



Published in final edited form as:

Behav Processes. 2014 May ; 0: 13–19. doi:10.1016/j.beproc.2014.02.012.

A fundamental role for context in instrumental learning and extinction

Mark E. Bouton and **Travis P. Todd**

University of Vermont

Abstract

The purpose of this article is to review recent research that has investigated the effects of context change on instrumental (operant) learning. The first part of the article discusses instrumental extinction, in which the strength of a reinforced instrumental behavior declines when reinforcers are withdrawn. The results suggest that extinction of either simple or discriminated operant behavior is relatively specific to the context in which it is learned: As in prior studies of Pavlovian extinction, ABA, ABC, and AAB renewal effects can all be observed. Further analysis supports the idea that the organism learns to refrain from making a specific response in a specific context, or in more formal terms, an inhibitory context-response association. The second part of the article then discusses research suggesting that the context also controls instrumental behavior before it is extinguished. Several experiments demonstrate that a context switch after either simple or discriminated operant training causes a decrement in the strength of the response. Over a range of conditions, the animal appears to learn a direct association between the context and the response. Under some conditions, it can also learn a hierarchical representation of context and the response-reinforcer relation. Extinction is still more context-specific than conditioning, as indicated by ABC and AAB renewal. Overall, the results establish that the context can play a significant role in both the acquisition and extinction of operant behavior.

The issue of the contextual control of behavior has been a focus of research in learning theory for many years (e.g., Balsam & Tomie, 1985). One reason is that contextual cues are thought to be essential in supporting memory retrieval, which predicts that if retention is tested in a context that is different from the context in which information is learned, there should be a decline in performance (e.g., Spear, 1978; Tulving & Thomson, 1972). A second reason is that a number of influential models of associative learning have given context a central role (e.g., Rescorla & Wagner, 1972; Pearce, 1994; Pearce & Hall, 1980; Wagner, 1978, 2003). For many years, our laboratory has therefore studied the role of context in Pavlovian conditioning, where organisms learn to associate a conditioned stimulus (CS) with a reinforcer or unconditioned stimulus (US). In recent years, however, we have begun

© 2014 Elsevier B.V. All rights reserved.

Correspondence: Mark E. Bouton, Psychology Dept., University of Vermont, Burlington, VT 05405-0134, Phone: 802 656-4164, Fax: 802 656-8783, mark.bouton@uvm.edu.

Publisher's Disclaimer: This is a PDF file of an unedited manuscript that has been accepted for publication. As a service to our customers we are providing this early version of the manuscript. The manuscript will undergo copyediting, typesetting, and review of the resulting proof before it is published in its final citable form. Please note that during the production process errors may be discovered which could affect the content, and all legal disclaimers that apply to the journal pertain.

to extend our analysis to the role of context in operant or instrumental learning, where organisms learn to associate their behavior with reinforcers or outcomes. The purpose of the present article is to review some of the work we have done to date on the latter problem, i.e., the contextual control of operant learning.

For some time, our laboratory has also been interested in *extinction*, an especially fundamental process of behavior change. Extinction learning is essential for the survival of organisms, because it allows them to adapt to changes in their environment. During extinction, responding declines when the contingencies in the environment change such that the US or reinforcer is no longer presented. In Pavlovian conditioning, the CS is repeatedly presented in the absence of the US, and the original conditioned response (CR) decreases. Although it has always been tempting to view extinction as a weakening or erasure of the original learning (e.g., Rescorla & Wagner, 1972), there is strong evidence that this is not the case. Instead, extinction results in new learning that is at least partly dependent on the context for expression (e.g., Bouton, 2002, 2004).

As suggested by our comments above, we have recently become interested in the extinction of instrumental, or operant, behavior. In operant extinction, responding declines when the reinforcer is no longer presented. Operant extinction learning might be a relevant model for understanding the suppression or inhibition of problematic voluntary behaviors in humans, such as overeating, gambling, and drug addiction. Further, operant extinction has become a powerful tool in studies of the neurobiology of drug addiction (for a review see Marchant, Li, & Shaham, 2013). In order to have a complete neurobiological theory of addiction, it is important to understand the behavioral principles of operant extinction (e.g., Todd, Vurbic, & Bouton, 2014a). Thus, the study of operant extinction learning can have implications for understanding the elimination of problem behaviors in humans as well as the neural mechanisms that underlie it.

In what follows, we summarize recent research from our laboratory that has examined the contextual control of operant behavior. The first section of the paper focuses on operant extinction, which began as an investigation of its possible parallels with Pavlovian learning. Although the research has uncovered similarities between Pavlovian and operant extinction, it has uncovered some important differences as well. The second part of the paper then reviews research that has examined the contextual control of operant behavior before it has been extinguished. As a whole, the research has uncovered a rather central role of the context in controlling both acquisition and extinction. It also tentatively suggests something specific about the mechanism of contextual control: In conditioning, the animal seems to learn to make a specific response in a specific context, and in extinction, it seems to learn to inhibit a specific response in a specific context.

Contextual control of operant extinction

As described above, it is now commonly understood that Pavlovian extinction results in new, context-dependent learning. One reason for this belief is that responding to an extinguished CS will return if the CS is tested outside the extinction context. This is known as the *renewal effect* (Bouton & Bolles, 1979). (In animal experiments, contexts are usually

defined as the chambers in which conditioning occurs; they typically differ in visual, tactile, and olfactory characteristics.) For example, after CS-US pairings in Context A, and CS alone presentations (extinction) in Context B, responding will return (renew) when the CS is subsequently tested in the original Context A (ABA renewal; Bouton & Bolles, 1979; Bouton & King, 1983; Bouton & Peck, 1989) or in a third, relatively neutral context (ABC renewal; Bouton & Bolles, 1979; Harris, Jones, Bailey, & Westbrook, 2000; Thomas, Larsen, & Ayres, 2003). Renewal is also observed when both conditioning and extinction occur in Context A, and the CS is then tested in Context B (AAB renewal; Bouton & Ricker, 1994; Laborda, Witnauer, & Miller, 2011; Tamai & Nakajima, 2000). The AAB and ABC forms of renewal are especially important at the theoretical level because they indicate that mere removal from the context of extinction is sufficient for renewal to occur. Thus, the extinction context somehow inhibits behavior, so that removal of the CS from that context can turn on responding to the CS again.

