

Context and Behavioral Processes in Extinction

Mark E. Bouton

Department of Psychology, University of Vermont, Burlington, Vermont 05405, USA

This article provides a selective review and integration of the behavioral literature on Pavlovian extinction. The first part reviews evidence that extinction does not destroy the original learning, but instead generates new learning that is especially context-dependent. The second part examines insights provided by research on several related behavioral phenomena (the interference paradigms, conditioned inhibition, and inhibition despite reinforcement). The final part examines four potential causes of extinction: the discrimination of a new reinforcement rate, generalization decrement, response inhibition, and violation of a reinforcer expectation. The data are consistent with behavioral models that emphasize the role of generalization decrement and expectation violation, but would be more so if those models were expanded to better accommodate the finding that extinction involves a context-modulated form of inhibitory learning.

Extinction is a well-known and important behavioral phenomenon that allows the organism to adapt its behavior to a changing environment. Nonetheless, it is difficult to find a clear answer to the question of why repeated presentation of the conditional stimulus (CS) alone after conditioning (CS–unconditional stimulus [US] pairings) causes behavior to return to essentially zero. One idea, surprisingly common in models of learning and memory (see Rescorla and Wagner 1972; McClelland and Rumelhart 1985; McCloskey and Cohen 1989), is that extinction involves the destruction of what was originally learned. However, there is ample evidence that this is not true; much of the original learning survives extinction (see Rescorla 2001; Bouton 2002; Myers and Davis 2002; Delamater 2004). In this article, I will selectively review results and theory from the behavioral literature in an effort to understand what is learned in extinction, and what causes the organism to learn it.

The first part of the article introduces several extinction phenomena that any adequate theory of extinction will need to explain and accommodate. They suggest that extinction does not destroy the original learning but instead involves new learning that is at least partly modulated by the context. In the second part of the article, I will discuss a set of behavioral phenomena that are theoretically linked to extinction, and ask what they further reveal about it. Extinction is just one example of a retroactive inhibition phenomenon in which new learning inhibits old, and the principles that explain it may have wide applicability. In the final part of the article, I will consider the question “If extinction is an example of new learning, what events ‘reinforce’ or cause it?” Several possibilities will be considered. In the end, the results may favor the view that extinction occurs because the omission of the US (1) causes generalization decrement and (2) violates the organism’s expectation of the US and therefore initiates new learning. These ideas have a long history in behavioral theories of extinction. Their fuller integration with a contextual analysis (part 1) may provide a more comprehensive behavioral account of extinction.

Extinction Learning Is Especially Context-Dependent

For many years, my colleagues and I have studied a number of experimental manipulations that can be conducted after extinction has taken place. In each of them, the extinguished response returns to performance. All of them therefore indicate that ex-

tingtion is not the same as unlearning, and because all of them can be seen as context effects (see Bouton 1993, 2002), they also support the idea that performance after extinction is context-dependent. Extinction involves new learning, and it therefore leaves the CS with two available “meanings” or associations with the US. As is true for an ambiguous word, the context is crucial in selecting between them.

The Renewal Effect

Perhaps the most fundamental of these effects is the renewal effect (see Bouton and Bolles 1979a; Bouton and King 1983). In this phenomenon, a change of context after extinction can cause a robust return of conditioned responding. Several versions of the renewal effect have been studied. In the most common one, “ABA renewal,” conditioning is conducted in one context (context A) and extinction is then conducted in a second one (context B). (The contexts are typically separate and counterbalanced apparatuses housed in different rooms of the laboratory that differ in their tactile, olfactory, and visual respects.) When the CS is returned to the original conditioning context (context A), responding to the CS returns (see Bouton and Bolles 1979a; Bouton and King 1983; Bouton and Peck 1989). In a second version, “ABC renewal,” conditioning is conducted in context A, extinction is conducted in context B, and then testing is conducted in a third, “neutral” context—context C. Here again, a renewal of responding is observed (see Bouton and Bolles 1979a; Bouton and Brooks 1993; Harris et al. 2000). In a final version, conditioning and extinction are both conducted in the same context (context A) and then the CS is tested in a second context (context B). Here again, conditioned responding returns (see Bouton and Ricker 1994; Tamai and Nakajima 2000), although there is currently less evidence of this “AAB renewal” effect in operant conditioning than in Pavlovian conditioning (see Nakajima et al. 2000; Crombag and Shaham 2002).

Several facts about the renewal effect are worth noting. First, it has been observed in virtually every conditioning preparation in which it has been investigated (for a review, see Bouton 2002). Second, it can occur after very extensive extinction training. In fear conditioning (conditioned suppression) in rats, Bouton and Swartzentruber (1989) observed it when 84 extinction trials followed eight conditioning trials. Other evidence suggests that it can occur after as many as 160 extinction trials (Gunther et al. 1998; Rauhut et al. 2001; Denniston et al. 2003), although a recent report suggests that it might not survive an especially “massive” extinction treatment (800 extinction trials after eight

E-MAIL Mark.Bouton@uvm.edu; **FAX** (802) 656-8783.

Article and publication are at <http://www.learnmem.org/cgi/doi/10.1101/lm.78804>.

conditioning trials; Denniston et al., 2003). Third, the role of the context is different from the one anticipated by standard models of classical conditioning (see Rescorla and Wagner 1972; Pearce and Hall 1980; Wagner 1981; Wagner and Brandon 1989, 2001). Those models accept the view that the context is merely another CS that is presented in compound with the target CS during reinforcement or nonreinforcement. It therefore enters into simple excitatory or inhibitory associations with the US. In the ABA renewal effect (for example), context A might acquire excitatory associations with the US, and context B might acquire inhibitory associations. Either kind of association would summate with the CS to produce the renewal effect (inhibition in B would reduce responding to the CS, whereas excitation in A would enhance it). However, a number of experiments have shown that the renewal effect can occur in the absence of demonstrable excitation in context A or inhibition in context B (see Bouton and King 1983; Bouton and Swartzentruber 1986, 1989). These findings, coupled with others showing that strong excitation in a context does not influence performance to a CS unless the CS is under the influence of extinction (described below; Bouton 1984; Bouton and King 1986), suggest that direct associations in a context are neither necessary nor sufficient for a context to influence responding to a CS. The implication (see Bouton and Swartzentruber 1986; Bouton 1991) is that the contexts modulate or “set the occasion” for the current CS–US or CS–no US association (see Holland 1992; Swartzentruber 1995; Schmajuk and Holland 1998). Put another way, they activate or retrieve the current relation of CS with the US.

A further important characteristic of the renewal effect is that it implies that extinction learning is more context-specific than is original conditioning. Notice that this must be true if one observes ABC and AAB renewal; in either case, conditioning transfers better to the final test context than extinction. But our experiments on renewal have often involved comparisons of groups that received extinction training in the context in which conditioning had occurred or in a discriminably different context. Strikingly, there was no measurable effect of switching the context after conditioning on responding to the CS (see Bouton and King 1983; Bouton and Peck 1989). In contrast, extinction itself was relatively context-specific, as the renewal effect itself suggests. Recent research suggests that both conditioning and extinction become somewhat context-specific after extinction has occurred (Harris et al. 2000). But there is little question that extinction is still more context-dependent than is the original conditioning. We have therefore emphasized the fact that extinction learning is especially context-dependent.

A final fact about the renewal effect is that it appears to be supported by many kinds of contexts. For example, when fear extinction was conducted in the interoceptive context provided by benzodiazepine tranquilizers chlordiazepoxide and diazepam, renewed fear was observed when the rat was tested in the original nondrug state (Bouton et al. 1990). Cunningham (1979) had reported compatible evidence with alcohol, and we have recently collected similar observations with the benzodiazepine midazolam (L. Pain, P. Oberling, and M.E. Bouton; unpubl.). State-dependent learning or retention can be conceptualized as the drug playing the role of context (see Overton 1985).

Spontaneous Recovery

The passage of time might also bring about changes in internal and external stimulation that provide a gradually-changing context. Pavlov (1927) first observed another well-known extinction effect. In spontaneous recovery, if time is allowed to pass following extinction, the extinguished response can recover. There are several available explanations of spontaneous recovery (for a discussion of alternatives, see Robbins 1990; Brooks and Bouton

1993; Devenport et al. 1997), and it seems likely to be multiply determined. However, we have argued (see Bouton 1988, 1991, 1993) that just as extinction is relatively specific to its physical context, so it may be specific to its “temporal context.” Spontaneous recovery can be seen as the renewal effect that occurs when the CS is tested outside its temporal context. Both are due to a failure to retrieve memories of extinction outside the extinction context. Consistent with this perspective, a cue that is presented intermittently during the extinction session can attenuate either spontaneous recovery or renewal if it is presented just before the final test (Brooks and Bouton 1993, 1994; Brooks 2000). The parallel results suggest that the two effects might be controlled by a common mechanism: a failure to retrieve a memory of extinction outside the extinction context. Interestingly, changing the physical context and temporal context together can have a bigger effect than changing either context alone, as if their combination creates an even larger context change (Rosas and Bouton 1997, 1998).

