

Contemporary Study of Pavlovian Conditioning

Robert A. Rescorla
University of Pennsylvania

Pavlov's first report on conditioning emphasized its role in allowing the animal to adjust to its environment. Contemporary theories have seen this adjustment in terms of developing accurate knowledge of the environment. Three aspects of that thinking are explored: how the animal acquires initial knowledge, how it changes its knowledge when conditions of the world change, and how it makes use of multiple knowledge representations.

Keywords: error correction, overexpectation, superconditioning, extinction, modulation, Rescorla-Wagner model

El primer informe de Pavlov sobre el condicionamiento subrayaba su función para la adaptación al ambiente. Las teorías contemporáneas han contemplado esta adaptación en términos del desarrollo de un conocimiento adecuado del medio ambiente. Se exploran tres aspectos de esa línea de pensamiento: cómo el animal adquiere el conocimiento inicial, cómo modifica su conocimiento cuando cambian las condiciones del medio y cómo hace uso de representaciones múltiples del conocimiento.

Palabras clave: corrección de errores, sobreespectativa, supercondicionamiento, extinción, modulación, modelo de Rescorla-Wagner

This research was supported by grants from the National Science Foundation and the National Institute of Health.

Correspondence concerning this article should be addressed to Robert A. Rescorla, Department of Psychology, University of Pennsylvania, 3815 Walnut Street, Philadelphia, PA, 19104.

A central theme in Pavlov's first public discussion of conditioning, which we are commemorating here, was that conditioning allows the organism to adapt to its world. In speaking of the unconditioned reaction he said, "Observing the normal activity of these glands, it is impossible not to be struck with the high degree in which they are adapted to their work. Give the animal some dry, hard food, and there is a great flow of saliva, but with watery food there is much less" (Pavlov, 1928, p. 48). He went on to emphasize the same property of conditioned reactions, "... these substances ... act exactly the same upon these glands ... when they are a certain distance from the dog. Dry food, even from a distance, produces much saliva; moist food, only a little" (Pavlov, 1928, p. 49). Thus, for Pavlov, one important purpose of conditioning was to allow the animal's behavior to reflect accurately the world in which it had to function.

For Pavlov, this reflection of the world occurred at the level of behavior itself. Hence his emphasis on what we would now call the issue of the similarity between the CR and the UR. Certainly discussions of adaptation at this level have been a continuing theme of students, especially in the middle part of the last century. However, I think it is fair to say that for many of us, conditioning represents adaptation at a more cognitive level. For many contemporary students of learning, conditioning provides a way in which the animal adjusts its knowledge so as to reflect the state of the world. Influenced not only by the physiology of Pavlov but also by the associationism of British philosophers, many see conditioning as a means by which organisms build up knowledge of the structure of the world. Like Pavlov, they see the organism as changing so as to reflect accurately the world; but unlike Pavlov they have seen that reflection less directly at the level of behavior and more at the level of knowledge representation.

This article touches on three aspects of this contemporary view of conditioning as generating an accurate knowledge representation: how the animal constructs that knowledge initially, how it modifies that knowledge when its experience with the world changes, and how it can maintain and use multiple knowledge representations as conditions demand. This discussion will illustrate something of the current state of work in conditioning, naturally enough using examples from work done in my own lab.

Constructing Initial Representations

Let me begin with the building of initial knowledge representations. Perhaps the key idea in our contemporary idea of how conditioning develops is that of error-correction. It is common to view the organism as comparing its current knowledge of the world with the new experiences that it has, evaluating the degree to which its knowledge is an accurate reflection of that experience. To the degree that the

knowledge of the world and experience do not agree, the organism detects an error. It then uses that error to correct its knowledge. The idea is that the animal builds an accurate understanding of the world by continually adjusting its knowledge in the light of its current experience.

A primitive, but still serviceable, version of such error correction was described by Rescorla and Wagner (1972) a quarter of a century ago. Although oversimplified, and even demonstrably incorrect in many ways, it has continued to provide a rough description of conditioning and to serve as the basis for a variety of more elaborate models. Figure 1 shows a version of this theory.

This theory follows Pavlov in attempting to characterize knowledge representation in terms of the formation of associations between a conditioned stimulus (CS) and an unconditioned stimulus (US), when the two are paired. As shown in the top equation in Figure 1, this model describes the changes in the strength of those associations that result

$$A \rightarrow US$$

$$\Delta V_A = k (\lambda - V_A)$$

$$AB \rightarrow US$$

$$\Delta V_A = k (\lambda - V_{AB})$$

$$\Delta V_B = k (\lambda - V_{AB})$$

$$V_{AB} = V_A + V_B$$

$$AB \rightarrow 0$$

$$\Delta V_A = k (0 - V_{AB})$$

$$\Delta V_B = k (0 - V_{AB})$$

$$V_{AB} = V_A + V_B$$

Figure 1. Error-correction equations for three kinds of conditioning trials, according to the Rescorla-Wagner model. In the first case, stimulus A is presented alone and followed by the US. The associative strength of A (V_A) is compared to the asymptote that the US will support (λ). Some percentage (k) of their difference governs changes in the associative strength of A (ΔV_A). In the second, case stimuli A and B are both followed by the US. In the third case, the AB compound is nonreinforced.

