Reflections on Dr. Fred Keller's

Learning: Reinforcement Theory

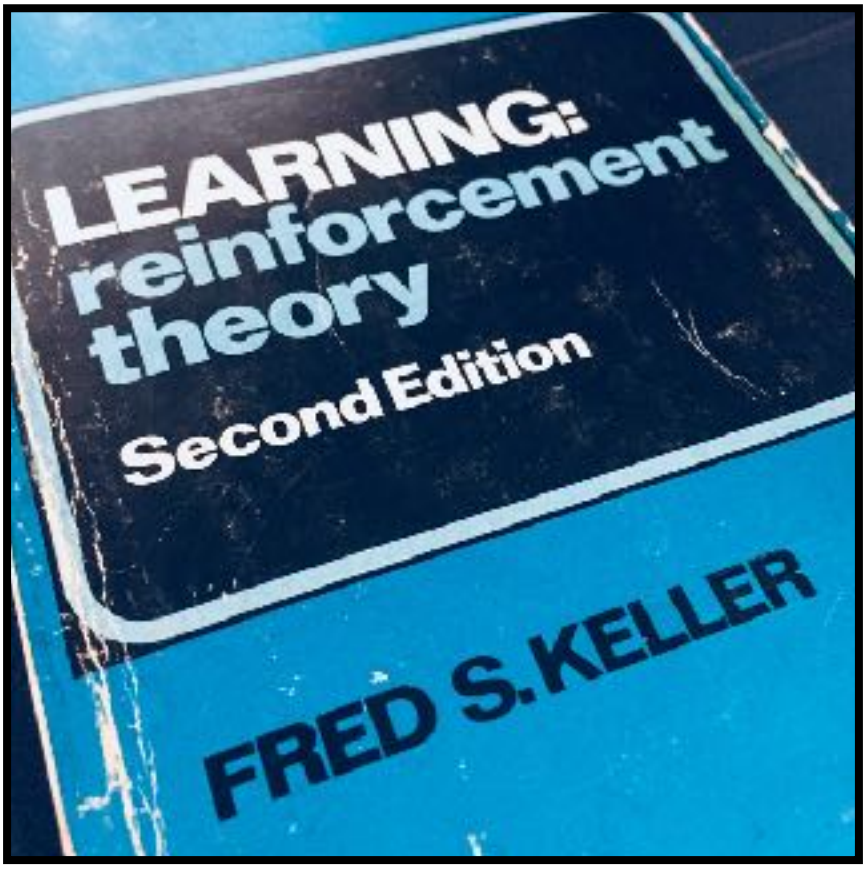

By Mary Hunter and Dr. Jesús Rosales-Ruiz

Introduction

Learning: Reinforcement Theory

In the 1950s, Dr. Fred Keller wrote a book called *Learning: Reinforcement Theory*.

The size of a small paperback novel, it's a short book. The whole thing is about 80 pages. However, within its pages, Dr. Keller explains the basics of operant and respondent conditioning, reinforcement, punishment, extinction, intermittent reinforcement, generalization, discrimination, chaining, and more.

For those who are not familiar with Dr. Keller, he was an accomplished behavior analyst and a contemporary of Dr. B. F. Skinner. Skinner and Keller met at Harvard in the 1930s and remained friends for life. Dr. Keller authored the textbook *Principles of Psychology* with Dr. William Shoenfeld and also developed the Personalized System of Instruction, a framework for teaching self-paced, mastery-based courses.

Learning: Reinforcement Theory is out of print and has largely been forgotten as other textbooks have replaced it. Still, it provides a solid introduction to the basic ideas of behavior analysis and contains many interesting examples and thought-provoking passages.

I re-read the book during the summer of 2019. During that time, Dr. Jesús Rosales-Ruiz and I published a series of short posts about the book on our Behavior Explorer Facebook page.

We have collected those 14 posts, edited them slightly, and assembled them in this PDF. We hope you will enjoy these short articles about Dr. Keller's book and hope they will help you think in new ways about how people and animals learn.