Rapid Reacquisition

A third effect further indicates that conditioning is not destroyed in extinction. In rapid reacquisition, when new CS–US pairings are introduced after extinction, the reacquisition of responding can be more rapid than is acquisition with a novel CS, indicating that the original learning has been “saved” through extinction (see Napier et al. 1992; Ricker and Bouton 1996; Weidemann and Kehoe 2003). Unfortunately, the early literature on rapid reacquisition was often difficult to interpret because many early designs were not equipped to rule out less interesting explanations (for a review, see Bouton 1986). To add to the complexity, studies of fear conditioning (conditioned suppression; Bouton 1986; Bouton and Swartzentruber 1989) and flavor aversion learning (Danguir and Nicolaidis 1977; Hart et al. 1995), have shown that reacquisition can be slower than acquisition with a new CS. (It is more rapid than initial acquisition with a CS that has received the same number of nonreinforced trials without conditioning [Bouton and Swartzentruber 1989].) In fear conditioning, slow reacquisition requires extensive extinction training; more limited extinction training yields reacquisition that is neither fast nor slow (Bouton 1986). At least part of the reason these preparations support slow reacquisition is that both typically involve very few initial conditioning trials. In contrast, procedures in which rapid reacquisition has been shown (rabbit nictitating membrane response (NMR) conditioning and rat appetitive conditioning) have usually involved a relatively large number of initial conditioning trials. Consistent with a role for number of trials, Ricker and Bouton (1996) demonstrated that slow reacquisition occurred in an appetitive conditioning preparation when the procedure used the number of conditioning and extinction trials that had been used in previous fear conditioning experiments. In rabbit NMR and heart rate conditioning, extensive extinction training has abolished rapid reacquisition, although slow reacquisition has yet to be observed (Weidemann and Kehoe 2003).

Ricker and Bouton (1996) suggested that rapid reacquisition may partly be an ABA renewal effect that occurs when the animal has learned that previous USs or conditioning trials are part of the original “context” of conditioning. That is, the animal might learn that recent CS–US pairings are part of the context of conditioning, whereas recent CS-only presentations are part of the context of extinction. When CS–US pairings are resumed after extinction, they would thus return the animal to the original conditioning context. The hypothesis is compatible with Capaldi’s (1967, 1994) sequential analysis of extinction, which has made excellent use of the idea that responding on a particular trial is determined by how the animal has learned to respond in

the presence of similar memories of previous trials (see below). Presumably, conditioning preparations that use a relatively large number of conditioning trials allow ample opportunity for the animal to learn that previous reinforced trials are part of the context of conditioning. Ricker and Bouton (1996) also reported evidence that high responding during the reacquisition phase was more likely after a reinforced than a nonreinforced trial, which presumably signaled conditioning and extinction, respectively.

In more recent experiments, Bouton et al. (2004) reasoned that if rapid reacquisition is caused by recent reinforced trials generating ABA renewal, then an extinction procedure that includes occasional reinforced trials among many nonreinforced trials should slow down rapid reacquisition by making recent reinforced trials part of the context of both conditioning and extinction. Consistent with this hypothesis, a very sparse partial reinforcement procedure in extinction slowed reacquisition in a final phase compared with a group that had received simple extinction. Such a result is consistent with the idea that rapid reacquisition is at least partly an ABA renewal effect. Because the partial reinforcement treatment involved many more CS–US pairings than did simple extinction, it is difficult to reconcile with the view that rapid reacquisition is a simple function of the strength of an association that remains after extinction (see Kehoe 1988; Kehoe and Macrae 1997).

Reinstatement

A fourth context-dependent extinction phenomenon is reinstatement. In this effect, the extinguished response returns after extinction if the animal is merely reexposed to the US alone (see Pavlov 1927; Rescorla and Heth 1975; Bouton and Bolles 1979b). If testing of the CS is contemporaneous with US delivery, then the USs may cause a return of responding because they were encoded as part of the conditioning context (as above; see Reid 1958; Baker et al. 1991; Bouton et al. 1993). On the other hand, in many studies of reinstatement, testing is conducted at an interval of at least 24 h after US re-exposure; here one still observes reinstatement compared with controls that were not re-exposed to the US (see Rescorla and Heth 1975; Bouton and Bolles 1979b). In this case, evidence strongly suggests that the effect is due to conditioning of the context. When the US is presented after extinction, the organism associates it with the context; this contextual conditioning then creates reinstatement. For example, if the reinstating USs are presented in an irrelevant context, there is no reinstatement when the CS is tested again (see Bouton and Bolles 1979b; Bouton and King 1983; Bouton 1984; Baker et al. 1991; Wilson et al. 1995; Frohardt et al. 2000). Independent measures of contextual conditioning also correlate with the strength of reinstatement (Bouton and King 1983; Bouton 1984). And if the animal receives extensive extinction exposure to the context after the reinstatement shocks are presented, reinstatement is not observed (Bouton and Bolles 1979b; Baker et al. 1991). These results indicate that mere re-exposure to the US is not sufficient to generate reinstatement. It is necessary to test the CS in the context in which the US has been re-exposed.

This effect of context conditioning is especially potent with an extinguished CS. For example, Bouton (1984) compared the effects of US exposure in the same or a different context on fear of a partially extinguished CS or another CS that had reached the same low level of fear through simple CS–US pairings (and no extinction). Although contextual conditioning enhanced fear of the extinguished CS, it had no impact on the nonextinguished CS (see also Bouton and King 1986). This result is consistent with the effects of context switches mentioned above: An extinguished CS is especially sensitive to manipulations of the context. One reason is that contextual conditioning may be another feature of the conditioning context; its presence during a test

may cause a return of responding after extinction because of another ABA renewal effect (Bouton et al. 1993).

In summary, a variety of research indicates that responding to an extinguished CS is susceptible to any of a number of recovery effects, suggesting that extinction is not unlearning. Indeed, based on the results of a number of tests that allow a specific comparison of the strength of the CS–US association before and after extinction (see Delamater 1996; Rescorla 1996), Rescorla (2001) has suggested that extinction involves no unlearning whatsoever; the original CS–US association seems to survive essentially intact. Extinction must thus depend on other mechanisms. The renewal effect, and the fact that extinction leaves the CS so especially sensitive to manipulations of context, is consistent with the idea that extinction involves new learning that is especially context-dependent. We have therefore suggested that extinction leaves the CS under a contextually modulated form of inhibition (see Bouton 1993): The presence of the extinction context retrieves or sets the occasion for a CS–no US association.

Other Phenomena With Theoretical Links to Extinction

Several behavioral phenomena have been linked theoretically with extinction, and it is worth considering them to see what insights they provide.

Counterconditioning and Other Interference Paradigms

In counterconditioning, a CS that has been associated with one US is associated with a second US, often incompatible with the first, in a second phase. Not surprisingly, performance corresponding to the second association replaces performance corresponding to the first. Counterconditioning is thus a paradigm that, similar to extinction, involves a form of retroactive interference.

The literature on counterconditioning is not as large as the literature on extinction. But there is evidence that similar principles may apply. For example, experiments in my own laboratory have demonstrated a renewal effect (Peck and Bouton 1990): If rats receive CS–shock pairings in one context and then CS–food pairings in another, the original fear performance returns (and replaces food performance) when the animals are returned to the original context. Complementary results were obtained when CS–food preceded CS–shock. Other experiments have demonstrated spontaneous recovery (Bouton and Peck 1992): In this case, after CS–shock and then CS–food, animals tested at a 1-d retention interval showed primarily appetitive performance, whereas animals tested 28 d later showed a recovery of fear performance (and a suppression of appetitive). A complementary pattern was observed when CS–food preceded CS–shock. Finally, we have observed reinstatement (Brooks et al. 1995): When CS–food follows CS–shock, a number of noncontingent shocks delivered in the same context (but not in a different context) can reinstate the original fear performance. Counterconditioning thus supports at least three of the effects suggesting that extinction involves context-dependent new learning.