from experiences such as the following of stimulus A by a US. According to the theory, the change in the associative strength of a signal A (V_A) is characterized by $\Delta V_A = k(\lambda - V_A)$, where λ is the asymptotic associative strength that the US is capable of producing. That is, on each trial, the organism calculates the discrepancy between its current associative strength, V_A , and the appropriate strength, λ , and then adjusts its current strength by some constant percentage of this discrepancy, or error. As can readily be seen, when the organism first experiences a US after a novel CS, the error will be substantial, because there is no existing associative strength; that is, initially V_A will be zero. As a result, the trial will produce a marked adjustment in V_A in order to bring it into line with λ . Repeated experiences of this sort will result in smaller and smaller error terms, with the result that the V_A will eventually become quite close to λ . Moreover, V_A will approach λ in a way that is characteristic of many conditioning processes, by a negatively accelerated path. That is, with experience, V_A provides an increasingly accurate representation.

Many error-correction models also apply this mechanism to situations that are somewhat more complicated, involving two alternative signals of the US. The second set of equations in Figure 1 displays the Rescorla-Wagner version of this application. According to such models, the animal adopts the same error detection and correction procedure for multiple signals, resulting in equations of the form $\Delta V_A = k(\lambda - V_{AB})$, and $\Delta V_B = k(\lambda - V_{AB})$, where $V_{AB} = V_A + V_B$. That is, the animal is seen as evaluating the error between the total strength of the AB stimulus (V_{AB}) and λ , allowing each stimulus to use that joint error term as a basis for correcting its strength. That correction will continue until the two stimuli together yield an accurate representation, a near- λ total associative strength.

Error signal models of this general sort are widely adopted not only within conditioning but more broadly in neural network and connectionist models, as well as in engineering applications. They provide a local learning rule that turns out to do a surprisingly good job of generating global representations of the world.

One early success of such models was to provide an account of some of the phenomena that Pavlov himself reported. For instance, Pavlov noted that the amount of conditioning that he observed to a stimulus, A, when it was paired with a US, would be diminished if another stimulus, B, were jointly presented with A. Pavlov called this phenomenon "overshadowing." It is as though the stimuli compete with each other for developing associative strength. Although there are multiple explanations for overshadowing, error correction models give one natural account: If an animal receives A+ trials, then the error must be corrected entirely by changes in V_A . However, if the animal receives AB+ trials, then changes in both V_A and V_B contribute to that correction. As a result, simple A+ trials will mean that V_A approaches λ , whereas AB+ trials will result in V_A approaching a value less than λ .

An extremely influential application of such models came in the phenomenon of blocking, first observed by Kamin (1968) but anticipated in the writings of Konorski (1948). In a blocking experiment, AB+ trials are preceded by separate A+ trials. Under these circumstances the overshadowing of B by A can be enhanced substantially, sometimes to the point where prior conditioning of A entirely prevents any conditioning of B. Error correction models naturally account for this finding because the prior conditioning of A results in a V_A that is close to λ , yielding only a small error term on the AB+ trials. Because the error is small on these trials, B undergoes little associative change; i.e., B is blocked despite its being followed by a US.

The important finding of conditioned inhibition, first reported by Pavlov, is viewed by error correction models as highly related to blocking. Pavlov noted that if he intermixed two kinds of conditioning trials, A followed by the US and AB followed by no US, not only did A become capable of evoking salivation, but B became capable of inhibiting that salivation. Any B stimulus treated in this way has two properties that led Pavlov to claim it was inhibitory: B can diminish the salivary response that other excitatory stimuli otherwise control, and B is difficult to convert into an exciter if it is then used to signal the US.

The third set of equations in Figure 1 shows the account that error correction models give of this learning. The separate A+ trials yield a V_A that is close to the λ appropriate to its reinforcement. However, on the AB- trials no reinforcement is given, with the result that the asymptote appropriate to such trials is quite low, perhaps zero. That means that once A has become conditioned, the error term appropriate on the AB- trials is $(0 - V_{AB})$ or $-\lambda$. That is to say that the AB- trials should produce negative change in both V_A and V_B . Since this is B's only learning experience, it develops an associative strength that is below zero, negative. Indeed, one can show that asymptotically V_B would approach $-\lambda$, whereas V_A would approach $+\lambda$. The B stimulus with negative associative strength would naturally subtract from those with positive strength and also be hard to train to have its own positive strength.

Error correction models explain these Pavlovian phenomena because they see associative change as governed not simply by the occurrence of US but rather by the discrepancy between the US that occurs and that which was expected on the basis of the associative strength of the antecedent stimuli. That is, error correction models are sensitive not to the absolute US but to the relation between that US and the associative strength of antecedent stimuli. That same feature has allowed error-correction models to predict new and surprising phenomena that had not been observed by Pavlov. To emphasize the power of such models to make novel predictions, I want to describe two such phenomena whose discovery was foretold by error correction models: overexpectation and superconditioning.

