

PAVLOV AND CLASSICAL CONDITIONING

Biographical Introduction

The father of modern learning theory is Ivan Petrovich Pavlov (1849-1936). Pavlov was born in Ryazan, Russia, the son of a poor village priest. Pavlov himself planned to become a priest until the age of 21, when he decided he was more interested in a scientific career. For many years he devoted his attention to physiological investigations, and in 1904 he won the Nobel Prize for his work on the digestive system. It was just a little before this time, when Pavlov was 50 years old, that he began his famous work on conditioned reflexes. This new interest came about through an accidental discovery about the nature of salivation in dogs. Ordinarily dogs salivate when food touches their tongues; this is an innate reflex. But Pavlov noticed that his dogs also salivated before the food was in their mouths; they salivated when they saw the food coming, or even when they heard approaching footsteps. What had happened was that the reflex had become conditioned to new, formerly neutral stimuli.

For a while Pavlov could not decide whether to pursue the implications of his new discovery or to continue with his earlier research. Finally, after a long struggle with himself, he began studying the

Learning Theory: Pavlov, Watson, and Skinner

conditioning process. Still, Pavlov believed that he was working as a physiologist, not a psychologist. In fact, Pavlov required that everyone in his laboratory use only physiological terms. If his assistants were caught using psychological language—referring, for example, to a dog's feelings or knowledge—they were fined (R. Watson, 1968, pp. 408-412).

Basic Concepts

The Classical Conditioning Paradigm. In a typical experiment (Pavlov, 1928, p. 104), a dog was placed in a restraining harness in a dark room and a light was turned on. After 30 seconds some food was placed in the dog's

mouth, eliciting the salivation reflex. This procedure was repeated several times—each time the presentation of food was paired with the light. After a while the light, which initially had no relationship to salivation, elicited the response by itself. The dog had been conditioned to respond to the light.

In Pavlov's terms (1927, lectures 2 and 3), the presentation of food was an unconditioned stimulus (US); Pavlov did not need to condition the animal to salivate to the food. The light, in contrast, was a conditioned stimulus (CS); its effect required conditioning. 1 Salivation to the food was called an unconditioned reflex (UR), and salivation to the light was called a conditioned reflex (CR). The process itself is called classical conditioning.

You might have noticed in this experiment that the CS appeared before the US; Pavlov turned on the light before he presented the food. One of the questions he asked was whether this is the best order for establishing conditioning. He and his students discovered that it is. It is very difficult to obtain conditioning when the CS follows the US (1927, pp. 27-28). Other studies have suggested that conditioning often occurs most rapidly when the CS is presented about one-half second prior to the US (see Schwartz, 1989, p. 83).

Pavlov discovered several other principles of conditioning, some of which we will briefly describe.

Extinction. A conditioned stimulus, once established, does not continue to work forever. Pavlov found that even though he could make a light a CS for salivation, if he flashed the light alone over several trials, it began to lose its effect. Drops of saliva became fewer and fewer until there were none at all. At this point, extinction had occurred (Pavlov, 1928, p. 297).

Pavlov also discovered that although a conditioned reflex appears to be extinguished, it usually shows some spontaneous recovery. In one experiment (Pavlov, 1927, p. 58), a dog was trained to salivate to the mere sight of food—the CS. (Previously, the dog would salivate only when food was in its mouth.) Next, the CS alone was presented at 3-minute intervals for six trials, and by

'Pavlov actually used the terms conditional and unconditional: they were translated conditioned and unconditioned, the terms psychologists now generally use.

Learning Theory: Pavlov, Watson, and Skinner

the sixth trial, the dog no longer salivated. The response appeared to have been extinguished. But, after a 2-hour break in the experiment, the presentation of the CS alone once again produced a moderate amount of salivation. Thus the response showed some spontaneous recovery. If one were to continue to extinguish the response, without periodically repairing the CS to the US, the spontaneous recovery effect would also disappear.

Stimulus Generalization. Although a reflex has been conditioned to only one stimulus, it is not just that particular stimulus that elicits it. The response seems to generalize over a range of similar stimuli without any further conditioning (Pavlov, 1928, p. 157). For example, a dog that has been conditioned to salivate to a bell of a certain tone will also salivate to bells of differing tones. The ability of the neighboring stimuli to produce the response varies with the degree of similarity to the original CS. Pavlov believed that we observe stimulus generalization because of an underlying physiological process he called irradiation. The initial stimulus excites a certain part of the brain that then irradiates, or spreads, over other regions of the cerebrum (p. 157).

Discrimination. Initial generalization gradually gives way to a process of differentiation. If one continues to ring bells of different tones (without presenting food), the dog begins to respond more selectively, restricting its responses to the tones that most closely resemble the original CS. One can also actively produce differentiation by pairing one tone with food while presenting another tone without food. This would be called an experiment in stimulus discrimination (Pavlov, 1927, pp. 118-130).

Higher-Order Conditioning. Pavlov showed, finally, that once he had solidly conditioned a dog to a CS, he could then use the CS alone to establish a connection to yet another neutral stimulus. In one experiment, Pavlov's student trained a dog to salivate to a bell and then paired the bell alone with a black square. After a number of trials, the black square alone produced salivation. This is called second-order conditioning. Pavlov found that in some cases he could also establish third-order conditioning, but he could not go beyond this point (p. 34).

Evaluation

In a sense, Pavlov's basic idea was not new. In the 17th century, Locke had proposed that knowledge is based on associations. Pavlov went beyond Locke,

however, and uncovered several principles of association through empirical experiments. He took the theory of learning out of the realm of pure speculation. Pavlov, as we shall see, did not discover everything there is to know about conditioning; in particular, his brand of conditioning seems restricted to a certain range of innate responses. Nevertheless, he was the first to put learning theory on a firm scientific footing.

186

Learning Theory: Pavlov, Watson, and Skinner

WATSON

Biographical Introduction

The man most responsible for making Pavlovian principles a part of the psychological mainstream was John B. Watson (1878-1958). Watson was born on a farm near Greenville, South Carolina. He said that in school "I was lazy, somewhat insubordinate, and so far as I know, I never made above a passing grade" (Watson, 1936, p. 271). Nevertheless, he went to college at Furman University and graduate school at the University of Chicago, where he began doing psychological research with animals. After earning his doctorate, he took a position at Johns Hopkins University in Baltimore, where he did his most productive work.