~ MH

Note: The quotes in this document come from the second edition of Learning: Reinforcement Theory*, which was published in 1969.*

Part 1: What is learning?

In the introductory chapter, Dr. Keller writes:

“Some form of the verb ‘to learn’ is used in many situations….

“Yet is is not easy to define learning….

“Is learning to skate like learning to be ashamed? Or like learning the value of money? Like learning to ski? Is learning to fear the dentist’s drill more like learning to avoid a flame than it is like learning to yell attention? Is there just one kind of learning? Are there two kinds? Seven? Or are there as many kinds as there are examples? What, exactly, is learning, anyhow?”

Take a minute and consider some of these examples. Are they alike? Are they not alike? What do you think?

Learning often seems like a mysterious process. How does a child learn to talk, read, or ride a bike? These three examples seem very different. Even if you know how to teach a child to ride a bike, this doesn’t seem to have very much in common with skiing or skating, never mind learning to read, learning to talk, or learning to avoid a hot stove.

Behavior analysts study the behavior of people and animals. They study how learning takes place and how behavior changes over time in relation to different environmental events. The really cool part is that their research has led to the development of basic principles of learning.

The more you understand these basic principles, the easier it will be for you to see parallels between teaching a toddler to talk, teaching a dog to come, and teaching an adult how to ice skate.

I hope you will enjoy our journey through Dr. Keller’s book and that it will help you understand even more about human and animal behavior and about the basic principles of learning.

Part 2: Neutral stimuli

In the chapter about respondent conditioning, Dr. Keller makes an interesting point about neutral stimuli.

He gives an example of sounding a buzzer just before delivering a brief shock to a person's hand. After some repetition of this, the person's heart rate changes when the buzzer is heard, before the shock is felt.

Most people would say that the buzzer was a neutral stimulus before this procedure started.

However, Dr. Keller cautions agains calling the buzzer a "neutral" stimulus.

He writes that the buzzer was " 'neutral' only in a relative sense."

He goes on to say, it "probably had some effect on behavior before the pairings, some subtle influence that you might not even be able to observe." He adds that it might have had "its own reflex response" and "was really eliciting in its own right."

We use lots of stimuli during teaching: words, sounds, gestures, pictures, body movements, and more. Whether we are doing respondent conditioning or operant conditioning, we often think of these stimuli as neutral or meaningless until our teaching procedure gives them meaning.

However, it may be wiser to think of them as unknown stimuli, rather than neutral stimuli. In some cases, you may want to probe how your learner reacts to a certain stimulus, before using it for teaching.

Part 3: Extinction doesn't equal zero

After the chapter on respondent conditioning, Dr. Keller discusses positive reinforcement, negative reinforcement, and extinction.

In the section on extinction, he writes, "If reinforcement is withheld, the response will ultimately return to its original unconditioned rate (sometimes called its operant level)."

People often discuss extinction as a way to get rid of a behavior or to completely eliminate the behavior. However, while extinction will greatly reduce the frequency of a behavior, the behavior will not likely go completely to zero.

Let's look at an example to help make more sense of this.

A hungry rat is placed in an operant chamber. A lever protrudes from one wall. At the beginning, no reinforcement is provided. As the rat explores the chamber, he occasionally presses the lever, perhaps once every 10 minutes or so.

After the rat has been in the chamber for an hour, the experimenter makes a change. Now when the rat presses the lever, the rat receives a food pellet. It doesn't take long for the rat to catch on to this change. He now presses the lever at a high rate.

However, the food container had only 50 pellets. After the rat receives the last pellet, pressing the lever no longer "pays off." Extinction has started. At first, the rat continues to press the lever. Then, his rate of pressing gradually decreases.

However, lever pressing does NOT go to zero. Instead, it eventually returns to the same level as before the behavior was reinforced, about once every 10 minutes. The rat returns to grooming, exploring, and the other activities he was doing at the beginning.

This example illustrates how extinction reduces the rate of a behavior. However, extinction may not completely eliminate the behavior.