Bouton (1993) reviewed the behavioral literature on a number of “interference paradigms” in conditioning. In these paradigms, a CS is associated with different outcomes in successive phases of the experiment. Importantly, performance in all such paradigms is sensitive to manipulations of context and time. For example, in addition to counterconditioning, the list includes discrimination reversal learning, in which two CSs (X and Y) are reinforced and nonreinforced (e.g., X+/Y–) before the relationship is reversed in a second phase (X–/Y+). The list also includes latent inhibition, in which a single CS is nonreinforced on a number of trials before it is paired with the US in a second phase. In both paradigms, if the first and second phase are conducted in separate contexts, a return to the phase 1 context can cause a

renewal of phase 1 performance (discrimination reversal: Spear et al. 1980; Thomas et al. 1981; Bouton and Brooks 1993; latent inhibition: Bouton and Swartzentruber, 1989; Maren and Holt 2000; Westbrook et al. 2000). Effects of time related to spontaneous recovery have also been observed (discrimination reversal: Gordon and Spear 1973; Thomas et al. 1984; Bouton and Brooks 1993; latent inhibition: Kraemer et al. 1991; Aguado et al. 1994; cf. De la Casa and Lubow 2000, 2002; Lubow and De la Casa 2002). Bouton (1993) argued that a retrieval account that accepts that both phases are learned and available, and that performance is therefore determined by which is retrieved, can go some distance in explaining all examples of interference (see also Spear 1981). At a broad level of analysis, then, extinction is just one example from a set of interference phenomena that all depend on context and time.

Conditioned Inhibition

Several theories of conditioning specifically attribute extinction to a build-up of inhibition (see Konorski 1948, 1967; Pearce and Hall 1980; Wagner 1981; Pearce 1994; see also Pavlov 1927). Although the Rescorla-Wagner model (1972) attributed extinction to unlearning rather than inhibitory learning, it provided a conceptualization of inhibition that has been extremely influential. Specifically, inhibition was viewed as a negative form of learning that develops when the summed strengths of all the CSs present on a conditioning trial “overpredict” the magnitude of the US that actually occurs on the trial. Most models now use a version of its error-correction learning mechanism, which essentially brings the overall expectation of the US (the summed values of the CSs present) into line with reality (the value of the US that is actually present); this means decrementing associative strengths until what is predicted equals what actually occurs.

According to this view, the main method for generating pure conditioned inhibition is the conditioned inhibition or feature-negative paradigm, in which the experimenter intermixes trials on which one CS (X) is paired with the US and other trials on which X is combined with another CS (Y) and presented without a US (X+, XY−). (X is sometimes called the “target” stimulus and Y the “feature” stimulus.) After early X+ trials, X begins to predict that the US will occur; when no US occurs on XY− trials, the error-correction mechanism therefore decrements both CSs to bring their strengths in line with no US. In the Rescorla-Wagner model, X will lose some excitatory strength on the negative trials, but it will not go below zero. But because Y starts with zero value, its associative strength becomes negative and thus becomes a pure conditioned inhibitor.

Other models actually predict inhibitory learning to both Y and X in this paradigm (see Pearce and Hall 1980; Wagner 1981; Wagner and Brandon 1989, 2001). For example, Wagner (1981) and Wagner and Brandon (1989, 2001) have presented a comprehensive real-time model that expands enormously on the Rescorla-Wagner model and is worth describing here. According to the model, known as SOP (for Sometimes-Opponent-Process), CS and US are represented as memory nodes that can become associated during conditioning. For the association between them to be strengthened, both nodes must be activated from inactivity to an active state, A1, at the same time. Once the association has been formed, the presentation of the CS activates the US node to a secondarily-active state, A2. This in turn generates the conditioned response. An inhibitory connection is formed between a CS and a US when the CS is activated to the A1 state and the US is activated to A2 rather than A1. These conditions are met in the feature-negative paradigm (X+, XY−). After some initial conditioning, X is able to activate the US node to A2. On XY− trials, the US node is therefore put into the A2 state at the same time that X and Y are in A1. This will create inhibition to both X

and Y. Because Y has no other association, it will become a pure conditioned inhibitor. For X, however, the inhibition is overlaid on its existing (and unchanged) excitatory association. The same thing occurs in simple extinction, because the CS also activates the US into A2 at that time. Because these processes occur in real time, during any nonreinforced trial, inhibition will accrue to the CS from the point in time at which the US node is first activated to A2 until the CS leaves the A1 state, which may not occur until the CS is turned off at the end of the trial (for further discussion, see the section Violation of Reinforcer Expectation below).

Because of these connections between inhibition and extinction, and the evidence that extinction is more context-specific than is simple conditioning, Bouton and Nelson (1994) and Nelson and Bouton (1997) asked whether pure inhibition acquired in the feature-negative paradigm was also context-specific. The designs of the experiments are sketched in Table 1. Rats were given different feature-negative discriminations in a series of intermixed sessions in two contexts (A and B), as shown at left. The design in the upper half of Table 1 asks whether inhibition conditioned to the feature in context A (CS Y) transfers to the other context (context B) in the final test; an inhibitor (Y) was tested in its original context and/or in another one. Surprisingly, its inhibition transferred without measurable disruption to the new context. However, we also performed experiments along the lines shown in the lower portion of Table 1. These asked whether any inhibition acquired by stimulus X, rather than the pure inhibitory Y, is lost with a context switch. Here there was clear evidence of a context effect: When switched to the alternate context, responding to X became more difficult to inhibit. Stimulus X, with its mixed history of both reinforcement and nonreinforcement, is similar to an extinguished CS: Its inhibitory association is context-specific. But Y’s pure inhibition is not.

Because inhibition was not generally context-specific, the results implied that extinction is not context-specific merely because it is a form of inhibition. A second reason why extinction might be context-specific is that it is the second thing the organism has learned about the CS. Notice that the same is true in other interference paradigms, such as counterconditioning: It is the second-learned association that seems dependent on the context and on time. And it is presumably also true of the target stimulus (X) in the feature-negative design; there must be excitation to X before the XY− trials can generate inhibition. Nelson (2002) confirmed the second-association hypothesis in a series of experiments. As usual, excitatory conditioning (tone-food conditioning) transferred undisturbed across contexts, unless the tone had first been trained as a conditioned inhibitor (stimulus Y in the feature-negative paradigm). Conversely, inhibition to a conditioned inhibitor also transferred across contexts unless the

Table 1. Designs Used by Bouton and Nelson (1994) and Nelson and Bouton (1997)

Training	Testing
Is inhibition of the feature (Y) context-specific?	
A: X+, XY−	A: X, XY
B: X+, XZ−	B: X, XY
Is inhibition of the target (X) context-specific?	
A: X+, XY−	A: X, XY
B: Z+, ZY−	B: X, XY

A and B are contexts; X, Y, and Z are CSs. + indicates reinforced; −, nonreinforced. During training, all rats received intermixed sessions in both contexts. Testing was then conducted in both contexts with either between-subject or within-subject methods.

CS had first been trained as a conditioned excitator (through initial tone–food pairings). Thus, regardless of whether the association was excitatory or inhibitory, the second thing learned was more context-specific than the first. Compatible data had been shown by Swartzentruber and Bouton (1992), who found that excitatory conditioning was relatively context-specific if it had been preceded by nonreinforced pre-exposure to the CS.

Thus, the evidence suggests that the learning and memory system treats the first association as context-free, but the second association as a kind of context-specific exception to the rule.¹ There may be functional reasons for this (Bouton 1994). A conditioning trial provides a sample from which an animal may make inferences about the state of the world (see Staddon 1988). Statistically, if the world is composed of two types of trials (CS–US and CS–no US), then the probability of sampling a particular type of trial will reflect its true prevalence in the world. Therefore, an early run of conditioning trials would reflect its high incidence in the population; a subsequent trial of another type might reflect an exception to the rule. Learning and memory may thus be designed to treat second-learned information as a conditional and context-specific. A mechanistic account has not been fully tested. One possibility is that when extinction begins, the omission of the US is surprising and boosts attention to the CS and context, as suggested by the attention rule of the Pearce-Hall model (Pearce and Hall 1980; see also Kaye and Pearce 1984). Interestingly, there is evidence that conditioned excitation may be context specific after one conditioning trial—when the US is also new and surprising (Hall and Honey 1990).

Inhibition Despite Continued Reinforcement

Other research has identified another inhibitory effect that deserves mention alongside extinction. Ayres et al. (1979; see also Vigorito and Ayres 1987) have shown in the conditioned suppression preparation that repeated conditioning trials may yield nonmonotonic learning curves in which conditioned fear reaches a peak and then declines despite continued pairings of the CS and shock. The process is poorly understood, although similar effects have been observed in several conditioning preparations (Kimmel and Burns 1975). Pavlov himself saw it often (1927), and attributed it to inhibition developing despite reinforcement.

Experiments in my laboratory have further documented the effect in the conditioned suppression preparation (M.E. Bouton, R.J. Frohardt, C. Sunsay, and J. Waddell, in prep.). We have also discovered a role for context. Rats received intermixed sessions in two contexts; in one context, one CS was repeatedly paired with footshock, and in the other, a second CS was similarly paired with the same shock. Nonmonotonic learning curves developed. However, when we then tested the CS in the alternate context, we saw a significant increase in fear of the CS. Over experiments, there was a strong correlation between the size of this increase and the degree of nonmonotonicity shown in the conditioning curve. The context switch thus attenuated the inhibitory process that led to the nonmonotonic learning curve.