Does operant extinction also result in context-dependent inhibitory learning? Until recently, the degree to which it did had been unclear. Although ABA renewal had been routinely demonstrated with either food or drug reinforcers (e.g., Bossert, Liu, Lu, & Shaham, 2004; Chaudri, Sahuque, & Janak, 2009; Crombag & Shaham, 2002; Hamlin, Clemens, & McNally, 2008; Hamlin, Newby, & McNally, 2007; Nakajima, Tanaka, Urushihara, & Imada, 2000; Welker & McAuley, 1978; Zironi, Burattini, Aircardi, & Janak, 2006), several reports had failed to demonstrate AAB renewal (see Bossert et al., 2004; Crombag & Shaham, 2002; Nakajima et al., 2000) and the evidence for ABC renewal was mixed (e.g., Zironi et al., 2006). The lack of evidence of AAB and ABC renewal left unanswered the crucial question of whether mere removal from the extinction context was sufficient to cause response recovery.

Recent research from our laboratory, however, has demonstrated all three forms of renewal after instrumental extinction (e.g., Bouton, Todd, Vurbic, & Winterbauer, 2011). In a representative experiment (Bouton et al., 2011, Experiment 1), rats first learned to lever-press for food pellets on a variable-interval 30 s (VI 30 s) schedule (pellets were made available on average every 30 s at which point the next lever press resulted in their delivery). After initial training in Context A, they then underwent extinction (in which lever presses no longer resulted in pellet delivery) in either Context A or B. Over the course of extinction, lever pressing declined and reached a very low rate. Finally, using a within-subject test procedure, all rats received a test session in both Contexts A and B (order counterbalanced) in which they could lever press, but no pellets were delivered. Both ABA and AAB renewal were observed. For rats trained in Context A and extinguished in B, lever pressing renewed when it was tested back in the original training context (A). Importantly, for rats trained and extinguished in Context A, responding also renewed when they were tested in Context B (see also Todd, Winterbauer, & Bouton, 2012a). In a separate experiment, renewal also occurred when training, extinction, and testing all occurred in different contexts (ABC renewal; see also Todd, Winterbauer, & Bouton, 2012b). Thus, for the first time the experiments provided clear evidence that testing outside the context of extinction is sufficient to cause renewal of operant behavior.

We also began to examine the behavioral mechanisms that might contribute to the phenomenon. Although the ABC and AAB effects suggest a role for some form of inhibition provided by the extinction context, our initial study (Bouton et al., 2011; Experiment 1) also found that the ABA renewal effect was significantly stronger than the AAB effect. This may have occurred because in the ABA design subjects are returned to a context that has been directly associated with the reinforcer. Other research has indicated that operant responding can be modulated after extinction by direct context-reinforcer associations (Baker, Steinwald, & Bouton, 1991). However, we found that extensive extinction exposure to Context A without the lever present prior to renewal testing did not reduce the strength of the ABA renewal effect (Bouton et al., 2011, Experiment 4). Exposure to Context A should have weakened any context-reinforcer associations that might be present. The fact that renewal was not weakened by extensive exposure to the renewal context suggested that the responding observed in the renewal context was likely not a function of direct association between the context and the reinforcer. Instead, Context A could have enabled strong ABA renewal by setting the occasion for the response-reinforcer relation or by directly eliciting the response-- two mechanisms we will return to in the second half of this paper.

In a subsequent set of experiments, Todd (2013) examined possible mechanisms of how the extinction context controls extinction. Theoretically, there were at least three possibilities. Perhaps the simplest was that during extinction the context might acquire a direct inhibitory association with the representation of the reinforcer (Polack, Laborda, & Miller, 2011; but see Bouton & King, 1983; Bouton & Swartzentruber, 1986, 1989). Mechanistically, this might occur because a strong excitatory association might be formed between the response and the reinforcer during acquisition. Then, during extinction, the surprising omission of the expected reinforcer would cause inhibition to accrue between the context and the reinforcer (cf. Rescorla & Wagner, 1972). A second possibility is that, instead of directly inhibiting the representation of the reinforcer, the extinction context might activate an inhibitory association between the response and the reinforcer (Bouton & Nelson, 1994, 1998; see also Swartzentruber & Rescorla, 1993). This type of mechanism, negative occasion setting (e.g., Holland, 1992), has been emphasized in Pavlovian extinction, primarily because a number of experiments have failed to detect simple inhibition in the extinction context (Bouton & King, 1983; Bouton & Swartzentruber, 1986, 1989). A reasonably well-accepted version of the mechanism is depicted in Figure 1 (e.g., Bouton & Nelson, 1994). Notice that the effect of the context is to influence an association of the CS with the US rather than enter into a direct association with the US. Finally, a third possible mechanism of operant extinction is that the context might come to directly inhibit the response. Rescorla (1993, 1997) has suggested this account of the extinction of instrumental responding in the presence of a discriminative stimulus, although he did not provide evidence that separated it from the occasion-setting account (see Rescorla, 1993, p. 335; Rescorla, 1997, p. 249). Nonetheless, according to any of these mechanisms, removal from the extinction context would cause renewal to occur due to a release from the context-specific inhibition learned during extinction.

To explore the different possibilities, Todd (2013) examined operant renewal when the extinction and renewal contexts were equated on their histories of reinforcement and nonreinforcement (e.g., Bouton & Ricker, 1994; Harris et al., 2000; Lovibond, Preston, &

Mackintosh, 1984; Rescorla, 2008). In Experiment 1, rats were reinforced for performing R1 in Context A, and R2 in Context B (R1 and R2 were lever press and chain pull, counterbalanced). This ensured that A and B were equally associated with the reinforcer during acquisition. Next, R1 underwent extinction in Context B, and R2 in Context A. This then ensured both contexts were equally associated with nonreinforcement. In a final renewal test, each rat was then tested in Context A and B with R1 or R2 present in separate sessions. There was a clear ABA (and BAB) renewal effect: Responding on R1 was high in A, but low in B, and responding on R2 was high in B, but low in A. The fact that renewal was observed when the contexts' direct associations with the reinforcer had been controlled indicated that differential context-reinforcer associations are not required for renewal. This is entirely consistent with the findings discussed earlier (Bouton et al., 2011, Experiment 4). Using analogous designs, Todd (2013) also demonstrated AAB and ABC renewal. Further, in all experiments, final tests with both R1 and R2 available simultaneously demonstrated that the context had an impact on *choice* behavior. The rats consistently preferred the response that was not in the context in which it had been extinguished. As a whole, the results suggested that the extinction context does not merely enter into a direct inhibitory association with the reinforcer. Instead, they suggested that the context can have a specific effect on the response, and thus pointed to a role for either the occasion-setting or inhibitory context-response mechanisms described above.

