—a response followed closely in time by the presentation of a stimulus—and the increased future responding that defines positive reinforcement. This two-term contingency is the fundamental building block for the selection of all operant behavior.

The stimulus that is presented as a consequence, and that is responsible for the subsequent increase in responding, is called a **positive reinforcer**, or, more simply, a reinforcer. Teacher attention in the form of positive praise was the reinforcer that increased Robbie's study behavior. Cold water flowing into a cup and the sight of a colorful bird are the reinforcers for the two behaviors shown in Figure 11.2.

It is important to remember that a reinforcer does not (and cannot) affect the response that it follows. When making this point, Skinner (1953) reminded us that reinforcement affects the operant response class.

> It is not correct to say the operant reinforcement "strengthens the response which precedes it." The response has already occurred and cannot be changed. What is changed is the future probability of responses in the same class. It is the operant as a class of behavior, rather than the response as a particular instance, which is conditioned. (p. 87)

Skinner (1938, 1966) used rate of responding as the fundamental datum for his research on reinforcement. To strengthen an operant is to make it occur more often.[2] However, rate is not the only dimension of behavior selected, shaped, and maintained by reinforcement. Reinforcement can also alter the duration, latency, interresponse time, magnitude, and/or topography of behavior. For example, if reinforcement follows only responses that fall within a range of magnitude—that is, above a minimum force but below a maximum force—responses within that range will occur more frequently. Reinforcement contingent on responses meeting multiple criteria will alter the response class meeting those criteria (e.g., responses by a golfer practicing 10-foot putts must fall within a narrow range of force and form to be successful).

Immediacy of Reinforcement

The direct effects of reinforcement involve "temporal relations between behavior and its consequences that are on the order of a few seconds" (Michael, 2004, p. 161). Although research with nonhuman animals suggests that as much as 30 seconds can elapse without critical loss of effect (e.g., Bradley & Poling,

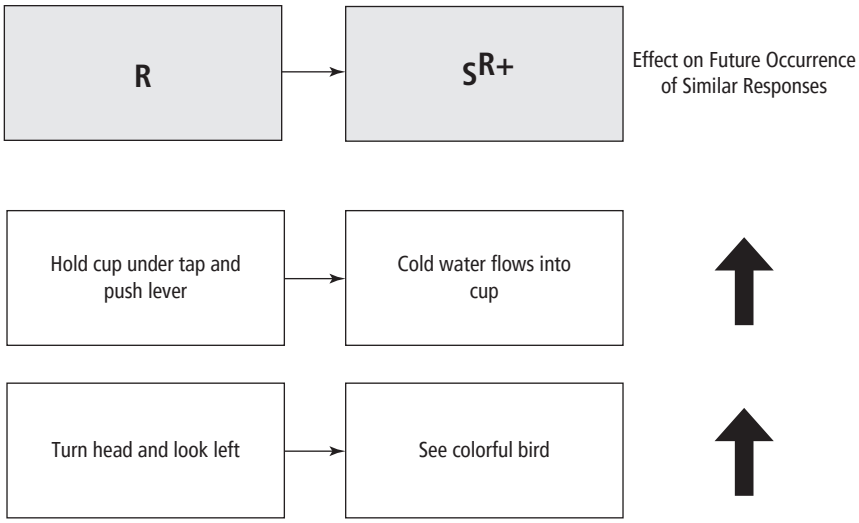

Figure 11.2 Two-term contingency illustrating positive reinforcement: A response (R) is followed closely in time by a stimulus change (S^{R+}) that results in an increased occurrence of similar responses in the future. Used with the permission of Skinner, B. F. (1953). Science and human behavior. New York: Macmillan.

2010; Byrne, LeSage, & Poling, 1997; Critchfield & Lattal, 1993; Lattal & Gleason, 1990; Wilkenfeld, Nickel, Blakely, & Poling, 1992), a response-to-reinforcement delay of as little as 1 second will be less effective than a reinforcer delivered immediately. This is because behaviors other than the target behavior can occur during the delay; the behavior temporally closest to the presentation of the reinforcer will be strengthened by its presentation. As Sidman (1960) described, "If the reinforcer does not immediately follow the response that was required for its production, then it will follow some other behavior. Its major effect will then be upon the behavior that bears, adventitiously to be sure, the closest prior temporal relationship to the reinforcement" (p. 371).