Dr. Ogden Lindsley would argue that there is not such a thing as a behavior going completely to zero. Behavior is a temporal phenomenon. Even if a behavior happens very, very rarely, perhaps once a decade, it still exists. Under the right conditions, it will return. This is why it is better to focus on teaching the behaviors you want, rather than using extinction to eliminate the behaviors you don't want.

Part 4: More about extinction

In the chapter on extinction, Dr. Keller gives an example of a pigeon that has been taught to peck a key.

He writes, “When, after long training, extinction is begun, the pigeon may respond as many as 7,500 times during the first hour, without any sign of let-up….

“Watching the bird, you might say that he was incurably addicted to key-pecking; you would wonder that he did not stop from sheer exhaustion! Ultimately, of course, he will no longer respond to the key, even when refreshed and still very hungry for the food that his pecking once produced, but an impatient observer might easily go away with the opinion that the habit was unbreakable.”

Imagine if you watched just a small piece of this process. You would see the pigeon pecking rapidly and never receiving any food. You would not know what was causing the behavior, and it might seem innate or automatic. Behavior can be difficult to interpret if the slice you are observing is too small.

There’s another point that is worth mentioning. The experimenter has put the behavior of pecking the key on extinction without giving the pigeon any alternative way to earn food. This is one reason why the behavior is so persistent.

Extinction is often suggested as a solution for dealing with unwanted behavior. However, extinction can be a slow and painful process if the learner has no other way to access the reinforcer. It’s much easier to change behavior if you teach the learner an alternative behavior for accessing that same reinforcer, rather than relying on just extinction.

Part 5: Extinction conditions

Soon, we'll be moving on to other topics, but we wanted to share one more idea about extinction. This quote comes from the chapter on intermittent reinforcement.

Dr. Keller writes, "An important factor in accounting for resistance to extinction is the similarity of the training conditions to the extinction conditions."

He discusses this in the context of schedules of reinforcement. For example, imagine you have one pigeon that receives grain every time it pecks a key. Meanwhile, a second pigeon receives grain when it pecks a key 100 times.

When extinction starts, the first pigeon will keep pecking for awhile. However, he will "realize" fairly quickly that something is different.

Extinction will take far longer for the second bird. This happens because, for him, the extinction condition "feels" much more similar to the training condition.

You can use this idea to your advantage.

Imagine you want to put an unwanted behavior on extinction and reinforce an alternative behavior. If you do this in the exact same context in which the unwanted behavior has been occurring, it will be more difficult for the learner to figure out the rules are different.

Instead, change something. Stand in a different position. Move to a different room. Wear a silly hat. Rearrange the furniture. Play some classical music. Do something so that the environment feels different to the learner, compared to the previous environment where the behavior was reinforced. This will lead to more rapid behavior change.

Part 6: DRL

In Chapter 9, Dr. Keller describes an experiment with a young boy. The boy is told to say random words. The experimenter gives him a penny every time he says the word "flower," as long as at least ten seconds have passed since he last said the word "flower."

This is what is called a DRL schedule of reinforcement. DRL stands for differential reinforcement of low rates of responding. This type of rule for reinforcement is used in some situations when the experimenter or practitioner wants a behavior to occur sometimes, but not too often.

Dr. Keller writes that, after some practice with this schedule, the boy "seldom responded, after reinforcement, in less than ten seconds, and seldom did he overshoot by more than two or three seconds. Yet he had no idea that time was involved in the experiment! He 'thought' he had to learn a series of responses, and, when the experiment was over, he apologized for his lack of success, saying: 'I'm sorry, but I couldn't figure out all of the words you wanted me to say.' "

Using PORTL, we have tested several types of time-based schedules, including differential reinforcement of low rates (DRL), differential reinforcement of other behavior (DRO), and non-contingent reinforcement. One benefit of PORTL is that our participant can tell us afterward what he or she thought.

With these types of schedules, we find similar results to what Dr. Keller describes. The participant experiments with different responses and eventually arrives at some chain or pattern of responses that meets the rule for reinforcement.