In 1913 Watson made a great impact on psychology by issuing a manifesto, "Psychology as the Behaviorist Views It." In this article he argued that the study of consciousness through introspection has no place in psychology as a science. Psychology should abandon "the terms consciousness, mental states, mind, content, introspectively verifiable, imagery and the like" (Watson, 1913, p. 166). Instead, its goal should be "the prediction and control of behavior" (p. 158). In particular, it should study only stimuli, responses, and the formation of habits. In this way psychology could become a science like the other natural sciences.

A year later he read the works of Pavlov and the Russians on conditioned reflexes and made Pavlovian conditioning the cornerstone of his thinking. Then, in 1916, Watson began research on young children, becoming the first major psychologist to apply principles of learning to the problems of development.

In 1929 Watson's academic career came to an abrupt end. His divorce from his wife became so widely and sensationally publicized that Johns Hopkins fired him. Watson remarried (Rosalie Raynor, a coworker) and entered the business world. In order to get a good sense of business, he worked for a while as a coffee salesman and a clerk at Macy's department store. He continued to

write, but now for magazines such as *Cosmopolitan*, *Harper's*, and *McCall's*, in which he advanced his ideas on child development.

Basic Concepts

Environmentalism. Watson was a behaviorist; he said we should study only overt behavior. He also was an environmentalist and made this famous proposal:

Give me a dozen healthy infants, well-formed, and my own specified world to bring them up in and I'll guarantee to take any one at random and train him to become any type of specialist I might select—doctor.

187

Learning Theory: Pavlov, Watson, and Skinner

lawyer, artist, merchant, chief, and yes, even begger-man and thief, regardless of his talents, penchants, tendencies, abilities, vocations, and race of his ancestors. (1924, p. 104)

In the next sentence Watson added that "I am going beyond my facts, and I admit it, but so have the advocates of the contrary and they have been doing it for many thousands of years" (p. 104).

Study of Emotions. One of Watson's major interests was the conditioning of emotions. He claimed that at birth there are only three unlearned emotional reactions—fear, rage, and love. Actually, all we observe are three different physical responses, but for the sake of simplicity we can call them emotions.

Fear, Watson said (1924, pp. 152-154), is observed when infants suddenly jump or start, breathe rapidly, clutch their hands, close their eyes, fall, and cry. There are only two unconditioned stimuli that elicit fear. One is a sudden noise; the other is loss of support (as when the baby's head is dropped). Yet older children are afraid of all kinds of things—strange people, rats, dogs, the dark, and so on. Therefore it must be that the stimuli evoking most fear reactions are learned. For example, a little boy is afraid of snakes because he was frightened by a loud scream when he saw one. The snake became a conditioned stimulus.

Rage is initially an unlearned response to the restriction of body movement. If we grab a 2-year-old girl, preventing her from going where she wants, she begins to scream and stiffens her body. She lies down stiff as a

rod in the middle of the street and yells until she becomes blue in the face (p. 154). Although rage is initially a reaction to one situation—being forcibly held—it later is expressed in a variety of situations; children become angry when told to wash their faces, sit on the toilet, get undressed, take a bath, and so on. Such commands elicit rage because they have been associated with physical restraint in these situations. The child becomes angry when told to get undressed because this order was initially associated with being forcibly held.

Love is initially a response that is automatically elicited by the stroking of the skin, tickling, gentle rocking, and patting. The baby responds by smiling, laughing, gurgling and cooing, and other responses that we call affectionate, good natured, and kindly. Although Watson had no use for Freud, he noted that such responses "are especially easy to bring about by the stimulation of what, for lack of a better term, we may call the erogenous zones, such as the nipples, the lips, and the sex organs" (p. 155).

Infants initially do not love specific people, but they are conditioned to do so. The mother's face frequently appears along with patting, rocking, and stroking, so it becomes a conditioned stimulus that alone elicits the good feelings toward her. Later, other people associated with the mother in some way

188

Learning Theory: Pavlov, Watson, and Skinner

also elicit the same responses. Thus tender or positive feelings toward others are learned through second-order conditioning.

Actually, much of Watson's writing on the emotions was speculation—and vague speculation at that. He said the three basic emotions become attached to a variety of stimuli and "we get marked additions to the responses and other modifications of them" (p. 165), but he said little about how these further developments occur. Where Watson did become specific was in his experimental work. His major experiment was on the conditioning of fear in an 11-month-old infant he called Albert B.

Conditioning Fear in Little Albert. Watson and Raynor (Watson, 1924, pp. 159-164) wanted to see if they could condition Albert to fear a white rat. At the beginning of the experiment Albert showed no such fear. Next, the experimenter on four occasions presented the rat and simultaneously pounded a bar behind Albert's head, producing a startle response. On the fifth trial Albert was shown the rat alone, and he puckered his face, whimpered, and withdrew. He had been conditioned to fear the rat. For good measure, the experimenter combined the rat and the pounding twice more, and on the next

trial, when the rat was again presented alone, Albert cried and tried to crawl away as rapidly as he could.

A few days later, Watson and Raynor tested for stimulus generalization. They found that although Albert played with many objects, he feared anything furry. He cried or fretted whenever he saw a rabbit, dog, fur coat, cotton wool, or a Santa Claus mask, even though he previously had not been afraid of these things. Albert's fear had generalized to all furry objects.

Practical Applications. One of Watson's major practical innovations was a method for deconditioning fears. He was not able to decondition Albert of his new fears, because Albert was an orphan who was adopted and taken out of town before this could be attempted. But Watson advised one of his colleagues, Mary Cover Jones, on procedures for eliminating the fears of another little boy, a 3-year-old called Peter.

Peter seemed active and healthy in every respect except for his fears. He was scared of white rats, rabbits, fur coats, feathers, cotton wool, frogs, fish, and mechanical toys. As Watson noted, "One might well think that Peter was merely Albert B. grown up, but Peter was a different child whose fears were 'home grown'" (1924, p. 173).

Jones tried a variety of methods, including having Peter watch other children play with a rabbit. But the procedure that she and Watson highlighted was the following. Peter was placed in his highchair and given a midafternoon snack. Then a caged white rabbit was displayed at a distance that did not disturb him. The next day, the rabbit was brought increasingly closer, until he showed a slight disturbance. That ended the day's treatment. The same thing was done day after day; the rabbit was brought closer and closer, with

189

Learning Theory: Pavlov, Watson, and Skinner

the experimenter taking care never to disturb Peter very much. Finally, Peter was able to eat with one hand while playing with the rabbit with the other. By similar means, Jones eliminated most of Peter's other fears as well.