We do not have a good understanding of this inhibitory process at the present point in time. In our experiments, the decline in fear over training is not a result of inhibition of delay, in which the animal learns the timing of the US (at the end of the

CS in most experiments), and therefore undergoes extinction to early parts of the CS (see Rosas and Alonso 1996). Interestingly, Rosas and Alonso (1997) have shown that inhibition of delay, developing over trials this way, is in fact attenuated with a context switch. However, in my laboratory the nonmonotonicity is still observed when shock occurs at unpredictable times in the CS, and there is again an increase in fear when the context is changed under these conditions. The result is also not due to an opioid process that helps the animal adapt to the US (see Vigorito and Ayres 1987); indeed, because our context switch experiments test the rat's reaction to the CS, not the US, the results indicate some sort of adaptation to the CS or to conditioned fear itself. Another possibility is suggested by SOP (see Wagner 1981): Repeated exposure to the CS would allow the animal to associate the CS with the context, which would permit the context to associatively activate the CS into the secondarily active state (A2). Putting the CS into that state might reduce its ability to evoke a response (see Hall and Honey 1990; Honey et al. 1993; Hall and Mondragón 1998). But perhaps contrary to this view, there is little correlation between the number of CSs presented in the conditioning procedure and the degree of nonmonotonicity. In addition, we have yet to observe nonmonotonic learning curves in appetitive conditioning, in which the animal also receives many opportunities to associate the context and CS. Although we do not yet understand the inhibition-despite-reinforcement phenomenon, there is a suggestion that, similar to extinction, it involves a context-specific inhibitory process that might be the second thing learned about the CS.

What Causes Extinction?

If we return specifically to extinction, another question is what event or behavioral process actually causes the loss of responding? Several ideas have been examined in recent experiments.

Discrimination of Reinforcement Rate

One possibility is that the animal eventually learns that the rate of reinforcement in the CS is lower in extinction than it was during conditioning. Gallistel and Gibbon (2000) have argued that the animal continually decides whether or not to respond in extinction by comparing the current rate of reinforcement in the CS with its memory of the rate that prevailed in conditioning. Because rate is the reciprocal of time, the animal computes a ratio between the amount of time accumulated in the CS during extinction and the amount of time accumulated in the CS between USs during conditioning. When the ratio exceeds a threshold, the animal stops responding.

This approach has been tested in several recent experiments. Haselgrove and Pearce (2003) examined the impact of varying the duration of the CS during extinction; when longer CSs are used in extinction, time in the CS accumulates more quickly, and the animal should stop responding after fewer trials. In some experiments, rats were given appetitive conditioning with a 10-sec CS and then given extinction exposures to a series of 10-sec or 270-sec presentations of the CS. When responding was examined at the start of each CS, there was an occasionally significant, but surprisingly small, effect of increasing the duration of the CS during extinction. For instance, by the 12th two-trial block, the 10-sec and 270-sec CS groups had similar nonzero levels of responding, even though they had accumulated a total of 4 and 108 min of exposure in the CS, respectively. On the other hand, responding did decline as a function of time within a single presentation of the 270-sec CS, perhaps reflecting generalization decrement resulting from the increasing difference between the current CS and the 10-sec CS employed in conditioning. Consistent with that view, when conditioning first occurred with a 60-sec CS, extinction of responding occurred more rapidly with a

¹The main exception to the second-association rule is latent inhibition, in which the first phase can be shown to exert a context-dependent influence on the second phase (see Hall and Channell 1985) despite the fact that it is arguably the first thing learned. Latent inhibition is unique, however, in that the CS is not paired with anything significant in the first phase. One possibility, therefore, is that it is in part encoded as a feature of the context, making it difficult to extract it from that context when it is paired with the US in phase 2 (cf. Gluck and Myers 1993).

10-s CS than with a 60-sec CS. Thus, either an increase or a decrease in the duration of the CS relative to conditioning accelerated the loss of responding. This effect of time was not anticipated by the rate-discrimination view (Gallistel and Gibbon 2000).

Drew et al. (2004) reported compatible results in experiments on autoshaping in ring doves. Doubling or halving the duration of the CS from the 8-sec value used in conditioning did not affect the number of trials required to stop responding. The fact that extinction was thus largely controlled by the number of CS presentations is consistent with experiments that have examined the effects of the number and duration of nonreinforced trials added to conditioning schedules (Bouton and Sunsay 2003). On the other hand, Drew et al. (2004) found that a more extreme increase in CS duration (from 8 to 32 sec) increased the rate of extinction. This was attributed to the animal learning to discriminate the longer nonreinforced CS presentations from the shorter reinforced CS presentations: When 8-sec CSs were presented again after extinction, birds extinguished with 4-sec and 32-sec CSs responded again. Animals are sensitive to time in the CS, but the number of extinction trials appears to be an important factor.

As noted by Gallistel and Gibbon (2000), the rate discrimination theory seems especially consistent with a well-known extinction phenomenon, the partial reinforcement effect (PRE; for a review, see Mackintosh, 1974). In this phenomenon, conditioning with partial reinforcement schedules (in which nonreinforced trials are intermixed with reinforced trials) creates a slower loss of responding in extinction than does conditioning with a continuous reinforcement schedule (in which every trial is reinforced). According to a rate-discrimination hypothesis (Gallistel and Gibbon 2000), the partially reinforced subjects have learned to expect the US after more accumulated time in the CS, and it thus takes more CS time in extinction to exceed the threshold of accumulated extinction time/expected time to each US. The more traditional approach, in contrast, has been to think that partially reinforced subjects have learned to expect the US after more trials than continuously reinforced subjects have. It therefore takes more trials to stop generalizing from conditioning to extinction (see Mowrer and Jones 1945; Capaldi 1967, 1994).

Contrary to the rate discrimination hypothesis, Haselgrove et al. (2004) and M.E. Bouton and A.M. Woods (in prep.) have shown that a PRE still occurs when partially and continuously reinforced subjects expect the reinforcer after the same amount of CS time. For example, both sets of investigators showed that a group that received a 10-sec CS reinforced on half its presentations (accumulated CS time of 20 sec) extinguished more slowly than did a continuously reinforced group that received every 20-sec CS presentation reinforced. M.E. Bouton and A.M. Woods (in prep.) further distinguished the "time-discrimination" account from the traditional "trial-discrimination" account (see Mowrer and Jones 1945; Capaldi 1967, 1994). Rats that had every fourth 10-sec CS reinforced extinguished more slowly over a series of alternating 10-sec and 30-sec extinction trials than rats that had received every 10-sec CS reinforced. This PRE was still observed when extinction responding was plotted as a function of time units over which the US should have been expected (every 40 sec for the PRF group but every 10 sec for the CRF group). In contrast, the PRE disappeared when extinction responding was plotted as a function of the trials over which the US should have been expected (every fourth trial for the PRF group and every trial for the CRF group). Ultimately, the PRE is better captured by trial-based theories (e.g., Capaldi 1967, 1994).

We have already seen that responding on a particular trial occurs in the context of memories of the outcomes of previous trials—that was the explanation provided earlier of rapid reac-

quisition as an ABA renewal effect (Ricker and Bouton 1996; Bouton et al. 2004). Interestingly, the recent finding that occasional reinforced trials in extinction (partial reinforcement) can slow down the rate of reacquisition (Bouton et al. 2004) is really just the inverse of the PRE: In the PRE, nonreinforced trials in conditioning allow more generalization from conditioning to extinction, whereas Bouton et al.'s finding suggests that reinforced trials in extinction allowed for more generalization of extinction to reconditioning. Either finding suggests the importance of considering recent trials as part of the context that controls performance in extinction.

In summary, there is little support for the idea that responding extinguishes when the US is omitted because the organism detects a lower rate of reinforcement in the CS. The number of extinction trials, rather than merely the accumulating time in the CS across trials, appears to be important to the extinction process. Time in the CS can have an effect: It appears to be another dimension over which animals generalize and discriminate (Haselgrove and Pearce 2003; Drew et al., 2004). But explanation of the PRE appears to be most consistent with a view that animals use their memories of the outcomes of preceding trials as a dimension over which they generalize and respond (for a more extended review, see also Mackintosh 1974).

Generalization Decrement

It is thus possible to claim that the animal stops responding in extinction at the point at which it stops generalizing between the stimuli that prevailed in conditioning than those that prevail in extinction (see Capaldi 1967, 1994). This idea has had a long and influential history in research on extinction, especially in research on the PRE. It is interesting to note that a generalization decrement theory of extinction does not imply destruction of the original learning in extinction, or indeed any new learning at all. However, there is still good reason to think that extinction also involves new learning. For instance, nonreinforcement of a food CS elicits measurable frustration, and this can be associated with stimuli present in the environment (Daly 1974). Nonreinforcement of the CS in a feature-negative paradigm also generates measurable new learning in the form of conditioned inhibition (see above). And there is also evidence for new learning in the renewal effect. For example, either ABC renewal or AAB renewal (see above) imply that context B (the extinction context) acquires an ability to modulate (suppress) performance to the CS. Such observations suggest that the animal has not merely stopped responding in extinction because of a failure to generalize. Instead, it appears to have learned that the CS means no US in the extinction context (see above).