First, overexpectation, a case in which error correction models predict that Pavlovian pairing of a CS with a US

will reduce, rather than augment, associative strength. An experiment conducted in my laboratory by Lattal and Nakajima (1998) provides a nice illustration. They used a popular contemporary Pavlovian preparation: the magazine approach. In this preparation, diffuse auditory and visual stimuli signal the availability of food in a particular location, the magazine. Rats given this Pavlovian treatment approach and place their heads in the magazine in anticipation of the food, something that can be measured by the interruption of a photocell. The top portion of Figure 2 shows the design of one of their experiments. Initially they conducted conditioning with a light, L, and two auditory stimuli, A and B. Then they presented L in compound with one of the auditory stimuli, B, and followed that compound by the same pellet that had followed each stimulus individually; the A was a nonpresented control. Finally, they tested responding to both A and B. Consider the expectations of such a manipulation from the point of view of an error correction model. The original training should result in both B and L having associative strengths near λ . When they are presented in compound, however, their total associative strength (V_{LB}) should greatly exceed λ ; indeed, the compound should have a value of 2λ . However, the error term applicable to a reinforced LB trial is $(\lambda - V_{LB})$. Since $V_{LB} (= 2\lambda)$ is initially greater than λ , the expectation of such models is that reinforcing the compound should not only fail to increase the strength of the B and L elements, it

should actually diminish them. In effect, the reinforcer is overexpected, even though physically it is still highly potent.

The bottom half of Figure 2 shows the results of their experiment. To the left are the final levels of responding during a pre-CS period and during the two auditory stimuli, A and B. Both stimuli augmented responding relative to the pre-CS rate, and to about the same extent. To the right are the results of presenting A and B again after the imposition of the LB+ trials. It is clear that although the control A continued to show substantial responding, the level of responding to the treated B was greatly diminished. That is, overexpectation was observed as predicted; reinforcing the B stimulus in the presence of another exciter reversed the normal effects of reinforcement. This result was confirmed in a variety of related experiments. It is a surprising, and relatively unique, prediction from error correction models. It highlights the fact that the impact of a US depends not on its absolute physical power, but rather on the error term it generates. That is, error correction allows the animal to accurately adjust its knowledge to the current conditions of the world.

The example of overexpectation shows how one can exploit an error correction mechanism to arrange for the reinforcing power of a US to be so diminished as to lead to associative decrements. The second example, superconditioning, is the converse, in which one uses an error correction mechanism to so enhance the reinforcing power of a US that it produces an extra measure of conditioning. A recent experiment in our laboratory provides an especially striking example (Rescorla, 2003). This experiment used an autoshaping procedure in pigeons. In this procedure localized visual stimuli are used to signal grain reinforcement; as conditioning proceeds, the pigeon comes to direct pecking at the visual stimulus. The design of this particular experiment is shown at the top of Figure 3. The birds first received a conditioning procedure of the form: A++, B++, AN-, BN-, in which A and B were visual stimuli, N was a diffuse auditory stimulus, and ++ indicates the immediate presentation of 5-sec of grain. As noted above, such a procedure will result in A and B becoming strong exciters and N becoming an inhibitor. Next, the animals were given NA and B trials, both followed by a reinforcer (+) that was weaker than the immediate grain originally used with A and B. That weaker reinforcer consisted of the delayed presentation of the 5-sec of grain; the interval between the stimuli and grain was filled by another visual stimulus. This delayed gap-filling procedure has been documented to yield stable, but reduced, associative strength, compared with the use of immediate grain.

The question of interest was what impact this reduced reinforcer would have on the associative strengths of A and B. The predictions of an error correction model are clear: B should lose strength because it is being followed by a diminished reinforcer. The asymptote that + can support is less than that which ++ can support, so the error term for




Figure 2. Design and results of an overexpectation experiment. Initially, three stimuli, L, A, and B, are each followed by a reinforcer (+). Then the LB compound is followed by the same reinforcer and the individual A and B stimuli are tested (after Lattal & Nakajima, 1998).

B is negative. However, A should gain strength from that same reinforcer because A is reinforced in the presence of the inhibitory N. Because the NA compound initially commands little strength it is quite discrepant from the asymptote that even the weaker reinforcer should produce. As a result, despite the fact that A is followed by a weaker reinforcer than it had received in the past, it should undergo increases in conditioning beyond its original strength.

The bottom half of Figure 3 shows the results of this experiment. The left-hand side shows the changes over the course of the treatment in which B and NA are both followed by the weaker reinforcer. It is clear that responding to B decreased, whereas that to NA increased, as expected. To the right are the results of the test in which A and B were both tested. Those results show that the reinforcement of A in the context of N had indeed augmented the power of the reinforcer, so that responding was greater to A than to B. That is, superconditioning had occurred. More generally, the same reinforcer could produce either decreases or increases in associative strength, depending on the error term that it generates.

These examples are intended to illustrate the considerable power of an error correction model. It has the ability to explain old phenomena, generate new phenomena, and provide an overall context in which to view many conditioning results. It is this that has made it a powerful




Figure 3. Design and results of a superconditioning experiment. Initially A and B are followed by immediate grain (++) when presented alone, but not when presented in compound with a noise (N). Then NA and B are both followed by delayed grain (+) and the individual A and B stimuli are tested (after Rescorla, 2003).

tool as the major learning mechanism used in neural network models. It implements the motivation behind Pavlov's ideas by providing a way to allow the animal to adjust accurately to the world.