However, afterward, the participant is confused about what we were trying to reinforce or reports a completely different rule than the rule the teacher was using. In addition, the participant often reports negative emotions.

Positive reinforcement is usually described as a way to create happy, confident learners. While this is often the case, it is too simplistic. If the learner cannot figure out the criteria for reinforcement, the learner often becomes confused or frustrated, even if the teacher is using a high rate of positive reinforcement.

Part 7: Bad at generalization

In the middle of the book, Dr. Keller has several chapters that discuss generalization, discrimination, and differentiation. He says this about generalization:

“In everyday life, examples of generalization are so common that they go unnoticed. They are most obvious, perhaps, in children, where they are often amusing. Parents smile at the child who calls out ‘Doggie!’ at the sight of a horse, cow, or some other four-legged creature; or they may laugh to hear a child say that soda-water ‘tastes like my foot’s asleep.’ They may fail to realize that the same principle is involved when the responses are much more common and undramatic. They may not see that one child’s ‘Chickie!’ in response to a robin is essentially the same as another child’s ‘Birdie!’ ”

We often overlook examples of generalization. In addition, we are sometimes perplexed when we expect generalization to happen, and then it doesn’t happen. This can lead us to blame the learner and label the learner as being “bad” at generalization.

Years ago, I was working with a young horse that was still fairly new to trailer loading. One day, the owner had the floor of the trailer covered with wood shavings because she had taken another horse somewhere the day before. The young horse refused to get into the trailer. He “failed to generalize” to this new situation.

One conclusion would be that this horse was “bad” at generalization. However, during the previous training sessions, he had already demonstrated that he could walk in the trailer when it was parked in several different places in the yard, with the windows open or shut, and even with a different person leading him.

He had already demonstrated generalization in several new conditions. The shavings just represented another variable, a different type of flooring, which was not part of our initial training.

Rather than thinking of generalization as something that an individual is good at or bad at, think instead about your teaching procedures. How can you design your teaching procedures so that generalization happens when you want it to and doesn’t happen when you don’t want it to?

Part 8: Difficult discriminations

What happens when an animal or human can't figure out when reinforcement will be provided? Dr. Keller discusses this in his chapter on discrimination. He writes:

"In a now-famous experiment, Pavlov once trained a dog to discriminate visually between a circle and an ellipse. Then, step by step, he brought the ellipse closer and closer to the shape of the circle. Ultimately the discrimination broke down, as you would expect. With continued demands upon his powers, the dog finally became "neurotic," to the degree that he had to be removed from the experiment and given a long rest for the sake of his health."

Dr. Keller gives a second example, involving a six-year-old Russian child who was taught to distinguish between different metronome beats. He says: "When the final discrimination between 144 and 132 beats per minute was attempted, the child became seriously upset, showing extreme rudeness, disobedience, excitement, and aggressive behavior, as well as sleepiness in the experimental situation."

During the final phase of the experiment, I imagine both learners sometimes still chose the correct answer by chance. However, neither learner could figure out the rule for reinforcement. As a result, both learners became increasingly unhappy.

This happened because the cue had become ambiguous with respect to the outcome, or consequence, that followed. This is similar to a poisoned cue situation, in which the cue has been followed sometimes by reinforcement and sometimes by a correction.

These types of ambiguous situations may occur accidentally during teaching. For example, a teacher may arrange a teaching situation that seems clear to him or her. However, the learner just can't figure it out. This may happen because the learner is missing prerequisite skills or because previously taught behaviors get in the way.

In other cases, this can occur when the teacher is inconsistent, sometimes reinforcing one response as correct and other times deciding to reinforce a different response under the same conditions.

Although the results may not always be as extreme as Dr. Keller's examples, there will be unwanted side effects when your learner can't figure out the rules of the game.

Part 9: The evolution of the word *shaping*

Chapter 13 of Dr. Keller's book is called "Differentiation (Shaping)."

