



Published in final edited form as:

J Exp Psychol Anim Learn Cogn. 2014 January ; 40(1): 92–105. doi:10.1037/xan0000002.

Contextual Control of Discriminated Operant Behavior

Mark E. Bouton, Travis P. Todd, and Samuel P. León

University of Vermont

Abstract

Previous research suggests that changing the context after instrumental (operant) conditioning can weaken the strength of the operant response. That result contrasts with the results of studies of Pavlovian conditioning, where a context switch often does not affect the response elicited by a conditioned stimulus. To begin to make the methods more similar, Experiments 1–3 tested the effects of a context switch in rats on a discriminated operant response (R, lever pressing or chain pulling) that had been reinforced only in the presence of a 30-s discriminative stimulus (S, tone or light). As in Pavlovian conditioning, responses and reinforcers became confined to presentations of the S during training. However, in Experiment 1, after training in Context A, a switch to Context B caused a decrement in responding during S. In Experiment 2, a switch to Context B likewise decreased responding in S when Context B was equally familiar, equally associated with reinforcement, or equally associated with the training of a discriminated operant (a different R reinforced in a different S). However, there was no decrement if Context B had been associated with the same response that was trained in Context A (Experiments 2 and 3). The effectiveness of S transferred across contexts, whereas the strength of the response did not. Experiment 4 found that a continuously reinforced response was also disrupted by context change when the same response manipulandum was used in both training and testing. Overall, the results suggest that the context can have a robust general role in the control of operant behavior. Mechanisms of contextual control are discussed.

Theories of learning and memory almost universally assume that contextual stimuli play a crucial role in determining performance. As just one example, memory retrieval is thought to depend on the match between the contextual cues present during learning and those that are present during testing (e.g., Spear, 1978; Tulving & Thomson, 1973). A clear implication is that a change of context should therefore weaken or disrupt memory or learned performance. However, although there are well-known results that are consistent with this idea (e.g., Godden & Baddeley, 1975; Tulving & Thomson, 1973), there are many exceptions to the rule (e.g., Bouton, 1993; see Rosas, Todd, & Bouton, 2013, for a recent review). For example, in experiments on human memory, context switches often fail to impair memory performance (e.g., see Smith, 1988; Smith & Vela, 2001, for reviews). And in experiments on Pavlovian conditioning in animals, a context switch after pairings of a conditioned stimulus (CS) and unconditioned stimulus (US) often causes surprisingly little decrement in the response evoked by the CS. In our laboratory, a context switch after

conditioning has had very little effect on conditioned responding in conditioned suppression (fear conditioning) (e.g., Bouton & King, 1983; Bouton & Swartzentruber, 1989; see also Bouton, Frohardt, Sunsay, Waddell, & Morris, 2008), appetitive conditioning (e.g., Bouton & Peck, 1989; Bouton et al., 2008; Nelson, 2002), or taste aversion learning (e.g., Rosas & Bouton, 1998). Similar results have been reported by other laboratories (e.g., Grahame, Hallam, Geier, & Miller, 1990; Hall & Honey, 1989, 1990; Harris, Jones, Bailey, & Westbrook, 2000; Lovibond, Preston, & Mackintosh, 1984; Rosas & Callejas-Aguilera, 2007). A context switch also has analogously little effect in human predictive learning (e.g., Nelson & Callejas-Aguilera, 2007; Rosas & Callejas-Aguilera, 2006). Although context switches occasionally disrupt performance in Pavlovian or predictive learning tasks (e.g., Bonardi, Honey, & Hall, 1990), Nelson and Callejas-Aguilera (2007) noted that a lack of an effect “appears to be the case more often than not” (p. 314).

In contrast, context switch effects readily occur after Pavlovian *extinction*, where a change of context typically causes a loss of extinction performance (and a renewal of conditioned responding) (e.g., Bouton & King, 1983; Bouton & Peck, 1989; Harris et al., 2000). In Pavlovian learning, the second thing learned about a CS appears to be more context-specific than the first (e.g., Nelson, 2002; see Bouton, 1994, 1997). One explanation has been emphasized by Rosas and colleagues (e.g., Rosas & Callejas-Aguilera, 2006; Rosas, Callejas-Aguilera, Ramos-Alvarez, & Abad, 2006). At the end of Phase 1 learning, the meaning of a predictive signal is known; when extinction begins, the outcome of each trial (no US) is surprising, and surprising outcomes (prediction error) might direct the animal’s attention to the context. What is learned in extinction is therefore more context-specific. This idea is consistent with the results of a number of experiments in both humans (e.g., Nelson, Lamoureux, & León, 2013; Rosas & Callejas-Aguilera, 2006) and animals (e.g., Rosas & Callejas-Aguilera, 2007). And it is consistent with theories of conditioning and attention, which predict increased attention to all stimuli present on a trial, including the context, when there is prediction error (e.g., Pearce & Hall, 1980; Pearce & Mackintosh, 2010; see Rosas et al., 2013, for more discussion).

The present experiments were designed to expand the analysis of context effects by assessing the effects of a context switch on instrumental (operant) conditioning. Research in this laboratory and in others has confirmed that extinction can be context-specific after operant extinction learning (Bouton, Todd, Vurbic, & Winterbauer, 2011; Bossert, Liu, Li, & Shaham, 2004; Crombag & Shaham, 2002; Hamlin, Clemens, & McNally, 2008; Nakajima, Tanaka, Urushihara, & Imada, 2000; Nakajima, Urushihara, & Masaki, 2002; Todd, 2013; Todd, Winterbauer, & Bouton, 2012a, b). For example, Bouton et al. (2011; see also Todd, 2013) reported that a lever-press response that was first reinforced on a variable interval (VI) 30-s reinforcement schedule in Context A and then extinguished in Context B was renewed during tests conducted in either Context A (ABA renewal) or in Context C (ABC renewal). When extinction and conditioning both occurred in Context A, renewal also occurred when tests were conducted in a second context, Context B (AAB renewal). The ABC and AAB renewal effects, like similar effects after Pavlovian extinction, suggest that extinction learning is more context-specific than the original conditioning.

However, unlike previous experiments on Pavlovian learning, the operant experiments also found results suggesting a role for context during conditioning (Bouton et al., 2011; Todd, 2013). Specifically, there was a substantial decrement in responding when the response was first tested in Context B after conditioning in Context A. We should note that Nakajima et al. (2000, Experiment 1) reported less effect than we did of switching the context on free operant responding, although a context switch did appear to weaken the response in a pilot experiment described on pp. 426–427 (see also Nakajima et al., 2000, Experiment 2). In addition, Bossert et al. (2004) and Crombag and Shaham (2002) reported less effect than we did of switching the context when lever pressing had been reinforced with intravenous injections of cocaine or heroin with cocaine. However, many aspects of the methods were different. The fact that Bouton et al. (2011) and Todd (2013) observed context switch effects after operant training suggests that the effects of switching the context after operant conditioning warrants further study and analysis.

The present experiments were therefore designed to extend and test a possible explanation of the previous findings. Bouton et al. (2011) and Todd (2013) studied the effects of a context switch after what we will refer to as “free-operant” training. In their method, the rat could press a lever (or pull a chain) for reinforcers at any time during the experimental session. In contrast to studies of Pavlovian conditioning, the context was the only stimulus present during performance of the response; there was no signal or proximate stimulus that could overshadow (e.g., Rescorla & Wagner, 1972) or “outshine” (Smith & Vela, 2001) learning about the context. The present experiments therefore studied the effects of a context switch after training with a *discriminated*-operant (or multiple schedule) procedure. In this procedure, the response was only reinforced in the presence of a discriminative stimulus (S) that had the same duration (30 s) as CSs that were used in previous context-switch experiments in appetitive Pavlovian conditioning (Bouton et al., 2008; Nelson, 2002). As in the Pavlovian experiments, the occurrence of responses and reinforcers became confined to the presence of the stimulus during training. However, when S was then tested for its ability to evoke responding in a second context, there was nevertheless a substantial drop in responding in the presence of S. Experiment 1 established this effect, and the remaining experiments tested factors that might contribute to it. Even when the associative histories of the contexts were controlled, the context switch consistently reduced the rate of the discriminated operant response. It did not, however, reduce the effectiveness of S when S could set the occasion for its usual response after the response had been trained (with a different S) in the switched-to context (Experiments 2 and 3).

Experiment 1

The first experiment provided an initial test of the effects of context change on discriminated operant responding. The design, which is illustrated in Table 1, combined discriminated operant training with the methods used by Bouton et al. (2011). Two groups learned to lever press in the presence of a light S in Context A. When the light was presented, responding was reinforced on a VI 30-s schedule; when it was off, the response was not reinforced. After discriminative responding was learned, it was tested over a series of extinction test trials in which the light was presented and responding was not reinforced. One group was tested for responding in the light in Context A, where training had occurred, and the other

group was tested in Context B. As in Bouton et al. (2011, Experiment 1), the rats had been exposed to Context B only once before during an initial magazine training session in which they received free food pellets there. In the previous free-operant experiments, a group tested in Context B showed less responding (e.g., Bouton et al., 2011). The question was whether a different pattern would emerge after discriminated operant training, where S bears a temporal relation to the reinforcer and the response that is more similar to that of a CS in Pavlovian conditioning.