Jones's technique, although anticipated by Locke, was quite innovative at the time. It is today known as a form of behavior modification called systematic desensitization (see Wolpe, 1969). The subject is relaxed and gradually introduced to the feared stimulus. The experimenter makes sure that the subject is at no time made to feel too anxious. Gradually, then, the subject learns to asso-

ciate relaxed feelings, rather than fear, to the object or situation.

Watson did not confine his advice to therapeutic procedures for eliminating fears. He also had much to say on child rearing, which he wanted to turn into a scientific enterprise. Watson recommended, among other things, that parents place babies on rigid schedules, and he insisted they refrain from hugging, kissing, or caressing their babies. For when they do so, their children soon associate the very sight of the parent with indulgent responses and never learn to turn away from the parent and explore the world on their own (Watson, 1928, p. 81). Watson's advice was quite influential in the 1930s, but it was too extreme to last. Under the influence of Spock, Bowlby, and others, parents relaxed their schedules and became more affectionate with their children. Nevertheless, Watson's more general goal—that of placing child training on the firm foundation of scientific learning principles—remains a vital part of child care in the United States.

Evaluation

Largely because of Watson's efforts, the classical conditioning paradigm became a cornerstone of psychological theory. It would seem that many of our reactions to objects and people develop through this conditioning process (see Liebert et al., 1977).

At the same time, we need to note that the model has certain limitations. For one thing, researchers have found it much more difficult to condition infants' responses than Watson implied. This seems particularly true during the first month of life (Lamb & Campos, 1982; Sameroff & Cavanaugh, 1979). Perhaps classical conditioning becomes easier once infants have developed what Piaget calls primary circular reactions. Once they can coordinate sensorimotor actions (e.g., look at what they hear), they might more readily learn to make various associations.

There also seem to be limitations to the kinds of conditioned stimuli humans will learn. When, for example, researchers attempted to classically condition infants to fear objects such as curtains and wooden blocks instead of rats, they had great difficulty. Perhaps humans are innately disposed to fear certain stimuli. There may be biological constraints on the kinds of stimuli we will associate with different responses (Harris & Liebert, 1984, pp. 108-109; Seligman, 1972).

190

Learning Theory: Pavlov, Watson, and Skinner

From a learning theory perspective, finally, classical conditioning seems

limited to certain kinds of responses. It seems to apply best to the conditioning of reflexes and innate responses (which may include many emotional reactions). It is questionable whether this kind of conditioning can also explain how we learn such active and complex skills as talking, using tools, dancing, or playing chess. When we master such skills, we are not limited to inborn reactions to stimuli, but we engage in a great deal of free, trial-and-error behavior, finding out what works best. Accordingly, learning theorists have developed other models of conditioning, the most influential of which is that of B. F. Skinner.

SKINNER AND OPERANT CONDITIONING

Biographical Introduction

B. F. Skinner (1905-1990) grew up in the small town of Susquehanna, Pennsylvania. As a boy, he liked school and enjoyed building things such as sleds, rafts, and wagons. He also wrote stories and poetry. After graduating from high school, he went to Hamilton College in New York. There, he felt somewhat out of place, but he graduated Phi Beta Kappa with a major in English literature.

Skinner spent the next two years trying to become a writer, but he eventually decided that he could not succeed because "I had nothing important to say" (1967, p. 395). Because he was interested in human and animal behavior, he enrolled in the graduate psychology department at Harvard, where he began doing research and formulating his ideas on learning. Skinner taught at the University of Minnesota (1936-1945), Indiana University (1945-1947), and Harvard University (1947 until his death in 1990).

Despite his successful career as a scientist, Skinner never completely abandoned his earlier interests. For one thing, he continued to display his boyhood enthusiasm for building things. When his first child was born, he decided to make a new, improved crib. This crib, which is sometimes called his "baby box," is a pleasantly heated place that does away with the necessity of excessive clothing and permits freer movement. It is not, as is sometimes thought, an apparatus for training babies. It is simply a more comfortable crib.

Skinner's literary interests also reemerged. In 1948 he published a novel, *Walden Two*, which describes a utopian community based on his principles of conditioning.

The Operant Model

Like Watson, Skinner was a strict behaviorist. He believed psychology should dispense with any references to intangible mental states (such as goals, desires, or purposes); instead, it should confine itself to the study of

Learning Theory: Pavlov, Watson, and Skinner

overt behavior. Like Watson, in addition, Skinner was an environmentalist; although Skinner recognized that organisms enter the world with genetic endowments, he was primarily concerned with how environments control behavior.

In contrast to Watson, however, Skinner's primary model of conditioning was not Pavlovian. The responses that Pavlov studied, Skinner said, are best thought of as respondents. These are responses that are automatically "elicited" by known stimuli. For example, the ingestion of food automatically elicits salivation, and a loud noise automatically elicits a startle response. Most respondents are probably simple reflexes.

A second class of behavior, which most interested Skinner, is called operant. In operant behavior, the animal is not harnessed in, like Pavlov's dogs, but moves freely about and "operates" on the environment. For example, in early experiments by Thorndike (1905), cats in a puzzle box would sniff, claw, and jump about until they hit upon the response—pulling a latch—that enabled them to get food. The successful response would then be more likely to recur. In such cases, we cannot always identify any prior stimulus that automatically elicits the responses. Rather, animals emit responses, some of which become more likely in the future because they have led to favorable consequences. Behavior, in Skinner's terms, is controlled by the reinforcing stimuli that follow it (Skinner, 1938, pp. 20-21; 1953, pp. 65-66). The two models, respondent and operant, are diagrammed in Figure 1.

To study operant conditioning, Skinner constructed an apparatus that is commonly referred to as a "Skinner box." This is a fairly small box in which an animal is free to roam about (see Figure 2). At one end there is a bar (lever) that, when pressed, automatically releases water or a pellet of food. The animal, such as a rat, at first pokes around until she eventually presses the bar, and then she gets the reward. As time goes on, she presses the bar more frequently. The most important measure of learning, for Skinner, is the rate of responding; when responses are reinforced, their rates of

Figure 1

Respondent and operant conditioning. In respondent (Pavlovian) conditioning, stimuli precede

responses and automatically elicit them. In operant conditioning, the initial stimuli are not always known; the organism simply emits responses that are controlled by reinforcing stimuli (S R s) that follow.