It is noteworthy that he chooses to call the chapter "Differentiation" and then includes the word *shaping* in parentheses. In the second paragraph of the chapter, he explains his choice of words.

"The word differentiation is not a very good one for our present purposes, since it is often used as if it meant discrimination, which it doesn't, at least in this account. Skill would be a better word if it weren't so inclusive, taking in both discrimination and differentiation and perhaps something more. Shaping is the term most commonly used, but it too raises problems since it is sometimes taken to mean a change in the stimulus control of some responses (for which a better term would be shifting). In this account, we'll stick to differentiation and try to make its meaning so clear that it won't be confused with anything else."

This passage is from the second edition of Dr. Keller's book, which was published in 1969. If you look in textbooks and articles from the 1930s, 1940s, and early 1950s, you will not see the word *shaping*. For example, the book *Principles of Psychology*, which Dr. Keller and Dr. William Shoenfeld published in 1950, talks only about differentiation.

This textbook from 1950 describes how, "Through a series of successive approximations to the desired reaction, the behavior is altered until it comes to bear little or no resemblance to the first-conditioned form." This certainly sounds like what we now call shaping!

It's fascinating to consider that the word *shaping* was not used at all by psychologists and behavior analysts 70 years ago. The basic concept of shaping was understood, but it was not identified by this label.

These quotes are a good reminder that terminology changes over time. The meaning of a certain term can shift, new words become part of the vocabulary, and other words fade away.

Part 10: Teach faster with smaller steps

In the chapter on shaping, Dr. Keller describes an experiment involving a pinball machine. The experimenter provides reinforcement when the subject pulls the plunger on the machine. At first, the experimenter has a broad criteria for reinforcement. The behavior of pulling the plunger is reinforced whether the subject pulls it a short distance or a very long distance.

The experimenter gradually narrows the requirement for reinforcement. By the end, the subject's behavior has changed considerably. The subject always pulls the plunger about one-third of the way back and his pulls fall in a fairly narrow range.

Dr. Keller writes that, if the experimenter had "jumped abruptly from the broad band to the narrow one, from the 'easy' to the 'hard,' he would have greatly increased the errors and time required for his subjects to solve their problem. By selectively reinforcing small advances in the right direction, he reduced the chances of failure and reached his goal without undue delay."

This quote contains an important lesson about shaping.

Teachers often make their teaching steps much too big. For example, the teacher envisions using a series of five steps to take the learner from where the learner is now to the goal behavior. However, it may actually be faster to use fifteen steps.

This seems counterintuitive to most people. How could fifteen steps be faster? However, with just five steps, some steps may be too big of a leap. The learner may miss important intermediate steps and may not spend time practicing essential component skills. As a result, the learner often makes errors that slow down the learning process. This approach may also produce unwanted behaviors and emotions that interfere with teaching.

In many situations, smaller teaching steps will help you get to your goal faster.

Part 11: It's ABC, not A—B—C

The relationships between behaviors and the environment can be described in terms of the ABCs. That is, Antecedents (cues) lead to particular Behaviors, which are followed by certain Consequences (reinforcers).

People often talk about antecedents, behaviors, and consequences as three completely separate categories that can be independently taught and analyzed.

For example, many textbooks diagram shaping as involving just Bs and Cs. The message is that the antecedent part of the equation is not important when you are gradually changing a response. Learning to respond to one antecedent and not to another involves discrimination, a completely separate process.

Toward the end of the chapter on differentiation (shaping), Dr. Keller writes: "The truth of the matter is that the processes of discrimination and differentiation march together, hand in hand, from the very beginnings of our behavioral development."

Sometimes, people talk about teaching a behavior, then adding an antecedent or cue for that behavior. But, learning doesn't happen in a vacuum. As Alexandra Kurland says, "Cues evolve during the shaping process." If you have taught a behavior, there are already environmental cues that go with that behavior!

This underscores the importance of carefully selecting or arranging the environment before you begin teaching. The environment you choose will have existing cues, and these cues will make certain behaviors more likely and other behaviors less likely.