Method

Subjects—Thirty-two female Wistar rats (Charles River Laboratories, Quebec, Canada), ranging in age from 75 to 90 days at the start of the experiment, were individually housed in suspended stainless steel cages in a room maintained on a 16:8-h light:dark cycle. Experimental sessions were run during the light portion of the cycle at the same time each day. The rats were food deprived and maintained at 80% of their baseline body weights throughout the experiment. Water was freely available in the home cage.

Apparatus—The apparatus consisted of two unique sets of four conditioning chambers (Med Associates, St. Albans, VT; model ENV-008-VP) housed in separate rooms of the laboratory. Each chamber was housed in its own sound attenuation chamber. All boxes measured 30.5 × 24.1 × 23.5 cm (l × w × h). The side walls and ceiling were made of clear acrylic plastic, while the front and rear walls were made of brushed aluminum. The floor was made of stainless steel grids (0.48-cm diameter). A recessed 5.1 × 5.1 cm food cup was centered in the front wall approximately 2.5 cm above the level of the floor. In both sets of boxes, a retractable lever (Med Associates model ENV-112CM) was positioned to the left of the food cup. The lever was 4.8 cm long and was 6.2 cm above the grid floor. It protruded 1.9 cm from the front wall when extended. A 28-V (2.8 W) panel light (2.5 cm in diameter) was mounted on the wall above the lever, 10.8 cm above the floor and 6.4 cm to the left of the food cup. The chambers were illuminated by two 7.5-W incandescent bulbs mounted to the ceiling of the sound attenuation chamber, 34.9 cm from the grid floor. Ventilation fans provided background noise of 65 dBA.

The two sets of boxes were fully counterbalanced and had unique features that allowed them to be used as different contexts. In one set of boxes, one side wall had black diagonal stripes, 3.8 cm wide and 3.8 cm apart. The ceiling had similarly spaced stripes oriented in the same direction. A distinct odor was continuously presented by placing 5 ml of Pine-Sol (Clorox Co., Oakland, CA) in a dish immediately outside the chamber. The grids of the floor were mounted in the same plane and were spaced 1.6 cm apart (center to center). The other set of boxes had no distinctive visual cues, and the grids of the floor were staggered such that odd- and even-numbered grids were mounted in two separate planes, one 0.5 cm above the other. An odor was provided by 1.5 ml of Lemon Extract (McCormick, Hunt Valley, MD) placed in a dish immediately outside the chamber. The reinforcer was a 45-mg food pellet (Traditional formula, Research Diets, New Brunswick, NJ). The apparatus was controlled by computer equipment located in an adjacent room.

Procedure—Throughout the experiment, there were two sessions each day separated by 2 hr. Training began after a week of daily handling and food restriction to attain the target body weights.

Magazine training: On Day 1, each rat was assigned to a box (Context A), with the restriction that boxes be equally represented in each group. All rats then received a 30-min session of magazine training in which pellets were delivered randomly on average every 30 s. After the session in the Context A, there was a similar session conducted in a box from the other set of boxes (Context B). The levers were retracted during these sessions.

Free operant training: On the next day, the rats received two 30-min sessions of training to press the lever in Context A. The levers were inserted 2 min after the rats were placed in the chambers. Lever presses were then reinforced on a variable interval (VI) 30-s schedule. With this method, widely-spaced responses emitted early in training were reinforced with a high probability. Perhaps because of this, as in other research in this laboratory (e.g., Bouton et al., 2011; Todd, 2013; Todd et al., 2012b), no additional response shaping was necessary.

Discriminated operant training: Over the next 6 days, the rats received a total of 12 sessions of discriminated operant training in Context A. Each session contained 16 trials in which the panel light was illuminated for 30 s and lever pressing was reinforced. The VI-30s schedule was in effect only during the light presentations; lever presses were not reinforced during the intertrial interval (ITI). The ITI was variable and averaged 30 s in the first two (16-min) sessions, 60 s in the third and fourth (24-min) sessions, and 90 s in the eight (32-min) sessions thereafter.

Testing: On the day that followed the last training session, rats were assigned to two groups matched on response rate in the final session. They were then given three 32-min sessions of extinction (the first two in one day and the third on the next day). For one group (Group A), the sessions were conducted in the original training context, Context A. For the other group (Group B), they were conducted in Context B. In each session, the rats received 16 presentations of the light with a variable ITI of 90 s. No pellets were delivered at any time.

Data analysis: To describe responding that was occasioned by each presentation of the light, we calculated an elevation score on each trial by subtracting the number of responses made during the 30 s immediately before each light presentation (the pre-S period) from the number of responses made during the 30-s light itself. The elevation scores and responding during the pre-S period were evaluated with parallel analyses of variance (ANOVAs) using a rejection criterion of $p.05$. For ANOVAs with more than one factor, we report η_p^2 as our measure of effect size; for comparisons between two means, we report η^2 . For either measure of effect size, we computed 95% confidence intervals using procedures described by Steiger (2004). First, we obtained the noncentrality parameter of the non-central F distribution. We then calculated the 95% confidence interval on the noncentrality parameter and used that to compute upper and lower boundaries of effect size estimates.

Results

The mean elevation scores for each session of discrimination training are summarized in the left side of the Figure 1. Stimulus control increased steadily over the sessions of training. A Group x Session ANOVA on the elevation scores revealed a main effect of Session, $F(11, 330) = 82.96$, $MSE = 6.90$, $\eta_p^2 = .73$, $CI [.68, .76]$, but no Group effect or a Group x Session interaction, $F_s < 1$. There was no difference between the elevation scores of the groups in the last session before testing, $F < 1$. Responding during the pre-S periods had mean values of 8.7 and 7.2 for Groups A and B (respectively) in the first session and 7.4 and 5.5 during the final session. A Group x Session ANOVA revealed a main effect Session, $F(11, 330) = 11.81$, $MSE = 6.77$, $\eta_p^2 = .28$, $CI [.18, .34]$, but no Group effect or interaction, larger $F(1, 10) = 3.48$, $MSE = 100.93$.

The results of the test phase are presented at right in Figure 1. A strong context-switch effect was evident in Group B. A Group x Session ANOVA found a main effect of Group, $F(1, 30) = 21.40$, $MSE = 19.04$, $\eta_p^2 = .42$, $CI [.14, .60]$, a main effect of Session, $F(2, 60) = 67.31$, $MSE = 5.76$, $\eta_p^2 = .69$, $CI [.54, .77]$, and a Group x Session interaction, $F(2, 60) = 11.23$, $MSE = 5.76$, $\eta_p^2 = .27$, $CI [.09, .42]$. A separate analysis of the first four trials of the first session also revealed a main effect of Group, $F(1, 30) = 19.20$, $MSE = 24.92$, $\eta^2 = .39$, $CI [.12, .58]$. Thus, the difference in discriminated operant responding between contexts was evident from the beginning of testing. The average pre-S responding for Groups A and B were 1.9 and 1.6 in the first session and 0.4 and 1.2 in the third. A Group x Session ANOVA on these data revealed a main effect of Session, $F(2, 60) = 42.14$, $MSE = 0.46$, $\eta_p^2 = .58$, $CI [.40, .68]$, but no Group effect or Group x Session interaction, $F_s < 1$. There was also no difference between the groups in pre-S responding on the first four trials, $F(1, 30) = 1.92$, $MSE = 17.17$, in which the mean scores for Groups A and B were 5.6 and 3.6.

Discussion

Lever pressing came under control of the S in both groups. However, the results of testing clearly indicated that a change of context caused a decrement in responding in the presence of the S. The results were consistent with our earlier findings with lever pressing trained as a free operant (Bouton et al., 2011), as well as previous findings with a discriminated operant reported by Nakajima et al. (2000, Experiment 2). The parallel between our results with free- and discriminated-operant procedures suggests that the earlier free-operant result cannot be attributed to the lack of overshadowing or outshining of the context by a proximal stimulus. Recall that previous experiments on appetitive Pavlovian conditioning have found no effect of a context switch on responding to a 30-s CS (Bouton et al., 2008; Nelson, 2002). In contrast, with the present methods, discriminated operant responding appears to be disrupted by context change.

Experiment 2

There are several reasons why the context switch could have disrupted responding in Experiment 1 and in the experiment reported by Nakajima et al. (2000, Experiment 2). The rats had received only one initial exposure to Context B during magazine training, so they were less familiar with Context B than Context A, and they had received many fewer

reinforcers in it. It is thus possible that responding in Context A was stronger because of its familiarity, or because of its direct association with the reinforcer, which could theoretically motivate or enhance responding (e.g., Rescorla & Solomon, 1967). In addition, only Context A had been associated with reinforcement of an operant response. The main purpose of Experiment 2 was therefore to ask whether a context switch effect still occurred when Contexts A and B were equated on these dimensions.