However, one should not think that the analysis of error correction models is in any sense complete. Indeed, some recent experiments in our laboratory have uncovered some potentially quite damaging flaws with one of the basic assumptions of error correction models (e.g., Rescorla, 2000). Such models have the strong implication that when A and B are reinforced together in compound, they are subjected to exactly the same error; the result is that they should show exactly the same associative change. Indeed, it is that feature of error correction mechanisms that makes neural network models so powerful. However, some recent data suggest that under some circumstances this implication may not be correct (e.g. Rescorla, 2000). It remains to be seen whether or not this will prove disastrous for such models.

It is important to realize that error correction models like this are not simply convenient frameworks within which to think about conditioning at the behavioral level. Of course, Pavlov saw the promise of conditioning as a way to investigate the brain. Contemporary error correction models have begun to pay off on that promise by guiding neurobiological research. For instance, recent work by Thompson and his collaborators (e.g., Kim, Krupa, & Thompson, 1998) has provided a detailed understanding of the neural circuitry responsible for eyelid conditioning in the rabbit. An essential part of that understanding has been the identification of a neural basis for the computation of the error term on which behavioral models depend. Thompson's laboratory has found evidence, for instance, of an inhibitory feedback loop acting on the pathway along which US information travels, a loop that allows a CS that has a strong excitatory association to diminish the impact of that US. In effect, as anticipated by behavioral models, an expected US has a diminished impact. Similarly, a great deal of evidence points to the involvement of dopamine neurons in the reinforcement process. Recent evidence reported by Schultz and his collaborators (e.g., Schultz, 1998) indicates that the output of these neurons is importantly governed by the kind of discrepancy that error term models identify. So, we are beginning to make serious progress in understanding the neural basis of such error correction mechanisms.

Changing Representations

Although understanding the initial development of knowledge about the world is important, it could be argued that it is even more important to understand how organisms adjust that knowledge when the world changes. In his first lecture on conditioning, Pavlov reported the earliest and most well studied of these adjustments to changed conditions, experimental extinction. Just as the following of a CS by a

US leads to one adjustment, the development of responding, so subsequently presenting the CS alone leads to another adjustment, the loss of responding. From the outset, Pavlov identified the importance of extinction and was fascinated by it. He describes (Pavlov, 1928, p 53) that "If you only show the dog food, and repeat this several times, at each repetition you get a weaker result, and finally no reaction whatever." Indeed, he used this observation as one of his primary arguments that the reaction to the sight of food was a conditioned, rather than an unconditioned, response.

In the continuation of that same section, Pavlov identified the first piece of evidence that led him to the important conclusion that extinction did not involve the erasure of original learning. He went on to say, "... if meat powder, having been repeatedly brought before the dog, fails to produce a flow of saliva, we may again make it active by giving it to the dog to eat [after showing it]." Pavlov suggested that reestablishment of the response was exceedingly rapid, implying preservation of some of the original learning. As Pavlov pursued the study of conditioning he was to come to describe several other phenomena that formed the principle factual basis on which he drew this conclusion: spontaneous recovery and disinhibition. These led him to suggest that the changed conditions that produced extinction did so not by changing the original knowledge but rather by superimposing some alternative knowledge. For Pavlov, conditioned inhibition provided this alternative knowledge.

Contemporary studies of extinction have stayed very much within the framework outlined by Pavlov, but have expanded it in two ways: (a) They have elaborated on the phenomena that support the conclusion that extinction does not remove original knowledge, and (b) They have provided new techniques that encourage the stronger conclusion of Pavlov that extinction left that original knowledge fully intact. Each point merits further comment.

Several techniques have been found, in addition to spontaneous recovery and disinhibition, which will restore responding to an extinguished stimulus, even without restoring the relation of the stimulus to the US. One of these is simple reexposure to the US itself in the absence of the CS. This US exposure will reinstate responding to the CS, often to a substantial degree (e.g., Bouton, 1984; Delamater, 1997; Rescorla & Heth, 1972). Actually, Pavlov himself reported this phenomenon but he made little of it. Another technique that restores responding involves moving the CS to a new context, different from that in which original training or extinction had taken place. Bouton (e.g., 1991) has reported many experiments in which such a procedure restores responding to the CS. Finally, one can reveal the presence of the original knowledge by testing the extinguished CS not for responding to itself, but for its ability to add to the responding of a different, nonextinguished, stimulus. Although Reberg (1972) reported this result many years ago, it has remained under-appreciated.

A recent example from our lab illustrates Reberg's finding. This experiment used the magazine approach in rats, giving initial conditioning of one stimulus, B and presenting another, C, without reinforcement. Then both B and C were presented repeatedly without reinforcement. As is clear in Figure 4, initially B showed substantial responding, but this declined over the course of 28 days of extinction. Indeed, after about the 10th day of 8 nonreinforced trials per day there was no difference in responding to B and C. At this point, spontaneous recovery from day to day had also disappeared. However, despite the similarity of performance to B and C, it is relatively easy to show that B has not lost its original conditioning. After the 28 days of extinction, B and C were both tested in compound with another stimulus, A, which had been conditioned to a moderate level. Figure 5 shows the results of that test. The moderately-trained A stimulus augmented responding relative to the pre-stimulus rate. However, when B was added to A, there was a great deal more responding; despite producing little response of its own, B greatly augmented the response to A. Moreover, this augmentation was indeed attributable to B's history of training and extinction, because the comparison stimulus, C, lacking any conditioning history, failed to have this effect. A summation test of this sort is apparently a powerful technique for revealing that even after extended extinction something of the original learning remains.