Malott and Shane (2014) discussed the importance of a brief response-to-reinforcement delay, as follows:

> [I]f you're trying to reinforce a response, don't push that 60-second limit, push the other end—the 0-second end. The direct effect of reinforcement drops off quickly as you increase the delay, even to 3 or 4 seconds. And even a 1-second delay may reinforce the wrong behavior. If you ask a young child to look at you and deliver the reinforcer 1 second after the response, you're liable to reinforce looking in the wrong direction. So, one problem with delayed reinforcement is that it reinforces the wrong response—the one that occurred just before the delivery of the reinforcer. (p. 4)

A common misconception is that delayed consequences can reinforce behavior, even when the consequences occur days, weeks, or even years after the responses occurred. "When human behavior is apparently affected by long-delayed consequences, the change is accomplished by virtue of the human's complex social and verbal history, and should not be thought of as an instance of the simple strengthening of behavior by reinforcement" (Michael, 2004, p. 36).

For example, suppose that a piano student practiced dutifully every day for several months in preparation for a statewide competition, at which she received a first-place award for her solo piano performance. Although some might believe that the award reinforced her persistent daily practice, they would be mistaken. Delayed consequences do not reinforce behavior directly. Delayed consequences can, when combined with language, *influence* future behavior through instructional control and rule following. A *rule* is a verbal description of a behavioral contingency (e.g., "Turnip seeds planted by August 15 will yield a crop before a killing freeze."). Learning to follow rules is one way that a person's behavior can come under the control of consequences that are too delayed to influence behavior directly. A statement by the piano teacher such as "If you practice your assignments every day for one hour between now and the competition, you could win first place" could have functioned as a rule that influenced the piano student's daily practice. The student's daily practice was evidence of rule-governed behavior if daily practice occurred because of her teacher's rule.[3] In other words, **rule-governed behavior** is behavior controlled by a rule (i.e., a verbal statement of an antecedent-behavior-consequence contingency) that enables human behavior to come under the indirect control of temporarily remote or improbable but potentially significant consequences. The following conditions provide strong indicators that behavior is the result of instructional control or is rule-governed behavior rather than a direct effect of reinforcement (Malott, 1988; Michael, 2004).

- No immediate consequence for the behavior is apparent.
- The response–consequence delay is greater than 30 seconds.
- Behavior changes without reinforcement.
- A large increase in the occurrence of the behavior occurs following one instance of reinforcement.
- No consequence for the behavior exists, including no automatic reinforcement, but the rule exists.

Reinforcement Is Not a Circular Concept

Another common misconception is that reinforcement is the product of circular reasoning and therefore contributes nothing to our understanding of behavior. Not so. Circular reasoning is a form of faulty logic in which the name used to describe an observed effect is mistaken as the cause for the effect. This confusion of cause and effect is circular because the observed effect is the sole basis for identifying the presumed cause.

In circular reasoning, the suspected cause is not independent of its effect—they are one and the same.

Here is an example of circular reasoning that occurs often in education. A student's persistent difficulties learning to read (effect) leads to a formal diagnosis of a learning disability, which is then offered as an explanation for the student's struggles with reading: "Yoshi's reading problem is due to his learning disability." How do you know Yoshi has a learning disability? Because he hasn't learned to read. Why hasn't Yoshi learned to read? Because his learning disability has prevented him from learning to read. And around and around it goes.

Similarly, it would be circular reasoning if we said that teacher attention increased Robbie's study behavior *because* it is a reinforcer. Instead, it is correct to say that because Robbie's study behavior increased when (and only when) it was followed immediately by teacher attention, teacher attention is a reinforcer. The difference is more than the direction of the relation, or some semantic sleight-of-hand. In circular reasoning, the suspected cause is not manipulated as an independent variable to see whether it affects the behavior. In circular reasoning, such experimental manipulation is impossible because the cause and effect are the same. Yoshi's learning disability cannot be manipulated as an independent variable because, as we used the concept in this example, it is nothing more than another name for the dependent variable (effect).