The design of the experiment is illustrated in Table 1. It employed two operant responses (R1 and R2, counterbalanced over lever press and chain pull) and two discriminative stimuli (S1 and S2, counterbalanced over a 30-s light and a 30-s tone). Four groups each received discriminated operant training in Context A in which one response (R1) was reinforced in the presence of one S (S1). In the final test, all groups were tested with the S1R1 combination in both Contexts A and B (order counterbalanced). The groups differed only in the training they received during conditioning sessions that were conducted in Context B (which were double-alternated with those in Context A). During those sessions, Group Exposure was merely placed in Context B with no lever or chain available or food pellets presented. The rats were thus exposed to the alternate context, but did not make a response or receive a stimulus or a reinforcer. Group No R also had no response available, but received presentations of a second stimulus (S2) that were paired with noncontingent food pellets delivered at a rate comparable to that received in Context A. Pellet deliveries were yoked to the pellets earned by another group member working for food in Context A. The third group, Group Diff R, likewise received food pellets during S2 presentations, but had to perform a different response (R2) in order to earn them. Context B was therefore similar to Context A in being home to discriminative operant training, but it was associated with a different combination of S and response. The final group, Group Same R, also received training with a new S (S2) in Context B. But in this case, the response required to earn pellets in B (R1) was the same as the response that was trained in Context A. Contexts A and B were thus home to training with different Ss but the same instrumental response.

As noted above, all rats were finally tested with S1R1 in both Contexts A and B. The question was whether R1 responding in the presence of S1 would transfer differentially in the different groups. If Experiment 1's contextual control of responding was due to differences in direct context-reinforcer associations, only Group Exposure should show weaker S1R1 responding in Context B (context-reinforcer associations were controlled in all other groups). If Experiment 1's contextual control was due to differences in the contexts' association with operant responding in general, Groups Exposure and No R, but not Groups Diff R and Same R, should show weaker responding in Context B. And if the contextual control were due only to the context-specificity of the operant response, we should find a context switch effect in Groups Exposure, No R, and Different R, but not Group Same R.

Method

Subjects—Thirty-two female Wistar rats from the same supplier were used. Their age and housing and maintenance conditions were the same as described in Experiment 1.

Apparatus—The apparatus was also the same as that used in Experiment 1, with two additions. First, a response chain (Med Associates model ENV-111C) was introduced to each chamber when needed. The chain was suspended from a microswitch mounted on top (outside) of the chamber's ceiling. The chain hung 1.9 cm from the front wall, 6.2 cm above the grid floor, and 3 cm to the right of the food cup. Although the chain and lever were never presented at the same time, they were positioned on opposite sides of the food cup. Second, presentation of a 30-s 3000 Hz tone (80 dBA; background was 65 dBA) provided a second S. The tone was presented through a speaker mounted to the ceiling of the sound attenuation chamber.

Procedure—As before, the experiment involved two daily sessions spaced 2 hr apart. During lever press sessions the chain was not present and during chain pull sessions the lever was not present. Training began after a week of daily handling and food restriction to attain the target body weights.

Magazine and free operant training: On Day 1, each rat received two 30-min magazine-training sessions in which pellets were delivered randomly on average every 30 s. One session occurred in Context A and the other in Context B (counterbalanced with respect to order). Neither the lever nor the chain was present during these sessions.

On each of the next two days, the rats received two 30-min sessions, one in Context A and the other in Context B (order counterbalanced), that involved free-operant VI-30 s training of the responses to be reinforced in each context. In Context A, all rats were trained on R1, which was lever press for half the rats and chain pull for the other half. During sessions in Context B, the treatment of the groups differed. For Group Exposure, the lever and chain were both absent and no reinforcers were delivered. For Group No R, the lever and chain were also absent, but a pellet was delivered whenever a yoked member of the same group earned a pellet while responding in Context A. For Group Diff-R, the alternate response to the one reinforced in Context A (R2, chain for half the rats and lever for the other half) was reinforced. Finally, for Group Same-R, the same response as the one trained in Context A (i.e., R1) was reinforced in Context B. As in Experiment 1, responses were always reinforced on a VI 30 s schedule, and no additional shaping was necessary.

Discriminated operant training: Over the next 10 days, the rats received a total of 20 sessions of discriminated operant training with the responses that had been conditioned during the pre-training phase. Half of the sessions were in Context A and the other half in Context B (double alternated beginning with AB). In Context A, all groups received 16 trials per session in which R1 was reinforced only during 30-s presentations of S1 (either panel light or tone, counterbalanced). The VI-30 s reinforcement schedule was in effect during S1, and no pellets were delivered during the ITI. The groups continued to differ in their treatment during the sessions in Context B. Group Exposure was merely placed in Context B with no response available, and with no presentations of pellets or S. For Group No R, responses were likewise unavailable, but each rat received 16 presentations per session of S2 (the alternate S from the one presented in Context A) on the usual trial schedule. Rats in Group No R continued to be yoked to a member of the same group that was being run in Context A, so that pellets were now delivered during S2 whenever the partnered rat earned

one during S1. Group Diff-R also received 16 trials with S2, but the rats were required to make the alternate response (R2) to earn pellets during S2 on a VI 30 schedule. Finally, Group Same R also received training with S2, but the *same* response trained in the Context A (R1) was available and reinforced. As in Experiment 1, the ITIs were variable, averaging 30 s in the first two (16-min) sessions, 60 s in the third and fourth (24-min) sessions, and 90 s in all (32-min) sessions thereafter.

Testing: On the final day, the rats received two 8-min extinction test sessions, one in Context A and the other in Context B (order counterbalanced). For all rats, only R1 was available, and there were four 30-s presentation of S1 spaced with the usual variable ITI of 90 s. No pellets were delivered at any time.

Results

Figure 2 shows the mean elevation scores for the responses emitted during the discrimination-training sessions in Context A (top panel) and Context B (bottom panel). Responding in the two contexts was analyzed separately. In Context A, where all groups received the same training with R1 in S1, a Group x Session ANOVA revealed a main effect of Session, $F(9, 252) = 44.36$, $MSE = 8.47$, $\eta_p^2 = .61$, $CI [.53, .68]$, but the main effect of Group and Group x Session interactions did not approach significance, $F_s < 1$. An analysis of the groups' elevation scores in the last acquisition session also found no difference among the groups, $F < 1$. In Context B, where only groups Diff R and Same R received operant training, a Group x Session ANOVA confirmed a main effect of Session, $F(9, 126) = 28.59$, $MSE = 5.16$, $\eta_p^2 = .67$, $CI [.55, .72]$, but no main effect of Group or Group x Session interaction, larger $F(9, 126) = 1.44$, $MSE = 5.16$. The groups also did not differ during the last session, $F < 1$.

Parallel analyses of pre-S responding were also performed. In Context A, the mean number of pre-S responses in the first training session were 9.6, 9.5, 11.2, and 8.2 for Groups Exposure, No R, Diff R, and Same R, respectively. In the last session, the means were 10.9, 5.2, 6.9, and 6.0. A Group x Session ANOVA revealed a main effect of Session, $F(9, 252) = 2.81$, $MSE = 8.84$, $\eta_p^2 = .10$, $CI [.01, .13]$, but no Group effect or interaction, larger $F(3, 28) = 1.49$, $MSE = 50.59$. In Context B, Groups Diff R and Same R had mean pre-S responses of 11.7 and 7.7 in the first session and 7.5 and 5.6 in the last. A Group x Session ANOVA revealed a main effect of Session, $F(9, 126) = 5.80$, $MSE = 4.82$, $\eta_p^2 = .29$, $CI [.12, .37]$, but no Group effect or interaction, larger $F(1, 14) = 4.57$, $MSE = 64.21$.

The test results are summarized in Figure 3, which shows the mean elevation scores of R1 responding in S1 for the groups during the four test trials in Context A and four in Context B. As suggested by the figure, responding decreased when the S1R1 combination was tested in a different context for all groups except Group Same R. A Group x Context ANOVA confirmed a main effect of Context, $F(1, 28) = 21.26$, $MSE = 7.35$, $\eta_p^2 = .43$, $CI [.15, .61]$, as well as a Group x Context interaction, $F(3, 28) = 3.79$, $MSE = 7.35$, $\eta_p^2 = .29$, $CI [.00, .46]$. Simple effect analyses revealed that the context switch decreased responding in Groups Diff R, No R, and Exposure, $F_s(1, 28) = 10.41, 5.74, \text{ and } 16.27$, respectively (all MSEs =

7.35, $\eta^2 = .17$, CI [.00, .40]). Importantly, there was no hint of a context switch effect on responding in Group Same R, $F < 1$.

The parallel analysis of pre-S responding revealed a main effect of Context, $F(1, 28) = 6.84$, $MSE = 2.64$, $\eta_p^2 = .20$, CI [.01, .42], and a Context x Group interaction, $F(3, 28) = 2.96$, $MSE = 2.64$, $\eta_p^2 = .24$, CI [.00, .42]. The mean number of pre-S responses in Context A were 4.8, 2.5, 5.0, and 6.8 for Groups Exposure, No R, Diff R, and Same R, respectively. In Context B, the corresponding scores were 3.2, 3.6, 3.3, and 4.8. To follow up on the Context x Group interaction, simple effect comparisons were conducted on the Context effect in each group. Somewhat surprisingly, Groups Same R and Diff R showed reliably less R1 responding in Context B than Context A, $F_s(1, 28) = 5.50$ and 4.47 , respectively, $MSEs = 2.64$, $\eta^2 = .14$, CI [.00, .36]. The differences are presumably due to sampling error. (Note that Group Same R had received equal training of R1 in both contexts; the result in elevation scores is replicated in Experiment 3 without a difference in pre-S rates.) The difference in Group Exposure approached significance, $F(1, 28) = 4.15$, $MSE = 2.64$, $p = .051$. All other $F_s = 1.16$, $ps = .21$.