CS v

>>

US-f:R

Respondent conditioning Operant conditioning

192

Learning Theory: Pavlov, Watson, and Skinner

Figure 2

A Skinner box. One side has been cut away to show the part occupied by the animal.

(From Skinner, B. F., *The Behavior of Organisms*, p. 49. Copyright 1938, renewed 1966. Reprinted by permission of Prentice-Hall, Inc.)

occurrence increase. In Skinner's apparatus, the bar presses are automatically registered on a graph, so the experimenter need not be present much of the time. The data are presented as a learning curve, illustrated in Figure 3.

Skinner believed that operant behavior, in comparison to respondent behavior, plays a much greater role in human life. When we brush our teeth, drive a car, or read a book, our behavior is not automatically elicited by a specific stimulus. The mere sight of a book, for instance, does not elicit reading in the same way a bright light automatically elicits an eyeblink. We may or may not read the book, depending on the consequences that have followed in

the past. If reading books has brought us rewards, such as high grades, we are likely to engage in this behavior. Behavior is determined by its consequences (Munn, Fernald, & Fernald, 1974, p. 208).

Figure 3

A typical learning curve.

(From Skinner, B. F. The Behavior of Organisms, Copyright 1938. Renewed 1966. Reprinted by permission of Prentice-Hall, Inc.)

Learning Theory: Pavlov, Watson, and Skinner

Principles of Conditioning

Reinforcement and Extinction. Skinnerians have performed numerous experiments showing that human behavior, beginning in infancy, can be controlled by reinforcing stimuli. For example, infants increase their rates of sucking when sucking results in sweet, as opposed to nonsweet, liquid (Lipsitt, 1975). Similarly, infants' rates of smiling and vocalization can be increased if the behavior leads to rewards such as the experimenter's smiles, caresses, and attention (Brackbill, 1958; Rheingold, Gewirtz, & Ross, 1959).

In such experiments, one is dealing with different kinds of reinforcers. Some reinforcers, such as food or the removal of pain, are primary reinforcers; they have "natural" reinforcing properties. Other reinforcing stimuli, such as an adult's smiles, praise, or attention, are probably conditioned reinforcers; their effectiveness stems from their frequent association with primary reinforcers (Skinner, 1953, p. 78).

Operant behavior, like respondent behavior, is also subject to extinction (p. 69). For example, because children do things "just to get attention" (p. 78), one can extinguish undesirable behaviors, such as excessive crying or temper tantrums, by consistently withdrawing one's attention whenever they occur (Etzel & Gewirtz, 1967; Williams, 1959).

Operant behavior that has apparently been extinguished may also show spontaneous recovery. For example, a little boy whose temper tantrums had been extinguished through the withdrawal of attention began having tantrums once again when placed in a new situation (Williams, 1959). The behavior had to be extinguished further.

Immediacy of Reinforcement. Skinner (1953, p. 101; 1959, p. 133) found he could initially establish responses at the highest rates when he reinforced them promptly. A rat will begin pressing a bar at a high rate only if she has promptly received a food pellet each time she has done so. As Bijou and Baer (1961, p. 44) point out, this principle has importance for child rearing. If a father shows pleasure immediately after his son brings him the newspaper, the boy is likely to repeat the behavior the next evening. If, however, the father

is so engrossed in something else that he delays reinforcing his son's behavior for a few minutes, the boy's behavior will not be strengthened. In fact, what gets strengthened is the boy's behavior at the moment of reinforcement. If he is building blocks at that moment, it is block-building, not newspaper-fetching, that gets reinforced.

Discriminative Stimuli. We have said that operant conditioning may be described without any reference to initiating stimuli. This is true, but it does not mean such stimuli are unimportant. Stimuli that precede responses may gain considerable control over them.

For example, Skinner (1953, pp. 107-108) reinforced a pigeon each time she stretched her neck. At this point Skinner had no knowledge of any initial

194

Learning Theory: Pavlov, Watson, and Skinner

stimulus; he simply waited for the pigeon to emit the response and then reinforced it. Next, however, he reinforced the response only when a signal light was on. After a few trials, the pigeon stretched her neck much more frequently when the light was flashing than when it was off. The flashing light had become a discriminative stimulus. The light controlled the behavior because it set the occasion upon which the behavior was likely to be reinforced.

Skinner (pp. 108-109) listed numerous examples to show how everyday behavior becomes attached to discriminative stimuli. In an orchard in which red apples are sweet and all others are sour, redness becomes a stimulus that sets the occasion upon which picking and eating will produce favorable outcomes. Similarly, we learn that a smile is an occasion upon which approaching another will meet with a positive response. When others frown, the same approach meets with aversive consequences, such as rebuffs. Insofar as this is true, the facial expressions of others become discriminative stimuli that control the likelihood that we will approach them.

Although discriminative stimuli do exert considerable control, it must be emphasized that this control is not automatic, as in the case of respondent conditioning. In Pavlov's experiments, prior stimuli automatically elicit responses; in operant conditioning, such stimuli only make responses more likely.

Generalization. In operant conditioning, as in respondent conditioning, there is a process of stimulus generalization (Skinner, 1953, p. 132).
Sup-

pose a little girl has been reinforced for saying "Da da" at the sight of her father, but not when she is looking at her mother or siblings. The father has become a discriminative stimulus. It is not unusual, however, to find the girl saying "Da da" when she sees any man at all, such as strangers on the street. The stimulus has generalized. Her parents must now teach her to make a finer discrimination. They might say, "That's right," when she utters "Da da" in the presence of her father, but not when she looks at any other man.

Similarly, we can observe response generalization. It has been shown, for example, that when children are reinforced for using one part of speech, such as plurals, they begin uttering new plurals—even though they haven't received reinforcement for those particular words. Reinforcement influences not only particular responses but those of the same general class (Lovaas, 1977, pp. 112-113).