However, Pavlov wanted to draw a conclusion stronger than such experiments permit. These procedures can only provide evidence that extinction left *something* of the original training intact. Pavlov, and others, have wished to conclude that extinction has left *everything* intact. It turns out that several more contemporary techniques have been developed that do provide relevant evidence. Typical results from one such technique are shown in Figure 6.




Figure 4. Extended extinction of a previously-reinforced (+) B stimulus, compared with presentation of a never-reinforced C stimulus.

That figure shows the results of a magazine approach experiment (Rescorla, 1996) in which rats had been exposed to 4 stimuli, two visual and two auditory. One stimulus within each modality had been paired with pellets and one with liquid sucrose. Subsequently, the animals were subjected to a devaluation procedure in which one of those outcomes was itself paired with LiCl, a substance that makes the animal ill. Previous experiments had shown that this treatment would dramatically reduce the value of the outcome so paired. More importantly, we knew it would reduce responding to any stimulus paired with that outcome. This is evidence for the specificity of the CS-US information in original learning. The data on the left-hand side of Figure 6 show an example of the impact of this devaluation. There is much less responding to the stimulus whose outcome had been devalued (DEV). But it is the data on the right-hand side that are of interest here. These come from a pair of stimuli that had also been paired with the different outcomes but were then subjected to extensive extinction before the devaluation. The point to notice is that the stimuli given extinction were as sensitive to the devaluation procedure as those not extinguished. This suggests that the information from acquisition that allowed the differential effect of devaluation remained fully intact despite extinction. Thus, modern experiments have sharpened Pavlov’s argument that extinction leaves original training intact. That is, they support

his view of what extinction is not: the removal of original learning.

Where contemporary work has been less successful is in identifying what extinction is. Pavlov expressed his own puzzlement in this original talk, “Among the above-mentioned facts there are some which at present are with difficulty explained from the physiological point of view. For example, why does a conditioned reflex, if repeated, lose its activity?” The best he could muster was “Probably we have to deal here with quite peculiar conditions for the excitation which travels along the temporary centripetal paths” (Pavlov, 1928, p. 57). Although some suggestions have been made (e.g., Rescorla, 2001), we are still a long way from a full understanding of extinction.

Use of Multiple Representations

If one supposes that changed conditions result both in the preservation of original knowledge and in the acquisition of new knowledge reflecting the new conditions, then an important question naturally arises: How do these two pieces of knowledge interact? This leads to the third point about adaptation to the environment. Contemporary work has found evidence that when two kinds of knowledge about the same CS exist side by side, other stimuli often play the role of selecting among them to control actual performance. Thus, if a CS is paired with a US in some circumstances but not in others, those circumstances can control the degree to which the CS evokes a CR. That is, other stimuli can gain conditional control over which association is expressed.




Figure 5. Summation testing of responding to the B and C stimuli whose extinction was shown in Figure 4. Stimulus A received moderate training prior to testing of A alone and in compound with B and C. Responding is also shown in the period just prior to stimulus presentation (Pre).

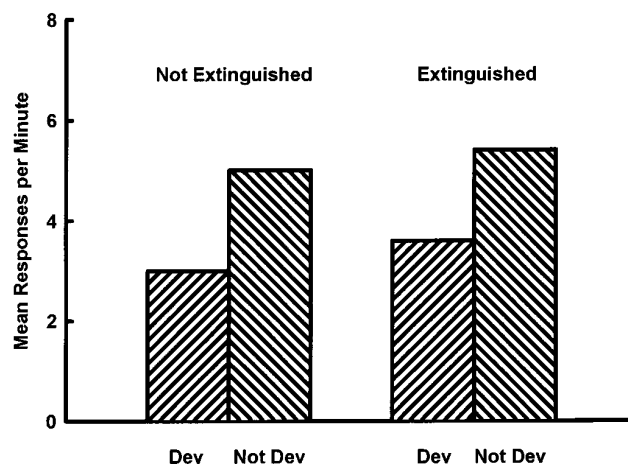


Figure 6. Responding to 4 stimuli, all of which had been given Pavlovian conditioning. Prior to the test, the outcome used with two of the stimuli had been devalued (DEV) by pairing with LiCl. Two of the stimuli had been extinguished prior to this devaluation (after Rescorla, 1966).

An example of this kind of control comes in contemporary studies of so-called occasion setting — in which one stimulus sets the occasion on which another will signal the US. Figure 7 shows two examples of such occasion setting in an autoshaping experiment conducted in my laboratory. In that experiment, pigeons received the illumination of two key lights, one red and the other green (A and B). Both key lights were sometimes reinforced with food and sometimes not followed by food, on a random half of their presentations. However, for each stimulus a diffuse white noise or houselight (X and Y) signaled when it would be followed by food. Key light A was reinforced except when accompanied by the diffuse stimulus, X. Key light B was nonreinforced except when accompanied by the diffuse stimulus, Y. Both X and Y came to have conditional control over responding to A and B. The right panel shows that the B-, YB+ positive occasion setting procedure resulted in little responding to B when presented alone but substantial pecking at B when it was accompanied by Y. The left panel shows the opposite pattern for X and A when an A+, XA- negative occasion setting procedure was used: Stimulus A controlled substantial responding except when accompanied by X. In each case, A and B have been associated with both reinforcement and nonreinforcement, and in each case X and Y signal whether A or B will be reinforced.