In another experiment, Todd (2013, Experiment 4) attempted to separate the occasion-setting and inhibitory context-response mechanisms. He reasoned that if the extinction context were acting as a negative occasion setter, then its occasion setting power should transfer and inhibit a second, similarly-trained target response. This type of transfer has been routinely observed in studies of operant occasion-setting by Peter Holland and colleagues (e.g., Holland & Coldwell, 1993; Morell & Holland, 1993). Essentially, the power of a negative occasion setter should transfer to a separately trained target if that target has been trained as a target of a different negative occasion setter. (Because this type of transfer is often not complete, even if the occasion setting mechanism were operating, renewal could have been detected in Todd's other experiments.) The design of the experiment is sketched in Table 1. Todd first trained two groups of rats on R1 in Context A and R2 in Context B. Next, both groups received extinction of R2 in Context A. The groups differed, however, with respect to the context in which R1 was extinguished. For one group (Ext-B), R1 was extinguished in Context B, potentially making it a negative occasion setter. For the other group (Ext-C), R1 was extinguished in an irrelevant Context C. In a renewal test, R2 was finally tested in both Context A (its extinction context) and in Context B (its renewal context). Renewal should have been weaker for Group Ext-B than Ext-C if Context B had acquired negative occasion setting properties. However, it was not; renewal was equivalent in the two groups. This result suggested that extinction of R1 in Context B had not given the context transferrable negative occasion-setting properties. Instead, the results are more consistent with the idea that the extinction context specifically inhibits a particular response. Renewal thus occurs upon removal from the extinction context because of a release from this context-specific response inhibition.

Todd, Vurbic, and Bouton (2014b) went on to seek additional evidence of an inhibitory context-response association developing in extinction. If such an association develops in extinction, they reasoned that it should weaken renewal of the same response if it were later tested in that context. To test the idea, they turned to a discriminated operant procedure. In the method, R1 and R2 (lever-pressing and chain-pulling, counterbalanced) were only reinforced in the presence of light or tone discriminative stimuli (S1 and S2, counterbalanced). They first demonstrated ABA renewal when Contexts A and B are equally associated with reinforcement and nonreinforcement (Experiment 1).

Todd et al. (2014b) went on to test for AAB renewal (Experiment 3), and showed that, consistent with the inhibitory context-response association hypothesis, it was eliminated when the response had been extinguished in the renewal context prior to testing. The design of the experiment is presented in Table 2. As illustrated there, a single group of rats learned four S/response combinations in three different contexts. S1R1 and S2R2 were trained in Context A. S3R1 was trained in Context B, and S4R2 was trained in Context C. Following acquisition, each S/response combination was extinguished in the context in which it was trained. Finally, S1R1 and S2R2 were tested for renewal in Contexts B and C. Conceptually, this would provide the opportunity for an AAB (AAC) renewal effect. However, because R2 had previously been extinguished in Context C, S2R2 was expected to be low in C but S1R1 was expected to be high. Conversely, because R1 had previously been extinguished in Context B, S3R1 was expected to be low but S2R2 was expected to be high. That was exactly the result obtained: Renewal was reduced when the same response had been previously extinguished in the renewal context. Further, by demonstrating AAB renewal, it is clear that removal from the extinction context is sufficient for the renewal of discriminated operant behavior.

Our studies of operant extinction and renewal thus confirm that operant extinction results in new, context-dependent, inhibitory learning. Although Pavlovian extinction also depends in part on context-dependent inhibition, the mechanism in operant extinction appears to be different. In Pavlovian extinction, the context appears to operate as a negative occasion setter and activates an inhibitory association between the CS and the US, as illustrated in Figure 1 (e.g., Bouton, 2004). However, our new evidence suggests that in operant extinction the context might simply and directly inhibit the response (Todd, 2013; Todd et al., 2014b). Such a simple, direct inhibitory context-response association can readily account for two key findings. First, there is no evidence that inhibition in a context transfers between different responses (Todd, 2013, Experiment 4), as negative occasion-setting should. Second, renewal is attenuated if the same, compared to a different, response has been inhibited (extinguished) in the renewal context before renewal testing (Todd et al., 2014b, Experiment 3). We should note, however, that there may also be an additional type of inhibition learned during extinction of discriminated operants. For example, some renewal can still be observed when both the extinction and renewal contexts are trained to inhibit the same response (Todd et al., 2014b, Experiment 2). In this case, renewal cannot be due to a release from context-specific response inhibition. Thus, the extinction context might also modulate some other association. One possibility is that, in addition to directly inhibiting the R, the context might also modulate an association between the S and the reinforcer. This mechanism is conceptually identical to the well-accepted mechanism of Pavlovian

extinction (e.g., Bouton, 2004; Bouton & Ricker, 1994). A complete understanding of this second mechanism will require further investigation, however.

In summary, the extinction of operant behavior appears to result in a direct inhibitory association between the context and the response. In effect, the animal learns not to make a particular response in a particular context. Renewal thus occurs when the context is changed after extinction because the response is released from context-specific response inhibition. In some cases, such as the extinction of discriminated operant behavior, an additional mechanism may also be operating that influences the effectiveness of the S.