Understanding the interconnectedness of the ABCs will help you arrange teaching environments that make it easier to teach the behaviors you want.

Part 12: Behavior chains

Chapter 14 is about chaining.

Dr. Keller writes: “We express this notion of chaining, today, in the simple statement that *one response may produce the stimulus for the next*. And we recognize, perhaps more than ever before, that it is the exceptional case in which responses do *not* occur in chains. It is seldom that a single response or stimulus-response connection does not lead to another or arise from something that has gone before.”

Many of the responses we call “behaviors” are actually chains of many smaller behaviors. For example, imagine a child washing her hands. You can think of this as one behavior. Or, you can think of it as many smaller behaviors that are connected by changes in the environment.

For example, the child first reaches for the faucet handle, grasps it, and twists. As a result, water begins to flow in the sink. Then, she moves her hands forward and down until she feels the water on them. Next, she reaches for the bar of soap, picks it up, and rubs it between her hands until she sees lather forming.

It’s worth noting that these three “steps” could actually be broken into even smaller units of responses and environmental changes.

If you can see the smaller sequences of actions that make up a larger behavior, this will help you when you are planning and teaching new behaviors. It will help you see approximations that you can reinforce and may help you troubleshoot when you get stuck.

Thinking in terms of chains also can help you when you are working on maintenance. Rather than “fading out” reinforcement, you can think in terms of building longer and longer chains of behavior.

Part 13: Cues can be conditioned reinforcers

Chapter 15 of Dr. Keller's book is called "Secondary Positive Reinforcement." He spends the first part of the chapter explaining the connection between conditioned reinforcers and chains.

Chains hold together because, as Dr. Keller explains, "every discriminative stimulus in a chain becomes a secondary reinforcer for the response that produces it."

Think again about the girl, from Part 12, who was washing her hands. At the beginning of the chain, the girl twists the handle. As a result, water flows into the sink. Next, she sticks her hands in the water.

Seeing and hearing the water serves as a conditioned reinforcer for twisting the handle. It tells the girl she has successfully completed this behavior.

At the same time, seeing the water tells the girl she can move to the next behavior in the chain. That is, it acts as a discriminative stimulus for the girl to put her hands under the water.

It is important to be careful when you present a discriminative stimulus or cue. If you consistently give a cue after an unwanted behavior, you may accidentally create an unwanted chain.

However, you can also use cues to your advantage. You can present a cue at a "reinforceable moment." The cue will signal the next behavior, but it also will reinforce what the learner was doing when the cue was given.

Part 14: Places to avoid

In the final pages of the book, Dr. Keller describes an experiment involving shock.

On the first day, a rat is placed in a white box that is connected to a black box. Shocks are delivered through the floor in the white box. Before long, the rat learns to escape to the black box. Every time the experimenter returns him to the white box, the rat runs quickly to the black box.

On the second day, the rat is again placed in the white box. This time, the door to the black box is closed, but the shocks are turned off. Even though no shocks are given, “urination, defecation, trembling, rapid breathing, and other classical indicators of fear are all to be seen in the rat’s behavior.” After an hour there is some improvement, but the rat still “continues to cower and quiver in obvious distress.”

On the third day, the shocks still remain off. This time, however, the door to the black box is open again. When the rat is placed in the white box, he quickly runs to the black box. This is repeated 60 times. By the end, the rat’s speed has slowed some. However, he still immediately leaves the white box for the comfort of the black box. Dr. Keller writes: “The white room, in spite of the absence of shock within it, is still something to be run away from. Its ‘removal’ is rewarding.”

These effects do not apply just to rats. Later in the chapter, Dr. Keller comments, “Human beings act like this on occasion. The place of emotional upset becomes, for them, a place never willingly to be reentered; they develop a ‘phobia' with respect to it and, through generalization, with respect to places that are like it.”

Shocks and other aversive stimuli can be used to change behavior. However, they can have significant side effects, even after they are no longer used. The place or person who is associated with the unwanted event becomes something to be avoided. In addition, the emotional effects can be long lasting.