This conditional signaling power of X and Y turns out not to be confined to the A and B stimuli. After this training, another key light, C, received both training and then partial extinction so as to endow it with both kinds of association. The question was whether X and Y would select from these two associations of C. Figure 8 shows responding to that trained and extinguished stimulus under four different conditions: When it was presented alone, when it was accompanied by the negative occasion setting X, when it was accompanied by the positive occasion setting Y, and when it was accompanied by both X and Y. It is clear that the yellow stimulus alone evoked a moderate level of performance, as expected because of its history of training and partial extinction. But X and Y modulated responding to it in different directions, presumably because they selected either the reinforcement or nonreinforcement association for expression. Of equal interest, when X and Y were both present, they tended to cancel each other out, again yielding an intermediate level of responding.

Experiments such as these suggest that stimuli signal not simply other events but relations among other events. The animal can maintain multiple associations between a CS and US and use other stimuli as modulators that signal when to use each of those associations. This is an exciting new area of research that is receiving a great deal of

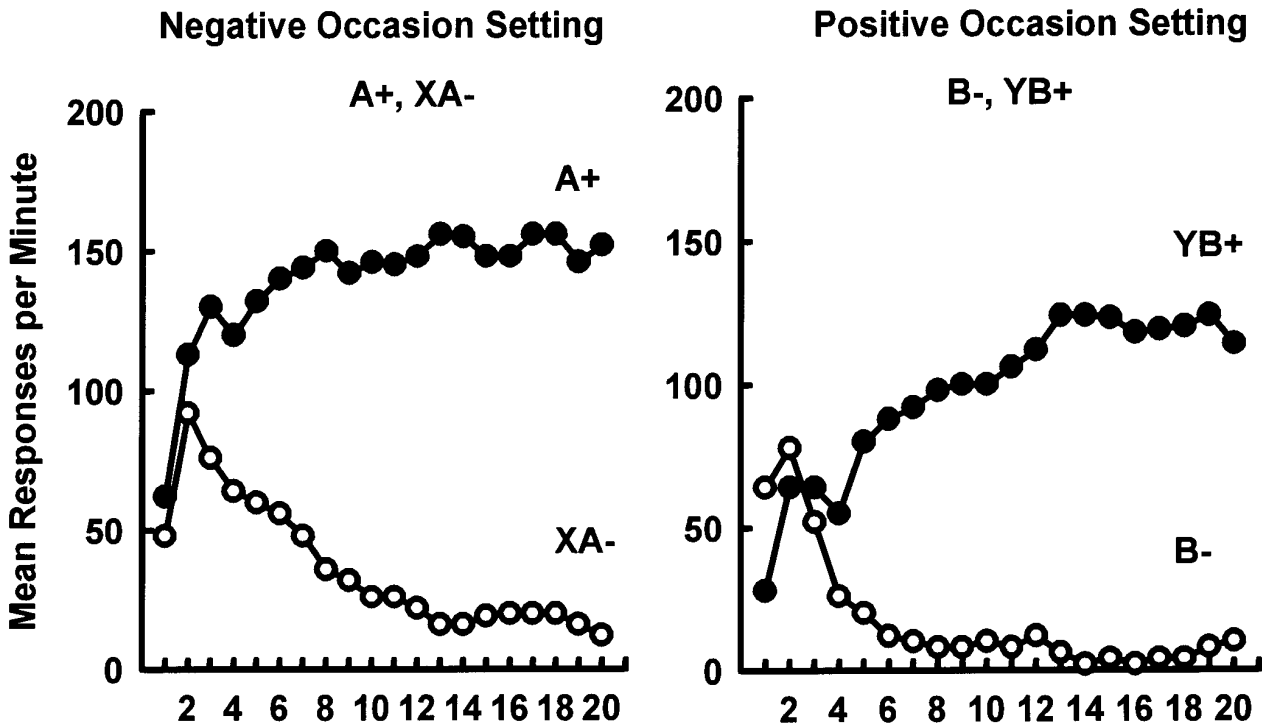


Figure 7. Two examples of occasion setting from an autoshaping situation. In negative occasion setting, key light A was followed by grain (+) except when accompanied by diffuse stimulus X. In positive occasion setting, key light B was nonreinforced except when accompanied by diffuse stimulus Y.

empirical and theoretical attention (see Schmajuk & Holland, 1998). It has implications not only for simple conditioning but for instrumental training, extinction, and indeed for associative theory in general. It illustrates a further way in which Pavlov's notion, that the role of conditioning is to help the animal adjust properly to the conditions of the world, is being played out in contemporary thinking.