Contextual control of non-extinguished operant behavior

Another interesting result that emerged in our studies of renewal in operant conditioning was that the context also controls the operant response before extinction has occurred. For example, in an experiment described earlier (Bouton et al., 2011, Experiment 1), different groups of rats received operant conditioning in Context A and then extinction of the response in either the same context or a different context (Context B). Rats that were being extinguished in Context B showed less responding from the outset than rats that were being extinguished in Context A. Todd (2013, Experiment 2) found the same context-switch effect in the experiments described above that controlled for the contexts' associative histories. That is, when R1 training in Context A was intermixed with R2 training in Context B, a switch of either response to the other context at the start of extinction also weakened the level of responding. Thus, the contextual control did not depend on differential reinforcement in the two contexts. The findings were unexpected, because many prior studies of Pavlovian conditioning have not produced the analogous effect. That is, if we conduct CS-US pairings in Context A and then extinguish the CS in either Contexts A or B, we (and others) have found that the context switch typically does not attenuate the response evoked by the CS (e.g., Bouton & King, 1983; Bouton & Peck, 1989; Hall & Honey, 1989; Harris et al., 2000; Nelson, 2002; Rosas & Bouton, 1998; see Rosas, Todd, & Bouton, 2013). In our laboratory, operant responding thus seems more specific than Pavlovian responding to the context in which it is learned.

One explanation of the difference between our operant and Pavlovian results was that the operant response had been intermittently reinforced with a VI 30-s reinforcement schedule. In contrast, the Pavlovian experiments have usually studied CSs that have always been reinforced or paired with the US. The difference is relevant, because Juan Rosas (e.g., Rosas & Callejas-Aguilera, 2006) has argued that surprising trial outcomes can direct attention to the context and thereby encourage the encoding of information with the context. In a VI schedule, the outcome of each response (i.e., whether it is reinforced or not) is never certain, and the resulting continued surprise after each response could therefore encourage attention to the context. To test this possibility, Bouton, Todd, and León (2014, Experiment 4) examined the effect of a context switch on operant lever-pressing and chain-pulling responses that had been continuously reinforced. The context switch still decreased the strength of the operant response. Thus, the decremental effect of the context switch does not depend on intermittent reinforcement.

Another difference between the operant and Pavlovian methods was that in the simple operant procedure, the only stimulus available for association with the response was in fact the background context. That is, there was no signal analogous to a CS that could compete with (i.e., block or overshadow) learning about the context. We therefore ran experiments that tested the effects of context change after *discriminated* operant training in which the response was only reinforced during a 30-s discriminative stimulus (S) (Bouton et al., 2014). [We had previously used a 30-s CS in studies testing the effects of a context switch after Pavlovian conditioning (e.g., Bouton, Frohardt, Sunsay, Waddell, & Morris, 2008; Nelson, 2002).] In this method, as training progresses, responses and reinforcers become confined to S. In an initial experiment (Bouton et al., 2014, Experiment 1), rats learned to press a lever that was only reinforced during S. Remarkably, when the context was changed, and responding was tested in the presence of S, there was still a substantial loss of responding, as we had seen in our simple operant methods. Thus, even in the presence of an S that controls responding, the context still plays a significant role.

Further experiments examined the context's mechanism of action in more detail. One of them (Bouton et al., 2014, Experiment 2) used the design illustrated in Table 3. Four groups first learned to perform a target response (R1, either lever pressing or chain pulling, counterbalanced) in the presence of an S (S1, either a light or a tone, counterbalanced) in Context A. In the final phase, each rat was tested for R1 responding in the presence of S1 in both Contexts A and B. The groups differed only in the treatments they received during training sessions in Context B that were intermixed with the ones in Context A. One group was simply exposed to Context B (with no response manipulandum or reinforcers presented) during those sessions; this group showed a substantial decrement in S1R1 responding when it was finally tested in Context B. A second group received response-noncontingent pairings of a second S (S2) in Context B with food pellets at a rate comparable to that received in S1 in Context A. These animals also showed a decrement in responding when S1R1 was first tested in Context B, even though the method had equated Contexts A and B on their Pavlovian association with the reinforcer. The third group (Diff R) received Context B training sessions in which a second response (R2) was reinforced in the presence of a second S (S2). Although this treatment equated Contexts A and B on their connections with operant as well as Pavlovian contingencies, there was still a substantial response decrement when S1R1 was first tested in Context B. The final group (Same R) received Context B training sessions in which a second S (S2) set the occasion for the same response that was being reinforced in Context A. In this case, when S1R1 was first tested in Context B, the response was not weakened. Notice that for this group, S1 was presented for the first time in Context B during the final tests. Analogous to a Pavlovian CS, S1's ability to control responding (R1) transferred perfectly across the two contexts.

The results with the "Same R" group just described were replicated in another experiment (Bouton et al., 2014, Experiment 3) that also tested a group given "Same S" training in which the same S1 set the occasion for R1 and R2 in Contexts A and B, respectively. Here, when S1R1 was tested for the first time in Context B, there was a strong decrement in responding. Thus, even when the rat was tested with an S that had set the occasion for responding in Context B, the strength of R1 decreased when it was switched and tested

there. The results clearly indicated that it is the strength of the response, and not the effectiveness of the discriminative stimulus, that is lost when the context is changed after discriminated operant training.

How are we to understand the context-switch effect theoretically? Bouton et al. (2014) suggested at least the two possibilities summarized in Figure 2. In the first, the context simply and directly evokes the response in the manner indicated in Panel a. The animal has a tendency to respond in the context in which the instrumental behavior is learned; any effect of the discriminative stimulus (via its direct association with R, its association with the reinforcing outcome (O), or its hierarchical connection with R-O) would be over and above that. In the second possibility, the context might work more hierarchically by signaling (or setting the occasion for) the R-O relation. Again, to account for the results, it is necessary to assume that the context operates on the response before the S can further modulate it—by any of the S-O, S-R, or S(R-O) mechanisms mentioned above.

More recent research has begun to uncover evidence of both the Context-R and hierarchical Context-(R-O) possibilities. Thraillkill and Bouton (in preparation) returned to the simple free-operant method (in which there are no explicit discriminative light or tone stimuli) and produced evidence consistent with a direct Context-R association (analogous to Figure 2's panel a.). They first found that the context switch effect occurs after an impressively wide range of training procedures. For example, in one experiment, rats received a context switch after training in Context A on a Random Ratio 15 schedule for either a small number of sessions (2 sessions and 120 reinforcers), an intermediate number of sessions (6 sessions and 360 reinforcers), or many sessions (18 sessions and 1080 reinforcers). Rats in each of these groups were also yoked to rats in other groups that could earn a reinforcer for the next response after the ratio-trained animal had earned one. This created three additional groups that received reinforcement with random-interval schedules with the same reinforcement rates. During final tests in which all rats were tested in both Contexts A and B, there was a decrement in operant responding in Context B regardless of the reinforcement schedule or the extensiveness of the training.