Discussion

The groups did not differ in their acquisition of SIR1 responding. However, important differences emerged during the tests of SIR1 in Contexts A and B. The key result was that all the groups but Group Same R showed evidence of a context switch effect; there was a decrease in responding when SIR1 was tested in Context B. The results help clarify the findings of Experiment 1. The fact that there was a context switch effect in Group Exposure, which received equivalent time in the two contexts during training, indicates that the context switch effect in Experiment 1 was not merely due to differential familiarity with the contexts. The fact that there was also a context switch effect in Group No R, for which context-reinforcer associations were controlled over A and B, suggests that Experiment 1's effect was not due to differential context-reinforcer associations. And the fact that there was a context switch effect in Group Diff R, which had equivalent discriminative operant training with different stimulus/response combinations in A and B, suggests that Experiment 1's effect did not result from the fact that the alternative context (B) had not been home to discriminative operant training. The explanation of why the context is important in discriminated operant training thus depends on another factor.

A clue is provided, of course, by the fact that Group Same R did not show a context switch effect during the test. For this group, Contexts A and B had equivalent familiarity, association with the reinforcer, and association with an S. Although the Ss used in the two contexts were different in this group, Group Same R was uniquely trained with the same operant action, R1, in both contexts. Evidently, that treatment was sufficient to support transfer of discriminated operant responding across the contexts. The results thus begin to isolate failure of the *response*, rather than the S, to transfer across contexts as the source of the context switch effect.

Experiment 3

As just noted, Group Same R in the preceding experiment showed no response decrement after the context switch despite the fact that S1 was being presented for the first time in the alternate context. This result implies that the response decrement observed in Group Diff R was due to the switch of the response, rather than the switch of the stimulus, during testing. Experiment 3 was designed to replicate and extend this result. As shown in Table 1, two groups both received discriminative operant training in Contexts A and B. Group Same R/Diff S received training like that of Group Same R in Experiment 2, so that the response, but not the stimulus, was common across contexts. A new group, Group Same S/Diff R, received a complementary experience in which the stimulus, rather than the response, was common across contexts: In Context A, R1 was reinforced in S1, whereas in Context B R2 was now reinforced in S1. In the final test, all rats were tested with the S1R1 combination in both contexts. If the response, rather than the stimulus, is more vulnerable to context change, then we should expect a bigger switch effect in Group Same S/Diff R than Group Same R/Diff S.

Method

Subjects and Apparatus—The subjects were 32 female Wistar rats purchased from the same vendor as those in the previous experiments and maintained under the same conditions. The apparatus was the same as that in Experiment 2.

Procedure

Magazine and free operant training: This phase was carried out following the procedure used in Experiment 2. On the first day, all rats received magazine training in both Contexts A and B. On each of the next two days, the rats received two sessions, one in Context A and the other in Context B (order counterbalanced). In Context A, half the rats in each group received training with lever press and half received training with chain pull. In Context B, Group Same R/Diff S were trained with the same response that was trained in Context A, while Group Diff R/Same S received training with the other response. A VI 30-s reinforcement schedule was used throughout, and no additional shaping was necessary.

Discriminated operant training: Over the next 10 days, discriminated operant training was carried out following a procedure similar to that of Experiment 2. For Group Same R/Diff S, R1 was reinforced in the presence of S1 (panel light or tone, counterbalanced) in Context A and in the presence of S2 in Context B. For Group Diff R/Same S, R1 was reinforced in the presence of S1 in Context A and R2 was reinforced in the presence of the same stimulus (S1) in Context B. The panel light and tone were counterbalanced. As usual, the ITI was 30 s in the first 2 sessions, 60 s in the next two, and 90 s in all sessions thereafter.

Testing: On the final day, all rats received two 8-min sessions, one in Context A and the other in Context B (order counterbalanced). In both contexts, there were four tests of R1 responding in the presence of S1; no pellets were delivered at any time. The ITI was variable with a mean of 90 s.

Results

Acquisition in Context A (top panel) and Context B (bottom panel) is shown in Figure 4. As before, responding was analyzed separately in each context. In Context A, where both groups received training with R1 in S1, a Group \times Session ANOVA revealed a significant main effect of Session, $F(9, 270) = 59.77$, $MSE = 7.20$, $\eta_p^2 = .67$, $CI [.59, .70]$, and a significant interaction between Group and Session, $F(9, 270) = 2.20$, $MSE = 7.20$, $\eta_p^2 = .07$, $CI [.00, .10]$. The main effect of Group approached significance, $F(1, 30) = 3.94$, $MSE = 81.51$, $p = .06$. Interestingly, training with the same response (R1) but different stimulus (S2) in Context B appeared to lead to lower S1R1 responding in Context A. However, an ANOVA on the last session of acquisition revealed no difference between the groups, $F < 1$. In Context B, a Group \times Session ANOVA revealed only a significant main effect of Session, $F(9, 270) = 61.62$, $MSE = 6.50$, $\eta_p^2 = .67$, $CI [.60, .71]$. Neither the main effect of Group nor the Group \times Session interaction were significant, largest $F(1, 30) = 2.43$. An ANOVA on the last session of acquisition in Context B also revealed no difference between groups, $F < 1$.

A parallel analysis was conducted on responding in the pre-S period during acquisition in Context A and Context B. In Context A, mean pre-S responding for Groups Same R/Diff S and Diff R/Same S was 9.8 and 9.4 in the first session and 8.3 and 5.7 in the last (respectively). A Group \times Session ANOVA revealed a main effect of Session, $F(9, 270) = 9.53$, $MSE = 8.31$, $\eta_p^2 = .24$, $CI [.14, .30]$, and a main effect of Group, $F(1, 30) = 8.86$, $MSE = 82.90$, $\eta_p^2 = .23$, $CI [.02, .44]$. There was also a significant interaction between Group and Session, $F(9, 270) = 2.01$, $MSE = 8.31$, $\eta_p^2 = .06$, $CI [.00, .09]$. An ANOVA on the final session of acquisition in Context A did not reveal a difference between the groups, $F(1, 30) = 2.11$, $MSE = 25.15$. In Context B, mean pre-S responding in Groups Same R/Diff S and Diff R/Same S was 11.3 and 9.7 in the first session and 8.2 and 5.9 in the last (respectively). A Group \times Session ANOVA revealed a main effect of Session, $F(9, 270) = 8.54$, $MSE = 8.44$, $\eta_p^2 = .22$, $CI [.12, .28]$, but neither the main effect of Group, $F(1, 30) = 3.12$, $MSE = 146.97$, nor the interaction, $F < 1$, were significant. An ANOVA on the final session of acquisition in Context B revealed no difference between the groups, $F < 1$.

Figure 5 presents the mean elevation score for the four test trials in each context for each group. There was little change in responding from Context A to Context B for Group Same R/Diff S, replicating the result of Experiment 2. However, there was a large decrement in responding in Group Diff R/Same S. A Group \times Context ANOVA revealed a significant main effect of Context, $F(1, 30) = 14.26$, $MSE = 18.16$, $\eta_p^2 = .32$, $CI [.07, .52]$. The Group factor was not significant, $F(1, 30) = 1.01$. However, the crucial Group \times Context interaction was significant, $F(1, 30) = 12.27$, $MSE = 18.16$, $\eta_p^2 = .29$, $CI [.05, .50]$. Follow up comparisons revealed that although there was no effect of Context in Group Same R/Diff S, $F < 1$, the effect of Context in Group Diff R/Same S was reliable, $F(1, 30) = 26.46$, $MSE = 18.16$, $\eta_p^2 = .47$, $CI [.19, .64]$. Further, responding in Context A was not different between the groups, $F < 1$, whereas Group Diff R/Same S responded significantly less than Group Same R/Diff S in Context B, $F(1, 30) = 12.54$, $MSE = 19.74$, $\eta_p^2 = .29$, $CI [.05, .50]$.

As usual, pre-S scores were analyzed in a parallel ANOVA. The mean number of pre-S responses for Group Same R/Diff S was 7.4 and 6.7 in Context A and B, respectively, and

2.4 and 3.7 for Group Diff R/Same S. A Group \times Context ANOVA found a main effect of Group, $F(1, 30) = 5.61$, $MSE = 45.25$, $\eta_p^2 = .16$, CI [.00, .38]. The effect of Context was not significant, $F < 1$, although the Group \times Context interaction approached significance, $F(1, 30) = 4.03$, $MSE = 3.73$, $p = .054$. The higher pre-S responding in Group Same R/Diff S is consistent with the pattern of responding that was initially observed during acquisition. However, the effect cannot account for the Group \times Context interaction observed in the elevation scores, because each group showed similar levels of pre-S responding in Contexts A and B. Within each group separately, the effect of Context was not significant, largest $F(1, 30) = 3.19$, $MSE = 3.73$.