Shaping. Operant behavior is not acquired in all-or-nothing packages. It is usually learned gradually, little by little. Even teaching a pigeon to peck a spot on the wall, Skinner (1953, p. 92) showed, must be gradually shaped. If we place a pigeon in a box and wait for her to peck the spot, we may have to wait days or even weeks. Much of the time, the pigeon doesn't even approach the spot. So we must shape her behavior. First, we give the bird food when she turns in the direction of the spot. This increases the frequency

195

Learning Theory: Pavlov, Watson, and Skinner

of this behavior. Next, we withhold food until she makes a slight movement in the right direction. We then keep reinforcing positions closer and closer to the spot, until the bird is facing it. At this point we can reinforce head movements, first giving food for any forward movement and finally reinforcing the bird only when she actually pecks the spot. Through this procedure we gradually shape the desired response. Shaping is also called the "method of approximations," because reinforcement is made contingent upon better and better approximations of the desired response.

We probably teach many human skills in this bit-by-bit shaping process. When we teach a boy to swing a baseball bat, we first say "Good" when he gets his hands into the right grip. We then say "Right" when he lifts his bat in the correct position over his shoulder. We then work on his stance, a level swing, and so on—gradually shaping the complete behavior.

Behavior Chains. Although behavior may be shaped bit by bit, it also develops into longer, integrated response chains. For example, batting in base-

ball involves picking up the bat, getting the right grip and stance, watching for the right pitch, swinging, running the bases, and so on. Skinnerians attempt to examine each step in terms of reinforcements and stimuli. Reaching for the bat is reinforced by obtaining it, which also serves as a stimulus for the next act, getting the right grip. Once the hands are placed on the bat, we get a certain "feel" that we recognize as the proper grip. This "feel" is a reinforcement, and it also signals the next action, pulling the bat over the shoulder. A little later, the sensation of the bat squarely striking the ball is a reinforcement for the swing, and it also signals the next action, running the bases. When a boy or girl has become a good hitter, the entire sequence is often performed in a smooth, integrated fashion (Schwartz, 1989).

Schedules of Reinforcement. Skinner (1953, p. 99) observed that our everyday behavior is rarely reinforced continuously, every time; instead, it is reinforced intermittently. We do not find good snow every time we go skiing or have fun every time we go to a party. Accordingly, Skinner studied the effects of different schedules of intermittent reinforcement.

Intermittent reinforcement may be set up on a fixed-interval schedule, such that the organism receives a reward for the first response after a specified period of time. For instance, a pigeon receives food after pecking a disc, but must wait 3 minutes before her next peck is rewarded, then 3 more minutes, and so on. The rate of responding on this schedule is generally low. Higher rates are produced by fixed-ratio schedules, as when the pigeon gets food after every fifth peck. On both schedules, however, there is a lull in responding immediately after reinforcement. It is as if the organism knows it has a long way to go before the next reinforcement (p. 103). Students often experience this effect immediately after completing a long term paper—it is difficult to get started on another assignment.

196

Learning Theory: Pavlov, Watson, and Skinner

The lulls produced by fixed schedules can be avoided by varying reinforcement in unpredictable ways. On variable-interval schedules, reinforcement is administered after an average length of time, but the intervals are mixed up. With variable-ratio schedules, we vary the number of responses needed to produce a reward. When put on these two schedules, organisms consistently respond at high rates, especially on variable-ratio schedules. They

keep responding because a reward might come at any time.

One of Skinner's most important findings is that intermittently reinforced behavior, in comparison to that which is continuously reinforced, is much more difficult to extinguish. This is why many of our children's undesirable behaviors are so difficult to stop. We might be able to resist a child's nagging or demanding behavior most of the time, but if we yield every once in a while, the child will persist with it (Bijou & Baer, 1961, p. 62).

If we wish to begin teaching a desirable form of behavior, it is usually best to begin with continuous reinforcement; this is the most efficient way to get the behavior started. However, if we also wish to make the behavior last, we might at some point switch to an intermittent schedule (Bijou & Baer, 1961, p. 62).

Negative Reinforcement and Punishment. So far we have been focusing on positive reinforcement. Reinforcement means strengthening a response (increasing its rate), and positive reinforcements strengthen responses by adding positive consequences such as food, praise, or attention. Responses may also be strengthened through negative reinforcement, by removing unpleasant or aversive stimuli. Basically, what is strengthened in this way is the tendency to escape, as when a girl standing on a diving board learns to escape the taunts of her peers by diving into the water (Skinner, 1953, pp. 73,173).

When we punish, in contrast, we do not try to strengthen behavior but to eliminate it. Punishment, Skinner said, is "the commonest technique of control in modern life. The pattern is familiar: If a man does not behave as you wish, knock him down; if a child misbehaves, spank him; if the people of a country misbehave, bomb them" (p. 182).

Punishment, however, does not always work. In an early experiment, Skinner (1938) found that when he punished rats for bar pressing (by having the bar swing back and smack them on the legs), he only temporarily suppressed the response. In the long run, punishment did not eliminate the response any faster than did extinction. Other studies (e.g., Estes, 1944) have obtained similar results, and the findings conform to everyday experience. Parents who hit their children get them to behave for a while, but the parents find that the misconduct reappears later on.

Skinner also objected to punishment because it produces unwanted side effects. A child who is scolded in school may soon appear inhibited and conflicted. The child seems torn between working and avoiding work because of

Learning Theory: Pavlov, Watson, and Skinner

the feared consequences. The boy or girl may start and stop, become distracted, and behave in other awkward ways (Skinner, 1953, pp. 190-191).

Some researchers believe Skinner overstated the case against punishment. In some instances punishment will in fact completely eliminate responses. This is especially true when the punishment is extremely painful. Also, punishment can be effective when it is promptly administered, and when the organism can make alternative responses that are then rewarded (Liebert et al., 1977, pp. 138-141). Nevertheless, the effects of punishment are often puzzling and undesirable.

Skinner recommended that instead of punishing children, we try extinction. "If the child's behavior is strong only because it has been reinforced by 'getting a rise out of' the parent, it will disappear when this consequence is no longer forthcoming" (1953, p. 192). Skinnerians often suggest that we combine extinction for undesirable behavior with positive reinforcement for desirable behavior. In one study, teachers simply ignored nursery school children whenever they were aggressive and gave them praise and attention whenever they were peaceful or cooperative. The result was a quieter classroom (P. Brown & Elliott, 1965).