Those results suggested that the dependence of operant behavior on context occurs over conditions that are supposed to engender goal-directed actions (ratio schedules or minimal training) and habits (interval schedules after extensive training) (e.g., Dickinson, 1994; Dickinson, Nicholas, & Adams, 1983). Habits and actions theoretically correspond to the two panels in Figure 2. Whereas goal-directed actions are thought to depend on a representation of the goal or reinforcer (as in panel b.), habits are thought to be evoked more directly by some antecedent stimulus (as in panel a.) Actions and habits are usually distinguished by their sensitivity to reinforcer devaluation effects (e.g., Dickinson et al., 1983). Actions, which depend on a reinforcer representation, are changed if the value of the reinforcer is changed. Habits, in contrast, are not. Thraillkill and Bouton (in preparation) therefore went on to test the effects of a context switch after first devaluing the reinforcer by repeatedly pairing it with illness induced by injections of lithium chloride. Responding was finally tested, in extinction, in the context where it had been trained and in a different context. With minimal initial operant training, responding was only partly reduced by reinforcer devaluation. The component that was sensitive to reinforcer devaluation

transferred across contexts. On the other hand, the component that remained after the reinforcer was devalued did not. With more extended training, the operant eventually became completely insensitive to reinforcer devaluation. But it was still weakened by a context switch. Overall, the results suggest that the context plays the role of the antecedent stimulus for habit, and (as in panel a. of Figure 2) controls the response directly, independently of the value of the reinforcer.

In other recent work, Trask and Bouton (in preparation) have produced evidence suggesting an additional role for the alternative Panel b possibility. They used an experimental design analogous to one used by Colwill and Rescorla (1990) in an investigation of the effects of discrete discriminative stimuli. Rats were initially trained to perform two free-operant responses (lever press or chain pull) in two contexts. The responses were reinforced by either sucrose pellets or grain pellets, with the actual contexts, responses, and pellet types being fully counterbalanced. In Context A, R1 led to O1 and R2 led to O2. In contrast, in Context B, R1 led to O2 and R2 led to O1. In a series of subsequent aversion-conditioning trials, O2 was separately paired with lithium chloride injections in both contexts until a complete conditioned aversion was established. Then, in extinction tests, R1 and R2 were tested simultaneously in both Contexts A and B. During testing, the rats performed R1 more than R2 in Context A but R2 more than R1 in Context B. The rats could not have performed this way unless they had learned about the specific combinations of context, response, and outcome, as implied by Figure 3's Panel b. Interestingly, suppression of the response associated with O2 in a given context was not complete, perhaps providing further evidence of additional control by a possible Context-R mechanism (Panel a).

In summary, our results clearly suggest that the context plays a fundamental role in controlling free-operant behavior (Bouton et al., 2011; Bouton et al., 2014; Todd, 2013). Over a range of conditions, a context change weakens the operant response. Evidence to date suggests that the context might often control through a direct association with the response (Thrailkill and Bouton, in preparation), as in Figure 2's panel a. But under certain training conditions, it can also set the occasion for the R-O relation (Trask & Bouton, in preparation), as in Panel b.

Conclusion

Our experiments on the role of context in operant learning have uncovered what appears to be a fundamental role for context after both acquisition and extinction. In the work investigating the effects of context switches after conditioning, the context controls responding when the response is occasion-set by a discriminative stimulus (Bouton et al., 2014) or when it is not (Bouton et al., 2011, 2014; Todd, 2013), and whether the response has been reinforced on ratio or interval schedules and after modest or extensive levels of training (Thrailkill & Bouton, in preparation). In all cases, the animal fundamentally appears to learn to make a specific response in a specific context. Our most recent experiments further suggest that the context can control the behavior by either evoking it directly (as in Panel a. of Figure 2; Thrailkill & Bouton, in preparation) or by hierarchically setting the occasion for the response-reinforcer relation (Panel b. of Figure 2; Trask & Bouton, in preparation). Further research will be required to pinpoint what training procedures are

required for each. It is worth noting, however, that the method used by Trask and Bouton (in preparation) uniquely involved explicit training of different R-O relations in the target contexts.

Up to a point, the role of context we have uncovered in operant *extinction* was anticipated by prior studies of Pavlovian extinction (e.g., Bouton, 2004). Most important, as in Pavlovian extinction, operant extinction performance depends at least partly on being in the context in which extinction has been learned. Further, we should emphasize that the fact that ABC and AAB renewal both occur after operant extinction (e.g., Bouton et al., 2011; Todd, 2013; Todd et al., 2014b) indicates that extinction must be more context-specific than the original operant learning. Moreover, our results point toward the idea that, in both the simple and discriminated-operant situations we have studied, the animal learns to refrain from making a specific response in a specific context. The idea is captured directly by the view that the organism learns an inhibitory association between the context and the response.

We have also noted that another perspective on the role of context is possible (e.g., Bouton et al., 2014). The context-specificity of instrumental responding might reflect the fact that the animal needs to learn what J. J. Gibson called an *affordance* (e.g., Gibson, 1977; see also Bolles, 1980). Without relevant experience, the rat cannot know that the lever or chain is “pressable” or “pullable” in a new context, just as a human cannot know whether a chair or bench in a new context will support her body weight or whether a door in the wall can be opened to permit movement into another room. In such cases, the organism must learn about the behaviors that a new stimulus will support. It remains to be seen whether future experiments will be able to distinguish this kind of approach from context-response or context-(response-outcome) associations noted above.

Acknowledgments

Preparation of the manuscript was supported by NIDA Grant RO1 DA033123. TPT is now at Dartmouth College.