Reinforcement is not a circular concept because the two components of the response–consequence relation can be separated, allowing the delivery of a consequence to be manipulated to determine whether it increases the occurrence of the behavior it follows. Epstein (1982) described it as follows:

> If we can show that a response increases in frequency because (and only because) it is followed by a particular stimulus, we call that stimulus a *reinforcer* and its presentation, *reinforcement*. Note the lack of circularity. *Reinforcement* is a term we invoke when we observe certain relations between events in the world. . . . [However,] if we say, for example, that a particular stimulus strengthens a response behavior *because* it is a reinforcer, we are using the term *reinforcer* in a circular fashion. It is *because* it strengthens behavior that we call the stimulus a *reinforcer*. (p. 4)

Epstein (1982) went on to explain the difference between using an empirically demonstrated principle such as reinforcement in a theoretical account of behavior and using a circular argument.

> In some of his writings, Skinner speculates that certain behavior (for example, verbal behavior) has come about through reinforcement. He may suggest, for example, that certain behavior is strong *because* it was reinforced. This use of the concept is not circular, only speculative or interpretive. Using the language of reinforcement in this way is reasonable when you have accumulated a large data base. . . . When Skinner attributes some everyday behavior to past reinforcers, he is making a plausible guess based on a large data base and principles of behavior established under controlled conditions. (p. 4)

Used properly, *reinforcement* describes an empirically demonstrated (or speculative, in a theoretical or conceptual analysis) functional relation between a stimulus change (consequence) immediately following a response and an increase in the future occurrence of similar responses. Table 11.1 shows restrictions and examples of appropriate use of the terms *reinforcer, reinforcing, reinforcement*, and *to reinforce*, as suggested by Catania (2013). Box 11.1 describes four mistakes commonly made when speaking and writing about reinforcement.

Relationship Between Reinforcement and Antecedent Stimulus Conditions

Reinforcement does more than increase the future occurrence of behavior; it also changes the function of stimuli that immediately precede the reinforced behavior. By virtue of being

TABLE 11.1 The Vocabulary of Reinforcement*

Term	Restrictions	Examples
reinforcer (noun)	A stimulus	Food pellets were used as reinforcers for the rat's lever presses.
reinforcing (adjective)	A property of a stimulus	The reinforcing stimulus was produced more often than the other, nonreinforcing stimuli.
reinforcement (noun)	As an operation, the delivery of consequences when a response occurs	The fixed-ratio schedule of reinforcement delivered food after every 10th key peck.
	As a process, the increase in responding that results from the reinforcement	The experiment with monkeys demonstrated reinforcement produced by social consequences.
to reinforce (verb)	As an operation, to deliver consequences when a response occurs; responses are reinforced and not organisms	When a period of free play was used to reinforce the child's completion of school work, the child's grades improved.
	As a process, to increase responding through the reinforcement operation	The experiment was designed to find out whether gold stars would reinforce cooperative play among first-graders.

*This vocabulary is appropriate if and only if three conditions exist: (1) A response produces consequences; (2) that type of response occurs more often when it produces those consequences than when it does not produce them; and (3) the increased responding occurs *because* the response has those consequences. A parallel vocabulary is appropriate to punishment (including *punisher* as a stimulus and *punish* as a verb), with the difference being that a punishing consequence reduces the likelihood of future responding.

Based on *Learning, Interim* (5th ed.) by A. C. Catania, 2013, p. 66. Cornwall-on-Hudson, NY: Sloan Publishing.

BOX 11.1

Common Mistakes in Talking and Writing About Reinforcement

A standard set of technical terms is prerequisite to the meaningful description of any scientific activity. Effectively communicating the design, implementation, and outcome of an applied behavior analysis depends on the accurate use of the discipline's technical language. The language of reinforcement includes some of the most important elements of the behavior analyst's vocabulary.

In this box we identify four mistakes made frequently by students of applied behavior analysis when describing reinforcement-based interventions. Perhaps the most common mistake—confusing negative reinforcement with punishment—is not discussed here. That terminology error was introduced in Chapter 2 and receives additional attention in Chapter 12.