Discussion

The results of this experiment replicate and extend those of Experiment 2. As before, when rats received training with the same response but different S in the two contexts, there was no effect of changing the context when the stimulus/response combination trained in A was tested in Context B (Group Same R/Diff S). However, when the animals received training with the same S but different responses in the two contexts (Group Same S/Diff R), there was a strong decrement when Context A's stimulus/response combination was tested in Context B. The results confirm that although S's ability to set the occasion for a response it has been trained with can transfer essentially perfectly across contexts (Group Same R/Diff S), the ability of a response to be augmented by an S it has been trained with does not (Group Same S/Diff R). It is the response, rather than the effectiveness of the stimulus, that depends on the context.

Experiment 4

The results of Experiments 1 – 3 add to those of Bouton et al. (2011) and Todd (2013) in suggesting that operant responses are influenced by changing the context. As noted in the Introduction, the effect of context change on operant behavior appears to contrast with that of Pavlovian conditioned responses, which often transfer effectively across contexts. However, a further difference between the operants and respondents studied so far is interesting theoretically. Specifically, in our studies of operant learning, the response has been reinforced on an intermittent (variable interval) schedule of reinforcement, which means that many responses go unreinforced. In contrast, in the typical Pavlovian experiment, the CS has always been paired with the reinforcer. The difference may be important, because results in both Pavlovian conditioning and predictive learning suggest that attention may increase to contextual cues during training when trial outcomes are surprising (e.g., Rosas & Callejas-Aguilera, 2006; see Rosas et al., 2013, for a recent review). Thus, the operant responses studied so far may have been sensitive to context change because intermittent reinforcement kept each delivery or omission of the reinforcer surprising.

Experiment 4 was therefore designed to test the effects of a context switch on an operant response that was always reinforced. Over a series of sessions, rats were continuously reinforced on R1 in Context A and R2 in Context B. Although the method equated Contexts A and B on their associations with the reinforcer and reinforcement of an operant response, according to Rosas et al. (2006), the continuous reinforcement schedule might encourage

less attention to the context than the VI schedule used in the previous experiments. The experiment used a free-operant method, because discriminated operant training would require nonreinforcing the response during the ITI, which would potentially increase attention to the context.

Experiment 4 also examined another variable that might make operant responding vulnerable to context change. Because each lever and chain is attached to a specific operant chamber, the rats in the previous studies were tested in Context B with a new lever or chain that was extremely similar to, but not the same as, the one on which responding had been reinforced. This change could introduce a source of generalization decrement if there were differences in subtle characteristics of the manipulandum, such as the “feel” of the microswitch that detected each press or pull. Alternatively, the precise way of activating a new and slightly different microswitch might be different, causing a decrease in response rate. To test the effect of such factors, Experiment 4 also examined responding on the same manipulandum in the different contexts during testing. For one group, Group Diff Manipulandum, the rat was tested (as usual) on the lever or chain that was affixed to each operant chamber. For Group Same Manipulandum, the specific lever or chain used in Context A was physically moved to Context B and responding on it was tested there. Since the manipulandum was the same, any response decrement caused by changing the context in Group Same Manipulandum could not be attributed to differences in its “feel” across contexts.

Method

Subjects and Apparatus—The subjects were 32 female Wistar rats purchased from the same supplier and maintained in an identical manner as the preceding experiments. The apparatus was the same as that used in Experiments 1–3.

Procedure

Magazine and operant training: On the first day of the experiment, all rats received two sessions of magazine training, one in Context A and the other in Context B (order counterbalanced). The procedure for this phase was identical to that in the previous experiments. The levers and chains were not present during magazine training.

Over the next 6 days, all rats received two daily sessions of instrumental training, one in Context A and one in Context B. Only one of the lever or chain was consistently available in a particular context (counterbalanced). For half the animals, lever pressing was trained in A and chain pulling was trained in B; the remaining half received the reverse. During each session, every response was reinforced. On the first day of acquisition, the two sessions were 30 min in duration. On the second day, they were 15 min. For Days 3 – 6, they were 10 min in duration.

Testing: On the last day of the experiment, each rat was tested in Context A and Context B in a counterbalanced order. The test sessions were 10 min long, and no pellets were delivered at any time. For all rats, the response that was originally trained in Context A was tested in both Context A and B. For Group Same Manipulandum, the actual lever or chain

on which responding had been trained in Context A was moved and responding on it was tested in Context B. For Group Diff Manipulandum, the manipulandum was not moved from A to B. Thus, as in Experiments 1–3, the lever or chain tested in Context B was not the same piece of equipment that had been used for training in Context A.

Results

Acquisition: The results of the acquisition phase are presented in Figure 6. Responding gradually increased over the course of acquisition. A Group x Context x Session ANOVA revealed a main effect of Session, $F(5, 150) = 650.88$, $MSE = 0.56$, $p < .001$, $\eta_p^2 = .95$, CI [.94, .96], and an interaction between Session and Context, $F(5, 150) = 32.40$, $MSE = 0.47$, $p < .001$, $\eta_p^2 = .52$, CI [.40, .59]. There were no other main effects or interactions, largest $F(5, 150) = 1.31$. A Group x Context ANOVA of the final session revealed that responding overall was higher in Context B than in Context A, $F(1, 30) = 5.62$, $MSE = 0.30$, $p < .05$, $\eta_p^2 = .16$, CI [.00, .38]. Importantly, neither the main effect of Group nor its interaction with Context was significant, $F_s < 1$. Although responding was significantly higher in B than A, the magnitude of this difference was not impressive. The mean responses per minute was 10.7 in Context B and 10.4 in Context A.

Testing: The results of the test day are summarized in Figure 7. Extinction was rapid, presumably due to the fact that the response had been continuously rather than partially reinforced, and analysis was therefore confined to the first 2 minutes of the test. A Group x Context ANOVA revealed that responding was higher in Context A than Context B, $F(1, 30) = 8.57$, $MSE = 125.66$, $p < 0.01$, $\eta_p^2 = .22$, CI [.02, .44]. No other main effects or interactions approached significance, largest $F(1, 30) = 1.07$. The nonsignificant difference between groups in Context A suggested by the figure was presumably due to sampling error. Regardless of whether the manipulandum was the same or different from the one that had been trained, responding decreased with the switch to Context B.

We also examined the latency to make the first response in each context. The time it takes the rat to make the first response might be even less affected by possible differences in the manipulandum's feel. These data are presented in Figure 8, which again strongly suggests the presence of a context-switch effect. A Group x Context ANOVA revealed a main effect of Context, $F(1, 30) = 12.77$, $MSE = 346.88$, $p < 0.01$, $\eta_p^2 = .30$, CI [.06, .50]. Neither the group main effect nor the interaction approached significance, $F_s(1, 30) < 1.05$.

Interestingly, when we then compared responses rates of the groups over the two minutes that followed performance of each rat's first response, the same pattern held: There was a significant Context main effect, $F(1, 30) = 4.36$, $MSE = 561.89$, $p < .05$, but no other main effect or interaction, $F_s < 1$. For Group Same Manipulandum, the mean rates were 25.1 and 17.4 in Contexts A and B; for Group Diff Manipulandum, the corresponding means were 20.5 and 15.8.

Discussion

The results of this experiment extend the previous results in two important ways. First, a context switch decreased the rate of responding after responding had been continuously reinforced. Thus, the context switch effect does not depend on using an intermittent

reinforcement schedule. Second, the context switch effect was also evident when the same lever or chain was used in Contexts A and B (Group Same Manipulandum). This result rules out possible effects of subtle differences in the “feel” of a specific response manipulandum. Indeed, whether the manipulandum was the same or different had no discernible impact on the results; the context effect did not interact with whether the group was tested with the same or different lever or chain. The role of the feel of the lever or chain may also be questioned by the fact that the rats had longer latencies to make the first response in Context B, before any difference in the microswitch could be detected. It is possible that the first R1 was delayed in Context B because the rats were searching for the absent lever or chain on which R2 had previously been reinforced there. However, in an unpublished experiment, we found the same difference in latency to the first response when a continuously reinforced response was first tested in a Context B in which no alternative behavior had been trained (the rats had received equal exposure to B with no lever or chain present). In any case, the results of Experiment 4 attest to the generality of the disruption of operant behavior by context change and suggest that it is not caused by difficulty making a subtly different response in Context B.

General Discussion

The results of the present experiments suggest that the context can have strong control over operant (instrumental) responding. Experiment 4 extended previous work suggesting that context change can disrupt the performance of free-operant responses (Bouton et al., 2011; Todd, 2013). The fact that the context switch effect still occurred despite training with continuous reinforcement indicates that it does not depend on attention being directed to the context because of the “surprising” delivery or omission of a reinforcer in intermittent reinforcement schedules (cf. Rosas et al., 2006). Moreover, the fact that the context switch still had an impact when the same lever or chain was used in Context B further indicates that the response decrement was not due to a difference in the “feel” of the different response manipulanda. In principle, such a factor could play a role in laboratories where the manipulanda differ more between operant chambers than they appear to in the present ones. However, the decremental effect of a context switch clearly does not depend on this factor. Experiment 4 thus strengthens and extends previous results suggesting that the context can play a fundamental role in free operant responding (Bouton et al., 2011; Todd, 2013).