References

- Balsam, P.D.; Tomie, A., editors. Context and learning. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc; 1985.
- Baker AG, Steinwald H, Bouton ME. Contextual conditioning and reinstatement of extinguished instrumental responding. *The Quarterly Journal of Experimental Psychology*. 1991; 43B:199–218.
- Bolles, R.C. Ethological learning theory. In: Gazda, G.M.; Corsini, R.J., editors. *Theories of learning: A comparative approach*. Itasca, IL: F. E. Peacock Publishers; 1980. p. 178-207.
- Bouton ME. Context, ambiguity, and unlearning: Sources of relapse after behavioral extinction. *Biological Psychiatry*. 2002; 52:976–986. [PubMed: 12437938]
- Bouton ME. Context and behavioral processes in extinction. *Learning & Memory*. 2004; 11:485–494. [PubMed: 15466298]
- Bouton ME, Bolles RC. Contextual control of the extinction of conditioned fear. *Learning and Motivation*. 1979; 10:445–466.
- Bouton ME, Frohardt RJ, Sunsay C, Waddell J, Morris RW. Contextual control of inhibition with reinforcement: Adaptation and timing mechanisms. *Journal of Experimental Psychology: Animal Behavior Processes*. 2008; 34:223–236. [PubMed: 18426305]

- Bouton ME, King DA. Contextual control of the extinction of conditioned fear: Tests for the associative value of the context. *Journal of Experimental Psychology: Animal Behavior Processes*. 1983; 9:248–265. [PubMed: 6886630]
- Bouton ME, Peck CA. Context effects on conditioning, extinction, and reinstatement in an appetitive conditioning preparation. *Animal Learning & Behavior*. 1989; 17:188–198.
- Bouton ME, Nelson JB. Context-specificity of target versus feature inhibition in a feature negative discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*. 1994; 20:51–65. [PubMed: 8308493]
- Bouton, ME.; Nelson, JB. Mechanisms of feature-positive and feature-negative discrimination learning in an appetitive conditioning paradigm. In: Schmajuk, N.; Holland, PC., editors. *Occasion setting: Associative learning and cognition in animals*. Washington, DC: American Psychological Association; 1998. p. 69-112.
- Bouton ME, Ricker ST. Renewal of extinguished responding in a second context. *Animal Learning & Behavior*. 1994; 22:317–324.
- Bouton ME, Swartzentruber D. Analysis of the associative and occasion-setting properties of contexts participating in Pavlovian discriminations. *Journal of Experimental Psychology: Animal Behavior Processes*. 1986; 12:333–350.
- Bouton ME, Swartzentruber D. Slow reacquisition following extinction: Context, encoding, and retrieval mechanisms. *Journal of Experimental Psychology: Animal Behavior Processes*. 1989; 15:43–53.
- Bouton ME, Todd TP, León SP. Contextual control of a discriminated operant behavior. *Journal of Experimental Psychology: Animal Learning and Cognition*. 2014; 40:92–105.
- Bouton ME, Todd TP, Vurbic D, Winterbauer NE. Renewal after the extinction of free operant behavior. *Learning & Behavior*. 2011; 39:57–67. [PubMed: 21279496]
- Bossert JM, Liu SY, Lu L, Shaham Y. A role of ventral tegmental area glutamate in contextual cue-induced relapse to heroin seeking. *The Journal of Neuroscience*. 2004; 24:10726–10730. [PubMed: 15564590]
- Chaudri N, Sahuque LL, Janak PH. Ethanol seeking triggered by environmental context is attenuated by blocking dopamine D1 receptors in the nucleus accumbens core and shell in rats. *Psychopharmacology*. 2009; 207:303–314. [PubMed: 19779700]
- Colwill RM, Rescorla RA. Evidence for the hierarchical structure of instrumental learning. *Animal Learning and Behavior*. 1990; 18:71–82.
- Crombag HS, Shaham Y. Renewal of drug seeking by contextual cues after prolonged extinction in rats. *Behavioral Neuroscience*. 2002; 116:169–173. [PubMed: 11895178]
- Dezfouli A, Balleine BW. Habits, action sequences and reinforcement learning. *European Journal of Neuroscience*. 2010; 35:1036–1051. [PubMed: 22487034]
- Dickinson, A. Instrumental conditioning. In: Mackintosh, NJ., editor. *Animal cognition and learning*. London: Academic Press; 1994. p. 4-79.
- Dickinson A, Nicholas DJ, Adams CD. The effect of the instrumental training contingency on susceptibility to reinforcer devaluation. *Quarterly Journal of Experimental Psychology*. 1983; 35B: 35–51.
- Gibson, JJ. The theory of affordances. In: Shaw, R.; Bransford, J., editors. *Perceiving, Acting, and Knowing*. Hillsdale, NJ: Erlbaum; 1977. p. 67-82.
- Hall G, Honey R. Contextual effects in conditioning, latent inhibition, and habituation: Associative and retrieval functions of contextual cues. *Journal of Experiment Psychology: Animal Behavior Processes*. 1989; 15:232–241.
- Hamlin AS, Clemens KJ, McNally GP. Renewal of extinguished cocaine-seeking. *Neuroscience*. 2008; 151:659–670. [PubMed: 18164822]
- Hamlin AS, Newby J, McNally GP. The neural correlates and role of D1 dopamine receptors in renewal of extinguished alcohol-seeking. *Neuroscience*. 2007; 146:525–536. [PubMed: 17360123]
- Harris JA, Jones ML, Bailey GK, Westbrook RF. Contextual control over conditioned responding in an extinction paradigm. *Journal of Experimental Psychology: Animal Behavior Processes*. 2000; 26:174–185. [PubMed: 10782432]