Internal Events:

Thoughts, Feelings, and Drives

Thoughts. It is sometimes said that Skinner proposed an "empty organism" theory. He examined only overt responses and ignored internal states. This characterization is accurate but slightly oversimplified. Skinner did not deny that an inner world exists. We do have inner sensations, such as the pain from a toothache. We also can be said to think. Thinking is merely a weaker or more covert form of behavior. For example, we may talk to ourselves silently instead of out loud, or we may think out our moves silently in a chess game. However, such private events have no place in scientific psychology unless we can find ways of making them public and measuring them (Skinner, 1974, pp. 16-17, and chap. 7).

Skinner was particularly distressed by our tendency to treat thoughts as the causes of behavior. We say we went to the store because "we got an idea to do so" or that a pigeon pecked a disc because she "anticipated" food. However, we are in error when we speak in this way. We go to stores, and pigeons peck discs, only because these actions have led to past reinforcements. Any discussion of goals or expectations is superfluous. Worse, it diverts us from the true explanation of behavior—the controlling effect of the environ-

ment (Skinner, 1969, pp. 240-241; 1974, pp. 68-71).

Feelings. Skinner acknowledged that we have emotions, just as we have thoughts. However, feelings do not cause behavior any more than thoughts do. We might say we are going to the movies because "we want

198

Learning Theory: Pavlov, Watson, and Skinner

to" or because "we feel like it," but such statements explain nothing. If we go to the movies, it is because this behavior has been reinforced in the past (Skinner, 1971, p. 10).

Emotional responses themselves can be explained according to learning-theory principles. In our discussion of Watson, we saw how emotional reactions might be learned through classical conditioning. Skinner believes an operant analysis is also useful. Many emotions are the by-products of different reinforcement contingencies. Confidence, for example, is a by-product of frequent positive reinforcement. When we learn to hit a baseball sharply and consistently, we develop a sense of confidence and mastery (Skinner, 1974, p. 58). Conversely, we become depressed and lethargic when reinforcements are no longer forthcoming. On certain fixed-ratio or fixed-interval schedules we find it difficult to get going after receiving a reward because further rewards will not be coming for some time (p. 59).

An operant analysis also helps us understand why various patterns of emotional behavior persist. If a little girl persistently behaves in an aggressive manner, it is important to know the consequences of this behavior. Do her actions succeed in getting attention or other children's toys? If so, her aggressiveness is likely to continue. Similarly, if displays of happiness, meekness, sympathy, fearfulness, and other emotional responses persist, it is because they have produced positive consequences (Bijou & Baer, 1961, pp. 73-74; Skinner, 1969, pp. 129-130).

Skinner believed, then, that we can understand emotions if we look at them as the products of environmental control. It is useless to consider emotions as intrapsychic causes of behavior, as the Freudians do. For example, a Freudian might talk about a man who fears sex because of anticipated punishment from an internal agency, the superego. To Skinner, such discussions get us nowhere. If we wish to understand why a person avoids sex, we must look at the past consequences of his sexual behavior (Skinner, 1974, chap. 10).

Drives. Skinner's refusal to look for causes of behavior within the organism led to certain difficulties. In particular, he had trouble with the concept of drive. Drives, such as hunger or thirst, would seem to refer to internal states that motivate behavior, and Skinner himself deprived his animals of food and water in order to make reinforcements effective.

Skinner argued that we do not need to conceive of drives as inner states, either mental or physiological. We simply specify the hours we deprive an animal of food or water and examine the effect of this operation on response rates (Skinner, 1953, p. 149).

Still, the drive concept has remained a thorn in the side of Skinnerians, and they have therefore searched for ways of conceptualizing reinforcement without reference to this concept. One interesting proposal has been made by Premack (1961), who suggests we think of reinforcement simply

199

Learning Theory: Pavlov, Watson, and Skinner

as the momentary probability of a response. Behavior that has a high probability of occurrence at the moment can serve as a reinforcer for behavior with a lower probability. If children are supposed to be eating their dinner but are busy playing instead, playing can be used as a reinforcer for eating. We simply say, "Eat some dinner and then play some more" (Hommel & Totsi, 1969). Conceptualized in this way, eating and drinking have no special status as reinforcers. Eating and drinking, like any other actions, may or may not be good reinforcers, depending on their probabilities of occurrence at a particular time.

Species-Specific Behavior

Skinner argued, then, that we need not look inside the organism for the causes of behavior. Behavior is controlled by the external environment. There do seem to be, however, certain limitations to environmental control. As we briefly mentioned in our evaluation of Watson, each species has a particular genetic endowment that makes it easier to teach it some things rather than other things. Operant research has found, for example, that it is hard to teach a rat to let go of objects, and it is hard to shape vocal behavior in nonhuman species (Skinner, 1969, p. 201). There are, as learning theorists increasingly say, biological "constraints" on what a species can learn.

In practice, Skinnerians often deal with species-specific behavior as

the topography of a response. That is, the experimenter maps out a description of the behavior he or she can work with—for example, vocal behavior in humans. The topography is merely a description and does not constitute the most important part of the analysis, which is the way reinforcements shape and maintain behavior. Nevertheless, the topography is essential (pp. 199-209).

In a larger sense, Skinner argued, even species-specific behavior is a product of environmental contingencies. For such behavior has become, in the course of evolution, part of the species' repertoire because it has helped that species survive in a certain environment. Thus environments selectively reinforce all behavior—not only that in an animal's lifetime but also that in its species' evolutionary past (pp. 199-209).

Practical Applications

Behavior Modification with Children with Autism. Skinner's research readily lends itself to practical applications. We have seen how Skinnerians might extinguish temper tantrums or get an unruly class to behave. The use of operant techniques to correct behavior problems is a branch of behavior modification. Operant techniques supplement the systematic desensitization procedures first employed by Watson and Jones.

200

Learning Theory: Pavlov, Watson, and Skinner

An impressive example of operant therapy is Lovaas's work with children with autism. Autism was first described by Kanner in 1943. It is a severe disorder in which children are extremely isolated. The children also engage in repetitive behavior such as spinning objects or flapping their hands over and over. Many are mute, and others are echolalic—they merely echo what one says. Some engage in self-injurious behavior, such as hitting themselves (Lovaas, 2003; Koegel & Koegel, 2006, p. 34).

Lovaas tries to gain control over the children's behavior so he can change it. He tries to eliminate socially inappropriate behavior and reinforce socially appropriate behavior. If a child engages in echolalia, repetitive behavior, or self-injurious behavior, Lovaas withdraws attention or punishes the child with a loud "No!" or a slap on the thigh. If the child does something more appropriate, such as emitting correct speech, Lovaas gives the child a reward, perhaps a bit of tasty cereal and the word "Good" (Lovaas, 1987).