The main purpose of the present experiments, however, was to test the effects of a context switch after lever pressing or chain pulling was trained as a *discriminated* operant. In Experiments 1 – 3, when pressing or pulling was reinforced only in the presence of an S, responding dropped when the S/response combination was tested in the alternate context (see also Nakajima et al., 2000, Experiment 2). That result indicated that the context’s role is not unique to free-operant procedures, which seemed possible because there is no salient proximal stimulus that could overshadow or block contextual conditioning. The decrement in discriminated operant responding was observed over a wide range of conditions. Switching from Context A to Context B caused a response decrement when Context B had only been associated with a single magazine training session and was less familiar than Context A (Experiment 1), as we have observed in free-operant procedures (e.g., Bouton et al., 2011). But the switch also decreased responding when Context B was equally familiar

(Group Exposure in Experiment 2), equally associated with the reinforcer (e.g., Group No R in Experiment 2), and equally home to the training of a different discriminated operant response (Groups Diff R and Diff R/Same S in Experiments 2 and 3, respectively). Thus, the context switch disrupted discriminated operant responding when the procedure gave the contexts comparable associations with a number of Pavlovian and operant contingencies. Furthermore, the fact that the context switch reduced responding regardless of whether Context B had been associated with a different response or not (e.g., Groups Diff R vs. Groups Exp and No R in Experiment 2) suggests that the effect does not depend on Context B cueing an incompatible response.

Experiments 2 and 3 also found that the response-modulating power of the S can transfer well across contexts (Groups Same R and Same R/Diff S in Experiments 2 and 3, respectively). The transfer of the S's ability to control the response across contexts is analogous to the transfer of a CS's response-eliciting power across contexts in Pavlovian conditioning experiments (e.g., Bouton & King, 1983; Bouton & Peck, 1989; Bouton & Swartzentruber, 1989; Nelson, 2002; Rosas & Bouton, 1998). However, the transfer of S's power was limited to the situation in which S1 could set the occasion for its usual response (R1) in a context in which R1 had been directly trained. When both R1 and S1 were tested for the first time in the new context, responding to the S1R1 combination decreased (Group Diff R in Experiment 2). And when R1 was tested with S1 in Context B for the first time after S1 had been trained there (to set the occasion for a different response, R2), responding also decreased (Group Same S/Diff R in Experiment 3). These observations suggest that the response, and not the effectiveness of S, is what is disrupted by changing the context after discriminated operant learning.

Over the last three decades, associative learning research has made considerable progress in understanding what is learned during discriminated operant learning (e.g., Balleine, 2005; Colwill, 1994; Colwill & Rescorla, 1986; de Wit & Dickinson, 2009; Dickinson, 1994). It is now accepted that discriminated operant training can allow the organism to learn associations between (1.) the response and the reinforcer, (2.) the stimulus and the reinforcer, and possibly (3.) the stimulus and the response. There is also evidence that the animal can learn (4.) a hierarchical "occasion setting" relation between the stimulus and the response-reinforcer association (e.g., Bradfield & Balleine, 2013; Colwill & Rescorla, 1990; Rescorla, 1992). Although research thus suggests that animals acquire a rich and varied set of associations and representations during discriminated operant learning, the present experiments may be the first to address the role of the context over and above the role(s) of the discriminative stimulus itself.

How can the present findings be integrated with the associative structures that are known to underlie discriminated operant learning? We might begin by considering the contextual control of free-operant behavior (Experiment 4; Bouton et al., 2011; Todd, 2013). Although the context might in principle play a role analogous to any of the roles of S enumerated above, in a series of experiments on free-operant learning, Todd (2013) found that when R1 is trained in Context A, R1 responding is weakened when it is tested in a second context (Context B) that has been associated with equivalent training of a second response, R2. Under those conditions, the decrement in responding caused by the context switch are not

likely due to differences in the strength of the contexts' associations with the reinforcer. Instead, one approach might be to accept the historically-popular idea that the animal learns a simple S-R, or Context-R, association during instrumental learning (e.g., Hull, 1943). An S-R approach is conceptually compatible with other results reported by Todd (2013), who found evidence consistent with the idea that the context controls extinction performance through a direct inhibitory S-R association. In any event, a straightforward account of the free-operant findings (Bouton et al., 2011; Todd, 2013) is that the animal learns a direct association between the context and the operant response (Figure 9, Panel A). When the context is changed, there is less support for the instrumental behavior.

The current results suggest that this mechanism might not change fundamentally for a discriminated operant response. Experiments 1–3 found that the context also has surprisingly strong control over discriminated operant responding. Thus, a Context-R association might also play a role in discriminated operant learning—and be added to the list of associations numbered above that are thought to be learned in a discriminated operant experiment. What, then, is the role of S? One possibility is that S might also simply elicit the response through a direct S-R association. This possibility would be consistent with the idea that S functions as a proximal or “local” context. There is some preliminary evidence for such an effect of S (cf. Colwill, 1993; Colwill & Rescorla, 1990). However, given the strong contextual control of the discriminated response observed in Experiments 1–3, we might instead suppose that S works at least partly by modulating the context-response association (Panel B in Figure 9). This mechanism can explain why S1R1 responding was weaker when there was a switch to a context in which R1 had not been trained; there would be less responding in Context B because the animal had not learned a B-R1 association that S1 could modulate. However, the mechanism cannot account for the results of Groups Same R (Experiment 2) and Same R/Diff S (Experiment 3). Those groups had the opportunity to associate R1 with Context B during training, and responding to S1R1 was accordingly similar in Contexts A and B. However, S1 had never been trained to modulate the B-R1 association; it had only augmented A-R1. Thus, the Panel B mechanism does not explain why S1 could be so effective at modulating responding in Context B, despite never having modulated an association with B before.

To account for that result, it is worth considering the associative structure shown in Figure 9's Panel C. This structure puts the role of S later in the information-processing sequence. Here S further modulates the response once it has already been evoked or enabled by the context with which the response has been associated. In this way, there is little transfer of S1R1 to a new context unless R1 has been associated with the transfer context. As illustrated by the figure, modulation of the response by S might work through any of the mechanisms previously identified (i.e., by activating the response, the association between the response and the reinforcer, or the representation of the reinforcer). The present results do not discriminate between these possibilities. But the point is that, as shown by the present results, the context may play a role in controlling the operant response that is at least as fundamental as that of S itself. If the response is not first enabled or evoked by the context, S cannot be effective.

Several other comments are in order. First, it is worth reiterating that the present results exclude the possibility that the context merely controls discriminated operant behavior through its direct association with reinforcer. This is because the context-switch effect still occurred in several groups for which Context A and B's associations with the reinforcer were controlled and equated in Experiments 2 and 3. Second, the results are also incompatible with the possibility that the context and S combine to form a unique configural stimulus that controls the response. This is because the ability of S to control responding appeared to transfer perfectly to a new context (i.e., there was perfect generalization between different context-S configurations, as suggested by Groups Same R and Same R/Diff S in Experiments 2 and 3, respectively). However, the present results cannot exclude the possibility shown in Panel D of Figure 9, where the context operates as an occasion setter modulating the response-reinforcer association. Although this approach seems less parsimonious than Panel C, it has the same advantage as Panel C: Once again, the S modulates the response once it has already been enabled by the context, which would not permit much transfer of S1R1 to a new context unless R1 has been reinforced in the new context. One challenge for either the Panel C or D mechanism is that response rates during the pre-S period (which occurred in the presence of the context alone) did not decrease when a response only trained in Context A was tested in Context B. The null results need to be interpreted with caution, however. Pre-S responding decreased over the course of discriminative acquisition training, suggesting that responding in the pre-S period was under the influence of extinction (inhibition) at the time of testing. This could allow an increase in responding on the switch from Context A to B (AAB renewal; Bouton et al., 2011), which might obscure any opposing tendency for responding to decline because of a Context-R or Context- (R-O) learning. Thus, the ideas embodied in panels C and D may remain viable as approaches for explaining the current results.

To summarize, the results of the present experiments suggest that the context can play a role in enabling instrumental behavior after either discriminated- or free-operant learning. We suggest that the context enables the response either through a direct context-response association or by setting the occasion for the response-reinforcer association. When testing occurs in the context in which the response has been trained, the response occurs and is available for modulation by the discriminative stimulus. But when testing occurs in a different context, it does not occur and is not available for modulation.

Although we have described the context's role in associationistic terms, it is possible to think of it in another way. The results are also consistent with the idea that the animal must learn that a particular context "affords" a particular instrumental response. The idea behind J. J. Gibson's *affordance* concept (e.g., Gibson, 1977) is that organisms learn that different items in the environment can support different behaviors or actions. The human must learn that a horizontal surface, such as a park bench or a particular floor, can be sat upon or walked upon, respectively. We must learn that the bench or floor can support our body weight and prevent us from falling through. We do not necessarily know that a new bench in a new context will provide the same support. The rat might similarly need to learn that a new lever or chain in a different context will support pressing or pulling, i.e., that those actions are associated with a positive outcome. Compared to the case with "reflexive"

Pavlovian responding, contexts may play a more pivotal role supporting instrumental behavior because the necessary affordances are inherently context-dependent.

Acknowledgments

This research was supported by Grants RO1 MH64847 and RO1 DA33123 from the National Institutes of Health to Mark Bouton. The participation of Samuel León (who visited the University of Vermont from the University of Jaén) was enabled by projects SEJ2007-267053/PSIC and BES-2008-003634 from the Spanish Ministry of Science and Innovation. We thank Eric Thraillkill for his consultation concerning the statistical analysis.