Role of Inhibition of the Response

Rescorla (2001) has recently suggested that extinction might involve learning to inhibit the conditioned response. For example, he has summarized evidence from instrumental (operant) conditioning experiments indicating that the effects of extinction can be specific to the response that undergoes extinction. For example, Rescorla (1993) reinforced two operant behaviors (lever pressing and chain pulling) with food pellets and then extinguished each response in combination with a new stimulus (a light or a noise). Subsequent tests of the two responses with both light and noise indicated that each response was more depressed when it was tested in combination with the cue in which it had been extinguished (see also Rescorla 1997). There is thus good reason to think that the animal learns something specific about the response itself during operant extinction: It learns not to perform a particular response in a particular stimulus. One possibility is that the animal learns a simple inhibitory S-R association (Colwill 1991). Another possibility, perhaps more consistent

with the context-modulation account of extinction emphasized above, is that the animal learns that S sets the occasion for a response–no reinforcer relationship. Rescorla (1993, p. 335; 1997, p. 249) has observed that the experiments do not separate the two possibilities. To my knowledge, no analogous experiments have been performed in the Pavlovian conditioning situation.

Instead, the main implication examined in Pavlovian conditioning is that extinction procedures should be especially successful at causing inhibitory S-R learning if they generate high levels of responding in extinction. This prediction may provide a reasonable rule of thumb (Rescorla 2001). For example, when a CS is compounded with another excitatory CS and the compound is extinguished, there is especially strong responding in extinction (due to summation between the CSs), and especially effective extinction as evidenced when the CS is tested alone (Wagner 1969; Rescorla 2000; Thomas and Ayres 2004). Conversely, when the target CS is compounded with an inhibitory CS, there is relatively little responding to the compound (excitation and inhibition negatively summate), and there is also less evidence of extinction when the target is tested alone (Soltysik et al. 1983; Rescorla 2003; Thomas and Ayres 2004). However, although these findings are consistent with the hypothesis that the effectiveness of extinction correlates with the degree of responding, they can also be interpreted in a different way. Either treatment also affects the degree to which the animal's expectation of the reinforcer is violated: The stimulus compound influences the size of the error term in the Rescorla-Wagner model and, in more cognitive terms, the extent to which the expectation of the US created by the compound is violated when the US does not occur. The results do not separate the response-inhibition hypothesis from an expectancy-violation hypothesis, which will be covered in the next section.

A recent eyeblink experiment by Krupa and Thompson (2003) manipulated the level of responding another way. During extinction, rabbits were given microinjections of the GABA agonist muscimol adjacent to the motor nuclei that control the conditioned response (the facial nucleus and the accessory abducens). The injection therefore eliminated the CR during extinction. However, when the subjects were then tested without muscimol, the CS evoked considerable responding, suggesting that evocation of the CR was necessary for extinction learning. Unfortunately, the muscimol microinjections also had robust stimulus effects. They caused complete inactivation of the ipsilateral facial musculature: "the external eyelids were flaccid, the left ear hung down unsupported, and no vibrissae movements were observed on the side of the infusion" (p. 10579). In effect, the rabbits received extinction in a context that was different from the one in which conditioning and testing occurred (the ordinary state without partial facial paralysis). There are thus strong grounds for expecting a renewal effect. The hypothesis that elicitation of the CR is necessary for extinction must await further tests.

There are also data suggesting that the number of responses or level of responding in extinction does not correlate with effective extinction learning. For example, Drew et al. (2004) noted that although animals given long CSs in extinction responded many more times in extinction than did animals given shorter CSs, extinction was mainly a function of the number of extinction trials. In fear conditioning experiments with mice, Cain et al. (2003) reported that extinction trials that were spaced in time produced a slower loss of freezing than did extinction trials that were massed in time. Nevertheless, there was less spontaneous recovery after the massed treatment, suggesting that extinction was more effective when the treatment involved less overall responding. Experiments in my own laboratory with different appetitive conditioning methods in rats (E.W. Moody, C. Sunsay,

and M.E. Bouton, in prep.) suggest a similar conclusion even though the results were different. Spaced extinction trials again yielded more responding in extinction than did massed trials, but the treatments caused indistinguishable amounts of extinction learning as assessed in spontaneous recovery and reinstatement tests. In related conditioned suppression experiments, M.E. Bouton, A. García-Gutiérrez, J. Zilski, and E.W. Moody (in prep.) compared the effects of extinction in multiple contexts on the strength of the ABA and ABC renewal effects. Rats received fear conditioning with a tone CS in context A, and then extinction of the tone for three sessions in context B, or a session in B, then C, and then D, before final renewal tests in the original context (context A) or a neutral fifth context (context E). Although the successive context switches in the BCD group caused more fear responding during extinction (due to renewal effects), the groups showed strikingly similar renewal in either context A or context E. Gunther et al. (1998) and Chelonis et al. (1999) have shown more favorable effects of extinction in multiple contexts on renewal. But higher responding in extinction does not guarantee better extinction learning. The results seem inconsistent with a response-inhibition hypothesis. Their impact on the expectancy violation hypothesis is perhaps less clear.

In summary, although animals that receive extinction after operant conditioning may in fact learn to refrain from performing a particular response in a particular context (see Rescorla 1993, 1997), the importance of response inhibition in Pavlovian extinction is not unequivocally supported at the present time. High responding in extinction does not guarantee more effective extinction learning. To date, the hypothesis has been difficult to distinguish from the violation-of-expectation hypothesis that is built into many mainstream models of classical conditioning (see Rescorla and Wagner 1972; Pearce and Hall 1980; Wagner 1981; Wagner and Brandon 1989, 2001) and is considered next.

Violation of Reinforcer Expectation

It is commonly thought that each CS presentation arouses a sort of expectation of the US that is disconfirmed on each extinction trial. For example, in the error-correction rule provided by Rescorla and Wagner (1972), the degree of unlearning (which we have seen can create inhibition) is provided by the difference in the overall associative strength present on a trial and the actual US that occurs on the trial. In the Pearce-Hall model (Pearce and Hall 1980), the discrepancy was conceptualized as an event that reinforced new inhibitory learning that is overlaid on the original excitatory learning (see also Daly and Daly 1982). Wagner's SOP model (1981) accepts a similar idea. One piece of evidence that seems especially consistent with the expectation-violation view is the "overexpectation experiment," in which two CSs are separately associated with the US and then presented together in a compound that is then paired with the US. Despite the fact that the compound is paired with a US that can clearly generate excitatory learning, the two CSs undergo some extinction (see Kremer 1978; Lattal and Nakajima 1998). The idea is that summation of the strengths of the two CSs causes a discrepancy between what the animal expects and what actually occurs, and some extinction is therefore observed. As mentioned above, the expectation-violation view is also consistent with the effects of compounding excitors and inhibitors with the target CS during no-US (extinction) trials (Wagner 1969; Soltysik et al. 1983; Rescorla 2000, 2003; Thomas and Ayres 2004).

One theoretical challenge has been to capture the expectancy violation in real time. Gallistel and Gibbon (2000) have emphasized the fact that traditional trial-based models such as the Rescorla-Wagner model have been vague about the precise point in time in a trial when the violation of expectation actually occurs. The issue is especially clear when trial-based models ex-

plain the extinction that occurs with a single extended presentation of the CS, as is the case for the context or background in conditioning protocols with very widely spaced conditioning trials. (Spaced trials are held to facilitate conditioning of the CS because long intertrial intervals allow more context extinction and thus less blocking by context.) There is good evidence that widely spaced trials do create less contextual conditioning than massed trials (see Barela 1999). To account for contextual extinction over long intertrial intervals, many trial-based models arbitrarily assume that the single long context exposure is carved into many imaginary trials, and that more imaginary trials occur and create more extinction in longer context exposures.

It is worth noting, however, that Wagner's SOP model (see Wagner 1981; Wagner and Brandon 1989, 2001) is relatively specific about when in time the process that generates extinction occurs. As already mentioned, extinction occurs because an inhibitory CS-US association develops when the CS is in the A1 state and the US is in the A2 state. After conditioning has occurred, the CS will continuously activate elements of the US node to the A2 state as long as it is present; therefore, extinction learning will occur continuously as long as the CS is on and no US occurs. A limiting factor, however, is the extent to which the CS itself is in the A2 state: The longer it remains on, the more likely the elements in the CS node will be in A2 rather than A1, making new learning about the CS more difficult. Nonetheless, extensions of the CS in extinction will have an effect, because elements in A2 eventually return to the inactive state, from which they will return to A1 because of the continued presence of the CS. SOP thus accounts for extinction in extended CSs without recourse to imaginary trials, and a recent extension of the model (Vogel et al. 2003) may also account for generalization decrement as a function of CS time (Haselgrove and Pearce 2003; Drew et al. 2004). Although a truly rigorous analysis of SOP requires computer simulations that are beyond the scope of the present article, the principles contained in the model are consistent with many of the facts of extinction reviewed here. From the current point of view, its most significant problem is that it underestimates the role of context in extinction and might not account for the negative occasion-setting function of context (see Bouton and King 1983; Bouton and Swartzentruber 1986, 1989; Bouton and Nelson 1998) that arguably provides the key to understanding renewal, spontaneous recovery, rapid reacquisition, and reinstatement phenomena (for a start at addressing occasion setting phenomena in terms of SOP, see Brandon and Wagner 1998; Wagner and Brandon 2001).