Frequently, appropriate behavior must be gradually shaped, as when the therapist teaches mute children to imitate words. At first, the therapist reinforces any vocalization the child makes, even blowing air out of the mouth. Once the child is regularly making sounds, the therapist says a word such as "baby" and rewards any sound that comes within the next 5 seconds. After that, rewards are made contingent on better and better approximations of "baby" (or other target words) (Lovaas, 1969,1977).

Initially, some children are so silent that the therapist must elicit sounds by tickling them or pressing their lips together and then letting the air out. These interventions are called manual prompts. Strictly speaking, these prompts violate Skinner's operant paradigm; operant conditioning reinforces freely emitted behavior—not behavior forced by the therapist's actions. Prompts are faded as soon as possible (Lovaas, 1977, pp. 36-37).

Lovaas's therapy is intensive. In his first major project, begun in the 1960s, Lovaas and his staff trained children seven hours a day, seven days a week, for one year in a residential treatment setting at UCLA. Most of the children were 5 to 8 years old (Lovaas, 1973,1977). Many made significant progress, but when they were discharged to state hospitals, they lost all they had gained. In his next major project, in the 1970s and 1980s, Lovaas avoided this discharge problem; he worked with children in their homes and taught parents to help train the children. He also worked with younger children—under the age of 4 years. The children were trained at least 40 hours a week. After 2 to 3 years, nearly half entered first grade as regular students—an achievement that would have once seemed impossible (Lovaas, 1987). A 6-year follow-up study found that almost all of these children still attended a regular school (McEachlin, Smith, & Lovaas, 1993).

Lovaas and his colleagues have described their treatment as applied behavior analysis, or ABA, a term they apply to any therapy that uses principles of learning in a measurable, scientific manner. Lovaas's ABA methods are still the most widely used, but there are new variations. Robert and Lynn Koegel and their colleagues (2006) have developed a program that focuses less on teaching

201

Learning Theory: Pavlov, Watson, and Skinner

discrete skills and more on broad areas such as motivation. The Koegels believe that children often have difficulty with Lovaas's tasks and sometimes lose their

enthusiasm for learning because reinforcement is infrequent. The Koegels therefore offer reinforcement more frequently—often just for the effort. (They call this loose shaping.) They also let the child choose toys and activities. And instead of giving cereal treats, they offer rewards that are important to the child in his or her everyday environment. If a girl wants a shovel to dig in the sand, the therapist asks her to say "Shovel," and if she does, she receives the shovel. If she says, "Swing," she gets a push on the swing. The Koegels have reported success, although their studies aren't nearly as extensive as those by Lovaas.

Compared to Lovaas, the Koegels give children more opportunities to take the lead. But the therapist still maintains control, as when the therapist requires the child to name an object such as a shovel before receiving it. Furthermore, if a child engages in highly repetitive behavior, the therapist actively intervenes to turn the child's attention to other activities. Sometimes the therapist actually stands in front of the object, such as a fan, that stimulates repetitive behavior (Koegel & Koegel, 2006, p. 221).

Programmed Instruction. Skinner contributed to the education of normal children through his invention of teaching machines and programmed instruction (Skinner, 1968). The teaching machine was a simple apparatus that permitted one to read a brief passage, answer questions, and then, by turning a knob, see if one was correct. Actually, the machine itself was less important than the programmed material it contained, and today the material is presented in simple booklet form or installed in a computer. To get an idea of how programmed instruction works, read the following material and pretend to fill in the blanks. As you do so, cover the answers on the left side with a piece of paper, sliding it down just far enough to check your answers.

1. Programmed instruction involves several basic principles of learning. One of these, called the principle of small steps, is based on the premise that

small new information must be presented in steps.

2. The learner gradually acquires more and more

small steps information, but always in

3. Because active readers generally acquire more knowledge than passive readers, programmed

instruction also is based on the principle of active participation. Writing key words as one is reading

active involves the principle of participation.

2 From Munn, N. L., Fernald, L. D., and Fernald, P. S., Introduction to Psychology, 3rd ed., Boston: Houghton Mifflin Co., 1974, pp. 249-250. By permission.

202

Learning Theory: Pavlov, Watson, and Skinner

active participation
small steps
active participation

knowledge

immediate
of results

immediate knowledge
of results

small steps
active participation
immediate knowledge
of results

4. While reading a book, an uninterested learner may

slip into a passive state and discover that he cannot recall what he has just "read." In using programmed instruction the learner is prompted to remain alert by writing the key words, thus utilizing the principle of_.

5. In these two techniques of programmed instruction, information is presented in_, and occasionally key words are missing thus requiring the learner's_to complete the statements.

6. A third principle, immediate knowledge of results, is illustrated when a professor returns quiz papers to his students at the end of the class in which they were written. These students receive almost immediate_of results.

7. If a student makes an incorrect response at any point in programmed instruction, he discovers his mistake because the correct answer may be seen immediately after the frame, before the next one is considered. Thus, in programmed instruction, the learner receives_knowledge_.

8. Notice that in programmed instruction, unlike the evaluation of term papers, "immediate" does not mean a week or even a day but rather a few seconds. The reader of the program is continuously informed concerning his progress; he receives_.

9. Let us review the three techniques of programmed instruction already considered. By means of

_, the reader learns new material,

which he acquires through_followed

by_ •

Programmed instruction embodies several Skinnerian principles. First, it proceeds in small steps, because Skinner has found that the best way to establish new behavior is to shape it bit by bit. Second, the learner is active, because this is the natural condition of organisms. (Recall how Pavlov's dogs, in contrast, were harnessed in and simply reacted to stimuli.) Third, feedback is immediate because Skinner found that learning is most rapid when promptly reinforced. (Reinforcement here is the knowledge that one's answer

is correct.)

A sample of programmed reading for children is found in Figure 4. In programmed instruction, students work independently and at their own pace. The instruction units are constructed so each student may begin at a level she can easily master. One does not want the student making many errors at first, for then she will lack positive reinforcement for learning. As with shaping, one

203

Learning Theory: Pavlov, Watson, and Skinner

Mat

Man

I am a ma

Figure 4

Programmed instruction for children.
(Adapted from Sullivan, M. W., Programmed learning in reading. In A. D. Calvin, Ed., Programmed Instruction: Bold New Venture. Bloomington: Indiana University Press, 1969, p. 111. By permission of the publisher.)