Reinforcing the Person

Although it is proper to speak of presenting a *reinforcer* to a learner (e.g., "The teacher delivered a token to Bobby each time he asked a question"), statements such as "The teacher reinforced Bobby when he asked a question" and "Chloe was reinforced with praise each time she spelled a word correctly" are incorrect. *Behaviors* are reinforced, not people. Bobby's teacher reinforced question asking, not Bobby. Of course, reinforcement acts on and affects the overall person, in that it strengthens behaviors within the person's repertoire. However, the procedural focus and the primary effect of reinforcement are on the behaviors that it follows.

Practice as Reinforcement for a Skill

Educators will sometimes say that students should practice a skill because "practicing reinforces the skill." The phrase poses no problem if the speaker is describing a common outcome of practice with the everyday language connotation of *reinforce*, as in "to make something stronger" (e.g., to reinforce concrete by embedding steel rods in it). Well-designed drill and practice on a skill usually yields stronger performance in the form of better retention, reduced latency, higher response rates, and/or increased endurance (e.g., Johnson & Layng, 1994; Swanson & Sachse-Lee, 2000). Unfortunately, a phrase such as "practicing reinforces the skill" is often misused and misinterpreted as technical usage of the language of operant conditioning.

Although a skill that has been practiced is often stronger as a result of the practice, the practice itself could not be a *reinforcer* for the behavior practiced. Practice refers to the form and manner in which the target skill is emitted (e.g., answering as many math problems as you can in 1 minute). Practicing is a behavior that could be reinforced with various consequences, such as a preferred activity (e.g., "Practice solving these math problems; then you can have 10 minutes of free time"). Depending on a learner's history and preferences, the opportunity to practice a certain skill may function as a reinforcer for practicing another skill (e.g., "Finish your math problems; then you'll get to do 10 minutes of repeated reading practice").

Artificial Reinforcement

A distinction between natural and artificial reinforcers is made sometimes, as in this statement: "As the students' success rates improved, we gradually stopped using artificial reinforcers, such as stickers and trinkets, and increased the use of natural reinforcers." Some authors have suggested that applications of the principles of behavior result in "artificial control" (e.g., Smith, 1992). A behavior–consequence contingency may be effective or ineffective as reinforcement, but none of its elements (the behavior, the consequence, or the resultant behavior change) is, or can be, artificial.

The reinforcement contingencies and stimuli used as reinforcers in any behavior change program are always contrived—otherwise, there would be no need for the program—but they are never artificial (Skinner, 1982). The meaningful distinction when talking about reinforcement contingencies is not between the natural and the artificial, but between contingencies that already exist in a given setting prior to a behavior change program and contingencies that are contrived as part of the program (Kimball & Heward, 1993). Although the ultimate effectiveness of a behavior change program may depend on shifting control from contrived to naturally occurring contingencies, there is no such thing as artificial reinforcement.

Reinforcement and Feedback as Synonyms

Some speakers and writers mistakenly use *reinforcement* and *feedback* interchangeably. The two terms refer to different operations and outcomes, though some of each term encompasses parts of the other term's meaning. *Feedback* is information a person receives about a particular aspect of his or her behavior following its completion (e.g., "Very good, Kathy. Two quarters equal 50 cents."). Feedback is most often provided in the form of verbal descriptions of performance, but it can also be provided by other means, such as vibration or lights (e.g., Greene, Bailey, & Barber, 1981). Because feedback is a consequence that often results in the increased future occurrence of behavior, it sometimes leads to the faulty assumption that reinforcement must involve feedback or that *reinforcement* is just a behaviorist's term for *feedback*.

Reinforcement always increases the future frequency of responding. Feedback may result in (a) an increase in the future occurrence of the student's performance as a reinforcement effect and/or as a prompt or instruction on how to respond next time (e.g. "Your handwriting is improving, Jason, but don't forget to cross your T's") and/or (b) a reduction in the occurrence of some aspect of the learner's performance as a function of punishment or instruction (e.g., "You dropped your elbow on that pitch. Don't do that."). Feedback may have multiple effects, increasing one aspect of performance and decreasing another. Feedback may also have no effect on future responding whatsoever.