Conclusions

Extinction is a highly complex phenomenon, even when analyzed at a purely behavioral level. It is worth noting that it is probably multiply determined. But according to the results reviewed here, it does not involve destruction of the original learning. Instead, the main behavioral factors that cause the loss of responding appear to be generalization decrement, and new learning that may be initiated by the violation of an expectation of the US. In SOP, perhaps the most powerful and comprehensive model of associative learning that is currently available, that expectation violation takes the form of the CS activating the US node into a secondarily-active (A2) state that potentially enables new inhibitory learning as long as the CS remains on and no US is presented. Importantly, this new inhibitory learning leaves the original CS-US association intact.

I have argued that the fact that extinction might leave the original learning intact means that the CS emerges from extinction with two available associations with the US. It therefore has properties analogous to those of an ambiguous word, and the

current performance depends on which of two associations is retrieved. Consistent with this idea, another fact that emerges from behavioral research on extinction is that it is relatively context-dependent. I have therefore suggested that the second (inhibitory) association of CS is especially dependent on the context for its activation or retrieval. The role of the context is modulatory; it activates or retrieves the CS's own second (inhibitory) association, much as a negative occasion setter might (see Holland, 1992). This hypothesis begins to integrate several facts about extinction and brings relapse effects such as the renewal effect, spontaneous recovery, rapid reacquisition, and reinstatement to center stage. Extinction is not the same as unlearning, and the context-dependence of extinction performance is a central part of that insight.

ACKNOWLEDGMENTS

I thank John Pearce for many useful discussions. Preparation of this article was supported by Grant RO1 MH64847 from the National Institute of Mental Health. It was written while the author was a visiting fellow at Cardiff University under the support of the Cardiff Visiting Fellows Scheme.

REFERENCES

- Aguado, L., Symonds, M., and Hall, G. 1994. Interval between preexposure and test determines the magnitude of latent inhibition: Implications for an interference account. *Anim. Learn. Behav.* **22**: 188-194.
- Ayres, J.J.B., Berger-Gross, P., Kohler, E.A., Mahoney, W.J., and Stone, S. 1979. Some orderly nonmonotonicities in the trial-by-trial acquisition of conditioned suppression: Inhibition with reinforcement? *Anim. Learn. Behav.* **7**: 174-180.
- Baker, A.G., Steinwald, H., and Bouton, M.E. 1991. Contextual conditioning and reinstatement of extinguished instrumental responding. *Q. J. Exp. Psychol. B* **43**: 199-218.
- Barela, P.B. 1999. Theoretical mechanisms underlying the trial spacing effect in Pavlovian conditioning. *J. Exp. Psychol. Anim. Behav. Processes* **25**: 177-193.
- Bouton, M.E. 1984. Differential control by context in the inflation and reinstatement paradigms. *J. Exp. Psychol. Anim. Behav. Processes* **10**: 56-74.
- . 1986. Slow reacquisition following the extinction of conditioned suppression. *Learn. Motiv.* **17**: 1-15.
- . 1988. Context and ambiguity in the extinction of emotional learning: Implications for exposure therapy. *Behav. Res. Ther.* **26**: 137-149.
- . 1991. Context and retrieval in extinction and in other examples of interference in simple associative learning. In: *Current topics in animal learning: Brain, emotion, and cognition* (eds. L. Dachowski and C.F. Flaherty), pp. 25-53. Lawrence Erlbaum, Hillsdale, NJ.
- . 1993. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychol. Bull.* **114**: 80-99.
- . 1994. Conditioning, remembering, and forgetting. *J. Exp. Psychol. Anim. Behav. Processes* **20**: 219-231.
- . 2002. Context, ambiguity, and unlearning: Sources of relapse after behavioral extinction. *Biol. Psychiatry* **52**: 976-986.
- Bouton, M.E. and Bolles, R.C. 1979a. Contextual control of the extinction of conditioned fear. *Learn. Motiv.* **10**: 445-466.
- . 1979b. Role of conditioned contextual stimuli in reinstatement of extinguished fear. *J. Exp. Psychol. Anim. Behav. Processes* **5**: 368-378.
- Bouton, M.E., and Brooks, D.C. 1993. Time and context effects on performance in a Pavlovian discrimination reversal. *J. Exp. Psychol. Anim. Behav. Processes* **19**: 165-179.
- Bouton, M.E., and King, D.A. 1983. Contextual control of the extinction of conditioned fear: Tests for the associative value of the context. *J. Exp. Psychol. Anim. Behav. Processes* **9**: 248-265.
- . 1986. Effect of context on performance to conditioned stimuli with mixed histories of reinforcement and nonreinforcement. *J. Exp. Psychol. Anim. Behav. Processes* **12**: 4-15.
- Bouton, M.E. and Nelson, J.B. 1994. Context-specificity of target versus feature inhibition in a feature-negative discrimination. *J. Exp. Psychol. Anim. Behav. Processes* **20**: 51-65.
- Bouton, M.E. and Nelson, J.B. 1998. Mechanisms of feature-positive and feature-negative discrimination learning in an appetitive conditioning paradigm. In *Occasion setting: Associative learning and cognition in animals* (eds. N. Schmajuk and P.C. Holland), pp. 69-112. American Psychological Association, Washington, DC.