Error correction at the level of knowledge

As noted earlier, Pavlov saw conditioning as allowing the animal to adjust to the world at the behavioral level. He illustrated this with the fact that the CS comes to evoke an anticipatory salivary response that is quite like that which the US will evoke when it is presented. But contemporary thinkers view conditioning as allowing the animal to match not its overt behavior but rather its internal knowledge to the world. One reason for this difference is that contemporary theorists sharply distinguish between learning and performance, frequently emphasizing that organisms may have a rich knowledge of the world that is quite imperfectly reflected in any particular behavioral measure.

There are many examples in which one is inclined to see conditioning as endowing the animal with accurate knowledge that does not immediately exhibit itself in its behavior. All of those cases encourage the contemporary

view that the animal is best viewed as adjusting its knowledge, not just its behavior. Figure 9 shows one such example from an experiment conducted in my lab some years ago. In this experiment pigeon subjects were given Pavlovian pairings of two different CSs with the same grain US. One of these CSs was a localized visual stimulus, an orange disk; the other was a diffuse auditory stimulus, an 1800 Hz tone. As a result of these pairings, the orange key light came to evoke directed pecking, a response quite like that evoked by the grain US. However, the tone showed no evidence of producing directed pecking, even though it received the same number of pairings with the US. The left-hand side of the figure shows the stark contrast in performance to these two stimuli. Here is plotted the percentage of trials on which each stimulus evoked pecking. Of course, an obvious possibility is that the animal learned about the key light but not about the tone, with its responding showing this learning at the behavioral level.

But it turns out that one can easily demonstrate that the animal has learned about the tone-grain pairings as well, despite the fact that this does not show in its pecking behavior. This can be shown by taking some measure of the animal's knowledge other than the ability of the tone to evoke UR-like behavior. For instance, one can use both the orange key light and the tone as reinforcers for other key lights in the manner of a second-order conditioning design. As Pavlov showed, a stimulus that has developed an

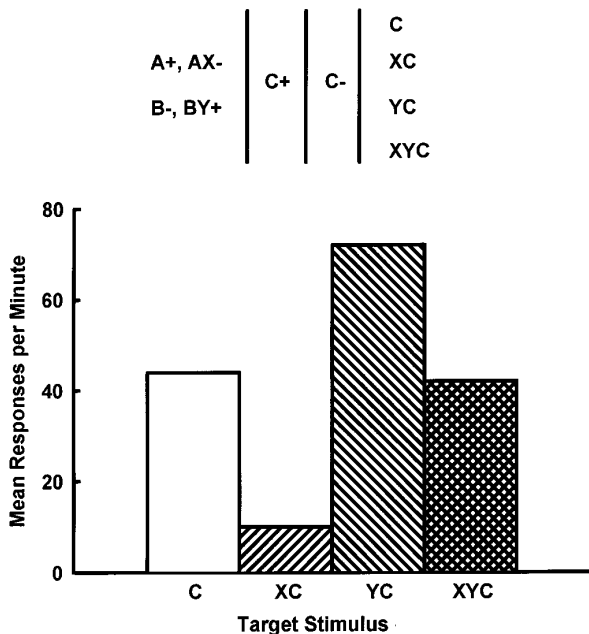


Figure 8. Transfer test of the X and Y stimuli whose training is shown in Figure 7. Stimulus C had received initial training and incomplete extinction. It was then tested alone, in combination with X and Y, and in combination with X and Y together.

| Acq | Test |
|-----|-------|
| V+ | X-->V |
| A+ | Y-->A |

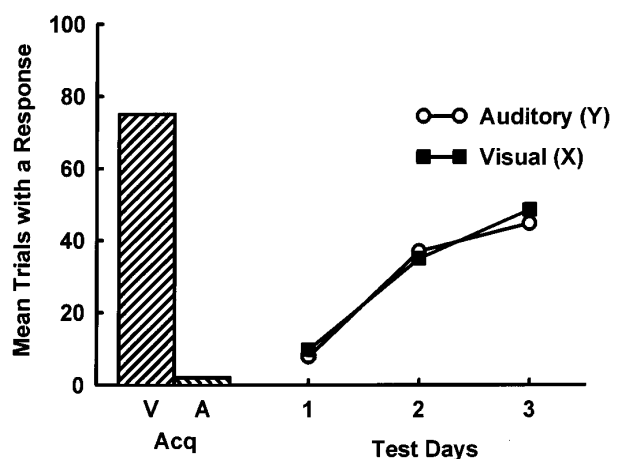


Figure 9. The results of pairing a visual (V) and an auditory (A) stimulus with grain in pigeon subjects. Despite substantial differences in their likelihood of evoking directed pecking, the two stimuli were quite similar in their ability to second-order condition two other key light stimuli (X and Y).

association with a US will not only evoke responses, it will also become capable of serving in lieu of the US for purposes of training other CSs. Indeed, we have frequently found the ability of stimulus A to condition stimulus B to be the best index of A's own association with the US. The right-hand side of Figure 9 shows the results of a second-order conditioning test in which the tone and light were each used as a reinforcer for other key lights. That figure shows that both of these stimuli were fully capable of establishing second-order conditioning.