References

- Balleine BW. Neural basis of food-seeking: Affect, arousal and reward in corticostriatal limbic circuits. *Physiology & Behavior*. 2005; 86:717–730. [PubMed: 16257019]
- Bonardi C, Honey RC, Hall G. Context specificity of conditioning in flavor-aversion learning: Extinction and blocking tests. *Animal Learning & Behavior*. 1990; 18:229–237.
- Bouton ME. Context, ambiguity, and classical conditioning. *Current Directions in Psychological Science*. 1994; 3:49–53.
- Bouton, ME. Signals for whether versus when an event will occur. In: Bouton, ME.; Fanselow, MS., editors. *Learning, motivation, and cognition: The functional behaviorism of Robert C. Bolles*. Washington, D. C: American Psychological Association; 1997. p. 385-409.
- Bouton ME. Context, time, and memory retrieval in the interference paradigms of Pavlovian learning. *Psychological Bulletin*. 1993; 114:80–99. [PubMed: 8346330]
- 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, Swartzentruber DA. Slow reacquisition following extinction: Context, encoding, and retrieval mechanisms. *Journal of Experimental Psychology: Animal Behavior Processes*. 1989; 15:43–53.
- 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]
- Bradfield LA, Balleine BW. Hierarchical and binary associations compete for behavioral control during instrumental biconditional discrimination. *Journal of Experimental Psychology: Animal Behavior Processes*. 2013; 39:2–13. [PubMed: 23316974]
- Colwill RM. Signaling the omission of a response-contingent outcome reduces discriminative control. *Animal Learning & Behavior*. 1993; 21:337–345.
- Colwill, RM. Associative representations of instrumental contingencies. In: Medin, DL., editor. *The psychology of learning and motivation*. Vol. 31. New York, NY: Academic Press; 1994. p. 1-72.
- Colwill, RM.; Rescorla, RA. Associative structures in instrumental learning. In: Bower, GH., editor. *The psychology of learning and motivation*. Vol. 20. New York, NY: Academic Press; 1986. p. 55-104.
- 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]

- de Wit S, Dickinson A. Associative theories of goal-directed behavior: a case for animal-human translational models. *Psychological Research*. 2009; 73:463–476. [PubMed: 19350272]
- Dickinson, A. Instrumental conditioning. In: Mackintosh, NJ., editor. *Animal learning and cognition*. London: Academic; 1994. p. 4-79.
- Gibson, JJ. The theory of affordances. In: Shaw, R.; Bransford, J., editors. *Perceiving, Acting, and Knowing*. Hillsdale, NJ: Erlbaum; 1977. p. 67-82.
- Godden DR, Baddeley AD. Context-dependent memory in two natural environments: On land and underwater. *British Journal of Psychology*. 1975; 66:325–331.
- Grahame NJ, Hallam SC, Geier L, Miller RR. Context as an occasion setter following CS acquisition and extinction or CS acquisition alone. *Learning and Motivation*. 1990; 21:237–265.
- Hall G, Honey R. Contextual effects in conditioning, latent inhibition, and habituation: Associative and retrieval functions of contextual cues. *Journal of Experimental 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]
- 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]
- Hull, CL. *Principles of behavior: An introduction to behavior theory*. New York: Appleton-Century-Crofts; 1943.
- Lovibond PF, Preston GC, Macintosh NJ. Context specificity of conditioning, extinction, and latent inhibition. *Journal of Experimental Psychology: Animal Behavior Processes*. 1984; 10:360–375.
- 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.
- Nakajima S, Urushihara K, Masaki T. Renewal of operant performance formerly eliminated by omission or noncontingency training upon return to the acquisition context. *Learning and Motivation*. 2002; 33:510–525.
- Nelson JB. Context specificity of excitation and inhibition in ambiguous stimuli. *Learning and Motivation*. 2002; 33:284–310.
- Nelson JB, Callejas-Aguilera JE. The role of interference produced by conflicting associations in contextual control. *Journal of Experimental Psychology: Animal Behavior Processes*. 2007; 33:314–326. [PubMed: 17620029]
- Nelson JB, Lamoureux JA, León SP. Extinction arouses attention to the context in a behavioral suppression method with humans. *Journal of Experimental Psychology: Animal Behavior Processes*. 2013; 39:99–105. [PubMed: 23316978]
- 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]
- Pearce, JM.; Mackintosh, NJ. Two theories of attention: A review and a possible integration. In: Mitchell, C.; LePelley, ME., editors. *Attention and Learning*. Oxford: Oxford University Press; 2010. p. 11-39.
- Rescorla RA. Associative relations in instrumental learning; The eighteenth Bartlett Memorial Lecture. *Quarterly Journal of Experimental Psychology*. 1991; 43:1–23.
- Rescorla RA, Solomon RL. Two process learning theory: Relationships between Pavlovian conditioning and instrumental learning. *Psychological Review*. 1967; 74:151–182. [PubMed: 5342881]
- 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, Callejas-Aguilera JE. Acquisition of a conditioned taste aversion becomes context dependent when it is learned after extinction. *The Quarterly Journal of Experimental Psychology*. 2007; 60:9–15. [PubMed: 17162504]
- Rosas JM, Callejas-Aguilera JE, Ramos-Alvarez MM, Abad MJF. Revision of retrieval theory of forgetting: What does make information context-specific? *International Journal of Psychology and Psychological Therapy*. 2006a; 6:147–166.
- Rosas JM, Todd TP, Bouton ME. Context change and associative learning. *Wiley Interdisciplinary Reviews: Cognitive Science*. 2013; 4:237–244. [PubMed: 23772263]
- Smith, SM. Environmental context-dependent memory. In: Davies, GM.; Thomson, DM., editors. *Memory in context: Context in memory*. New York: John Wiley & Sons; 1988. p. 13-34.
- Smith SM, Vela E. Environmental context-dependent memory: A review and meta-analysis. *Psychonomic Bulletin & Review*. 2001; 8:203–220. [PubMed: 11495110]
- Spear, NE. *The processing of memories: Forgetting and retention*. Hillsdale, NJ: Erlbaum; 1978.
- Steiger JH. Beyond the *F* test: Effect size confidence intervals and tests of close fit in the analysis of variance and contrast analysis. *Psychological Methods*. 2004; 9:164–182. [PubMed: 15137887]
- Todd TP. Mechanisms of renewal after the extinction of instrumental behavior. *Journal of Experimental Psychology: Animal Behavior Processes*. 2013 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]
- Tulving E, Thomson DM. Encoding specificity and retrieval processes in episodic memory. *Psychological Review*. 1973; 80:352–373.




Figure 1. Results of Experiment 1. Mean elevation scores of the groups during acquisition in Context A (left) and testing in Context A or B (right; $\pm SEM$).




Figure 2. Mean elevation scores of the groups during acquisition in Context A (top panel) and Context B (bottom panel) in Experiment 2.




Figure 3. Mean elevation scores of the groups' R1 responding in S1 in Contexts A and B during testing in Experiment 2.