- Bouton, M.E. and Peck, C.A. 1989. Context effects on conditioning, extinction, and reinstatement in an appetitive conditioning preparation. *Anim. Learn. Behav.* **17**: 188–198.
- Bouton, M.E. and Peck, C.A. 1992. Spontaneous recovery in cross-motivational transfer (counter-conditioning). *Anim. Learn. Behav.* **20**: 313–321.
- Bouton, M.E. and Ricker, S.T. 1994. Renewal of extinguished responding in a second context. *Anim. Learn. Behav.* **22**: 317–324.
- Bouton, M.E. and Sunsay, C. 2003. Importance of trials versus accumulating time across trials in partially-reinforced appetitive conditioning. *J. Exp. Psychol. Anim. Behav. Processes* **29**: 62–77.
- Bouton, M.E. and Swartztruber, D. 1986. Analysis of the associative and occasion-setting properties of contexts participating in a Pavlovian discrimination. *J. Exp. Psychol. Anim. Behav. Processes* **12**: 333–350.
- Bouton, M.E. and Swartztruber, D. 1989. Slow reacquisition following extinction: Context, encoding, and retrieval mechanisms. *J. Exp. Psychol. Anim. Behav. Processes* **15**: 43–53.
- . 1992. Spontaneous recovery in cross-motivational transfer (counterconditioning). *Anim. Learn. Behav.* **20**: 313–321.
- Bouton, M.E., Kenney, F.A., and Rosengard, C. 1990. State-dependent fear extinction with two benzodiazepine tranquilizers. *Behav. Neurosci.* **104**: 44–55.
- Bouton, M.E., Rosengard, C., Achenbach, G.G., Peck, C.A., and Brooks, D.C. 1993. Effects of contextual conditioning and unconditional stimulus presentation on performance in appetitive conditioning. *Q. J. Exp. Psychol.* **46B**: 63–95.
- Bouton, M.E., Woods, A.M., and Pineño, O. 2004. Occasional reinforced trials during extinction can slow the rate of rapid reacquisition. *Learn. Motiv.* (in press).
- Brandon, S.E. and Wagner, A.R. 1998. Occasion setting: Influences of conditioned emotional responses and configural cues. In *Occasion setting: Associative learning and cognition in animals* (eds. N.A. Schmajuk and P.C. Holland), pp. 343–382. American Psychological Association, Washington, DC.
- Brooks, D.C. 2000. Recent and remote extinction cues reduce spontaneous recovery. *Q. J. Exp. Psychol.* **53B**: 25–58.
- Brooks, D.C. and Bouton, M.E. 1993. A retrieval cue for extinction attenuates spontaneous recovery. *J. Exp. Psychol. Anim. Behav. Processes* **19**: 77–89.
- . 1994. A retrieval cue for extinction attenuates response recovery (renewal) caused by a return to the conditioning context. *J. Exp. Psychol. Anim. Behav. Processes* **20**: 366–379.
- Brooks, D.C., Hale, B., Nelson, J.B., and Bouton, M.E. 1995. Reinstatement after counterconditioning. *Anim. Learn. Behav.* **23**: 383–390.
- Cain, C.K., Blouin, A.M., and Barad, M. 2003. Temporally massed CS presentations generate more fear extinction than spaced presentations. *J. Exp. Psychol. Anim. Behav. Processes* **29**: 323–333.
- Capaldi, E.J. 1967. A sequential hypothesis of instrumental learning. In *Psychology of learning and motivation*, Vol. 1 (eds. K.W. Spence and J.T. Spence), pp. 67–156. Academic Press, New York.
- . 1994. The sequential view: From rapidly fading stimulus traces to the organization of memory and the abstract concept of number. *Psychonomic Bull. Rev.* **1**: 156–181.
- Chelonis, J.J., Calton, J.L., Hart, J.A., and Schachtman, T.R. 1999. Attenuation of the renewal effect by extinction in multiple contexts. *Learn. Motiv.* **30**: 1–14.
- Colwill, R.W. 1991. Negative discriminative stimuli provide information about the identity of omitted response-contingent outcomes. *Anim. Learn. Behav.* **19**: 326–336.
- Crombag, H.S. and Shaham, Y. 2002. Renewal of drug seeking by contextual cues after prolonged extinction in rats. *Behav. Neurosci.* **116**: 169–173.
- Cunningham, C.L. 1979. Alcohol as a cue for extinction: State dependency produced by conditioned inhibition. *Anim. Learn. Behav.* **7**: 45–52.
- Daly, H.B. 1974. Reinforcing properties of escape from frustration aroused in various learning situations. In *The psychology of learning and motivation*, Vol. 8 (ed. G.H. Bower), pp. 187–232. Academic Press, New York.
- Daly, H.B. and Daly, J.T. 1982. A mathematical model of reward and aversive nonreward: Its application in over 30 appetitive learning situations. *J. Exp. Psychol. Gen.* **111**: 441–480.
- Danguir, J. and Nicolaidis, S. 1977. Lack of reacquisition in learned taste aversions. *Anim. Learn. Behav.* **5**: 395–397.
- De la Casa, L.G. and Lubow, R.E. 2000. Super-latent inhibition with conditioned taste aversion testing. *Anim. Learn. Behav.* **28**: 389–399.
- . 2002. An empirical analysis of the super-latent inhibition effect. *Anim. Learn. Behav.* **30**: 112–120.
- Delamater, A.R. 1996. Effects of several extinction treatments upon the integrity of Pavlovian stimulus-outcome associations. *Anim. Learn. Behav.* **24**: 437–449.
- . 2004. Experimental extinction in Pavlovian conditioning: Behavioural and neuroscience perspectives. *Q. J. Exp. Psychol.* **57B**: 97–132.
- Denniston, J.C., Chang, R.C., and Miller, R.R. 2003. Massive extinction treatment attenuates the renewal effect. *Learn. Motiv.* **34**: 68–86.
- Devenport, L., Hill, T., Wilson, M., and Ogden, E. 1997. Tracking and averaging in variable environments: A transition rule. *J. Exp. Psychol. Anim. Behav. Processes* **23**: 450–460.
- Drew, M.R., Yang, C., Ohyama, T., and Balsam, P.D. 2004. Temporal specificity of extinction in autoshaping. *J. Exp. Psychol. Anim. Behav. Processes* **30**: 163–176.
- Frohardt, R.J., Guarraci, F.A., and Bouton, M.E. 2000. The effects of neurotoxic hippocampal lesions on two effects of context after fear extinction. *Behav. Neurosci.* **114**: 227–240.
- Gallistel, C.R. and Gibbon, J. 2000. Time, rate, and conditioning. *Psychol. Rev.* **107**: 289–344.
- Gluck, M. and Myers, C.E. 1993. Hippocampal mediation of stimulus representation: A computational theory. *Hippocampus* **3**: 492–516.
- Gordon, W.C. and Spear, N.E. 1973. Effect of reactivation of a previously acquired memory on the interaction between memories in the rat. *J. Exp. Psychol.* **99**: 349–355.
- Gunther, L.M., Denniston, J.C., and Miller, R.R. 1998. Conducting exposure treatment in multiple contexts can prevent relapse. *Behav. Res. Ther.* **36**: 75–91.
- Hall, G. and Channell, S. 1985. Differential effects of contextual change on latent inhibition and on the habituation of an orienting response. *J. Exp. Psychol. Anim. Behav. Processes* **11**: 470–481.
- Hall, G. and Honey, R.C. 1990. Context-specific conditioning in the conditioned-emotional-response procedure. *J. Exp. Psychol. Anim. Behav. Processes* **16**: 271–278.
- Hall, G. and Mondragón, E. 1998. Contextual control as occasion setting. In *Occasion setting: Associative learning and cognition in animals* (eds. N. Schmajuk and P.C. Holland), pp. 199–222. American Psychological Association, Washington, DC.
- Harris, J.A., Jones, M.L., Bailey, G.K., and Westbrook, R.F. 2000. Contextual control over conditioned responding in an extinction paradigm. *J. Exp. Psychol. Anim. Behav. Processes* **26**: 174–185.
- Hart, J.A., Bourne, M.J., and Schachtman, T.R. 1995. Slow reacquisition of a conditioned taste aversion. *Anim. Learn. Behav.* **23**: 297–303.
- Haselgrove, M. and Pearce, J.M. 2003. Facilitation of extinction by an increase or a decrease in trial duration. *J. Exp. Psychol. Anim. Behav. Processes* **29**: 153–166.
- Haselgrove, M., Aydin, A., and Pearce, J.M. 2004. A partial reinforcement extinction effect despite equal rates of reinforcement during Pavlovian conditioning. *J. Exp. Psychol. Anim. Behav. Processes* **30**: 240–250.
- Holland, P.C. 1992. Occasion setting in Pavlovian conditioning. In *The psychology of learning and motivation*, Vol. 28 (ed. D.L. Medin), pp. 69–125. Academic Press, San Diego, CA.
- Honey, R.C., Hall, G., and Bonardi, C. 1993. Negative priming in associative learning: Evidence from a serial-conditioning procedure. *J. Exp. Psychol. Anim. Behav. Processes* **19**: 90–97.
- Kaye, H. and Pearce, J.M. 1984. The strength of the orienting response during Pavlovian conditioning. *J. Exp. Psychol. Anim. Behav. Processes* **10**: 90–109.
- Kehoe, E.J. 1988. A layered network model of associative learning: Learning to learn and configuration. *Psychol. Rev.* **95**: 411–433.
- Kehoe, E.J. and Macrae, M. 1997. Savings in animal learning: Implications for relapse and maintenance after therapy. *Behav. Ther.* **28**: 141–155.
- Kimmel, H.D. and Burns, R.A. 1975. Adaptational aspects of conditioning. In *Handbook of learning and cognitive processes: Conditioning and behavior theory*, Vol. 2 (ed. W.K. Estes), pp. 99–142. Erlbaum, Hillsdale, NJ.
- Konorski, J. 1948. *Conditioned reflexes and neuron organization*. Cambridge University Press, Cambridge, UK.
- . 1967. *Integrative activity of the brain*. University of Chicago Press, Chicago.
- Kraemer, P.J., Randall, C.K., and Carby, T.J. 1991. Release from latent inhibition with delayed testing. *Anim. Learn. Behav.* **19**: 139–145.
- Kremer, E.F. 1978. The Rescorla-Wagner model: Losses in associative strength in compound conditioned stimuli. *J. Exp. Psychol. Anim. Behav. Processes* **4**: 22–36.
- Krupa, D.J. and Thompson, R.F. 2003. Inhibiting the expression of a classically conditioned behavior prevents its extinction. *J. Neurosci.* **23**: 10577–10584.
- Lattal, K.M. and Nakajima, S. 1998. Overexpectation in appetitive Pavlovian and instrumental conditioning. *Anim. Learn. Behav.* **26**: 351–360.
- Lubow, R.E. and De la Casa, L.G. 2002. Superlatent inhibition and spontaneous recovery: Differential effects of pre- and