What these data suggest is that the conditioning procedure has allowed the animal to build an accurate representation of the relation to grain of both the tone and the orange key light, but that the representations can play out in behavior quite differently. The visual stimulus showed the knowledge in pecking and the auditory stimulus did not. Indeed, one of the major advances that has occurred in thinking about conditioning is the realization that the way in which knowledge is exhibited is determined not only by the US, as Pavlov readily saw, but also by the CS (e.g., Holland, 1977). As illustrated here, different CSs, paired with the same US, come to control quite different learned responses. Yet both show their knowledge in the ability to produce second-order conditioning. Clearly, any one behavioral measure can give a highly misleading picture of the product of conditioning. It is data like these that have led many to emphasize conditioning as a means of the animal's adjusting its knowledge, rather than its behavior. This realization greatly expands the importance of Pavlov's discovery and brings it into more complete contact with the rest of psychology.

Conclusion

So, how would Pavlov feel about what has happened in the 100 years since his first announcement of conditioning? From his point of view, the good news is that his ideas have had a tremendous impact on psychology, arguably a greater impact than those of any other scientist. Conditioning has come to be seen as the basic learning process underlying a vast array of phenomena. Most organisms capable of benefiting from experience have proven to use conditioning as a major means by which they do so. Moreover, conditioning applies to a wide range of learning situations, from simple associative ones, to perceptual situations, to judgments of causality. Indeed, we have realized one of Pavlov's own claims—to have placed associationism on a firm empirical footing. The ideas that blossomed in the area of conditioning have proven central to the broad range of network learning models that have been successfully applied to many complex processes. They have proven basic to many applied clinical phenomena, something that Pavlov himself foresaw.

But from Pavlov's point of view the bad news is that this tremendous impact has been principally among

psychologists. Moreover, much of the impact has been because psychologists have used conditioning in a more abstract, conceptual way, rather than as a direct influence on behavior. However, even the bad news has a silver lining. In recent years psychologists have brought conditioning back to physiologists to help neuroscientists guide the study of learning at the neural level.

Clearly, Pavlov would only partly recognize the child to whom he gave birth. Like many parents he would surely have wished that the child had matured somewhat differently. But I am confident that he would also be proud of the independent and highly influential adult that we see in conditioning today.

References

- Bouton, M. E. (1984). Differential control by context in the inflation and reinstatement paradigms. *Journal of Experimental Psychology: Animal Behavior Processes*, 10, 56-74.
- Bouton, M. E. (1991). Context and retrieval in extinction and in other examples of interference in simple associative learning. In L. W. Dachowski & C. F. Flaherty (Eds.), *Current topics in animal learning: Brain, emotion, and cognition* (pp. 25-53). Hillsdale, NJ: Erlbaum.
- Delamater, A. R. (1997). Selective reinstatement of stimulus-outcome associations. *Animal Learning and Behavior*, 25, 400-412.
- Holland, P. C. (1977). Conditioned stimulus as a determinant of the form of the Pavlovian conditioned response. *Journal of Experimental Psychology: Animal Behavior Processes*, 3, 77-104.
- Kamin, L. J. (1968). Attention-like processes in classical conditioning. In M. R. Jones (Ed.), *Miami symposium on the prediction of behavior: Aversive stimuli* (pp. 9-32). Coral Gables, FL: University of Miami Press.
- Kim, J. J., Krupa, D. J., & Thompson, R. F. (1998). Inhibitory cerebello-olivary projections and blocking effect in classical conditioning. *Science*, 279, 570-573.
- Konorski, J. (1948). *Conditioned reflexes and neuron organization*. Cambridge, UK: Cambridge University Press.
- Lattal, K. M., & Nakajima, S. (1998). Overexpectation in appetitive Pavlovian and instrumental conditioning. *Animal Learning and Behavior*, 26, 351-360.
- Pavlov, I. P. (1928). *Lectures on conditioned reflexes*. New York: International.
- Reberg, D. (1972). Compound tests for excitation in early acquisition and after prolonged extinction of conditioned suppression. *Learning and Motivation*, 3, 246-258.
- Rescorla, R. A. (1996). Preservation of Pavlovian associations through extinction. *Quarterly Journal of Experimental Psychology*, 49B, 245-258.
- Rescorla, R. A. (2000). Associative changes in excitors and inhibitors differ when they are conditioned in compound. *Journal of Experimental Psychology: Animal Behavior Processes*, 26, 428-438.

- Rescorla, R. A. (2001). Experimental extinction. In R. R. Mowrer & S. B. Klein (Eds.), *Handbook of contemporary learning theories* (pp. 119-154). Mahwah, NJ: Erlbaum.
- Rescorla, R. A. (in press). Superconditioning from a reduced reinforcer. *Quarterly Journal of Experimental Psychology: Comparative and Physiological Psychology*.
- Rescorla, R. A., & Heth, C. D. (1975). Reinstatement of fear to an extinguished stimulus. *Journal of Experimental Psychology: Animal Behavior Processes*, 1, 88-96.
- Rescorla, R. A., & Wagner, A. R. (1972). A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In A. Black & W. F. Prokasy (Eds.), *Classical conditioning II* (pp. 64-99). New York: Appleton-Century-Crofts.
- Schmajuk, N. A., & Holland, P. C. (1998). *Occasion setting: Associative learning and cognition in animals*. Washington, DC: American Psychological Association.
- Schultz, W. (1998). Predictive reward signal of dopamine neurons. *Journal of Neurophysiology*, 80, 1-27.

Received June 20, 2003

Revision Received September 30, 2003

Accepted October 8, 2003