- Holland, PC. Occasion setting in Pavlovian conditioning. In: Medin, DL., editor. *The psychology of learning and motivation*. Vol. 28. New York: Academic Press; 1992. p. 69-125.
- Holland PC, Coldwell SE. Transfer of inhibitory stimulus control in operant feature-negative discrimination. *Learning and Motivation*. 1993; 24:345–375.
- Laborda MA, Witnauer JE, Miller RR. Contrasting AAC and ABC renewal: the role of context associations. *Learning & Behavior*. 2011; 39:46–56. [PubMed: 21264561]
- Lovibond PF, Preston GC, Mackintosh NJ. Context specificity of conditioning, extinction, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*. 1984; 10:360–375.
- Marchant NJ, Li X, Shaham Y. Recent developments in animal models of drug relapse. *Current Opinion in Neurobiology*. 2013 in press.
- Morell JR, Holland PC. Summation and transfer of negative occasion setting. *Animal Learning & Behavior*. 1993; 21:145–153.
- Nakajima S, Tanaka S, Urushihara K, Imada H. Renewal of extinguished lever-press responses upon return to the training context. *Learning & Motivation*. 2000; 31:416–431.
- Nelson JB. Context specificity of excitation and inhibition in ambiguous stimuli. *Learning and Motivation*. 2002; 33:284–310.
- Pearce JM. Similarity and discrimination: A selective review and a connectionist model. *Psychological Review*. 1994; 101:587–607. [PubMed: 7984708]
- Pearce JM, Hall G. A model for Pavlovian learning: Variations in the effectiveness of conditioned but not of unconditioned stimuli. *Psychological Review*. 1980; 87:532–552. [PubMed: 7443916]
- Polack CW, Laborda MA, Miller RR. Extinction context as a conditioned inhibitor. *Learning & Behavior*. 2011; 40:24–33. [PubMed: 21786176]
- Rescorla RA. Inhibitory associations between S and R in extinction. *Animal Learning & Behavior*. 1993; 21:327–336.
- Rescorla RA. Response inhibition in extinction. *The Quarterly Journal of Experimental Psychology*. 1997; 50B:238–252.
- Rescorla RA. Within subject renewal in sign tracking. *The Quarterly Journal of Experimental Psychology*. 2008; 61:1793–1802. [PubMed: 18609366]
- Rescorla, RA.; Wagner, AR. A theory of Pavlovian conditioning: Variations in the effectiveness of reinforcement and nonreinforcement. In: Black, AH.; Prokasy, WF., editors. *Classical conditioning II: current research and theory*. New York: Appleton-Century-Crofts; 1972. p. 64-99.
- Rosas JM, Bouton ME. Context change and retention interval can have additive, rather than interactive, effects after taste aversion extinction. *Psychonomic Bulletin & Review*. 1998; 5:79–83.
- Rosas JM, Callejas-Aguilera JE. Context switch effects on acquisition and extinction in human predictive learning. *Journal of Experimental Psychology: Learning, Memory, and Cognition*. 2006; 32:461–474.
- Rosas JM, Todd TP, Bouton ME. Context change and associative learning. *Wiley Interdisciplinary Reviews: Cognitive Science*. 2013; 4:237–244. [PubMed: 23772263]
- Spear, NE. *The processing of memories: Forgetting and retention*. Hillsdale, NJ: Erlbaum; 1978.
- Swartzentruber D, Rescorla RA. Modulation of trained and extinguished stimuli by facilitators and inhibitors. *Animal Learning & Behavior*. 1993; 22:309–316.
- Tamai N, Nakajima S. Renewal of formerly conditioned fear in rats after extensive extinction training. *International Journal of Comparative Psychology*. 2000; 13:137–146.
- Thomas BL, Larsen N, Ayres JJB. Role of context similarity in ABA, ABC, and AAB renewal paradigms: implications for theories of renewal and for treating human phobias. *Learning and Motivation*. 2003; 34:410–436.
- Thraillkill E, Bouton ME. Contextual control of instrumental actions and habits. in preparation.
- Todd TP. Mechanisms of renewal after the extinction of instrumental behavior. *Journal of Experimental Psychology: Animal Behavior Processes*. 2013; 39:193–207. [PubMed: 23627796]
- Todd TP, Vurbic D, Bouton ME. Behavioral and neurobiological mechanisms of extinction in Pavlovian and instrumental learning. *Neurobiology of Learning and Memory*. 2014a; 108:52–64. [PubMed: 23999219]

- Todd TP, Vurbic D, Bouton ME. Mechanisms of renewal after the extinction of discriminated operant behavior. *Journal of Experimental Psychology: Animal Learning and Cognition*. 2014b in press.
- Todd TP, Winterbauer NE, Bouton ME. Contextual control of appetite: Renewal of inhibited food-seeking behavior in sated rats after extinction. *Appetite*. 2012a; 58:484–489. [PubMed: 22200411]
- Todd TP, Winterbauer NE, Bouton ME. Effects of the amount of acquisition and contextual generalization on the renewal of instrumental behavior after extinction. *Learning & Behavior*. 2012b; 40:145–157. [PubMed: 22002545]
- Trask S, Bouton ME. Contextual control of operant behavior: Evidence for hierarchical associations in instrumental learning. in preparation.
- Tulving E, Thomson DM. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*. 1973; 80:352–373.
- Wagner, AR. Expectancies and the priming of STM. In: Hulse, SH.; Fowler, H.; Honig, WK., editors. *Cognitive processes in animal behavior*. Hillsdale, NJ: Lawrence Erlbaum Associates, Inc; 1978. p. 177-209.
- Wagner AR. Context-sensitive elemental theory. *Quarterly Journal of Experimental Psychology B: Comparative and Physiological Psychology*. 2003; 56B:7–29.
- Welker RL, McAuley K. Reductions in resistance to extinction and spontaneous recovery as a function of changes in transportational and contextual stimuli. *Animal Learning & Behavior*. 1978; 6:451–457.
- Zironi I, Burattini C, Aircardi G, Janak PH. Context is a trigger for relapse to alcohol. *Behavioural Brain Research*. 2006; 167:150–155. [PubMed: 16256214]

Highlights (for review)

- The article reviews research on the contextual control of instrumental (free-operant) behavior.
- Context change after extinction reduces extinction performance
- Evidence suggests that an inhibitory context-response association develops in extinction
- Context change also reduces instrumental behavior after conditioning (before extinction)
- Evidence suggests that the context might be associated with the response and/or hierarchically signal the response-reinforcer relation