- postconditioning CS-alone presentations after long delays in different contexts. *Anim. Learn. Behav.* **30**: 376–386.
- Mackintosh, N.J. 1974. *The psychology of animal learning*. Academic Press, New York.
- Maren, S. and Holt, W. 2000. The hippocampus and contextual memory retrieval in Pavlovian conditioning. *Behav. Brain Res.* **110**: 97–108.
- McClelland, J.L. and Rumelhart, D.E. 1985. Distributed memory and the representation of general and specific information. *J. Exp. Psychol. Gen.* **114**: 159–188.
- McCloskey, M. and Cohen, N.J. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *The psychology of learning and motivation*, Vol. 24 (ed. G.H. Bower), pp. 109–165. Academic Press, San Diego, CA.
- Mowrer, O.H. and Jones, H.M. 1945. Habit strength as a function of the pattern of reinforcement. *J. Exp. Psychol.* **35**: 293–311.
- Myers, K.M. and Davis, M. 2002. Behavioral and neural analysis of extinction. *Neuron* **36**: 567–584.
- Nakajima, S., Tanaka, S., Urshihara, K., and Imada, H. 2000. Renewal of extinguished lever-press responses upon return to the training context. *Learn. Motiv.* **31**: 416–431.
- Napier, R.M., Macrae, M., and Kehoe, E.J. 1992. Rapid reacquisition in conditioning of the rabbit's nictitating membrane response. *J. Exp. Psychol. Anim. Behav. Processes* **18**: 182–192.
- Nelson, J.B. 2002. Context specificity of excitation and inhibition in ambiguous stimuli. *Learn. Motiv.* **33**: 284–310.
- Nelson, J.B. and Bouton, M.E. 1997. The effects of a context switch following serial and simultaneous feature-negative discriminations. *Learn. Motiv.* **28**: 56–84.
- Overton, D.A. 1985. Contextual stimulus effects of drugs and internal states. In *Context and learning* (eds. P.D. Balsam and A. Tomie), pp. 357–384. Erlbaum, Hillsdale, NJ.
- Pavlov, I.P. 1927. *Conditioned reflexes*. Oxford University Press, Oxford, UK.
- Pearce, J.M. 1994. Similarity and discrimination: A selective review and a connectionist model. *Psychol. Rev.* **101**: 587–607.
- Pearce, J.M. and Hall, G. 1980. A model for Pavlovian conditioning: Variations in the effectiveness of conditioned but not unconditioned stimuli. *Psychol. Rev.* **87**: 332–352.
- Peck, C.A. and Bouton, M.E. 1990. Context and performance in aversive-to-appetitive and appetitive-to-aversive transfer. *Learn. Motiv.* **21**: 1–31.
- Rauhut, A.S., Thomas, B.L., and Ayres, J.J.B. 2001. Treatments that weaken Pavlovian conditioned fear and thwart its renewal in rats: Implications for treating human phobias. *J. Exp. Psychol. Anim. Behav. Processes* **27**: 99–114.
- Reid, R.L. 1958. The role of the reinforcer as a stimulus. *B. J. Psychol.* **49**: 202–209.
- Rescorla, R.A. 1993. Inhibitory associations between S and R in extinction. *Anim. Learn. Behav.* **21**: 327–336.
- . 1996. Preservation of Pavlovian associations through extinction. *Q. J. Exp. Psychol.* **49B**: 245–258.
- . 1997. Response inhibition in extinction. *Q. J. Exp. Psychol.* **50B**: 238–252.
- . 2000. Extinction can be enhanced by a concurrent exciter. *J. Exp. Psychol. Anim. Behav. Processes* **26**: 251–260.
- . 2001. Experimental extinction. In *Handbook of contemporary learning theories* (eds. R.R. Mowrer and S.B. Klein), pp. 119–154. Erlbaum, Mahwah, NJ.
- . 2003. Protection from extinction. *Learn. Behav.* **31**: 124–132.
- Rescorla, R.A. and Wagner, A.R. 1972. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In *Classical conditioning II: Current research and theory* (eds. A.H. Black and W.K. Prokasy), pp. 64–99. Appleton-Century-Crofts, New York.
- Rescorla, R.A. and Heth, C.D. 1975. Reinstatement of fear to an extinguished conditioned stimulus. *J. Exp. Psychol. Anim. Behav. Processes* **1**: 88–96.
- Ricker, S.T., and Bouton, M.E. 1996. Reacquisition following extinction in appetitive conditioning. *Anim. Learn. Behav.* **24**: 423–436.
- Robbins, S.J. 1990. Mechanisms underlying spontaneous recovery in auto-shaping. *J. Exp. Psychol. Anim. Behav. Processes* **16**: 235–249.
- Rosas, J.M. and Alonso, G. 1996. Temporal discrimination and forgetting of CS duration in conditioned suppression. *Learn. Motiv.* **27**: 43–57.
- . 1997. The effect of context change upon long-term memory of CS duration. *Behav. Processes* **39**: 69–76.
- Rosas, J.M. and Bouton, M.E. 1997. Additivity of the effects of retention interval and context change on latent inhibition: Toward resolution of the context forgetting paradox. *J. Exp. Psychol. Anim. Behav. Processes* **23**: 283–294.
- . 1998. Context change and retention interval can have additive, rather than interactive, effects after taste aversion extinction. *Psychonomic Bull. Rev.* **5**: 79–83.
- Schmajuk, N.A. and Holland, P.C., eds. 1998. *Occasion-setting: Associative learning and cognition in animals*. American Psychological Association, Washington, DC.
- Soltysik, S.S., Wolfe, G.E., Nicholas, T., Wilson, J., and Garcia-Sanchez, J.L. 1983. Blocking of inhibitory conditioning within a serial conditioned stimulus-conditioned inhibitor compound: Maintenance of acquired behavior without an unconditioned stimulus. *Learn. Motiv.* **14**: 1–29.
- Spear, N.E. 1981. Extending the domain of memory retrieval. In *Information processing in animals: Memory mechanisms* (eds. N.E. Spear and R.R. Miller), pp. 341–378. Erlbaum, Hillsdale, NJ.
- Spear, N.E., Smith, G.J., Bryan, R., Gordon, W., Timmons, R., and Chiszar, D. 1980. Contextual influences on the interaction between conflicting memories in the rat. *Anim. Learn. Behav.* **8**: 273–281.
- Staddon, J.E.R. 1988. Learning as inference. In *Evolution and learning* (eds. R.C. Bolles and M.D. Beecher), pp. 59–78. Erlbaum, Hillsdale, NJ.
- Swartztruber, D. 1995. Modulatory mechanisms in Pavlovian conditioning. *Anim. Learn. Behav.* **23**: 123–143.
- Swartztruber, D. and Bouton, M.E. 1992. Context sensitivity of conditioned suppression following preexposure to the conditioned stimulus. *Anim. Learn. Behav.* **20**: 97–103.
- Tamai, N. and Nakajima, S. 2000. Renewal of formerly conditioned fear in rats after extensive extinction training. *Int. J. Comp. Psychol.* **13**: 137–147.
- Thomas, B.L. and Ayres J.J.B. 2004. Use of the ABA fear renewal paradigm to assess the effects of extinction with co-present fear inhibitors or exciters: Implications for theories of extinction and for treating human fears and phobias. *Learn. Motiv.* **35**: 22–51.
- Thomas, D.R., McKelvie, A.R., Ranney, M., and Moye, T.B. 1981. Interference in pigeons' long-term memory viewed as a retrieval problem. *Anim. Learn. Behav.* **9**: 581–586.
- Thomas, D.R., Moye, T.B., and Kmose, E. 1984. The recency effect in pigeons' long-term memory. *Anim. Learn. Behav.* **12**: 21–28.
- Vigorito, M. and Ayres, J.J.B. 1987. Effect of naloxone on conditioned suppression in rats. *Behav. Neurosci.* **101**: 576–586.
- Vogel, E.H., Brandon, S.E., and Wagner, A.R. 2003. Stimulus representation in SOP, II: An application to inhibition of delay. *Behav. Processes* **62**: 27–48.
- Wagner, A.R. 1969. Stimulus selection and a "modified continuity theory." In *The psychology of learning and motivation*, Vol. 3 (eds. G.H. Bower and J.T. Spence), pp. 1–41. Academic Press, New York.
- . 1981. SOP: A model of automatic memory processing in animal behavior. In *Information processing in animals: Memory mechanisms* (eds. N.E. Spear and R.R. Miller) pp. 5–47. Erlbaum, Hillsdale, NJ.
- Wagner, A.R. and Brandon, S.E. 1989. Evolution of a structured connectionist model of Pavlovian conditioning (AESOP). In *Contemporary learning theories: Pavlovian conditioning and the status of traditional learning theory* (eds. S.B. Klein and R.R. Mowrer), pp. 149–189. Lawrence Erlbaum Associates, Hillsdale, NJ.
- . 2001. A componential theory of Pavlovian conditioning. In *Handbook of contemporary learning theories* (eds. R.R. Mowrer and S.B. Klein), pp. 23–64. Erlbaum, Mahwah, NJ.
- Weidemann, G. and Kehoe, E.J. 2003. Savings in classical conditioning in the rabbit as a function of extended extinction. *Learn. and Behav.* **31**: 49–68.
- Westbrook, R.F., Jones, M.L., Bailey, G.K., and Harris, J.A. 2000. Contextual control over conditioned responding in a latent inhibition paradigm. *J. Exp. Psychol. Anim. Behav. Processes* **26**: 157–173.
- Wilson, A., Brooks, D.C., and Bouton, M.E. 1995. The role of the rat hippocampal system in several effects of context in extinction. *Behav. Neurosci.* **109**: 828–836.