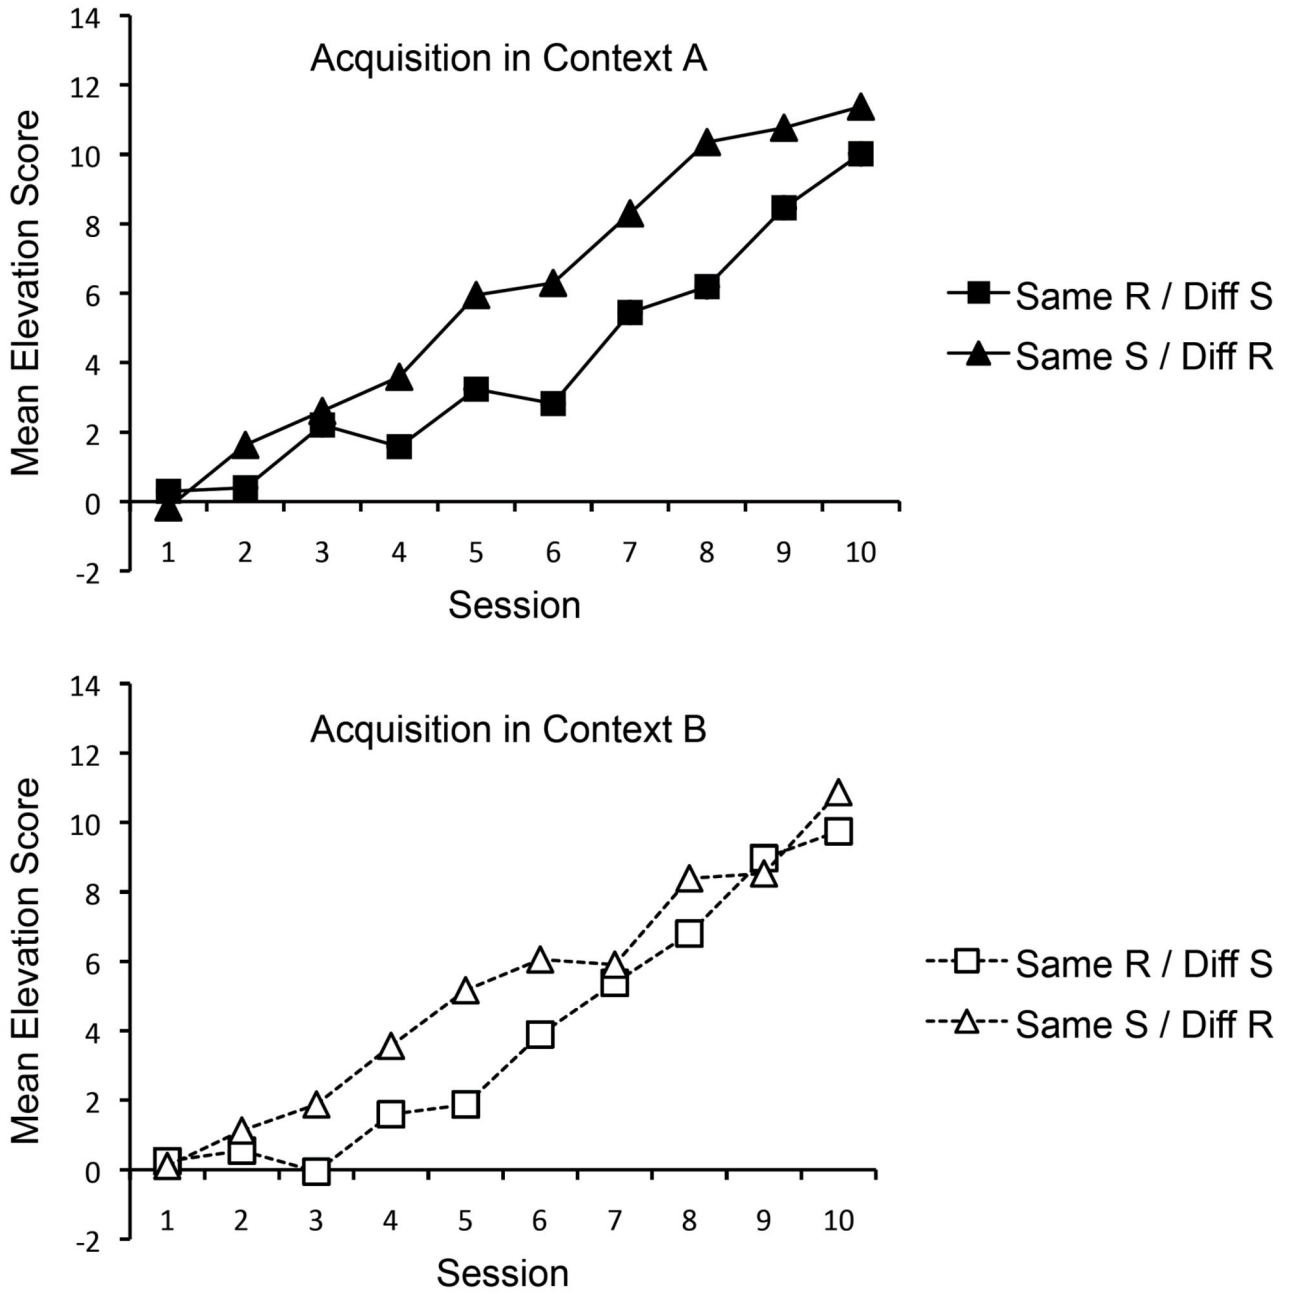


Figure 4. Mean elevation scores of the groups during acquisition in Context A (top panel) and Context B (bottom panel) in Experiment 3.




Figure 5. Mean elevation scores of the groups' R1 responding in S1 in Contexts A and B during testing in Experiment 3.




Figure 6. Mean responses per minute for Group Same Manipulandum and Group Diff Manipulandum in Contexts A and B during acquisition in Experiment 4. The sessions were 30-min (1), 15-min (2), and 10-min (3–6) in duration.




Figure 7. Mean responses during the first two minutes of testing for Group Same and Diff in Contexts A and B in Experiment 4.




Figure 8. Mean latency to first response for Groups Same and Diff in Contexts A and B during testing in Experiment 4.

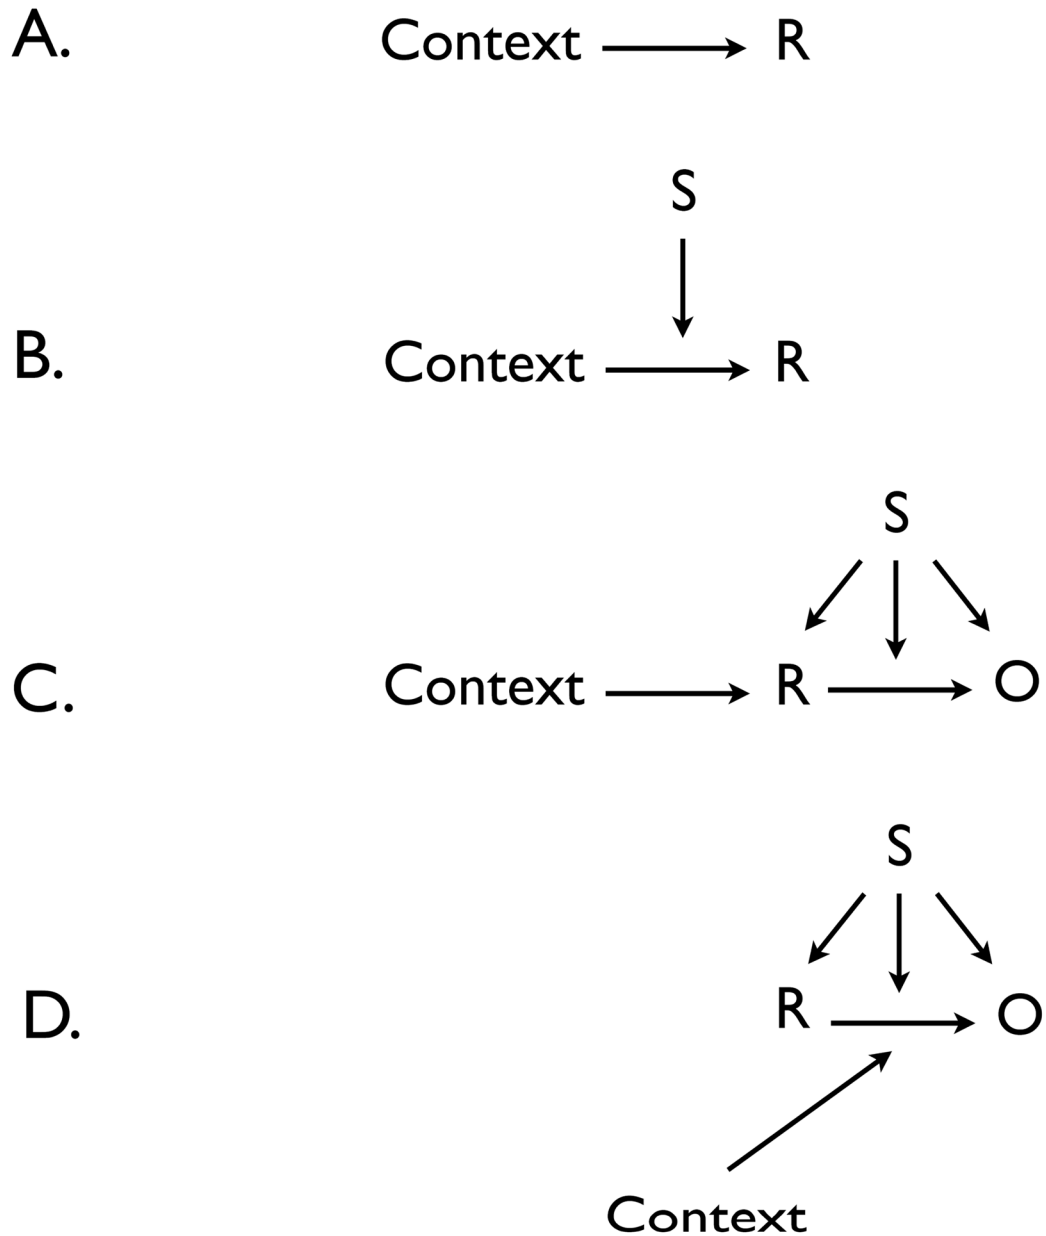


Figure 9.

Four possible roles for the context in operant learning. Panel A. Direct association between the context and the response. Panel B. Direct association between the context and the response that is modulated by S. Panel C. 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 D. 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.

Table 1

Designs of Experiments 1–3

Group	Acquisition	Test
Experiment 1		
A	A: S1R1+	A: S1R1–
B	A: S1R1+	B: S1R1–
Experiment 2		
Exp	A: S1R1+	A: S1R1–
	B: ---	B: S1R1–
No R	A: S1R1+	A: S1R1–
	B: S2+	B: S1R1–
Diff R	A: S1R1+	A: S1R1–
	B: S2R2+	B: S1R1–
Same R	A: S1R1+	A: S1R1–
	B: S2R1+	B: S1R1–
Experiment 3		
Same R/Diff S	A: S1R1+	A: S1R1–
	B: S2R1+	B: S1R1–
Same S/Diff R	A: S1R1+	A: S1R1–
	B: S1R2+	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. The schedule of delivery of 16 pellets to Group No R during S2 was yoked to a member of the same group that was being run in Context A, so that pellets were now delivered during S2 whenever the partnered rat earned one during S1.