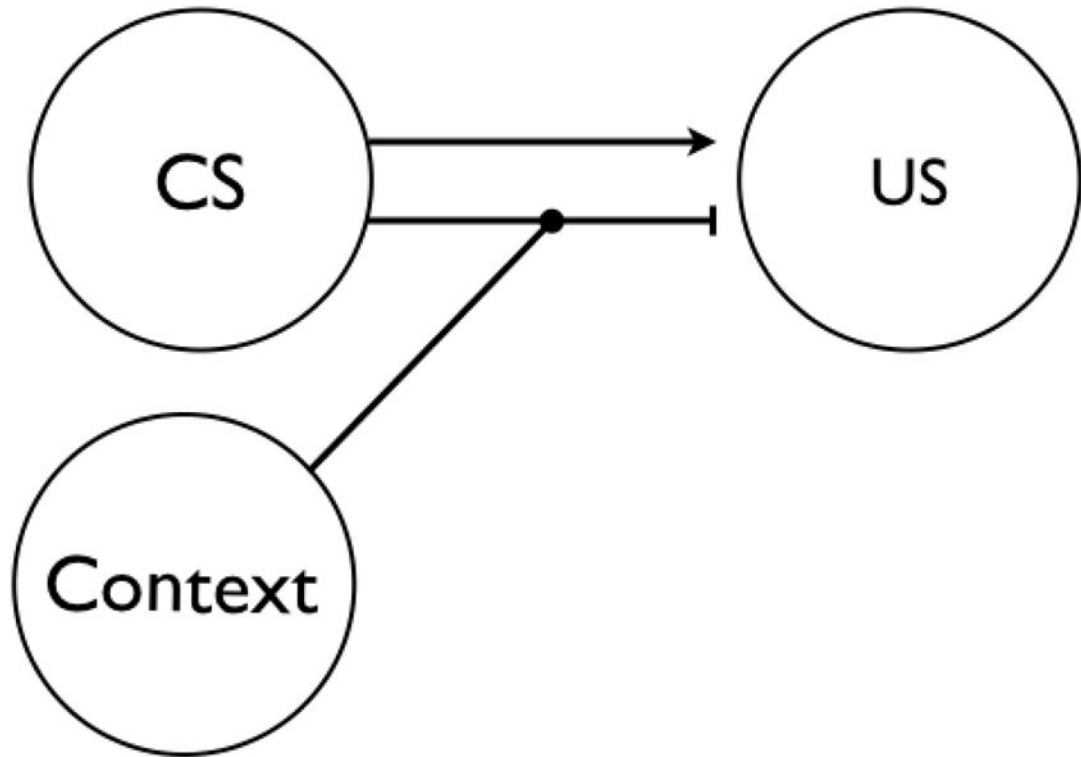


Figure 1.

Hypothetical mechanism of Pavlovian extinction. Arrow indicates excitatory association between the CS and US; blocked line indicates inhibitory association learned during extinction. Inhibition requires input from both the Context and the CS.

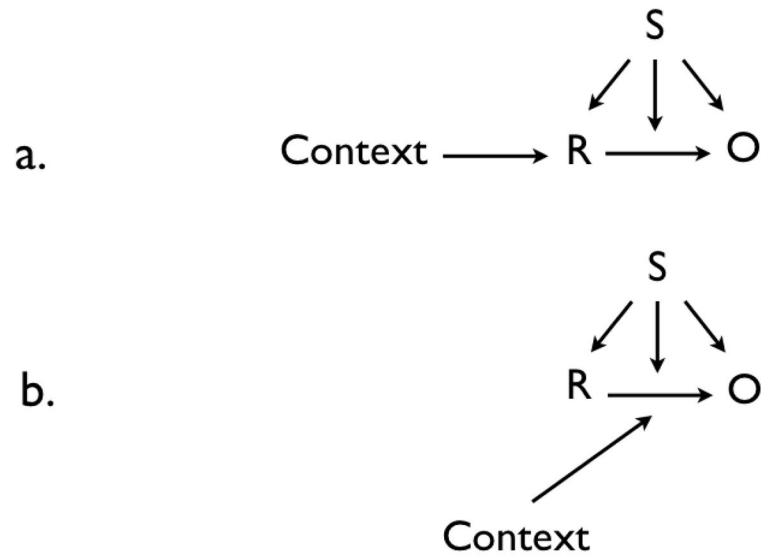


Figure 2.

Possible roles for the context in operant learning. Arrows represent excitatory associations. Panel A. Direct association between the context and the response that is then influenced by S via its possible S-R, S-(R-O), and S-O associations. Panel B. Context as a modulator of the association between the response and the outcome which is then modulated by S via its possible S-R, S-(R-O), and S-O associations. Adapted from Bouton, Todd, & León (2013).

Table 1

Design of an experiment testing the role of negative occasion setting by an extinction context (Context B; Todd, 2013, Experiment 4).

Group	Acquisition	Extinction	Test
Ext-B	A: R1+ B: R2+ C: --	A: R2- B: R1- C: --	A: R2- B: R2-
Ext-C	A: R1+ B: R2+ C: --	A: R2- B: -- C: R1-	A: R2- B: R2-

Note. A, B, and C refer to counterbalanced contexts. R1 and R2 refer to lever press and chain pull, + designates reinforcement, - designates nonreinforcement (extinction). - designates simple exposure to the context. In a subsequent test not shown, the groups were also tested with A: R2- and C: R2-.

Table 2

Design of an experiment demonstrating context-specific response inhibition arising in extinction (Todd, Vurbic, & Bouton, 2014b, Experiment 3).

Acquisition	Extinction	Test
A: S1R1+, S2R2+ B: S3R1+ C: S4R2+	A: S1R1-, S2R2- B: S3R1- C: S4R2-	A: S1R1- v. B: S1R1- A: S2R2- v. B: S2R2- and A: S1R1- v. C: S1R1- A: S2R2- v. C: S2R2-

Note. A, B, and C refer to counterbalanced contexts. R1 and R2 refer to lever press and chain pull, counterbalanced. S1 and S2 refer to discriminative stimuli (light and tone, counterbalanced), S3 and S4 (flashing light and click, counterbalanced). + designates reinforcement, - designates nonreinforcement (extinction).

Table 3

Design of an experiment demonstrating contextual control of non-extinguished operant behavior (Bouton, Todd, & León, 2014, Experiment 2).

Group	Acquisition	Test
Experiment 2		
Exp	A: S1R1+ B:---	A: S1R1- B: S1R1-
No R	A: S1R1+ B: S2+	A: S1R1- B: S1R1-
Diff R	A: S1R1+ B: S2R2+	A: S1R1- B: S1R1-
Same R	A: S1R1+ B: S2R1+	A: S1R1- B: S1R1-

Note. A and B refer to contexts. R1 and R2 refer to lever press and chain pull, counterbalanced. S1 and S2 refer to discriminative stimuli (light and tone, counterbalanced). + designates reinforcement, - designates nonreinforcement (extinction). --- = context exposure without the lever or chain present.