Mat

I am a ma

begins by reinforcing responses that are within the student's behavioral repertoire and gradually building up from there.

On a technical level, programmed instruction has run into some difficulties. For example, students sometimes rush through the programs without fully mastering the material (Munson & Crosbie, 1998). But the underlying principles are important and make efforts to solve the problems worthwhile.

Surprisingly, the principles underlying programmed instruction overlap somewhat with Montessori's. Both Skinner and Montessori wanted to make learning an individualized, self-paced activity that begins at the student's own level and builds skills gradually. For both, the goal is not to tear down, through criticism or punishment, but to make learning a consistently positive experience.

But the two approaches also differ. For one thing, programmed instruction involves material that young children read (see Figure 4), whereas Montessori materials are largely physical. Even when learning to read, Montessori children begin with sandpaper letters, metal insets, and so on. Montessori thought young children find such physical activities more natural.

More fundamentally, there is the difference in the extent to which the child's work is free from adult direction. Montessori allowed children to choose their own tasks and work on them while the teacher steps into the background. She wanted children to discover for themselves how something is out of place, how cylinders fit, how water is poured, and whatever else is important to them. In programmed instruction, in contrast, adult direction is pervasive. Although it might seem that children work

204

Learning Theory: Pavlov, Watson, and Skinner

independently on the booklets, in fact an adult (the program developer) has structured each small response. The child follows the adult's lead,

repeatedly checking with this social authority to see if she is right. Children probably derive less sense that they are making their own discoveries about the world.

Nevertheless, it is important not to overlook the similarities between the two methods—especially the way both try to make learning a positive experience. One can even imagine Skinner approving Montessori's physical tasks, albeit in his own terms. He would say they work not because they allow for spontaneous discoveries, but because they allow children to make responses that readily result in positive feedback from the physical environment.

EVALUATION

Skinner considerably widened the scope of learning theory. After noting the limitations of classical conditioning, he explored the nature of operant behavior, where the organism acts freely and is controlled by the consequences of its actions. In a brilliant series of studies, Skinner showed how such control is exerted—by schedules of reinforcement, shaping, the influence of discriminative stimuli, and other factors. Furthermore, Skinner amply demonstrated the practical importance of his ideas.

In the process, Skinner stirred up controversies on many fronts. To some, his work lends itself to authoritarian practices—for he suggests ways to control, manipulate, and program others' behavior. Skinner's (e.g., 1974, p. 244) reply was that environments do, in fact, control behavior, and how we use our knowledge of this fact is up to us. We can create environments that suit humane purposes, or we can create ones that do not.

Developmentalists, too, often enter into heated, value-laden debates with Skinnerians. Developmentalists cringe at talk of controlling and changing children's behavior, when we should, instead, try to understand children and give them opportunities to grow on their own. To many Skinnerians, such sentiments are romantic and naive, for children chiefly develop through the molding influence of the external environment.

In a more objective vein, there are essentially three ways in which Skinner and writers in the developmental tradition disagree. First, developmental theorists often discuss internal events. Piaget described complex mental structures, even though he did not expect to find direct evidence for all of them in any individual case. Freudians discuss internal events, such as unconscious fantasies, that we cannot directly observe at all. Skinner believed such concepts divert us from scientific progress, which is made when we confine ourselves to the measurement of overt responses and

Learning Theory: Pavlov, Watson, and Skinner

environmental stimuli. But on this point, Skinner is now generally considered too extreme. Since the 1960s, there has been a dramatic new interest in cognition, and even growing numbers of learning theorists have been considering internal, cognitive events, even if the events cannot be directly measured.

Second, developmental theorists and Skinnerians disagree on the meaning and importance of developmental stages—periods when children organize experience in very different ways. In Piaget's theory, for example, a child's stage is a crucial variable; it is a predictor of the kind of experience the child can learn from. A child at the sensorimotor level will not learn tasks that involve language, nor will a child beginning to master concrete operations learn much from lectures covering abstract theory.

Skinnerians doubt the validity of stages as general, distinct ways of thinking or behaving; they believe the environment shapes behavior in a gradual, continuous manner (Bijou, 1976, p. 2; Skinner, 1953, p. 91). Skinner did acknowledge that one must note the child's age in any experiment, just as one must note an animal's species and characteristic behavior (Skinner, 1969, p. 89).

Age contributes to the "topography" of behavior; it helps describe the behavior that the experimenter sets about to shape or maintain. However, such information is still merely descriptive; it is secondary to environmental variables that control behavior. The question is whether the child's developmental status deserves this secondary role.

A third issue dividing Skinner and developmental theorists is the most important of all. This issue concerns the source of behavioral change. Developmentalists contend that in crucial instances a child's thoughts, feelings, and actions develop spontaneously, from within. Behavior is not exclusively patterned by the external environment. Gesell, for example, believed children stand, walk, talk, and so on from inner maturational promptings. Piaget was not a maturationist, but he also looked primarily to inner forces underlying developmental change. In his view, children's behavior is not structured by the environment but by children themselves. Children, out of a spontaneous interest in moderately novel events, construct increasingly complex and differentiated structures for dealing with the world.

Consider, for example, a baby girl who drops a block, hears the sound, and drops it again and again, making this new and interesting sound last. In Skinner's theory, the sound is a reinforcer that controls her behavior. But

this reinforcer will soon lose its effectiveness, for she will soon become interested in more complex outcomes (Kohlberg, 1969a). She may, for instance, begin listening for different sounds as she drops objects from different heights. For Piaget, we cannot look to external reinforcements as the determinants of behavior, for these often vary with the child's developing interests. For him, the main variable is the child's spontaneous curiosity about increasingly complex events.

206

Learning Theory: Pavlov, Watson, and Skinner

Developmental theorists, then, try to conceptualize ways in which children grow and learn on their own, somewhat independent of others' teachings or external reinforcements. At the same time, no one can deny that environments also reinforce and control behavior to a considerable extent, and often in ways Skinner described. Skinner's theory and research, moreover, have a clarity and elegant simplicity that others would do well to emulate. It is clear that Skinner's enormous contribution to scientific method and theory will be a lasting one.